




IBM Nways RouteSwitch Implementation Guide

*John Parker, Corky Camin, Jocelyn Cecire, Matthias Enders
Milind Gupte, Richard Shaw, Kevin Treweek, ,*



International Technical Support Organization

<http://www.redbooks.ibm.com>



International Technical Support Organization

SG24-4881-01

IBM Nways RouteSwitch Implementation Guide

September 1998

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix D, "Special Notices" on page 365.

Second Edition (September 1998)

This edition applies to Version 3, Release 2 of Nways RouteSwitch Software Program, Program Number 5697-B70 (8274) and 5697-B69 (8273).

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. HZ8 Building 678
P.O. Box 12195
Research Triangle Park, NC 27709-2195

When you send information to IBM, you grant IBM a non-exclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© **Copyright International Business Machines Corporation 1997, 1998. All rights reserved.**

Note to U.S. Government Users — Documentation related to restricted rights — Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Contents

Preface	vii
The Team That Wrote This Redbook	vii
Comments Welcome	viii
 Chapter 1. IBM 8274 and 8277 RouteSwitch - Functional Overview	 1
1.1 1996 Announcement	1
1.2 1998 Announcement	1
1.3 Functional Overview	1
1.3.1 The 8274 LAN RouteSwitch	2
1.3.2 8274 Description	2
1.4 Hardware Components	4
1.4.1 RouteSwitch Architecture	4
1.4.2 Backplane	5
1.4.3 MPMs	6
1.4.4 Hardware Routing Engine (HRE)	6
1.4.5 Frame-to-Cell Switching Module (FCSM)	7
1.4.6 RouteCell Modules	7
1.4.7 ATM Access Switching Modules	7
1.4.8 Switching Modules	7
1.5 Functional Overview of IBM 8277 RouteSwitch	8
1.5.1 8277 Description	8
1.5.2 The 8277 Ethernet RouteSwitch Model 524	9
 Chapter 2. Basic RouteSwitch Setup	 11
2.1 Nways RouteSwitch Software Program	11
2.2 Basic Setup and Operation of the RouteSwitch	12
2.2.1 Getting Started	12
2.3 Switch Configuration Using Console	14
2.3.1 Basic Console Setup	14
2.3.2 SNMP Setup	18
2.3.3 Listing the Files Loaded in the MPM	19
2.3.4 Copying and Deleting Configuration Files	20
2.3.5 File Names of Switching Modules	21
2.4 Management Processor Module (MPM)	22
2.4.1 Management Processor Module Memory	23
2.4.2 Management Processor Module Redundancy	24
2.4.3 MPM Changeover Procedure	25
2.4.4 Module Replacement	26
2.4.5 Configuring the Modem Port on the MPM	28
2.5 Installing Software on the MPM	30
2.5.1 Using an FTP Server	30
2.5.2 Using an FTP Client	31
2.5.3 Using ZMODEM	32
 Chapter 3. VLAN and Mobile Groups Concepts	 37
3.1 Traditional LAN Design	37
3.2 VLAN Design	38
3.3 RouteSwitch Terminology	40
3.3.1 RouteTracker Policies	40
3.3.2 Port Rules	41
3.3.3 MAC Address Rules	42

3.3.4 Protocol Rules	44
3.3.5 Network Address Rules	45
3.3.6 User-Defined Rules	46
3.3.7 Port Binding Rule	47
3.3.8 DHCP Port Rules	47
3.3.9 DHCP MAC Address Rules	47
3.3.10 Non-Mobile Groups	47
3.3.11 Mobile Groups	50
3.4 Ports, MAC Devices and VLAN Timers - Non-Mobile Groups	52
3.5 RouteSwitch Packet Processing for Non-Mobile Groups	53
3.5.1 The Switching Process for Non-Mobile Groups	54
3.5.2 Frame Flooding for Non-Mobile Groups	54
3.5.3 VLAN Leakage for Non-Mobile Groups	58
3.5.4 Optimized Device Switching Ports	59
3.6 RouteSwitch Packet Processing for Mobile Groups	59
3.6.1 Broadcast Frames Processing for Mobile Groups	60
3.6.2 Unicast Frames Processing for Mobile Groups	63
Chapter 4. Basic VLAN Configuration	67
4.1 Non-Mobile Group and VLAN Configuration	67
4.1.1 Non-Mobile Group Configuration Using the Console	67
4.1.2 VLAN Configuration Using the Console	71
4.1.3 The Precedence of the Policies	80
4.2 Basic Mobile Group Configuration	80
4.2.1 Mobile Group Configuration Using the Console	81
4.2.2 Configuration of the RouteSwitch	81
Chapter 5. Bridging, Routing, Switching and Trunking	89
5.1 Bridging Methods	89
5.1.1 Transparent Bridging	90
5.1.2 Source Route Bridging	94
5.1.3 Translational Bridges	95
5.1.4 Source Route Transparent Bridges	96
5.2 The Spanning Tree	96
5.2.1 The Spanning Tree in Source Route Bridges	96
5.2.2 The Spanning Tree in Transparent Bridging	97
5.2.3 The 802.1d Spanning Tree Protocol	98
5.2.4 Transparent Bridges and Network Topology Changes	102
5.2.5 Setting the Parameters That Control the Spanning Tree	103
5.2.6 Summary of the IEEE 802.1d Spanning Tree Algorithm	112
5.3 Configuring Source Route Bridging	112
5.3.1 Virtual Token-Rings	120
5.3.2 Invalid Configuration	121
5.3.3 Source Route Bridging between RouteSwitches	122
5.3.4 Summary of Bridge Types	123
5.4 VLAN and Group Routing	125
5.4.1 RouteSwitch Handling of Router Ports	130
5.5 Advanced Routing	130
5.5.1 Advanced Routing Example 1	131
5.5.2 Advanced Routing Example 2	140
5.6 Next Hop Resolution Protocol (NHRP)	142
5.6.1 ELAN Support	144
5.6.2 NHRP Example 1	146
5.7 Mixed Media	150
5.7.1 Any-to-Any Switching	150

5.7.2 Theory of Operation	151
5.7.3 Default Translation Tables	163
5.8 Example of Bridging, Routing and Switching	165
5.9 Trunking	169
5.9.1 ATM Trunking	170
5.9.2 Group Multiplexing on FDDI	185
Chapter 6. RouteSwitch ATM LAN Emulation	189
6.1 LAN Emulation Benefits	189
6.2 Emulated LANs (ELANs)	189
6.3 ATM Addresses	191
6.3.1 ATM Addresses of LAN Emulation Components	191
6.4 Overview of ILMI Functions	192
6.4.1 LAN Emulation Components	192
6.4.2 LAN Emulation VC Connections	195
6.4.3 LE Service Operation	196
6.4.4 LAN Emulation Summary	197
6.5 LAN Emulation on the RouteSwitch	197
6.5.1 SAR Option	198
Chapter 7. ATM Cell Switching	201
7.1 Cell Switching Modules	201
7.1.1 Frame to Cell Switching Module	201
7.2 Cell Switching Matrix	202
7.3 Traffic Management	203
7.3.1 Buffer Management	203
7.3.2 Traffic Contract Parameters	203
7.3.3 Class of Service	204
7.4 Private Network-to-Network Interface (PNNI)	205
7.4.1 Understanding PNNI Networks	205
7.4.2 PNNI Network Initialization	205
7.4.3 PNNI Data Path	206
7.4.4 PNNI and IISP	206
7.4.5 PNNI Implementation in the 8274	207
Chapter 8. LECS and LES/BUS Functionality	213
8.1 LANE Services Configuration	213
8.2 LSM Configuration Examples	214
8.2.1 Auto-Creating an ATM Environment	214
8.2.2 LSM Example: Routing between ELANs	221
8.2.3 Example of Redundant LES/BUS	230
8.2.4 Example of LES/BUS Redundancy with IP Routing	247
Chapter 9. Wide Area Networking	257
9.1 Interfaces	257
9.2 Compression	258
9.3 Virtual Circuits and DLCIs	259
9.4 Congestion Control	260
9.4.1 FECN and BECN	260
9.4.2 Comitted Information Rate and Burst Sizes	261
9.5 WSM Routing and Interoperability Examples	263
9.5.1 Back-to-Back Bridging/Routing Using a Frame Relay Group	263
9.5.2 Frame Relay Routing with an IBM 2210	268
Chapter 10. Network Configuration Examples	273

10.1 Fast Ethernet Interconnected RouteSwitches	273
10.1.1 Network Topology	273
10.1.2 RouteSwitch Configuration	274
10.2 RouteSwitch Using Original Port Rule	283
10.3 RouteSwitch Interconnection Using MSS and 8260	284
10.3.1 Network Topology	285
10.3.2 RouteSwitch Configuration	286
10.4 IP Routing at the MSS	291
10.4.1 Network Topology	291
10.4.2 RouteSwitch Configuration	292
10.5 IP Routing at the MSS Using Multiple LECs	300
10.5.1 Logical Network Topology	300
10.5.2 RouteSwitch Configuration	301
10.6 Configuration Example of the 8274 as an ATM Switch	305
10.7 Configuration of Two 8274s over MSS and Using DHCP Protocol Rule	308
10.7.1 Logical View of the Network	308
10.7.2 RouteSwitch Configuration	309
 Appendix A. Correcting ARI/FCI Issues	 335
A.1 Read Me File from the TRDRVFIX Diskette	335
 Appendix B. Ether-Types and SAP Listings	 341
B.1 Ether-Type Listing	341
B.2 SAP Listing	343
 Appendix C. Sample GATED Configuration File	 345
 Appendix D. Special Notices	 365
 Appendix E. Related Publications	 367
E.1 International Technical Support Organization Publications	367
E.2 Redbooks on CD-ROMs	367
E.3 Other Publications	367
 How to Get ITSO Redbooks	 369
How IBM Employees Can Get ITSO Redbooks	369
How Customers Can Get ITSO Redbooks	370
IBM Redbook Order Form	371
 List of Abbreviations	 373
 Index	 375
 ITSO Redbook Evaluation	 377

Preface

This redbook describes the features, models and operation of the IBM Nways RouteSwitch family. This book will help you to understand basic switching concepts such as VLANs, ELANs and groups.

Procedures for hardware and software installation, setup and configuration related to the 8273 and 8274 are also documented. The primary focus of this book is on Ethernet, token-ring and ATM environments.

The unique any-to-any switching capabilities of the RouteSwitch are discussed and explored in several chapters. Particular emphasis is placed on those areas of interoperability that have to do with any-to-any configuration issues.

Several practical examples document the different policies/types of VLANs by providing detailed configuration scripts. Examples include TCP/IP routing, bridging, and trunking, including ATM ELANs. With minor modifications, these scripts can be used as models for many customer environments.

Some knowledge of LAN architecture, switching, ATM, and bridging is assumed.

The Team That Wrote This Redbook

This redbook was produced by a team of specialists from around the world working at the Systems Management and Networking ITSO Center, Raleigh.

The authors of this redbook were:

John Parker is an Advisory ITSO Specialist in Campus LAN at the Systems Management and Networking ITSO Center, Raleigh. He writes extensively and teaches IBM classes worldwide on all areas of Campus LAN. Before joining the ITSO, John worked in Availability Services, USA as an I/T Services Specialist.

Corky Camin is an Advisory Networking Strategist at RTP, NC providing direct network support for IBM and external customers. He has been in networking for 14+ years and LANs for 11+ years. His areas of expertise include bridges, switches, and routers for Ethernet and token-ring networks. He also has a background in ATM networking.

Jocelyn Cecire is a Senior Networking Systems Specialist in Canada. He has 25 years of experience in the telecommunication/networking field. He has worked at IBM for 25 years. His areas of expertise include hubs, bridges, switches as well as analyzing traces and performing LAN Doctor services. He is also the author of 3174INFO, 82XXINFO and the Advanced LAN Problem Determination course. He is a token-ring CNX.

Matthias Enders is a Networking Specialist in Germany. He has seven years of experience in the networking field. His areas of expertise include TCP/IP, campus LAN products and LAN protocol analysis as well as multiprotocol network design. He has written extensively on the fifth edition of the redbook *TCP/IP Tutorial and Technical Overview*.

Milind Gupte is a Technical Manager - Networking Services in India. He has 11 years of experience in networking. His areas of expertise include designing and implementing campus/building networks as well as WAN. This is his first book on the topic.

Richard Shaw is a Network Specialist in South Africa. He works extensively in the pre-sales role, specializing in network design and switching. He presents extensively. This is his first book on the topic.

Kevin Treweek is an Advisory IT Specialist in South Africa. He has eight years of experience in the networking field. He works in the pre-sales environment and his areas of expertise are Campus and LAN network design. This is his first book on the topic.

Thanks to the following people for their invaluable contributions to this project:

Volkert Kreuk
International Technical Support Organization, Raleigh Center

Matt Darlington
IBM, Raleigh

Wes Wegner
Xylan California

Robert Roohparvar
Xylan California

Comments Welcome

Your comments are important to us!

We want our redbooks to be as helpful as possible. Please send us your comments about this or other redbooks in one of the following ways:

- Fax the evaluation form found in "ITSO Redbook Evaluation" on page 377 to the fax number shown on the form.
- Use the electronic evaluation form found on the Redbooks Web sites:

For Internet users	http://www.redbooks.ibm.com/
For IBM Intranet users	http://w3.itso.ibm.com/

- Send us a note at the following address:

redbook@us.ibm.com

Chapter 1. IBM 8274 and 8277 RouteSwitch - Functional Overview

The 8274 and 8277 RouteSwitch provide connectivity and functionality never before found in an IBM LAN switch. In this chapter we review the RouteSwitch announcements and general capabilities of these unique LAN switches.

1.1 1996 Announcement




Figure 1. 8274 RouteSwitches Models W93 and W53

The 8274 RouteSwitches and the associated software were originally announced on September 12th 1996. By announcing these switches, IBM is bringing its switched virtual networking concept to the LAN/VLAN as well as the ATM/ELAN level. The 8274 LAN RouteSwitch provides you with powerful, inexpensive, transparent switching at wire speeds, at the same time providing extensive VLAN support.

1.2 1998 Announcement

The 8277 RouteSwitch can provide drop-in solutions for many requirements in today's environment. It provides with the fixed configuration, 10/100Base-T auto-sensing switched ports and can consolidate a cluster of 24 workstations into a large network, providing each workstation with its own dedicated port.

1.3 Functional Overview

Outlined in the next three sections are the general characteristics of the IBM Nways RouteSwitch. The hardware features and general functions are explained.

1.3.1 The 8274 LAN RouteSwitch

The following list highlights some of the features and functions available in the 8274 RouteSwitch:

- Provides low-cost, high-function LAN and ATM switching
- Improves network availability with hot-swappable modules and power supplies
- Connects to network segments, file servers, or individual workstations
- Supports Ethernet, token-ring, FDDI, CDDI, Fast Ethernet, frame relay (WAN), and ATM switching to the desktop and the backbone
- Protects your investments and positions you for future expansion with the 13.2 Gbps ATM cell switching backplane (Models 513, 913, W53, and W93)
- Uses a flexible, failure-resistant architecture

The 8274 LAN RouteSwitch incorporates a scalable, hot-swappable, modular chassis. The LAN RouteSwitch combines an innovative hardware architecture with sophisticated features and configuration software. The Nways 8274 LAN RouteSwitch is available in three models providing you with the unique combination of LAN switching and ATM switching to the desktop and to the backbone. The IBM 8274 Nways LAN RouteSwitch is for customers who need functionality and higher port density.

1.3.2 8274 Description

The IBM 8274 Nways LAN RouteSwitch is a flexible, powerful, highly reliable new switching platform. It combines an innovative hardware architecture with a sophisticated feature set.

IBM 8274 LAN RouteSwitch is uniquely versatile. It supports routing of IP and IPX. It connects to network segments, file servers, or individual workstations. It supports any combination of Ethernet, FDDI, CDDI, Fast Ethernet, token-ring, and ATM at wire speed with automatic any-to-any translation.

The 8274 LAN RouteSwitch is designed for situations where customers need functionality of virtual LANs and a higher port density. The 8274 provides customers with extensive VLAN support. An MPM or Management Processor Module is required in all models of the 8274 RouteSwitch. The originally announced MPM II will support environments where packet-only switching is needed.

The 8274 models 513, 913, W33, W53, and W93 with the 13.2 Gbps ATM cell switching backplane matrix will include all components required to support the Cell Switching Modules (CSM) in addition to all existing framed-based modules. The MPM modules that will support the cell-based ATM switching backplane include the MPM II-16, MPM 1G or the MPM 1GW. The IBM 8274 Models W53 and W93 require the MPM 1GW, which will also support the ATM cell switching function. Additionally, if the CSMs are installed in one of these RouteSwitches, a Frame-to-Cell Switching Module (FCSM) will be required. The FCSM is required for the MPM to do call setup for the new ATM switching function.

The 8274 IBM LAN RouteSwitch is the most reliable product in its class. There are no active components in the IBM LAN RouteSwitch backplane. Every module has its own processing power. You can add redundancy to any critical component.

Some of the reliability and serviceability features include:

- Multiple SPARC RISC processor architecture
- Hot-swappable modules
- Redundant power supplies
- Hot-swappable power supplies
- Dual inputs
- Redundant cooling fans
- Temperature alarm
- Flash memory
- Extensive LED indicators
- Global agency listings for safety and emissions
- Small profile wall mount/standard 19-inch rack mount

1.3.2.1 8274 Model Descriptions

The 8274 LAN RouteSwitch is available in three models: 300 series, 500 series, and 900 series. The 300 series includes the W33 model. The 500 series includes the 500, 513 and W53 models. The 900 series includes the 900, 913 and W93 models. The 300 series requires MPM module in slot one and would have two slots available if using frame (LAN) switching. Frame and cell switching are not supported together in the W33, consequently if the W33 is to be used as an ATM switch, then an MPM, frame to cell switching module (FCSM) and one cell switching module (CSM) would be required. The 500 series has four slots available for LAN modules with a Management Processor Module (MPM) installed, while the 900 series has eight available slots with the MPM inserted. Below is a brief description of the 8274 models:

- 8274 Model 513 (discontinued)
 - Five slots
 - 250 watt AC power supply
 - 13.2 Gbps capable, cell-based backplane
 - MPM II-16 and MPM 1G only
- 8274 Model 913 (discontinued)
 - Nine slots
 - 500 watt power supply
 - 13.2 Gbps capable, cell-based backplane
 - MPM II-16 and MPM 1G only
- 8274 Model W33
 - Three slots 1-1/2 inch wide slots
 - 150 watt AC power supply
 - 13.2 Gbps capable, cell-based backplane
 - Requires MPM 1GW
- 8274 Model W53
 - Five 1-1/2 inch wide slots

- 250 watt AC power supply
- 13.2 Gbps capable, cell-based backplane
- Requires MPM 1GW
- 8274 Model W93
 - Nine 1-1/2 inch wide slots
 - 500 watt AC power supply
 - 13.2 Gbps capable, cell-based backplane
 - Requires MPM 1GW

1.4 Hardware Components

The following sections describe the hardware features that are available for the 8274.

1.4.1 RouteSwitch Architecture

The RouteSwitch's architecture is based on a store-and-forward technology. This allows the RouteSwitch to filter out runts, packets that exceed the maximum allowed length, CRC-flawed packets, misaligned packets, etc. Another advantage of store-and-forward is the rate conversion, allowing connectivity to high-speed servers and backbones.

Figure 2 on page 5 is a representation of the architecture. Terms included in the figure are:

- MBUS - Management Bus
- VBUS - Frame path
- CAM - Content Addressable Memory (MAC addresses, VLAN membership flags)
- MPM - Management Processor Module
- ESM - Ethernet Switching Module
- HSM - High-Speed Switching Module (FDDI, token-ring, ATM, etc.)
- NSM - Network Switch Module (ESM, HSM, etc.)

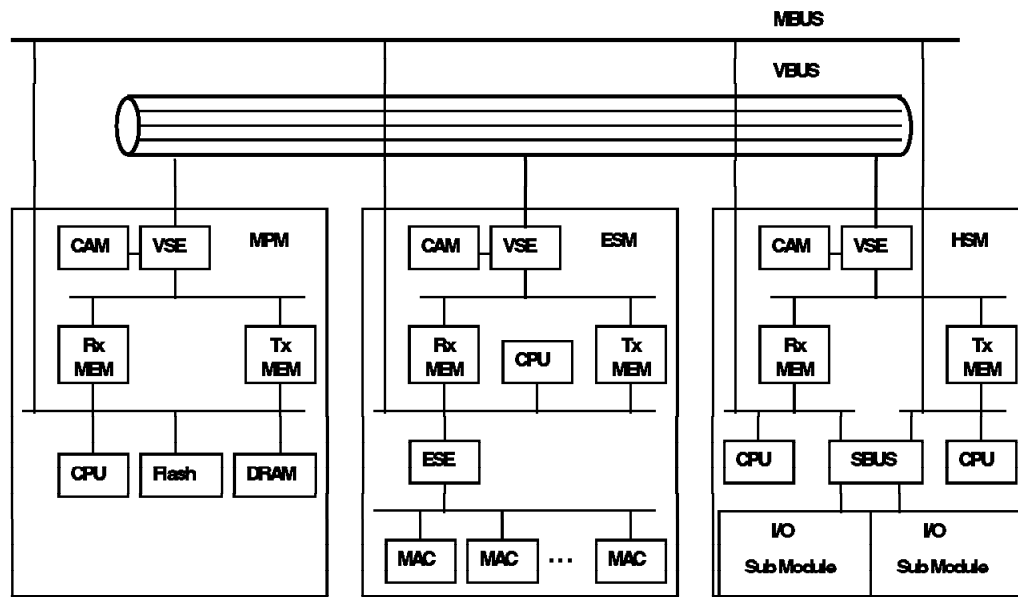


Figure 2. RouteSwitch Architecture

Each switching module, as well as the Management Processor Module have multiple SPARC RISC processors, supported by an advanced memory architecture. Processor power is added with the addition of each switching module. Switching tasks are distributed over all processors. A sophisticated operating system gives the RouteSwitch the ability to rapidly integrate powerful new features.

The RouteSwitch architecture uses a resilient switching fabric called the VLAN Bus (VBUS). The VBUS supports high-performance switching while maintaining separation between up to 96 groups within any one RouteSwitch and up to 65,000 groups in a network. Within each non-mobile group up to 32 VLANs can be configured in any one switch.

Management of the switching modules is provided via a separate management bus, the MBUS. The switching functionality is distributed across both the Management Processor Module and all of the switching modules, resulting in a scalable, high-capacity capability.

1.4.2 Backplane

IBM 8274 frame bus operates at 640 or 960 Mbps, depending on the MPM used and the switching modules revision levels. The MPM 1G supports a frame bus capacity of 960 Mbps while MPM and MPM II supports 640 Mbps. Management of the switching modules is provided via a separate management bus which operates at 120 Mbps.

IBM 8274 Models W33, 513, W53, 913 and W93 provide an ATM cell matrix in addition to the frame and management bus. This ATM cell matrix operates at up to 13.2 Gbps.

1.4.3 MPMs

The following three MPMs were announced on March 11th 1997:

- MPM II-16
- MPM-1G
- MPM-1GW

The first two, the MPM II-16 and the MPM-1G, can be used in the all models except the wide chassis types. The MPM-1GW is supported in the 8274 models W33, W53 and W93 wide chassis only. These MPMs include 16 MB of DRAM, 4 MB flash memory and connect to the management bus as well as the LAN/frame bus. Unlike the prior MPMs, these MPMs are capable of operating the frame bus at 960 Mbps versus the original speed of 640 Mbps.

1.4.4 Hardware Routing Engine (HRE)

The HRE performs the IP and IPX routing functions previously done in software on a separately priced and orderable module for the currently offered and newly announced 8274 models. The MPM-1G or the MPM-1GW is a prerequisite of this HRE feature. This new feature is a hardware add-on daughter card for the MPM. The HRE vastly improves the IP and IPX routing performance of the RouteSwitch. No additional software is required in the RouteSwitch for the HRE. The HRE Plus has the same capacity for routing frames, and also offers HRE redundancy when two MPMs are installed in the same RouteSwitch. The HRE and HRE Plus do not require any additional software in the MPM. Before installing the HRE some jumpers need to be removed from the MPM.

Refer to Figure 3 for the location of the jumpers. Remove all of the jumpers in location E30 to E35 before installing the HRE or the HRE Plus. The HRE Plus also requires the installation of two modules in location U37 and U55.

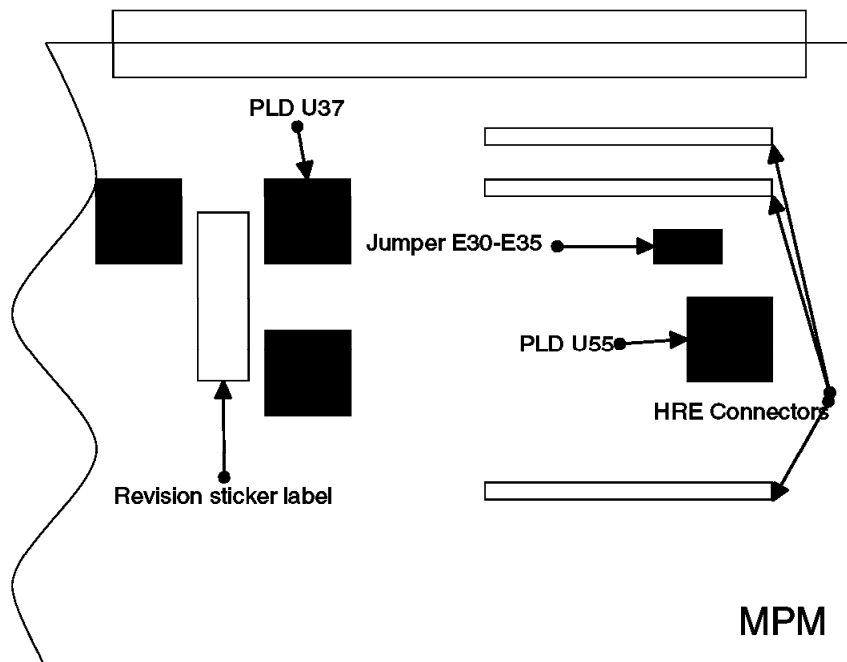


Figure 3. Jumpers and Modules Location on MPM for HRE and HRE Plus

1.4.5 Frame-to-Cell Switching Module (FCSM)

When an IBM 8274 is used as a hybrid LAN frame and cell switch, or as a stand-alone ATM switch the frame-to-cell switching module (FCSM) serves as an internal linkage between the IBM 8274's frame and cell bus. By performing the critical SAR function (segmentation and reassembly), the FCSM converts between LAN frames and cells enabling hybrid frame/cell switch operation in the IBM 8274 without external ATM connections. Additionally, in current RouteCell-only configurations, this module also performs certain ATM management and control functions.

1.4.6 RouteCell Modules

The RouteCell modules are a subgroup of 8274 modules that are optimized for ATM cell switching and are supported in the following 8274 models only:

- 513, 913
- W33, W53, W93

These RouteCell modules only have connections to the 13.2 Gbps cell matrix and are intended to provide the ATM connections when these 8274 models are used as an ATM switch (ATM-to-ATM). Each of these RouteCell modules occupies one slot in any of the currently offered 8274 models. RouteCell module requires FCSM module installed in the 8274

1.4.7 ATM Access Switching Modules

ATM access switching modules allow you to connect the IBM 8274 to ATM servers, backbones and switches without having frame-to-cell switching modules (FCSM). These modules are suited for connecting the switch to an ATM backbone or directly to the ATM server. These ASM modules have connection to the frame bus on the backplane. The SAR function (segmentation and reassembly) is included in each of the ASM modules.

Through the use of point-to-point bridging one can extend all LAN traffic over the ATM backbone. Several IBM 8274s can be connected over one or more backbones. In such a configuration, one combines the flexibility of the 8274's any-to-any switching with the power and speed of the ATM backbone without the use of an ATM backbone switch.

If the IBM 8274 is connected directly to the ATM server using an ASM module, then all non-ATM devices in the LAN can communicate with the high-speed of the ATM server through the 8274.

1.4.8 Switching Modules

Switching modules are available for Ethernet, token-ring, ATM, FDDI, CDDI and WAN interfaces. A variety of connectors, speeds, and signaling options are available for each network interface type.

Each switching module port is assigned a dedicated amount of bandwidth. For example, a 10 Mbps Ethernet module contains ports that each provide the full 10 Mbps of bandwidth. Likewise, an ATM OC-3 module contains ports that each provide the full 155 Mbps of bandwidth

Since the IBM 8274 employs a distributed architecture, each switching module that is added, increases the processing and memory power of the entire switch. The Management Processor Module (MPM) passes off much of the processing

and memory functions to individual switching modules. Switching modules perform software filtering, translations between dissimilar network interfaces (for example, token-ring and Ethernet, Ethernet and ATM, ATM and frame relay), and hardware-based switching.

Each switching module contains at least one RISC processor, RAM for software storage, ASICs for performing hardware-based switching and content addressable memory (CAM) for storing the MAC addresses of source devices. A MAC address for a single source device only needs to be stored once in the CAM of the switching module that received the original frame. Each module's CAM is capable of storing up to 1,024 MAC addresses and one can optionally add another CAM chip to boost the total addresses stored by the module to 2,048 or 4,096, depending on the module's requirement and the maximum CAM supported by the module.

All switching modules provide front panel LEDs that give a quick view of the status of the module, ports, connections and traffic. All switching modules may be hot swapped as long as the same type of module is re-inserted.

1.5 Functional Overview of IBM 8277 RouteSwitch

The following sections describe the newly announced 8277 RouteSwitch. Details on its description and a list of features are listed.

1.5.1 8277 Description

The 8277 is equipped with twenty-four 10/100 Base-T Ethernet ports and can consolidate a cluster of 24 workstations with its own dedicated port. If the higher performance is needed, the 8277 offers smooth migration from 10 Mbps to 100 Mbps without changing the 8277 switch. IBM 8277 Ethernet RouteSwitch is ideal for small-to-mid-size workgroups. It uses advanced hardware architecture to provide high-speed switching between a fixed number of Ethernet ports, and an uplink for high-speed server and backbone access. It supports a very large number of hubs and segments throughout the network by connecting multiple 8277s together over 100Base-TX, 100Base-FX or ATM backbone links. As soon as it is connected to the network it automatically learns the locations of over 1000 devices and starts passing traffic to the network. A hub can be connected to the 8277 switched port.

The IBM Nways 8277 RouteSwitch Software Program (NRSP) is the main controlling software program that is preloaded on the 8277 Ethernet RouteSwitch. The 8277 provides multiple networking protocols between like-to-like entities. The base software, NRSP, offers a wide variety of advanced functions, including layer three switching, IP/IPX routing, Remote Network Monitoring (RMON), policy-based virtual local area networks (VLANs) and LANE client. The Nways RouteSwitch Advanced Routing Software Program (NRAR), Nways RouteSwitch LANE Software Module (NRLS), and Nways RouteSwitch Next Hop Resolution (NHRP) are optional and offer additional functions. The base software, NRSP, must be in place prior to installing any of these optional programs.

1.5.2 The 8277 Ethernet RouteSwitch Model 524

The following list highlights some of the features and functions available in the 8277 RouteSwitch:

- 24-port 10/100 auto-sensing full-duplex Ethernet
- Internal IP and IPX routing
- Layer 3 switching
- HRE, GeB capable
- Full range of options for high-speed connections.
- LAN Emulation, Multiprotocol Encapsulation over ATM and Classical IP over ATM

Chapter 2. Basic RouteSwitch Setup

This chapter describes the Management Processor Module (MPM) and its required software. Its function is to control the RouteSwitches as well as the additional software used to configure the RouteSwitches and to create and manage VLANs.

There are three basic software programs that are used to control and manage the RouteSwitch hardware. They are:

- Nways RouteSwitch Software Program
- Nways RouteTracker Manager
- Nways RouteSwitch Network Manager

For powerful enterprise management of an entire network including equipment from different vendors, RouteSwitch Manager applications can be integrated with a variety of enterprise network management platforms. These platforms include:

- HP OpenView for Windows
- HP OpenView for UNIX
- SunSoft SunNet Manager
- IBM NetView for AIX

2.1 Nways RouteSwitch Software Program

The operating system capabilities of the RouteSwitches are controlled and managed by software. This software is mandatory when ordering an 8274 RouteSwitch. The software controls all the hardware in the switch as well as performing microcode functionality for all the routing, bridging, and setup of VLAN configurations. These functions include:

- IP routing
- IPX routing
- Frame switching
- Backup management
- IEEE 802.1D Spanning Tree Transparent Bridging
- Source route
- Source route transparent
- Any-to-any switching
- Optimized device switching
- Virtual ring extensions
- ATM cell switching

This RouteSwitch Software Program can be thought of as similar in function to the software required for products such as the IBM Nways 2210 and its required software, the Multiprotocol Routing Services (MRS).

2.2 Basic Setup and Operation of the RouteSwitch

This section outlines the setup and basic operation of the RouteSwitch. It discusses basic console functions and commands. The operation of the RouteSwitch is described in relation to basic network configurations and Virtual LANs (VLAN). Further technical studies and examples of VLAN concepts can be found in Chapter 3, "VLAN and Mobile Groups Concepts" on page 37.

2.2.1 Getting Started

After the physical installation of the 8274 the switch is ready to be utilized. Various switching modules could have been installed within the 8274 chassis.

2.2.1.1 Powering on the Switch and Power Diagnostic Test

The 8274 models take up to two power supply modules and have ON/OFF switches. When two power supplies are installed in an 8274, they act as redundant power sharing sources. During normal operation, the power supplies share the load. In the event of failure of one of the power supplies, the other takes over the load of the failed unit. The takeover does not affect the operation of the RouteSwitch.

After power on, the 8274 and its modules go through a power diagnostic check. This is called the Power-On Self Test (POST) cycle. On the 8274, while the diagnostics are running, the MPM OK2 LED will blink amber. When the diagnostics are completed and the modules are operating correctly, the OK1 LED on all the modules should be solid green and the OK2 LED should be blinking green. The RouteSwitch will take between two to five minutes to power-up depending of the hardware and software configurations.

The correct condition on the power display LEDs on the front of all models after power on is:

- 8274 - PS1 LED - Steady green
- 8274 - PS2 LED - steady green if two modules are installed

2.2.1.2 Basic Operation of the RouteSwitch

After powering up, without any software configuration entered the 8274 RouteSwitch will perform switching (transparent bridging) between ports of the same type and any-To-any switching between ports of a different type.

At Version 3.2.3 and higher all of the ports are in optimized mode. Upon activation, the port will not participate in the spanning tree protocol and will start forwarding frames immediately. If more than one MAC address is detected on a port, then the port will go to bridge mode and participate in the spanning tree protocol. Prior to V3.2.3, the ports were open in bridge mode, participating in the spanning tree and would take up to 30 seconds before forwarding frames. Optimized ports were a configurable parameter.

In a switch-based network, a broadcast domain or *group* can be within one physical switch or it can span across multiple switches. As was the case with legacy hub-based networks of the past, a broadcast domain within the LAN switching environment is not limited by:

- A single network interface (for example, Ethernet or token-ring)
- Specific geographic location

An unconfigured RouteSwitch contains one default group called GROUP #1. Within this group, there is also a default VLAN called VLAN #1. The default group and VLAN are comprised of all the physical ports of the switch. All workstations and servers connected to the switch are initially members of the default group and default VLAN before any configuration changes are made to the switch. When a new switching module is added to the RouteSwitch, these additional ports are immediately added to the default group and VLAN.




Figure 4. Basic Operation of RouteSwitch

Any device or workstation that is attached to either Switch A or Switch B will simply acquire network connectivity automatically on power up. The inter-switch link whether it is ATM, Fast Ethernet, Ethernet, FDDI or token-ring will provide connectivity between the two switches with no configuration of the inter-switch port necessary. This operation is sometimes referred to as *plug and play*. In essence what is created by linking multiple switches together is one large, flat, switched network. Any port can be switched from one to another.

2.2.1.3 Any-to-Any Switching

The RouteSwitch with its innovative chassis design provides the option of configuring different media types. Frames from these media types are transported across the switch backplane fabric. For example, an Ethernet frame can be transported onto an FDDI ring or a token-ring frame can be transported onto a Fast Ethernet link. This feature of the RouteSwitch is called *any-to-any switching*.

Workstations and servers connected to the switch communicate by means of established routing protocols but the broadcast domains or groups handled by the RouteSwitch's internal router are no longer limited by specific media type. The switch therefore transforms the frame type of one media into a frame type of another media in such a way that it is acceptable to the routing protocols.

Any-to-any switching is automatically performed by the RouteSwitch software.

This transformation process is discussed in greater detail in 5.7.1, “Any-to-Any Switching” on page 150.

2.3 Switch Configuration Using Console

Several basic definitions are necessary to allow access from the graphical applications of the RouteSwitch. These basic steps must be defined from the console. Once these definitions are in place, then one can use either the console or RouteTracker Manager or RouteSwitch Network Manager for further configuration. Examples of specific configurations for both console mode and the graphical interfaces are documented in Chapter 7, “ATM Cell Switching” on page 201.

These basic definitions may include:

- System name
- System contact
- System location
- TCP/IP address
- Community name

2.3.1 Basic Console Setup

Access to the switch management software console can be gained through a serial (RS-232C) female 9 pin (DB-9) port labelled Console on the front panel of the MPM. The interface cable that is required should be of the straight through pin-to-pin variety and conform to the IBM AT serial port specification.

Flow control on the ASCII terminal emulation software must be set to none.

The terminal emulation software on the RouteSwitch will allow you to change emulation settings. The default factory settings are:

- 9600 baud
- 8 data bits
- No parity
- 1 stop bit

The default user ID is admin and the default password is password.

The user ID of diag and password of switch is also available. It is used to run online diagnostics on the modules installed in the RouteSwitch.

The password will not display on the screen. It is recommended for security purposes that the password be changed on a regular basis by the designated network administrator.

With V3.2 up to four users can now log on as admin or diag.

Only the first user will have write privilege.
The other user will be able to view only.

The following message will appear when a user does not have write privilege:

You are logged in as 'admin' without the WRITE privilege.
The WRITE privilege is currently in use by another user.

Any admin or diag user can see who else is logged on and can possibly gain write privilege.

The command who is used to see who else is logged on as admin (or diag if the command is issued from a diag user).

The command write is used to send a message to another user. The session number displayed by the who command is required. Typing write 0 What is your phone number? will display the message at the console port.

The command kill will grant the admin or diag user the write privilege. The session number displayed by the who command is required. As an example kill 0 will remove the write privilege from the admin user on the console port.

```
Welcome to the IBM Corporation LAN RouteSwitch! (Serial # 64999721)
login   : admin
password:

*****

IBM Corporation LAN RouteSwitch - Copyright (c) 1994, 1995, 1996, 1997
All rights reserved.
      System Name:      no_name
Command      Main Menu
-----
File          Manage system files
Summary       Display summary info for VLANs, bridge, interfaces, etc.
VLAN          VLAN management
Networking    Configure/view network parameters such as routing, etc.
Interface     View or configure the physical interface parameters
Security      Configure system security parameters
System        View/set system-specific parameters
Services      View/set service parameters
Switch        Enter Any to Any Switching Menu
Help          Help on specific commands
Diag          Display diagnostic level commands
Exit/Logout   Log out of this session
?             Display the current menu contents

/ %
```

Figure 5. Logging into the RouteSwitch

Once logged into the RouteSwitch, basic information may be changed. Type syscfg at the command line. The following figure shows a y entered on the Change any of the above? line and the entry of the new information.

```
/% syscfg

System Contact           : Unset
System Name              : no_name
System Location          : Unset
System Description       : DESCRIPTION NOT SET.
Duplicate MAC Aging Timer : 0 (not configured)
Change any of the above {Y/N}? (N) : y
System Contact (Unset)   : Jocelyn Cecire
System Name (no_name)   : ibm8274a
System Location (Unset) : LAB
System Description (DESCRIPTION NOT SET.) : RouteSwitch at IBM
Duplicate Mac Aging Timer (0) :
/%
```

Figure 6. Changing the Basic System Information

The next screen shows the configuration steps for changing the default IP address of the switch. First you will see that the modvl command was issued to modify the default group and default VLAN. The following screen shows the current values.

```
/ % modvl 1
Current values associated with GROUP 1.1 are as follows:

1) GROUP Number          - 1:1
2) Description            - Default GROUP (#1)
IP parameters:
3) IP enabled             - Y
4) IP Network Address     - 192.168.10.1
5) IP Subnet Mask         - 255.255.255.0
6) IP Broadcast Address   - 192.168.10.255
7) Router Description     - GROUP #1.0 IP router vport
8) RIP Mode               - Silent
   {Active(a), Inactive(i), Deaf(d), Silent(s)}
9) Routing disabled       - N
10) NHRP enabled          - N
11) Default Framing       - Ethernet II
   {Ethernet II(e), Ethernet 802.3(8), fddi(f),
    token-ring(t), source route token-ring(s)}
IPX parameters:
12) IPX enabled           - N

(save/quit/cancel)
:
```

Figure 7. Default RouteSwitch IP Address

Then as Figure 8 on page 17 shows, use option 4 and 5 to change the default IP address.

```

1) GROUP Number          - 1:1
2) Description           - Default GROUP (#1)
IP parameters:
3) IP enabled            - Y
4) IP Network Address    - 192.168.10.1
5) IP Subnet Mask        - 255.255.255.0
6) IP Broadcast Address  - 192.168.10.255
7) Router Description    - GROUP #1.0 IP router vport
8) RIP Mode              - Silent
                        {Active(a), Inactive(i), Deaf(d), Silent(s)}
9) Routing disabled      - N
10) NHRP enabled         - N
11) Default Framing      - Ethernet II
                        {Ethernet II(e), Ethernet 802.3(8), fddi(f),
                        token-ring(t), source route token-ring(s)}
IPX parameters:
12) IPX enabled          - N

(save/quit/cancel)
: 4=9.24.105.99
New IP address generates new subnet and broadcast addresses.
Enter '?' to view the changes.
: 5=255.255.255.0
New mask caused change in broadcast address.

: ?

1) GROUP Number          - 1:1
2) Description           - Default GROUP (#1)
IP parameters:
3) IP enabled            - Y
4) IP Network Address    - 9.24.105.99
5) IP Subnet Mask        - 255.255.255.0
6) IP Broadcast Address  - 9.24.105.255
7) Router Description    - GROUP #1.0 IP router vport
8) RIP Mode              - Silent
                        {Active(a), Inactive(i), Deaf(d), Silent(s)}
9) Routing disabled      - N
10) NHRP enabled         - N
11) Default Framing      - Ethernet II
                        {Ethernet II(e), Ethernet 802.3(8), fddi(f),
                        token-ring(t), source route token-ring(s)}
IPX parameters:
12) IPX enabled          - N

: save

```

Figure 8. Changing a RouteSwitch IP Address

Enter save to save the changes. The options of quit and cancel are also available.

If the RouteSwitch is to be placed on a different IP subnet behind an existing router, an IP static route must be defined. This provides the path back from the RouteSwitch through the router to the management station or telnet client. Figure 9 on page 18 shows the commands necessary to implement the static route. In the following example any unresolved IP addresses will be forwarded to the 9.24.105.1 gateway.

```
/ % a1sr

Do you want to see the current route table? (y or n) (y) : y

IP ROUTING TABLE

Network          Mask          Gateway        Metric    Group VLAN
-----
9.24.105.0       255.255.255.0  9.24.105.99    1         1:1
127.0.0.1        255.255.255.255 127.0.0.1      1         LOOPBACK
-----

Destination IP address of host or network : 0.0.0.0
IP address of next hop                    : 9.24.105.1
Route successfully added
```

Figure 9. Adding an IP Static Route

Once the initial IP address is set up using the console, the RouteSwitch Network Manager and RouteTracker Manager software can access the switch.

2.3.2 SNMP Setup

Figure 10 on page 19 shows the setup necessary to enable SNMP on the 8274. First option 5 is set to enabled and even though option 6 does not appear on the screen it must be set. Option 6 is to configure the TCP/IP address of the SNMP management station.

Note

If SNMP is going to be used, the default TCP/IP address of group 1 should be changed to an IP address that complies with your network. The SNMP process always uses the TCP/IP address defined for the default group/VLAN as its reporting IP address, regardless of whether IP is enabled or not in the default VLAN. If IP is disabled in the default group, then the same IP address can now be used in a policy-based VLAN for another group.

```

/ % snmpc
SNMP current configuration:

1) Set Community Name - public
2) Get Community Name - public
3) Trap Community Name - public
4) Broadcast Traps - disabled
5) 0 Unicast Traps - disabled

(save/quit/cancel)
: 5=enabled
(save/quit/cancel)
: ?
1) Set Community Name - public
2) Get Community Name - public
3) Trap Community Name - public
4) Broadcast Traps - disabled
5) 0 Unicast Traps - enabled

(save/quit/cancel)
: 6=192.168.26.188
    Enter trap mask words 0:1 (ffffffff:ffffffff):
    Enter trap mask words 2:3 (ffffffff:ffffffff):
    Enter destination port (162):
    NMS state (on):
    Special Access? (no):
(save/quit/cancel)
: ?
1) Set Community Name - public
2) Get Community Name - public
3) Trap Community Name - public
4) Broadcast Traps - disabled
5) 1 Unicast Traps - enabled
6) NMS IP address - 192.168.26.188 /162 -- bffffffff:ffffffff
                                -- fffffffff:ffffffff (on)

(save/quit/cancel)
: save

```

Figure 10. Enabling SNMP

2.3.3 Listing the Files Loaded in the MPM

The files loaded in the MPM can be listed using the ls command.

```

/ % ls

mpm.img                1563137 01/01/70 00:00
e12.img                26862 01/01/70 00:00
mpm.cmd                 18 01/01/70 00:02
mpm.cnf                 32768 01/01/70 00:00
mpm.cfg                 1024 01/01/70 00:00
mpm.log                 18072 06/02/98 14:20
oem.res                 122 06/02/98 16:02
fesm2.img              30948 06/03/98 08:00
asm.img                805254 06/03/98 08:01
diag.img               157148 06/03/98 08:06
dni.img                17936 06/03/98 08:06

                               998027 bytes free

/ %

```

Figure 11. Listing File in the MPM

The file name is case-sensitive. The RouteSwitch only loads files when the file name is in lowercase.

2.3.4 Copying and Deleting Configuration Files

Occasionally, it may be necessary to alter the configuration settings of the RouteSwitch for the following reasons:

- To reset the RouteSwitch to its factory default settings
- Duplication of files for backup purposes
- Configuration changes implemented on the network

All the configuration settings are saved in two files: mpm.cfg and mpm.cnf.

These files can be erased in order to reset the switch to its factory default settings. The commands to do this are `rm mpm.cfg` and `rm mpm.cnf`.

A file can also be duplicated and the duplicate given a new name thereby creating a backup file. The command to do this is `cp`. An example would be `cp mpm.cfg bak.cfg`. The result would be two identical files: one with the file name `mpm.cfg` and the other `bak.cfg`.

The following console screen captures display how the `mpm.cfg` file has changed.

```

/ % cp mpm.cfg bak.cfg

/flash/mpm.cfg -> /flash/bak.cfg : 100%
/ % ls
  mpm.img                1563137 01/01/70 00:00
  e12.img                26862 01/01/70 00:00
  mpm.cmd                 18 01/01/70 00:02
  mpm.cnf                32768 01/01/70 00:00
  mpm.cfg                1024 01/01/70 00:00
  mpm.log               18072 06/02/98 14:20
  oem.res                122 06/02/98 16:02
  fesm2.img             30948 06/03/98 08:00
  asm.img               805254 06/03/98 08:01
  diag.img             157148 06/03/98 08:06
  dni.img              17936 06/03/98 08:06
  bak.cfg               1024 06/03/98 08:40
                        996795 bytes free. 32768 4/ 1/97 09:12
/ %

```

Figure 12. Creating a Backup File

When the `cp mpm.cfg bak.cfg` is complete, you can see that the file called `bak.cfg` is created.

2.3.5 File Names of Switching Modules

Flash memory on the MPM holds the RouteSwitch's executable images and configuration data for each image file of a switching module. When a switching module comes online, the MPM downloads the appropriate image file for that module into the module's memory. Image files have the extension `.img`. The image files contain executable code for the different modules.

The following table identifies the relevant image file to its hardware switching module.

Table 1 (Page 1 of 2). File Names of Switching Modules		
File Name	Switching Module	Module Type
asm.img	ATM	ASM-155-Fx,ASM-155-C, ASM-DS3, ASM-E3 ASM-DS3,ASM-E3
asmce.img	Circuit emulation	ASM-CE, CSM-AB-CE
cell.img	Cell	CSM-155, CSM-622, CSM-A25, FCSM
diag.img	Diagnostics	MPM
dni.img	Diagnostics	MPM
ds3e3drv.img	DS3	ASM-DS3, ASM-E3, CSM-AB-DS3, CSM-AB-E3 CSM-AB-E3
esm.img	Ethernet	ESM-C-8,ESM-U
e12.img	Ethernet	ESM-C-12,ESM-T-12, ESM-F-8 ESM-F-8

<i>Table 1 (Page 2 of 2). File Names of Switching Modules</i>		
File Name	Switching Module	Module Type
fesm2.img	Fast Ethernet	ESM-100-C, ESM-100-C-FD, ESM-100-Fx-FD, ESM-100C-5, ESM-100CFx-5 ESM-100-Fx-FD, ESM-100C5, ESM-100CFx-5
fsm2.img.	FDDI	FSM-M, FSM-S, FSM-C, FSM-M-C, FSM-SH FSM-MC, FSM-SH
gated.img	Advanced routing functions	MPM
lsm.img	LES/BUS functions	MPM
mesm.img	Ethernet	ESM-C-100-12, ESM-C-16, ESM-C-32, ESM-100FM-8 ESM-100FM-8
mpm.cfg	Basic configuration files	MPM
mpm.cnf	Network configuration files	MPM
mpm.cmd	Startup file	MPM
mpm.img	Base software	MPM
oem.res	IBM unique file	MPM
t1e1drv.img	E1/T1	WSM-FT1/E1, ASM-CE, CSM-AB-T1, CSM-AB-E1 CSM-AB-T3
tsm.img	token-ring	TSM-C6,TSM-F6,TSM-CD6
tsm.pga	token-ring	TSM-F6
wsm.img	WAN	WSM-S

Not all image files loaded in flash memory are required. Only those files needed for the RouteSwitch hardware configuration need to be loaded. Any files that are not required can be removed by using the rm command. For example, if no FDDI module is installed in the RouteSwitch, then the command `rm fsm2.img` could be used to delete the FDDI image file.

The RouteSwitch will alter flash memory contents for the following reasons:

- When a software command requests a configuration change
- When a remote administrator downloads a new executable image
- When the switch fails, a record of the failure is written to flash memory

2.4 Management Processor Module (MPM)

The MPM is the core of the distributed management function of the Nways LAN RouteSwitch. It provides such systems services as:

- Maintenance of user configuration information
- Downloading of switching module software

- Basic bridge management functions
- Basic routing functions
- SNMP management agent
- Access to the software console interface

The MPM contains additional hardware logic to support future routing enhancements to the Nways RouteSwitch and provides clocking for the RouteCell backplane.

2.4.1 Management Processor Module Memory

The standard memory configuration for the MPM II-16, 1G and 1GW is 4 MB of flash and 16 MB of DRAM. Version 3 of the RouteSwitch requires 4 MB of flash and 16 MB of DRAM.

The flash memory is used to store the configuration and the executable files of the RouteSwitch.

The DRAM is used as a working memory to store IP Route entries, PVCs, LE ARP cache and also when the Advanced Routing function is loaded.

Of the 4 MB of flash available, only 3.45 MB can be utilized for the system files. Free memory is required to expand the mpm.cfg and to store the dump files (if one gets created).

The maximum number of files in flash cannot exceed 32. This limit also includes the dump files.

The maximum memory required by V3.2 and above exceeds 4 MB. It is still possible to use V3.2 if all of the required files do not exceed the limit of 3.45 MB.

See Table 1 on page 21 for a list of files that would be required for a particular configuration. The releases notes will contain the size of each file.

In some cases, it may be necessary to add more DRAM memory. There are currently three options available for the MPM-1G and MPM-1GW: 32/4 MB, 32/8 MB and 64/4 MB. The 4/32 and 4/64 will still only have 4 MB flash for system files. The 8/32 option adds an additional 4 MB to the original flash. Some of the system files can now be loaded in the new 4 MB SIMM. This extra 4 MB will appear as a subdirectory in the RouteSwitch file system. The mpm.img, mpm.cmd, mpm.cfg and mpm.cnf files must reside in the /flash directory. The other files can be loaded in either the /flash or the /simm directory. The MPM II-16 is limited to 16 MB of DRAM.

The following scenarios provide basic examples to show when more DRAM memory is required.

Scenario 1: The customer is planning to use:

- Basic switching technology, installing only the RouteSwitch Base Software (without installing any optional Software modules, such as advanced routing or cell switching)
- Less than 2,000 PVCs.
- The internal RouteSwitch routing functionality with less than 20,000 IP route entries

If the customer's requirements match any combination of the above items, they should purchase an MPM 1G or MPM 1GW, which comes supplied with 16 MB DRAM and 4 MB flash memory as standard. This scenario also applies to customer owned MPM-II-16MB modules.

Scenario 2: The customer is planning to use:

- Any of the optional software modules, in addition to the RouteSwitch Base Software.
- More than 2,000 PVCs
- The internal RouteSwitch routing functionality with up to 45,000 IP route entries

If the customer requires any combination of the above conditions, they should purchase an MPM 1G-32MB or MPM 1GW-32MB, which comes supplied with 32 MB DRAM and 4 MB flash memory as standard. Existing customers can upgrade their existing MPMs to 32 MB DRAM.

Scenario 3: The customer is planning to use:

- Basic switching technology, installing only the RouteSwitch Base Software (without installing any optional software modules, such as advanced routing or cell switching)
- Less than 2,000 PVCs
- The internal RouteSwitch routing functionality with up to 70,000 IP route entries

If the customer requires a minimal configuration with a very high number of IP route entries, as defined in the bullet items above, they should purchase an MPM 1G-64MB or MPM 1GW-64MB Management Processor Module, which comes supplied with 64 MB DRAM and 4 MB flash memory as standard. Existing customers can upgrade their existing MPMs to 64 MB DRAM.

Note:

Only IP route entries are referenced in the scenarios above. The reason for this is that the MPM's IPX processing capabilities is restricted to about 2,000 IPX route entries and 2,000 SAP entries, which can be supported by any DRAM configuration, including the minimum 16 MB option. Therefore, increasing the size of DRAM has no effect in an IPX routing environment.

2.4.2 Management Processor Module Redundancy

To provide greater reliability and resilience within the 8274, redundant MPMs are supported. Depending on how they are assigned, they will either be primary or secondary. If the primary MPM fails, the secondary MPM takes over the management functions of the RouteSwitch automatically without any operator intervention.

If two MPMs are implemented within a single chassis, they must be installed in slots 1 and 2 of the switch. Only one processor can be active at a time. The roles of the MPMs are the following:

- Primary - This processor is active and processing commands.
- Secondary - This processor is in a standby and monitoring state.

It is recommended that for full redundancy, the secondary MPM is at the same software version as the primary MPM and the network configuration on the secondary is the same as the primary. In this state, the secondary MPM is able to take over the primary role at any time.

A physical indication of the state of the processors is displayed on the front panel of the MPMs. The PRI LED is green on the primary and the SEC LED is green on the secondary.

2.4.3 MPM Changeover Procedure

The secondary MPM continuously monitors the primary MPM. This monitoring serves two purposes:

- To notify the the secondary MPM that the primary is up and processing
- To update the configuration and keep the two MPMs in-sync

If the secondary MPM detects that the primary is no longer operational, it begins take over as primary instantaneously. When a secondary MPM becomes primary, it resets all other modules in the chassis and performs a primary MPM initialization which takes approximately 30 seconds. For a primary/secondary configuration to be in a full redundant state, the relationship between the two MPMs must meet the conditions shown in the table below:

Table 2. MPM Status	
MPM State	Requirement for State
Redundant	Both MPMs are running the same version of software and the configurations are in-sync.
Configuration Fallback	Both MPMs are running the same version of software but the configurations are different.
Software Fallback	The MPMs are running different versions of software and their configurations may be the same or different.
None	There is only one MPM installed in the chassis.

The current configuration state of the MPM can be viewed from the console by typing the command `slot`.

```

/ % slot

```

Slot	Module-Type	Adm-Status	HW	Board	Mfg	Firmware-Version
Part-Number	Oper-Status	Rev	Serial #	Date	Base-MAC-Address	
1*	MPM-II	Enabled	1.00	64999721	12/17/96	<u>2.3.33</u>
	5012013	Operational				00::20::da::75::8e::d0
2	MPM-II	Enabled	1.00	64999795	12/21/96	<u>2.3.33</u>
	5012013	Redundant				00::20::da::75::e1::c0

Figure 13. Redundant MPM

As can be seen in the preceding screen, slot 1 contains the primary MPM and slot 2 contains the secondary MPM. This screen also displays an indication of

the Firmware-Version level of the MPM. This is highlighted in the furthestmost right-hand column of the screen console. In this case, it is shown that the version of firmware that this switch is at is 2.1.3.

Another useful command to display the status of the processors is `mpm`.

```
/ % mpm
```

```
Currently this slot 1 holds the Primary MPM and slot 2 holds the  
Secondary
```

Figure 14. MPM Slot Allocation

2.4.4 Module Replacement

Even though all modules including the MPM processors are *hot-swappable*, that is, they can be removed from the chassis without powering down the hardware, it is recommended that one of the following procedures be followed when either reseating or changing a module:

1. Procedure One:

- Initiate the swap command from the console.
- Change the swap state to on.
- Change the swap timeout to the time required to perform swapping function. (The default is five minutes.)
- Some additional steps are required when hot swapping an ESM-C-32, M-Ether/12, or the ESM-100F-8, refer to the next procedure.
- Remove or reset the specific module.
- Replace the module.
- Turn swap state to off.

2. Procedure Two

- Initiate the swap command from the console.
- Change the swap state to on.
- Change the swap timeout to the time required to perform swapping function. (The default is five minutes.)
- Enter `reset x reset` where `x` is the actual slot location of the module being swapped.
- Remove or reset the specific module.
- Replace the module.
- Enter `reset x enable` where `x` is the actual slot location of the module being swapped.
- Turn swap state to off.
- After entering the swap command, enter `reset x reset` where `x` is the physical slot location of the module being reset.

If the swap state of the system is not on, the system may halt or restart when a module is replaced. While the swap state is on, a special function called REFERENCE CHECKING takes place. In this state, performance of the RouteSwitch may decrease. Therefore, it is recommended that the swap state only be turned on when you want to swap modules. After completion of the swapping procedure, the swap state should be switched to off. If this is not done, the system will automatically turn the swap state to off after the specified timeout period expires.

Only modules of the same type can replace modules that have been removed. For example, if an ATM switching module has been removed, it can only be replaced with a similar ATM module.

Modules can only be hot swapped when the MPM's OK2 light is in its normal flashing green state. When the module is removed, the OK2 light will flash amber one or two times and then return to normal. On inserting a module, the OK2 light will flash amber again for several seconds and then return to normal. If a module is removed or inserted when the OK2 light is flashing amber, this can cause the system to reset.

If after swapping modules the OK2 light continues to flash amber for more than 10 seconds, a reset of the switch is needed.

An example of the swap command is in the following figure:

```
/ % swap
Swap is OFF, timeout is 5 minutes
usage swap { ON { minutes } | OFF { minutes } }
/ % swap on
Swap is ON for 5 minutes
/ % swap off

Swap time expired without change to chassis THU APR 03 07:08:31 1997.

Swap is OFF, timeout is 5 minutes
```

Figure 15. Initiating the Swap Command

2.4.4.1 Installing a New Module

A new module might be installed into an operational RouteSwitch for one the following reasons:

- Replacing an existing module with a different media module type, for example, an ATM module to replace a FDDI module
- A new module is to be installed in a vacant slot of the RouteSwitch chassis

The RouteSwitch must be rebooted after the installation of the new module or the newly installed module will not be recognized by the MPM and will not work.

If after the reboot the new module is still not operational, the RouteSwitch must be powered off and then powered on in order for the MPM to recognize the new hardware configuration.

2.4.5 Configuring the Modem Port on the MPM

The modem and console ports can be operated simultaneously.

The two serial ports have a physical component and a logical component associated with them.

- Physical console port
 - 9 pin D-Shell female connector
 - Requires a straight-through cable from a terminal emulator
 - Requires a crossover cable from a modem
- Physical modem port
 - 9 pin D-Shell male connector
 - Requires a crossover cable from a terminal emulator
 - Requires a straight-through cable from a modem

The logical definitions can be assigned to either the console port or the modem port. The logical definitions are:

- Console
- Slip
- Auxiliary Console (Version 3.2 and higher)
- Down

By default the physical console port is set to the logical console mode and the physical modem port is set to the down mode.

The 8274 RouteSwitch will display the messages during the boot process to the logical console port.

In the auxiliary console mode, the serial port will be available to access the MPM console even if the other serial port is in use.

<i>Table 3. Port Configuration Status</i>	
Console Port	Modem Port
Console	Down, SLIP or AuxConsole
Down	Console, SLIP
SLIP	Console, Down
AuxConsole	SLIP, Console

The procedure to configure the modem port to be in auxiliary console mode is shown in Figure 16 on page 29.

```

/ % ser

Port to configure? {(C)onsole,(M)odem} (Console) : m
Current Modem port configuration:
  9600 bps, 8 data bits, None parity, 1 stop bit, running Down
Speed {1200/9600/19200/38400} (9600) :
Data size {7/8} bits (8) :
Parity {(N)one/(E)ven/(O)dd} (None) :
Stop bits {0/1/2} (1) :
Mode {(D)own,(C)onsole,(A)uxConsole,(S)LIP} (D) : a
Set (and save) these settings? {(S)ave/(Q)uit} (Save) : s

```

Figure 16. Auxiliary Console Setup for Modem Port

Additional configuration parameters are required when defining SLIP mode on one of the physical port.

```

/ % ser

Port to configure? {(C)onsole,(M)odem} (Console) : m
Current Modem port configuration:
  9600 bps, 8 data bits, None parity, 1 stop bit, running Down
Speed {1200/9600/19200/38400} (9600) :
Data size {7/8} bits (8) :
Parity {(N)one/(E)ven/(O)dd} (None) :
Stop bits {0/1/2} (1) :
Mode {(D)own,(C)onsole,(A)uxConsole,(S)LIP} (A) : s

Current SLIP configuration:

SLIP not running on any ports, do you want to configure it?
Yes, No {Y/N} (Y) : y

Configuring SLIP device sl0:
Local IP address (0.0.0.0) : 1.2.3.4
Remote IP address (0.0.0.0) : 1.2.3.5
Set (and save) these settings? {(S)ave/(Q)uit} (Save) : s

/ %

```

Figure 17. SLIP Setup for Modem Port

Note:

In order for the new serial port mode to take effect the MPM must be rebooted.

When using a modem, it is recommended to configure the modem port for SLIP as this will allow access to the user interface software through a modem SLIP connection.

2.5 Installing Software on the MPM

The Nways RouteSwitch Software Program (NRPS) comes preloaded on the MPM. You do not have to reload unless one of the following conditions occurs:

- Upgrading
- Backing up
- Reloading due to file corruption

There are different transfer methods for loading software onto the MPM. The transfer method that is used depends on the hardware configuration and the operational condition of RouteSwitch. These methods are:

- FTP Server - The RouteSwitch has a built-in FTP server.
- FTP Client - The RouteSwitch can also be an FTP client.
- ZMODEM - Software can be loaded directly through the serial port with any terminal emulator that supports the ZMODEM protocol.

Existing files should be erased before attempting to load new files in the MPM. Existing files can be removed individually using the command `rm` and the file name or the complete image can be removed with the command `imgcl`. The `imgcl` command will not remove the `mpm.cfg`, `mpm.cnf` and `mpm.cmd` files.

When loading software, the versions of software on all the modules must be at the same level. Mixing earlier versions of software with current versions can cause the RouteSwitch to reset or hang.

It is strongly recommended to keep a backup copy of the `mpm.cfg`, `mpm.cnf` and `mpm.cmd` files. This way if an MPM is replaced, the RouteSwitch configuration can be reloaded quickly.

Note:

The latest level of RouteSwitch software can be found at:
<http://www.networking.ibm.com/nes/neshome.html>

2.5.1 Using an FTP Server

The RouteSwitch is an FTP server. Software can be loaded to the RouteSwitch using FTP-compatible software. General instructions in using FTP to transfer software to the RouteSwitch are listed below:

- An IP address must be configured in the RouteSwitch. If this has not been done, refer to Figure 8 on page 17.
- Use your FTP client software just as you would with any FTP server. When you connect to the switch you will be able to see the files contained in the flash directory. It is the only directory in the switch.
- Because of the organization of files in the switch, any time a file is deleted, the flash memory is compacted. Depending on the number of files in the switch and where they are located in memory, this compaction can take anywhere from a few seconds to several minutes.
- Before a file can be transferred to the switch, the old file with the same name must first be deleted. When the old file is deleted, compaction takes place and the transfer can begin. You may not see anything happening for

approximately 2 minutes due to the file compaction procedure. After compaction, the file will be transferred.

- The file name of each file must be in lowercase. The RouteSwitch will not load a file name in uppercase or in mixed-case.

2.5.2 Using an FTP Client

Follow the steps below to start the FTP client:

- Log on to the switch and type `ftp`.
- The system will prompt for a host. The host must be a valid and active FTP server.
- The system will prompt for a user name. If the username `anonymous` is used with no password, the message `Login failed.` will appear immediately after the message `230 User anonymous logged in.` The user `anonymous` is logged in anyway.
- If the system prompts you for a password, enter your password.
- After logging onto the system you will receive the `ftp>` prompt.

Type a question mark (?) to review the FTP commands. The following screen displays:

```
ftp> ?  
Supported commands:  
  ascii    binary    bye      ed         delete  
  dir      get      help     hash       ls  
  put      pwd      quit     remotehelp user  
lpwd  
ftp>
```

Figure 18. FTP Commands

Following is a brief description of each command.

- | | |
|-------------------|--|
| ascii: | Set transfer type to ASCII (7-bit). |
| binary | Set transfer type to binary (8-bit). |
| bye | Close gracefully. |
| cd | Change to a new directory on the remote machine. |
| delete | Delete a file on the remote machine. |
| dir | Obtain a long listing on the remote machine. |
| get | Retrieve a file from the remote machine. |
| hash | Print the # for every block of data transferred. This command toggles hash enabling and disabling. |
| ls | Summary listing of the current directory on the remote host. |
| put | Send a file to the remote machine. |
| pwd | Display the current working directory on the remote host. |
| quit | Close gracefully. |
| remotehelp | List the commands that the remote FTP server. |

user	Send new user information.
lpwd	Display the current working directory on the local host.

If you lose communications while running FTP, you may receive the following message Waiting for reply (Hit Ctrl C to abort). You may press Ctrl+C to abort the FTP or wait until the communication failure is resolved and the FTP transfer will continue.

2.5.3 Using ZMODEM

Normally, it is advisable to use FTP to transfer files to the RouteSwitch. FTP is faster than using the serial port of the RouteSwitch, but the following situations may arise where it is necessary to use ZMODEM:

- No access to an FTP client or server program.
- Image software files on the switch have been deleted or corrupted.

To use ZMODEM, you must have a terminal emulator that supports the ZMODEM protocol. If a file you are transferring already exists in the switch, it must be removed before transferring the new file via ZMODEM.

2.5.3.1 Using ZMODEM from the Command Line

If the RouteSwitch is up and working, log onto the switch and type the command `load` to start the ZMODEM transfer.

```

/File % load

The Console (DCE) port is currently running at 9600 baud
Type 'y' to start ZMODEM download. 'q' to quit (y) : y

Upload directory: /flash
ZMODEM ready to receive file, please start upload or
(send 5 Ctrl X's to abort)
**B000000023be50

```

Figure 19. ZMODEM Load Command

When the transfer is completed, the file or files that have been loaded can be seen by entering `ls` to list the transferred files.

2.5.3.2 Using ZMODEM from the Boot-Line Prompt

If you encounter the situation where you have deleted some or all of the files in your switch, it is necessary to load files through the boot line prompt. This load procedure is done before the switch has finished the boot process. If there is no software available in the switch, then it cannot boot until the software has been reloaded. Loading software through the boot prompt should only be done when the switch is offline and not during the normal working process. The procedures for utilizing ZMODEM through the boot line prompt are:

1. Connect a terminal to the console port. Set the terminal to the last values set in the switch. If the `mpm.cfg` has been deleted or corrupted, the console port values will default to the factory settings.
2. Switch the RouteSwitch on.

- a. If the software cannot load, the RouteSwitch will stop at the boot prompt.
- b. If the software can load, the following message will be displayed.

```
System Boot

Press any key to stop auto-boot...
2
```

Figure 20. System Boot

The number 2 shown counts down to 0. To stop the boot, press any keyboard key before the number counts down to 0.

3. Type ? to view a list of commands available from the boot prompt.

```
:[Boot:]: ?

?                - print this list
@                - boot (load and go)
p                - print boot params
c                - change boot params
l                - load boot file
g adrs           - go to adrs
d adrs:[,n:]     - display memory
m adrs           - modify memory
f adrs, nbytes, value - fill memory
t adrs, adrs, nbytes - copy memory
e                - print fatal exception
n netif          - print network interface device address
L                - list ffs files
P                - Purge system: removes ALL ffs files
R file :[files:] - remove ffs file(s)
S                - save boot configuration
V                - display bootstrap version
$dev(0,procnum)host:/file h=# e=# b=# g=# u=usr :[pw=passwd:] f=#
                  tn=targetname s=script o=other

Boot flags:
0x02 - load local system symbols
0x04 - don't autoboot
0x08 - quick autoboot (no countdown)
0x20 - disable login security
0x40 - use bootp to get boot parameters
0x80 - use tftp to get boot image
0x100 - use proxy arp
0x1000 - factory reset

available boot devices: sl ffs zm
:[Boot:]:
```

Figure 21. Boot Prompt Commands

4. Type c to select a different boot device. The default boot device is the flash memory in the RouteSwitch. Change the boot device to zm. This will tell the system to load files from a ZMODEM connection instead of flash memory.

The default file the RouteSwitch is expected to receive is mpm.img as shown in Figure 22 on page 34. If you are loading more than one file, the file name can then be changed when changing the boot mode.

```

:[Boot:]: c

'.' = clear field; '-' = go to previous field; ^D = quit

Boot device      : ffs zm
Boot file        : /flash/mpm.img
Local SLIP addr   : 0.0.0.0
Local SLIP hostname :
Remote SLIP addr  : 0.0.0.0
Remote SLIP hostname :
SLIP gateway addr :
User             :
Remote password   :
Startup script    : /flash/mpm.cmd
Console params    : 9600,n81c
Modem params      : 9600,n81a
Boot flags        : 0xb
Other             : dvip:ibm8274a,9.24.105.99,255.255.255.0,9.24.105.2

:[boot:]:

```

Figure 22. Changing the Boot Parameters

Press the Return key to accept all the defaults. Enter a period followed by the return key to clear the entry from a field.

5. When this is completed, the system will return to the boot prompt. Type in the @ command to load the boot parameters.

```

:[Boot:]: @

Boot device      : zm
Boot file        : /flash/mpm.img
Local SLIP addr   : 0.0.0.0
Remote SLIP addr  : 0.0.0.0
Startup script    : /flash/mpm.cmd
Console params    : 9600,n81c
Modem params      : 9600,n81a
Boot flags        : 0xb
Other             : dvip:ibm8274a,9.24.105.99,255.255.255.0,9.24.105.2

Attaching network interface lo0... done.
Disk load or Boot load {D/B/Q}? -> d

Upload directory: /flash
ZMODEM ready to receive file, please start upload (or send 5 CTRL-X's to
**B000000023be50

```

Figure 23. Loading the Boot Parameters

6. At the Disk load or Boot load D/B/Q prompt type in d. This tells the system to load from disk and accept the ZMODEM transfer.

```

Upload directory: /flash
ZMODEM ready to receive file, please start upload (or send 5 CTRL-X's to
abort)
**B0I00000023be50

```

Figure 24. Loading from Disk

7. Activate the ZMODEM transfer through the terminal emulation software.
8. When the transfer is completed, use the L command (case-sensitive) to list the files that have been loaded.
9. Repeat this procedure for every file that needs to be uploaded.

Chapter 3. VLAN and Mobile Groups Concepts

This chapter looks briefly at traditional LANs and their evolution to switching. It examines the past ways of segmenting a LAN using bridges and routers, and moves on to the way a switch is capable of achieving this using Virtual LANs (VLANs). It then looks briefly at the policies that the IBM Nways RouteSwitch can use to set up these VLANs.

3.1 Traditional LAN Design

When local area networks were first implemented, they were traditionally very small and consisted of a server and a couple of PCs. Due to the fact that early PCs had limited processing capacity, and could only sustain limited bursts of data on and off the network, the process worked quite well.

LANs have grown at an incredible rate over the last couple of years, and with this growth has also come more powerful PCs as well as more graphical-intensive software. Client/server computing is also present and transferring large files across the network is required.




Figure 25. Example of a Traditional LAN

To enable this growth, the LAN had to be segmented by bridges or routers or both. Bridges, by their nature, propagate traffic (broadcasts) to all stations. This works well in smaller networks, since broadcasts are normally a very small percentage of the total traffic. But in large networks, the number of devices becomes large enough that the broadcasts start overloading the network.




Figure 26. Example of a Traditional Bridged Network

Routing, both in file servers and in external units, solved this problem. Broadcasts are stopped at the router and thus user traffic between LAN segments is controlled. Routers do this by analyzing the layer three header of the protocol, which allows them to filter broadcasts intelligently; however, it does add latency in the network, which occurs in the router. Another disadvantage of routers is that the layer three network address assignment is done on a per port basis. That means if a network station is moved to another building for example, its protocol address has to be changed accordingly.




Figure 27. Example of a LAN That Has Been Segmented Using a Router

So with the incredibly rapid growth in desktop computers in the organizations and the increased processor power of these PCs, we are being driven to a new era of networking: switching. A switch port also provides dedicated bandwidth to the segment or device that is attached to it. Thus, a server attached directly to a port on the switch will be on its own collision domain with dedicated bandwidth.

Network managers are also facing new problems. Continuous moving within the company results in people changing floors or buildings only to find that they can no longer access their servers. The cost to keep segmenting the LAN is very high in a router-based network, due to the cost of a port for the router. In comparison, a switch port is inexpensive.

The physical constraints of the segments in a router and bridged environment limit the mobility of devices. An example of this is a salesman can't just walk into the office, find a free desk, plug into the network and start working. This is because he or she is physically on the wrong segment for the configuration of his or her computer. Thus he or she has to have a specific area to work and this space is wasted when he or she is not in the office.

With the Virtual LAN concept that switching provides, devices can be plugged in on any segment without a problem, and all resources used, for example file servers, printers, etc. will be available. This is because the device is automatically connected to its Virtual LAN by the switch.

3.2 VLAN Design

Virtual LANs are an inherent outgrowth of LAN switching and to fully understand the reasoning and concepts of Virtual LANs, it is necessary to also look at the need to use LAN switches.

Following on from the topics discussed in 3.1, “Traditional LAN Design” on page 37, LANs have continued to grow and so has the traffic on them. With everyone within an organization using a networked PC, the growth of users, and demands for network bandwidth to run applications, a switch may be an attractive alternative.

Flat bridged networks have limitations. They flood the network with broadcasts, multicasts and even unicasts so as to learn paths to destination addresses. In larger networks, this can have an impact on point-to-point traffic. Therefore, it is important to keep broadcast traffic limited in broadcast domains.

A router is a very effective device in controlling broadcasts due to the routing tables and lookups it performs. It also has segmentation of the network on its ports and does not have to flood the entire network when it does broadcast, but rather only the segment(s) involved. The segments are essentially broadcast domains.

A network based on LAN switches also needs to provide this function of controlling broadcasts on the network, hence the concept of the Virtual LAN. The Virtual LAN is a combination of fast MAC layer switching and broadcast control based on analysis of frames to be forwarded.

Switches achieve high performance levels due to the fact that they “switch” data at layer two (of the OSI Model). This differs from a router that has to do things such as address resolution, adaptation of MAC headers, path determination, etc. which is done in software and thus at layer three of the OSI Model. Today’s switches function similarly to transparent bridges and thus are typically much faster than a router.

The idea of controlling broadcasts is not new. What is new is that physically, the broadcast domain is not restricted to one area, but can be scattered around the network. This is because in the past, router networks segmented LANs on a per port basis. Now in a switched environment, logical Virtual LANs based on other policies can be built.

VLANs are logical rather than physical collections of devices. In router-based networks, users are identified by their physical location in the network. In the old model, the location of users would be in one building or a section of that building. A Virtual LAN on the other hand need not be constrained to a physical area, but may be spread out to selected users in different buildings or even across a WAN. Thus VLANs make it possible to build large switched networks in which the VLANs control the broadcasts. VLANs don’t even have to consist of only one single LAN type; there may be Ethernet, token-ring and ATM devices for example, all members of one VLAN based on a common IP subnet. This possibility of having different MAC types within one Virtual LAN allows connections of central resources such as file servers, application servers, or mainframes to run at speeds greater than those of the workstations using switching rather than routing, which is typically faster.

3.3 RouteSwitch Terminology

In order to make a decision on how the RouteSwitch could fit in your environment and to be able to do the proper configuration, the following new terms have to be understood quite well:

- VLAN and mobile group policies
- Non-mobile groups and VLAN
- Mobile groups
- Ports, MAC devices and VLAN timer

Note: The non-mobile groups in Release 3.2.X and later, are exactly the same as groups in Release 3.1.X and lower.

3.3.1 RouteTracker Policies

The RouteTracker policies are used to define membership criteria that must be met in order for a device to join a VLAN in a non-mobile group, or to join a mobile group.

1. Non-mobile groups

In order to limit flooding of broadcasts within a non-mobile group to the necessary ports only, policies must be in place. Policies are used to group devices with common characteristics dynamically together in order to build smaller broadcast domains. Since the assignment to VLANs is based on network-wide policies, devices can be moved within the entire RouteSwitch network without any reconfiguration.

In this section we look at the types of policies and the implications of using each. Keep in mind that a port or device is included in a VLAN if it matches any one VLAN rule. The types of policies that are available for producing VLANs in a RouteSwitch environment are as follows:

- Port rules
- MAC address rules
- Protocol rules
- Network address rules
- User-defined rules
- Binding rules
- DHCP port rules
- DHCP MAC rules

2. Mobile groups

The same RouteTracker policies are used to define membership criteria for mobile groups. The policies are defined in VLAN 1 of the mobile group and no other VLANs can be configured in the mobile group. When a policy is matched, the port and MAC address are moved to the mobile group whose policy has been matched. Multiple policies can be defined for a mobile group, but only one policy has to be matched in order to be dynamically added to the mobile group.

These defining policies are now examined in further detail, in respect to non-mobile groups.

3.3.2 Port Rules

Port-based VLANs are the simplest form of virtual LANs. Being the simplest they also offer the most control and security. They are configured by assigning physical ports to that particular Virtual LAN. Basically this type of VLAN simply groups ports together. There are some disadvantages using port-based VLANs:

- They offer only limited flexibility when relocating stations from one switch to another. The proper port assignment always has to be proven at the new location prior to connecting a station.
- Port policies can cause problems in a multi-switch environment because they are switch-specific and not network-wide.
- In Nways RouteSwitch software Release 2.1 and later, port policies are no longer used to learn VLAN assignments for traffic received on such configured ports. Thus devices not matching any logical policies stay in the default VLAN and are not moved to the port-based VLAN.

Note: To change the port policy so that it works as in Release 2.0 and earlier, add the following line to the mpm.cmd file `reg_port_rule=1`.

Port policies are, in fact, not useful in the creation of consistent VLAN membership across multiple switches, so why use port policies at all? Port policies should be added as a second rule to logical VLANs in these situations:

- If a device does not transmit any traffic during its initialization phase that matches an appropriate VLAN (such as some network printers), the port to which the device is connected never gets assigned to a VLAN, thus other stations that have already joined a VLAN can not communicate with that device.
- VLANs and their internal virtual router are not activated until at least one port is assigned to that VLAN. The RouteSwitch as mentioned earlier only assigns ports to VLANs if a frame is received at a particular port that matches any one VLAN policy. This means, in some network situations, a port policy may be needed to force a VLAN active, thus allowing the internal router to be able to participate in routing updates.
- All RouteSwitch ports that are used for backbone connections should be added to a VLAN by an additional port policy in order to enable traffic from that VLAN to flow out onto the backbone.

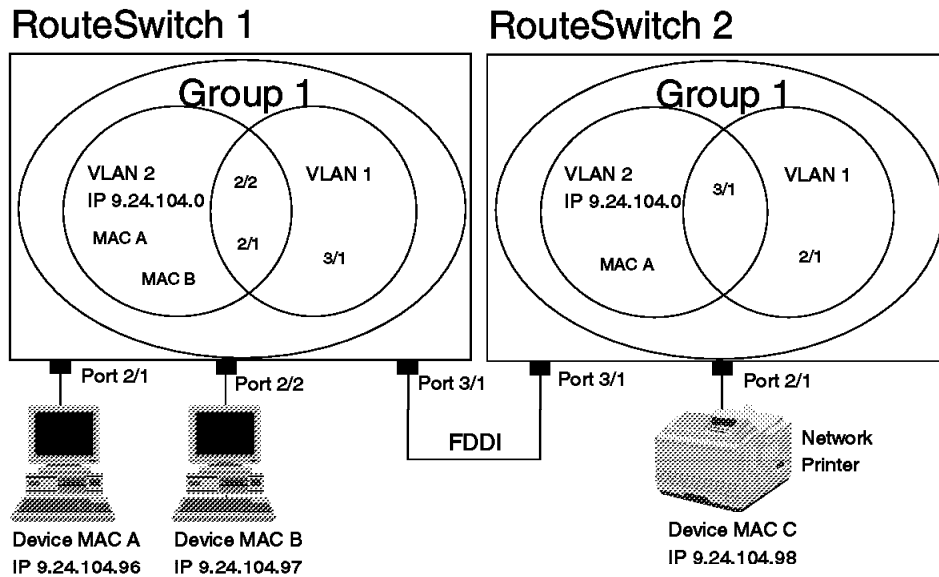


Figure 28. The Usefulness of Port Policies

In Figure 28, initially just the network policy-based VLAN 2 was configured for the IP network 9.24.104.0 in each of the RouteSwitches. Assume device C is the only station at switch 2 that could join VLAN 2 and it did not send a frame during initialization. As soon as device A communicates with B, their MAC addresses are moved from default VLAN 1 to VLAN 2 and they are deleted from VLAN 1. Their ports also become members of VLAN 2, and the ports 2/1 and 2/2 are now members of VLAN 1 and VLAN 2. Port 3/1 of switch 1 remains in VLAN 1 since no traffic was received on this port. On the other hand, port 3/1 of switch 2 becomes a member of VLAN 2 in that switch because of the very first ARP broadcast sent by A, which was flooded to all ports of group 1 including port 3/1. If device A now wants to communicate with device C, it sends out an ARP request in order to resolve device C's MAC address. This broadcast, according to the flooding rules discussed in 3.5.2, "Frame Flooding for Non-Mobile Groups" on page 54, stays within VLAN 2 and thus is not forwarded to switch 2 because port 3/1 is not a member of VLAN 2. This is the current status of the VLAN assignment in Figure 28. To solve this problem, port 3/1 was added to VLAN 2 by configuring a port rule as a second rule to each VLAN 2, which forces port 3/1 to VLAN 2. Also port 2/1 of switch 2 was added to VLAN 2 in order to force the quiet network printer to that VLAN. Now the ARP request arrives at C and thus, both devices are able to communicate. With this configuration, a network-wide VLAN membership consistency is achieved.

3.3.3 MAC Address Rules

A MAC address VLAN is a grouping of MAC addresses that builds a broadcast domain (VLAN). Configuration is done by assigning MAC addresses to the VLAN.

MAC address Virtual LANs can take a long time to implement; imagine configuring a network with 1,000 users and each of these users must be

assigned to a VLAN. In addition, MAC addresses are inherently cryptic, making them a bit more difficult to configure.

MAC address-based VLANs provide a little more flexibility than port-based VLANs because VLAN assignment is done by the source MAC address thus a station could be connected to any port of the switch.

MAC address-based VLANs provide a larger degree of control and security if universal adapter addresses are used because by their very nature, they can not be changed. Thus, a device won't be able to join a MAC-based VLAN if its MAC address is not predefined in the switch.

But MAC-based policies have several disadvantages:

- In a multi-switch environment every MAC address has to be defined in all switches in order to have a consistent network-wide VLAN membership.
- Every time a new device is added to the network, the new MAC address has to be added to all switches.
- If a network adapter card fails and has to be replaced, the new address has to be added to all switches if universal addresses were used, as well as deleting the old adapter card address, if it is known.
- More effort is needed to gather all addresses and keep track of them.
- Some communication problems may occur in multi-switch networks since the ports used to link the switches will not become members of a MAC-based VLAN unless they are forced by an additional port rule.

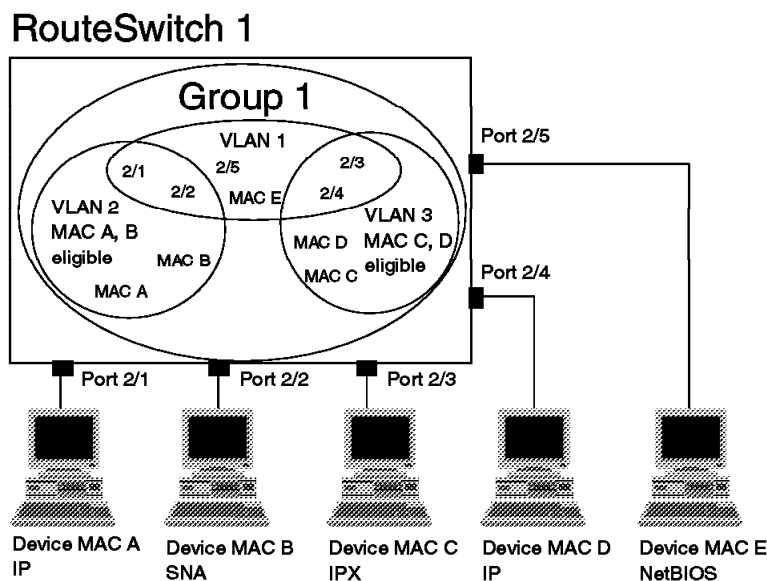


Figure 29. MAC Address-Based VLAN Example

In Figure 29, MAC devices A and B and MAC devices C and D will be able to communicate respectively. Since universal MAC addresses that are not configurable at the devices were used, station A has no chance to become a member of VLAN 3 and thus will not be able to communicate to any station in

VLAN 3. Since MAC E is not defined, device E remains in VLAN 1 and gets isolated.

Note: If a hub is connected to a port, all devices connected to this hub will be members of a MAC-based VLAN if only one device's MAC address matches a defined MAC rule.

3.3.4 Protocol Rules

Protocol-based VLANs provide a convenient way of grouping devices together based on a common layer three protocol. Since traffic is separated by protocol, this can significantly reduce the amount of unnecessary packets an adapter has to receive and forward to the software for further processing. Later, these packets are discarded by the upper layer software because a particular protocol is not active or even installed in that device.

Furthermore, protocol-based VLANs are well-suited to configure consistent VLAN assignment of devices in a multi-switch environment.

The protocol policies are a set of given fields within a frame that are common to all traffic designated to participate within a protocol-based Virtual LAN. Because the intelligent policies specify a field within the frame, the user does not have to be concerned with the media from which the frame was received, or the offset into the frame. Following is a list of the set of protocol policies the user can define to govern traffic within a Virtual LAN:

- All IP protocol traffic
- All IPX protocol traffic
- All DECnet protocol traffic (DECnet Phase IV only)
- All AppleTalk traffic
- All traffic of a specific Ether-Type
- All traffic with a specified source and destination SAP header
- All traffic with a specified SNAP type

The SNAP header consists of five bytes. The first three bytes of the SNAP header is the vendor code, generally the same as the first three bytes of the source MAC address, although it is sometimes set to zero. Following the vendor code is a two-byte field that typically contains an Ether-Type for the frame. Here is where the backwards compatibility with Ethernet Version II is implemented.

Please refer to Appendix B, "Ether-Types and SAP Listings" on page 341 for valid Ether-Type or SAP values or look at http://www.optimized.com/tech_cmp/index.html for further explanations.

Protocol rules should be used to define VLANs for all protocols that are not currently routed by the RouteSwitch in order to isolate broadcasts to the participating devices.

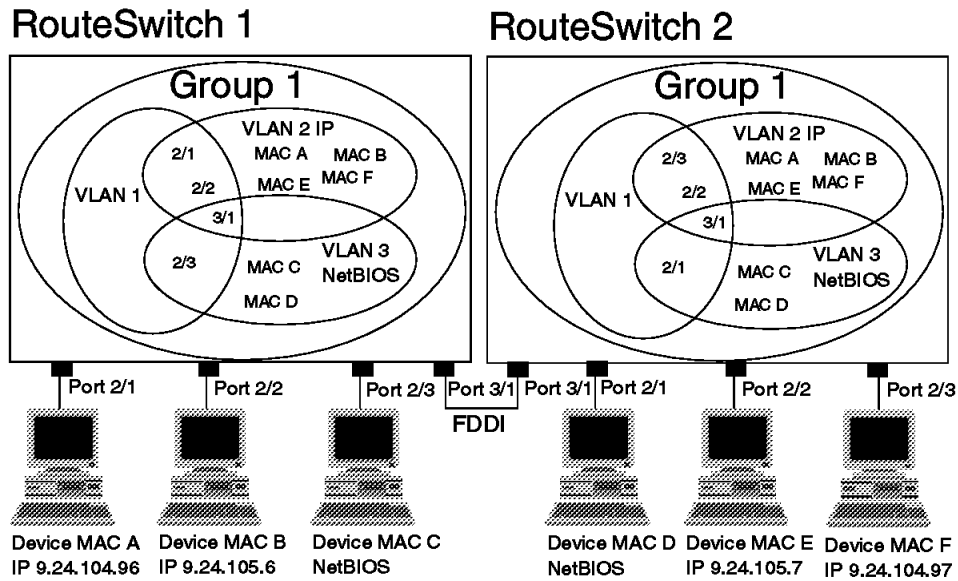


Figure 30. Protocol-Based VLAN Example

In Figure 30, the VLANs are determined by the protocol that a device is using, therefore all devices using NetBIOS make up one VLAN (in this case VLAN 3), and all devices using IP make up a separate VLAN (in this case VLAN 2). NetBIOS relies heavily on broadcasting. Therefore, two broadcast domains based on protocol type were built, thus no unnecessary NetBIOS broadcasts will be forwarded to IP stations and vice versa.

Note: A port rule to both protocol-based VLANs in both switches was added in order to force port 3/1 in VLAN 2 and 3.

3.3.5 Network Address Rules

Network address rules allow the network administrator to build broadcast domains based on the logical layer three network topology. This is the exact criteria that network administrators have been using to create broadcast domains with routers, but with the difference that devices can be moved to another location without the need of any change. This is the most effective way to limit unnecessary broadcast flooding within the network. The following list is a set of the network address policies the network administrator can define to govern traffic within a virtual LAN:

- IP network address and IP network mask
- IPX network number and encapsulation type

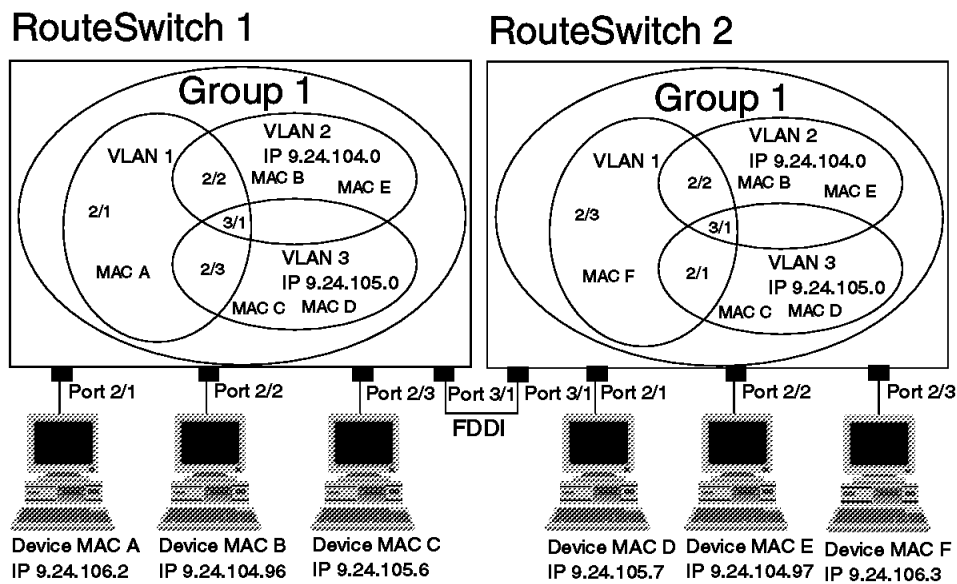


Figure 31. Network Address-Based VLAN Example

The VLAN policy in Figure 31 has been defined to split the devices into two VLANs. Devices with IP addresses in the 9.24.104.0 subnet are defined to be in VLAN 2, and devices with IP addresses in the subnet 9.24.105.0 are members of VLAN 3. A subnet mask of 255.255.255.0 is assumed. The advantage of using network rules is that all ARP requests stay within the appropriate VLAN once the MAC addresses have been learned by the RouteSwitch.

Notes:

1. A port rule to both protocol-based VLANs in both switches was added in order to force ports 3/1 in VLAN 2 and 3.
2. Do not activate IPX routing on the default VLAN (VLAN 1) if IPX network-based VLANs have been configured. This is because when the internal router sends out RIP and SAP broadcasts on VLAN 1, the broadcasts are flooded out of all ports in the group. Thus all servers receive the broadcasts and respond to this with a router configuration error.

3.3.6 User-Defined Rules

Network managers may also define their own policies for configuring Virtual LANs. User-defined policies provide the definition of VLAN membership on the basis of a specified pattern at a specified location within a frame. All devices that originate frames containing the specified pattern at the specified location are included in the VLAN. User-defined policies allow the definition of custom policies for a specific application not covered by the RouteSwitch's network or protocol policies. To create a user-defined policy, the following components must be defined:

- An offset, to define the location of the pattern in the frame
- A value, to define the pattern

- A mask, to define the bits in the value that should be recognized

3.3.7 Port Binding Rule

This rule allows you to:

1. Bind an IP address to a port and a MAC address.
2. Bind a MAC address to a protocol and a port. The protocols that the MAC address can be bound to are the same as the protocols in 3.3.4, “Protocol Rules” on page 44.

3.3.8 DHCP Port Rules

This rule is the same as the port rule (see 3.3.2, “Port Rules” on page 41) but applies only to ports on the switch on which DHCP clients are attached.

Note: For non-mobile groups, the port must be moved to the appropriate group using `addvp` before enabling the DHCP port policy.

3.3.9 DHCP MAC Address Rules

This rule is the same as the MAC address rule in 3.3.3, “MAC Address Rules” on page 42, but also only applies to MAC addresses of DHCP clients that are attached to the switch.

3.3.10 Non-Mobile Groups

Version 3.2 introduced new configuration options for the RouteSwitch. This section details the original method used for defining groups and VLANs in the RouteSwitch.

1. Non-mobile group

Non-mobile groups make it possible to physically subdivide the network on a per port basis. The RouteSwitch supports network-wide Non-mobile group identifiers ranging from 1 to 65,535. One non-mobile group of physical ports can be spread over different RouteSwitches throughout the entire network. Each switch chassis can support up to 32 different non-mobile groups. Each non-mobile group is viewed as a separate entity and traffic from one non-mobile group cannot be forwarded to another non-mobile group without the use of routing. To summarize:

- A non-mobile group is a broadcast domain.
- A non-mobile group is a collection of physical ports.
- Non-mobile groups can span switches.
- Non-mobile groups may not overlap.
- Ports may belong to only one non-mobile group.
- Frames are routed between non-mobile groups.
- Frames are routed between VLANs in the same non-mobile group.

2. RouteTracker VLAN

In prior releases, RouteSwitches supported port-based VLANs only. These VLANs were equivalent to today’s non-mobile groups. In order to provide for multiple VLANs in the same physical network, the paradigm of port-based Virtual LANs was expanded and the VLAN assignment is now based on logical criteria called policies rather than physical ports. The software

component that provides this functionality is called RouteTracker, thus policy-based VLANs are RouteTracker VLANs.

Note: In the switch configuration software IBM Nways RouteTracker, policy-based VLANs are called RouteTracker VLANs but at the console, the term AutoTracker is used instead of RouteTracker. Both terms refer to the same function within the RouteSwitch.

In summary a RouteTracker VLAN:

- RouteTracker VLAN is a smaller portion of a non-mobile group's broadcast domain.
- Broadcasts as well as frames with unknown destination addresses stay within a VLAN, provided the broadcast and unknown frames originate inside the Virtual LAN.
- RouteTracker VLAN is a logical association of virtual ports.
- VLANs may overlap, which means that one port or MAC device may be member of different VLANs.
- VLANs are defined by rules or policies using the RouteTracker capabilities.
- The assignment of a MAC device to a VLAN(s) is done by a complete analysis of the first frame a station sends to the RouteSwitch, while broadcasts and multicasts are always checked by the MPM.
- Upon this analysis, a station becomes a member of all VLANs that have matching policies.
- Up to 32 policies of each type may be defined per non-mobile group. Since there are five types of policies, there is a maximum of 160 policies per non-mobile group.
- A VLAN first becomes activated if at least one device or port is assigned to it by the RouteTracker.
- LAN frames matching a rule or a set of rules are logically grouped together.
- The rules for all VLANs are contained in the RouteTracker rules database.
- VLAN numbers must be unique within a non-mobile group.
- Switching occurs within a VLAN.
- Routing occurs between VLANs.
- Only one virtual router port with only one IP and/or IPX network address may be assigned to a VLAN.
- There can be up to 32 VLANs within a non-mobile group.

3. Multicast VLAN

Multicast VLANs (MVLANs) provide flooding control of multicast traffic in a network. Multicast destination addresses are used by several applications to deliver information from one source to multiple recipients. For example, an application is a server distributing CNN Newscasts. The goal of multicast VLANs is to reduce CPU cycles because a defined multicast is only processed once by the MPM and then delegated to the switch modules. The other advantage is that the multicasts propagation could be limited to predefined stations, which really need this kind of information. In summary:

- MVLANs operate independently of RouteSwitch VLANs.
- Up to 32 MVLANs may be created per non-mobile group.
- MVLANs are defined by a multicast address.
- The members of an MVLAN are specified ports or MAC addresses.
- There is no default MVLAN.
- Routing may not be configured for MVLANs.

4. Default Group and Default VLAN

Since the default group as well as the default VLAN have special characteristics, the following rules apply:

- Initially, the RouteSwitch has a default group (Group #1), with one default VLAN (VLAN #1) inside of it and all physical ports are part of this non-mobile group and VLAN.
- Whenever a new non-mobile group is created, a default VLAN for this non-mobile group is also created.
- The default group is the pool from which ports are selected to become members of other non-mobile groups.
- If RouteSwitches are interconnected, all default groups are bridged by default while all additionally created non-mobile groups can be trunked in order to share one link between RouteSwitches with different non-mobile groups. The traffic remains separated between non-mobile groups even if it flows through one single link.
- A non-mobile group must exist prior to VLAN creation.
- Neither the default group nor the default VLAN may be deleted.
- VLAN #1 is the "home of the homeless" which means that any traffic that has not been tagged as belonging to a policy-based VLAN goes here.
- If the default VLAN is turned off by using the command `defvl off`, frames from devices are automatically dropped when they do not match any VLAN policy, which means they are isolated from the network.
- If the default VLAN is off, all types of traffic existing in the network should be covered by VLAN policies, otherwise devices may not be able to communicate.
- For maximum control of LAN traffic, turn default VLAN off.
- No policies can be added to the default VLAN.
- All physical ports always remain members of the default VLAN, but they can also become members of other VLANs.
- All MAC devices are initially part of the default VLAN.
- Individual MAC devices are removed from the default VLAN if they match any one defined policy and moved to a non-default VLAN.

5. Virtual Port

The RouteSwitch assigns at least one virtual port to every physical port for internal processing. Virtual ports are numbered sequentially within the entire RouteSwitch, thus each virtual port number is unique. Ports for which a service may be configured, such as FDDI or ATM, may have multiple virtual ports. Each configured service has its own virtual port. An ATM port,

for example, may have three virtual ports if there are three LECs configured for it.

By default, the description field of all virtual ports contains the virtual port number, which can be displayed by using the `via` command; however, this could be changed to describe the actual configuration. If there are multiple different services configured for one port for example, the service should be described here for easier administration in the future. However, these virtual port numbers are not of importance because console commands never refer to them; they always relate to group number, module slot, service, instance, and physical port numbers.

Note: There is no virtual port number assigned to the internal virtual router ports.

3.3.11 Mobile Groups

Mobile groups allow ports to be dynamically assigned to the group based on RouteTracker policies. Both mobile and non-mobile groups can coexist in the same network and even the same RouteSwitch, this is shown in Figure 32. Only Ethernet and token-ring ports can be dynamically assigned to groups.

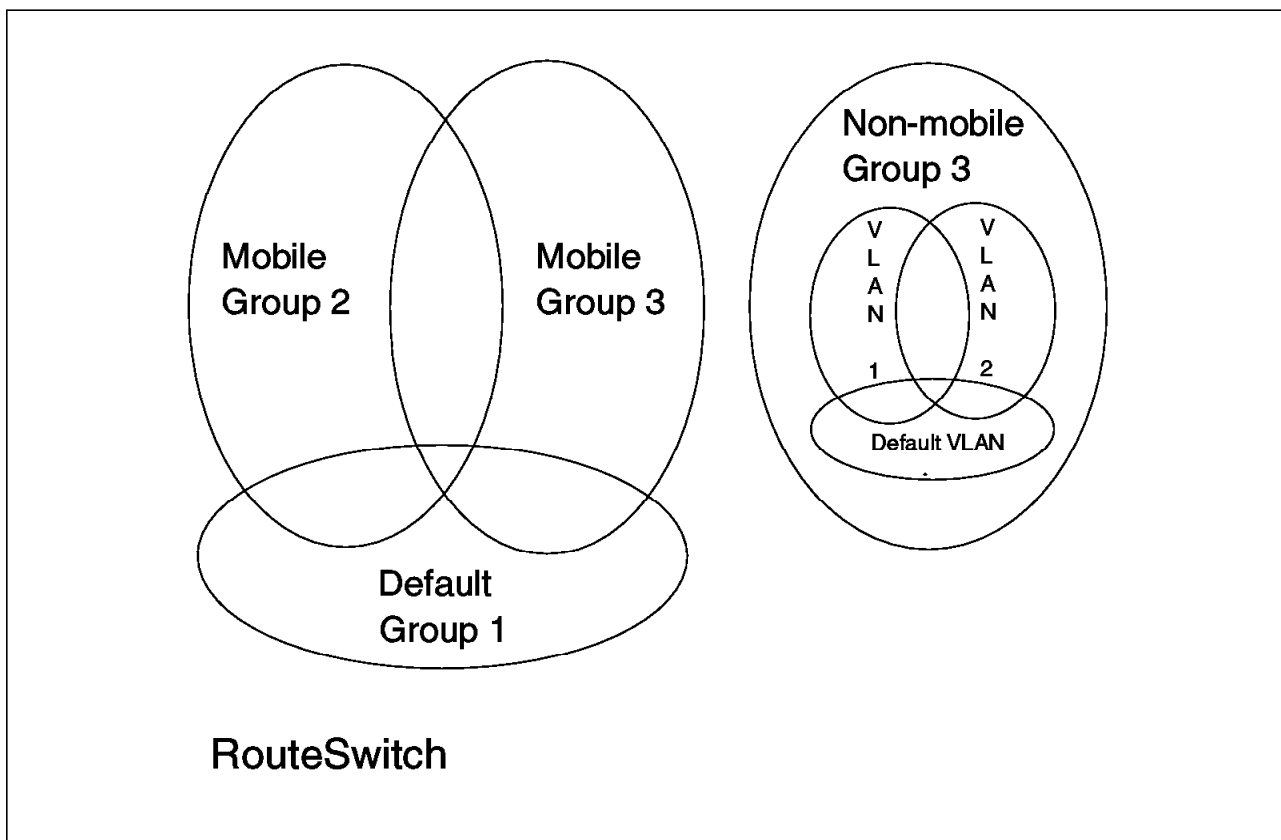


Figure 32. Coexistence of Mobile and Non-Mobile Groups

In the following list the terms used in the mobile group environment will be defined. Some of the terms from the non-mobile group environment are the same, however their function may differ in a mobile group configuration.

1. Mobile groups

The dynamic method of assigning ports to groups is called group mobility, thus the terminology we use is:

- Mobile groups - Groups for which group mobility has been enabled.
- Non-mobile groups - Groups for which group mobility is disabled.

Important

By default, the mobile group function is disabled in the RouteSwitch. The command `gmcfg` is used to enable the mobile group function within a RouteSwitch.

There are no RouteTracker VLANs within the mobile groups. However the RouteTracker policies are defined to the default VLAN within the mobile group, and these policies apply to the entire mobile group in which they are defined. To summarize:

- Physical ports are dynamically assigned to mobile groups.
- Ports can belong to multiple mobile groups.
- Frames are routed between mobile groups.
- Frames are switched within mobile groups.
- Whenever a new mobile group is created, a default VLAN for this mobile group is created.

2. RouteTracker Policies

RouteTracker policies are defined in the default VLAN in the mobile group. These policies apply to the entire mobile group and membership is granted if a policy is matched within the mobile group.

- Mobile groups have to be defined prior to creating RouteTracker policies.
- Policies can also be defined during the configuration on the mobile group.

3. Default Group

The following characteristics apply to the default group with regard to mobile groups.

- Physical ports must belong to the default group so that they can be assigned to mobile groups once a policy has been matched.
- If ports are removed from the default group and assigned to a non-mobile group, they can not be dynamically assigned to a mobile group.

4. Dynamic Port Assignment

When a device on a port matches a RouteTracker policy defined in a mobile group, the port is dynamically added to that mobile group.

To summarize:

- Each port is initially assigned to the default group.
- RouteTracker examines traffic on the ports to see if policies can be matched.
- If more than one device is coming in on a port, the port can belong to more than one mobile group.
- Only ports in the default group and other mobile groups are eligible for dynamic port assignment.

- Only token-ring and Ethernet ports can be dynamically assigned to mobile groups.
- LANE service ports carrying token-ring or Ethernet traffic can also be dynamically added to a mobile group.
- If the same RouteTracker policy is defined on two separate mobile groups, a port will be dynamically added to the mobile group with the lowest group number.
- Once a port has been added to a mobile group it will stay in that group, even if the MAC address of the device times out of the RouteTracker VLAN. This variable can be changed globally per RouteSwitch.

3.4 Ports, MAC Devices and VLAN Timers - Non-Mobile Groups

Ports and MAC addresses are treated differently in the RouteSwitch in order to limit the broadcasts and unicasts with unknown destination addresses to the VLAN of the source device. This section explains how this is achieved by the RouteSwitch.




Figure 33. Port and MAC Address Handling

Initially all ports and MAC devices belong to the default VLAN 1. As soon as the first frame of device A is received at a RouteSwitch port, the frame is flooded out of all ports within the same group and in parallel it is completely analyzed by the MPM. Assuming the frame contents matches the defined policy for VLAN 2, then following will occur:

- The MAC address of device A is removed from VLAN 1 and moved to VLAN 2. Therefore, device A is no longer a member of the default VLAN 1.
- The port to which device A is connected to (port 2/1 in this case) will be added to VLAN 2 but also remain in VLAN 1 because all ports are always members of the default VLAN.

- VLAN 2 is activated because port 2/1 is the first port assigned to that VLAN.

If device A sends a subsequent broadcast frame, this frame is only forwarded to ports that are in the same circle, just VLAN 2, and not flooded out to the rest of the RouteSwitch ports.

Note: If MAC A was assigned to multiple VLANs, the broadcast would be forwarded to all ports that have VLANs in common with MAC A.

The RouteSwitch uses a timeout mechanism to remove inactive devices and ports from their logical VLAN(s). All learned MAC addresses and their VLAN membership information is stored in the Content Addressable Memory (CAM). Each switching module maintains its own CAM. Only source MAC addresses of frames that were received at a particular port of a switch module are stored in that module's CAM. There is a timer in place in order to delete outdated MAC addresses and their VLAN membership information. If the switch does not receive a subsequent frame from a device within a certain amount of time, the device's MAC address and VLAN membership flags are removed from CAM. The CAM timeout for inactive devices varies depending on available CAM storage and is initially set to the maximum value of 300 seconds. Further information can be found in 1.4.1, "RouteSwitch Architecture" on page 4.

This means, if device A is inactive for more than five minutes, the RouteSwitch deletes its MAC address from VLAN 2 and MAC A becomes an unknown device again.

The port timeout is different and initially set to 1200 seconds. Therefore, VLAN 2 membership of device A is given for at least 20 minutes based on the port membership.

Important

If a station is connected to a dedicated port at the RouteSwitch and does not send a frame out that matches a particular VLAN policy within a time period of 20 minutes, its MAC address and port assignment is deleted from that VLAN by the RouteTracker. Thus, this station is no longer available to other devices in that VLAN until it sends out a matching frame again.

For example, the above described RouteSwitch behavior may cause communication problems for network printers that are quiet for more than 20 minutes.

Please refer to 3.3.2, "Port Rules" on page 41 for a solution to such problems.

3.5 RouteSwitch Packet Processing for Non-Mobile Groups

The RouteSwitch switching process combines hardware-based switching with software-based management. The result of this design is that the standard switching is performed by hardware while the RouteSwitch software provides the opportunity to examine each frame.

3.5.1 The Switching Process for Non-Mobile Groups

The following list describes the steps the MPM and NSMs use in determining the switching of frames.

1. A frame is received on a network interface.
2. Hardware looks up the source address to ensure that the source has been learned and that it is in the correct VLAN.
3. If the source address is unknown (not in CAM), the NSM tags the frame indicating that it is intended for the MPM, and forwards the frame to the VBUS.
4. The MPM examines the entire frame to determine which VLAN(s) the source address qualifies for membership in and adds the source address to the forwarding table. The MPM then relays information concerning VLAN membership to the NSM via the MBUS.
5. The NSM now adds an entry into its CAM and modifies the filtering database with VLAN membership information.
6. Broadcasts and multicasts are always sent to the MPM for a complete analysis.
7. If the source address is known to belong to another VLAN, the management system is notified and the port is automatically reconfigured to join the correct VLAN. (This applies to optimized device switched ports.)
8. Simultaneously, all NSMs copy the frame from the VBUS where it is examined by the hardware forwarding engine of each module.
9. If the destination address is recognized, the frame will be picked up by the designated switching module.
10. This NSM will check the filtering database to determine if the destination device belongs to a VLAN that is also associated with the virtual port that received the frame from the source device. If so, the NSM will do any necessary translations that may be required and forward the frame out the correct port; otherwise the frame will be dropped. The NSM also sends a signal claiming the frame and all other NSMs abort.
11. If no module recognizes the destination address, and the source address is known, the frame is forwarded by all switching modules within common VLANs with the source address.
12. If no module recognizes the destination address, and the source address is unknown, the frame is forwarded by all switching modules to all ports within the same group. RouteSwitches do not flood unrecognized frames to optimized device switching ports, further conserving the dedicated bandwidth at that port.

3.5.2 Frame Flooding for Non-Mobile Groups

Flooding occurs when a frame received by the RouteSwitch is destined for a MAC station that has never been heard from before. In a typical bridged environment, the frame will be forwarded out all ports. The situation is different in the RouteSwitch because of the existence of RouteSwitch non-mobile groups that lead to segmentation of the network into broadcast domains. Frames are treated differently under different conditions. A RouteSwitch non-mobile group consists of a set of ports which is the maximum set of ports out of which a broadcast frame will be forwarded. The following tables list the criteria for flooding traffic within a RouteSwitch non-mobile group.

<i>Table 4. Unicast Processing</i>		
MAC Address	Known Destination	Unknown Destination
Known Source	Frame is switched to destination port if source and destination addresses have at least one VLAN in common. Frame is not analyzed by the MPM and no additional VLAN assignment is made.	Frame is forwarded out of all ports that have at least one VLAN in common with the source MAC address except optimized device switching ports. Frame is not analyzed by the MPM and no additional VLAN assignment is made.
Unknown Source	The frame is completely analyzed by MPM and VLAN assignment is updated in the NSM's CAM. Then the frame is switched to the destination port if the source and destination address have at least one VLAN in common.	Frame is flooded out all ports that are in the same non-mobile group as the source port except optimized device switching ports. The frame is completely analyzed by MPM and VLAN assignment is updated in the NSM's CAM.

<i>Table 5. Broadcast and Multicast Processing</i>	
MAC Address	Processing
Known Source	Frame is forwarded out of all ports that have at least one VLAN in common with the source MAC address. The frame is completely analyzed by MPM and VLAN assignment is updated in the NSM's CAM.
Unknown Source	Frame is flooded out all ports that are in the same non-mobile group as the source port. The frame is completely analyzed by MPM and VLAN assignment is updated in the NSM's CAM.

Please refer to Chapter 7, “ATM Cell Switching” on page 201 for further information about how these rules are applied in a network.

3.5.2.1 VLAN Example 1 - Non-Mobile Groups

The following example provides a better understanding of how VLANs are implemented for non-mobile groups in the RouteSwitch. Please refer to 3.5, “RouteSwitch Packet Processing for Non-Mobile Groups” on page 53 in order to see how received frames are actually processed in the RouteSwitch.

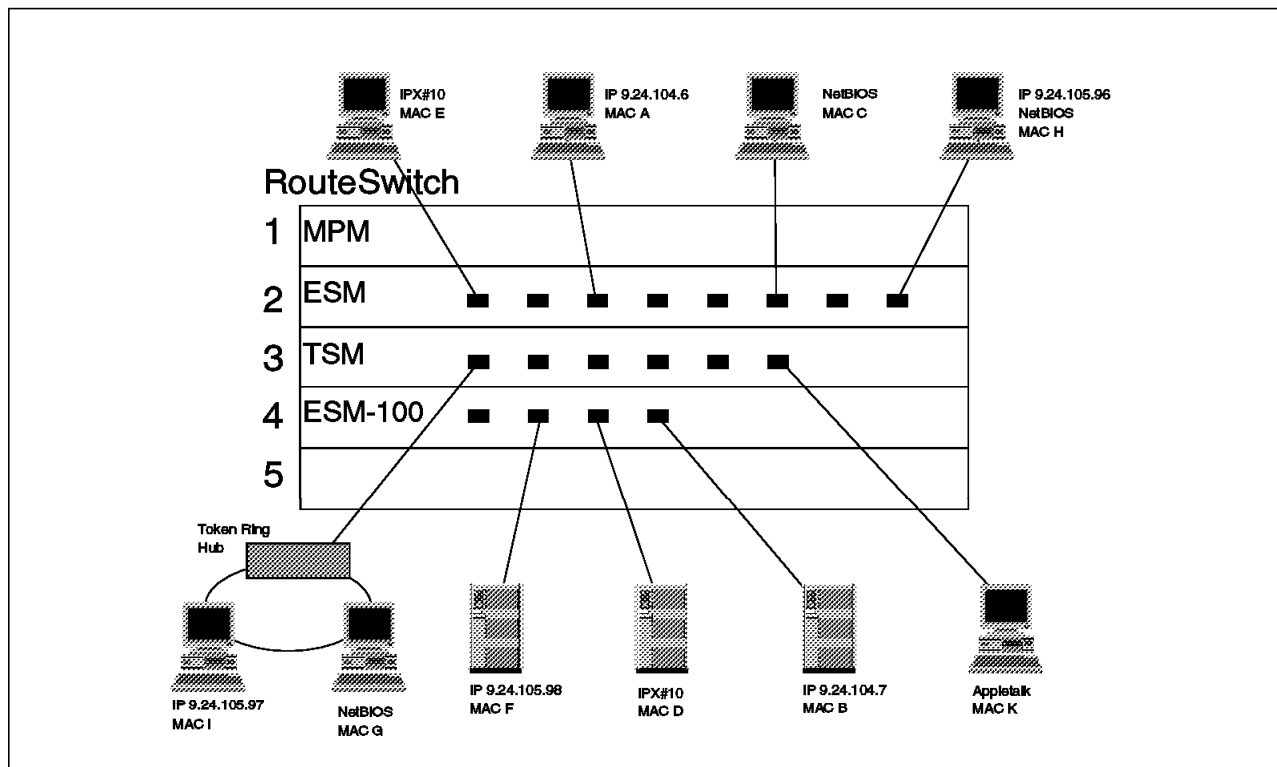


Figure 34. Physical View of a RouteSwitch Network

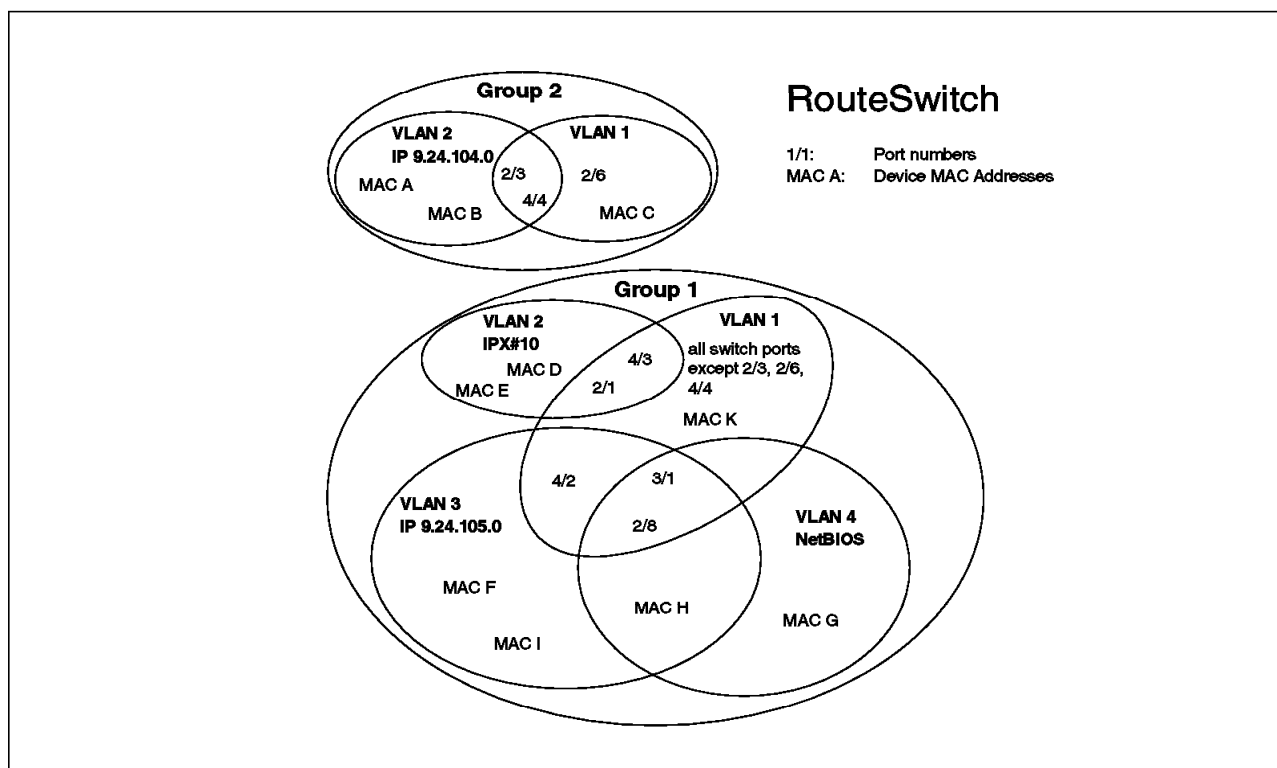


Figure 35. Logical View of RouteSwitch VLANs

Configuration of the RouteSwitch In Figure 34:

1. A second non-mobile group (group 2) was created and the 10Base-T ports 2/3 and 2/6 and the 100Base-TX port 4/4 were assigned respectively.
2. The RouteSwitch automatically created the default VLAN 1 in non-mobile group 2.
3. The network policy-based VLAN 2 was created in non-mobile group 2 and the IP subnetwork address of 9.24.104.0 with the subnetwork mask of 255.255.255.0 was created.
4. The network policy-based VLAN 2 in group 1 was created and the IPX network number 10 was assigned.
5. A second IP VLAN was created in group 1 in order to group all devices with the IP subnetwork number of 9.24.105.0 together. The same subnetwork mask as above was used.
6. Finally, the protocol-based VLAN 4 for NetBIOS traffic was created.
7. No further configuration was done.

Details on creating VLANs can be found in Chapter 4, “Basic VLAN Configuration” on page 67.

VLAN Assignment: In Figure 35 on page 56, if every station has sent at least one frame to the RouteSwitch, all devices will be known by the RouteSwitch and assigned to the appropriate VLAN.

Communication will work as follows:

- Communication between devices in different non-mobile groups is impossible because non-mobile groups physically separate the network. If an external router or the internal virtual router would be connected between these two non-mobile groups, communication would be possible.
- All ports are at least members of the default VLAN 1 and they remain there even if the port becomes a member of another VLAN.
- Initially all MAC addresses and ports belong to the particular default VLANs. As a result of the first received frame of each station, the MAC addresses were moved from VLAN 1 to the appropriate VLAN or VLANs according to the defined policies. (For example, MAC H is member of VLAN 3 and VLAN 4.) The corresponding port is added to the VLAN or VLANs.
- The AppleTalk (MAC K) and NetBIOS (MAC C) devices haven’t matched any defined policy; their ports and MAC addresses remain in the default VLAN.
- Devices G and I share a port at the RouteSwitch and use different protocols; port 3/1 becomes a member of both VLANs (VLAN 3 and VLAN 4) and of course still remains in VLAN 1. Thus, all broadcasts of both protocols will be forwarded to the token-ring hub on port 3/1. Device I is only a member of VLAN 3 and device G is only member of VLAN 4 even though they are connected to the same switch port 3/1. This shows clearly the different treatment of ports and MAC addresses.
- Station H has two network protocols active, thus the station’s MAC address is added to both VLANs. Station H is a member of two broadcast domains.
- VLAN membership is independent from the MAC layer protocol. Stations F, H and I are members of the IP network VLAN 9.24.105.0 even if they are connected at different speeds and MAC layer protocols.

- IPX RIPs and SAPs of Novell server D in VLAN 2, group 1 are only forwarded to port 2/1, thus preserving bandwidth at all other ports.
- If MAC I sends an ARP broadcast, this broadcast will be forwarded to ports that are in the same circle of MAC I, thus VLAN 3 only. In this case, the ARP will show up at ports 2/8, 3/1, and 4/2 only.
- The VLAN assignment is checked by the RouteSwitch network switch module's (NSM) hardware on which a frame was received to always ensure a proper VLAN membership. If a device is moved from one port to another or even to another switch, its MAC address remains in the same VLAN. The only thing that changes is the assigned port number.
- The actual port and MAC address assignments can be checked by the following commands:
 - vivl - Displays the VLAN membership of virtual ports
 - fwvvl - Displays the VLAN membership of MAC devices

3.5.3 VLAN Leakage for Non-Mobile Groups

Since the assignment of devices to VLANs is done dynamically through learning unknown source addresses, some frames may be forwarded under certain circumstances to ports that belong to a different VLAN. This is called VLAN leakage. VLAN leakage occurs if at least one of the MAC addresses (source or destination) within a received frame is unknown to the switch (not in CAM). The frame is then flooded within the non-mobile group according to the rules described in 3.5.2, "Frame Flooding for Non-Mobile Groups" on page 54, in order to provide connectivity with unknown devices. This may lead to a device that is only a member of a NetBIOS VLAN and directly connected to a switch port, receiving IP ARP broadcast frames.

Following are some examples that discuss the problems that may arise because of the VLAN leakage.

3.5.3.1 Multiple Novell Servers Using the Same Encapsulation Type

Assume there are two Novell servers connected to a RouteSwitch network. Both servers have different IPX network numbers configured but they are using the same encapsulation type. There are two IPX network address-based VLANs configured, one for each IPX network configured at the servers. Furthermore an IP network address-based VLAN is configured. As soon as both servers send out their first RIP frame, the switch assigns them to the appropriate VLAN. Now all subsequent RIPs and SAPs of both servers stay within their VLANs and everything is fine.

Assume IP is bound to both servers as a second protocol to do backup with ADSM. Now both servers also become members of the common IP VLAN. According to Table 5 on page 55 broadcasts from known sources are forwarded to all ports that have at least one VLAN in common; now both servers receive the RIPs and SAPs of the other server because of the common IP VLAN. Both servers will show up routing configuration errors.

In summary:

- If devices that are assigned to different VLANs join an additional common VLAN, the broadcast domain is expanded to the common VLAN.

- The very first frame of all devices is always forwarded to all ports within the non-mobile group regardless of VLAN membership.
- If an already known MAC address is timed out and removed from CAM, the next broadcast from such a device is forwarded to all ports within the non-mobile group regardless of VLAN membership.
- Unicasts are forwarded to all ports within a non-mobile group if source and destination addresses were timed out and removed from CAM.

3.5.4 Optimized Device Switching Ports

When the RouteSwitch detects a single device attached to a port, it automatically puts that port into optimized device switching mode. In this mode, the port only receives traffic specifically destined for it as well as broadcasts and multicasts. Unicast floods (unknown destinations) and BPDU frames will not be forwarded to this port, reducing extraneous traffic. The feature is especially valuable when the device is a server, in order to maximize its throughput. The RouteSwitch default setting is auto optimized for all ports. When the RouteSwitch detects a single MAC address on the port, the port is automatically placed in optimized switching mode. Should the RouteSwitch detect multiple devices attached to a port, the port is automatically placed into transparent bridging mode.

3.6 RouteSwitch Packet Processing for Mobile Groups

The following list describes the steps the MPM and NSMs use in determining the switching of frames:

1. A frame is received on a network interface.
2. Hardware looks up the source address to ensure that the source has been learned and that it is in the correct mobile group.
3. If the source address is unknown (not in CAM), and the physical port belongs to the Default Group, MPM picks up the frame.
4. The MPM examines the entire frame to see if policies can be matched and determines which mobile group the source address qualifies for membership in and adds the source address to the forwarding table. The MPM then relays information concerning mobile group membership to the Network Switching Module (NSM) via the MBUS.
5. Once the policies are matched, the port and MAC address are moved to the mobile group whose policy has been matched. The NSM now adds an entry into its CAM and modifies the filtering database with mobile group membership information.
6. MPM drops the frame after moving the port to the respective group. If there are no matching policies to the frame, then the frame is forwarded to all the ports in the default group.
7. If the port is moved to a different mobile group other than default group, then the second frame from the source is sent to all the ports in the respective mobile group.
8. Broadcasts and multicasts are always sent to the MPM for a complete analysis.
9. Once a port has been added to a mobile group it will stay in that group even if the MAC address of the device times out of the RouteTracker VLAN.

10. If the destination address is recognized, the frame will be picked up by the designated switching module.
11. This NSM will check the filtering database to determine if the destination device belongs to a mobile group that is also associated with the virtual port that received the frame from the source device. If so, the NSM will do any necessary translations that may be required and forward the frame out the correct port; otherwise the frame will be dropped. The NSM also sends a signal claiming the frame and all other NSMs abort.
12. If no module recognizes the destination address, and the source address is known, the frame is forwarded to all ports within a common mobile group with the source address.
13. If no module recognizes the destination address, the source address is unknown to be in a particular mobile group (also does not match any policies), then the frame is forwarded to all the ports in the default mobile group. RouteSwitches do not flood unrecognized frames to optimized device switching ports, further conserving the dedicated bandwidth at that port.

3.6.1 Broadcast Frames Processing for Mobile Groups

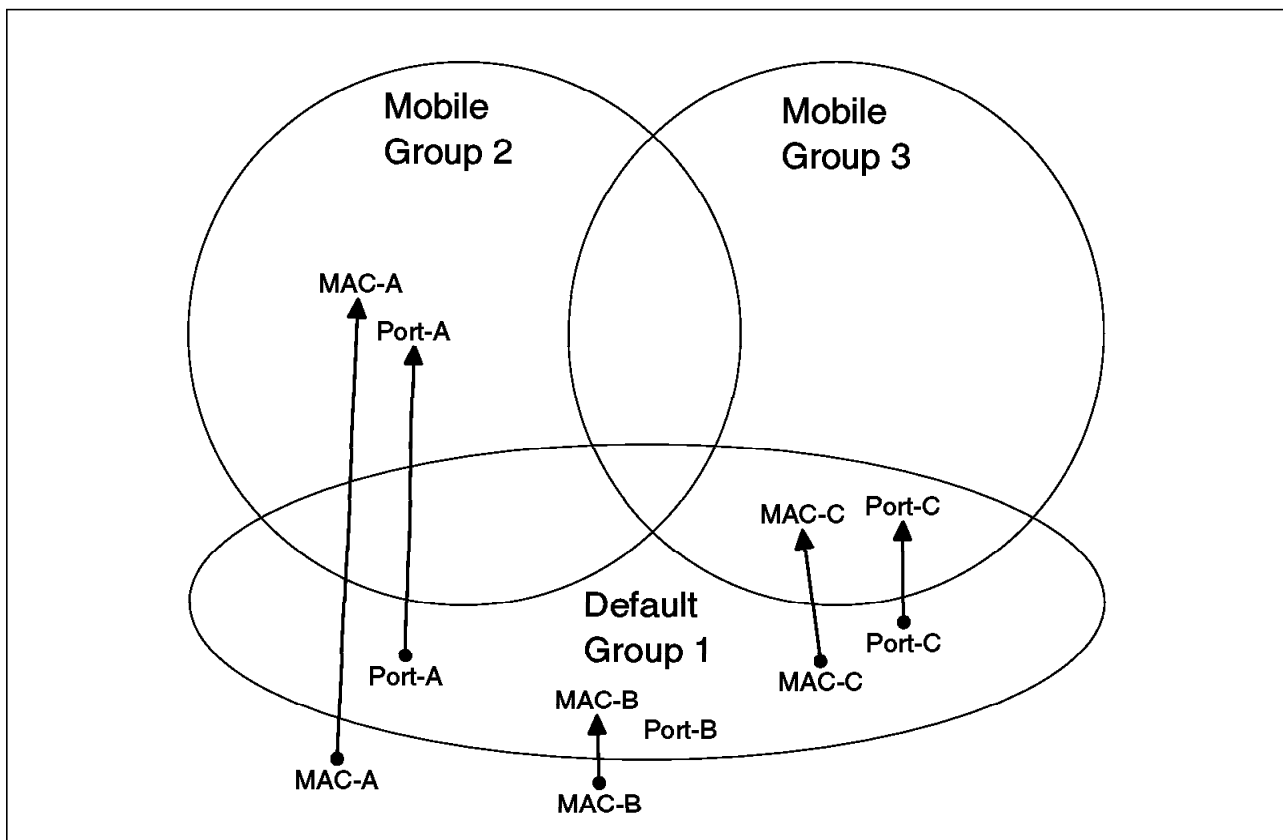


Figure 36. Broadcast from the Default Group

Figure 36 illustrates three possible scenarios for a broadcast frame.

1. MAC-A, Port-A
 - MAC-A is unknown and the broadcast frame matches the policy for group 2.
 - Port-A is in the default group.

- The first broadcast frame will be forwarded to all ports that are a member of the default group.
- The frame is analyzed by the MPM.
- Port-A will become a member of group 2.
- MAC-A will be added to the address table for group 2.
- Subsequent broadcast frames matching the policy for group 2 will be forwarded to all ports that are members of group 2.
- In this case it does not matter if the default group is enabled or disabled.

2. MAC-B, Port-B

- MAC-B is unknown and the broadcast frame does not match any policy.
- Port-B is in the default group.
- The first broadcast frame will be forwarded to all the ports that are a member of the default group.
- The frame is analyzed by the MPM.
- Port-B will remain a member of the default group.
- MAC-B will be added to the address table for the default group.
- Subsequent broadcast frames will be forwarded to all ports that are a member of the default group.
- With the default group enabled, MAC-B will be able to send unicast frames to any devices that are members of the default group.
- With the default group disabled, broadcast frames will still be forwarded to all ports in the default group. MAC-B will not be added to the default group's address table and unicast frames will not be switched to any ports.

3. MAC-C, Port-C

- MAC-C is known and the broadcast frame matches the policy for group 3. (This is possible if MAC-C has sent an earlier broadcast frame not matching any policy. See the details for MAC-B).
- Port-C is in the default group.
- The first broadcast frame will be forwarded to all the ports that are members of the default group 1.
- The frame is analyzed by the MPM.
- Port-C will remain a member of the default group and also become a member of group 3.
- MAC-C will remain in the address table of the default group and will be added to the address table of group 3.
- The second (and subsequent) broadcast frames matching the policy for group 3 will be forwarded to all ports that are members of group 3. Any other broadcast frame, not matching any policy will be forwarded to all ports that are members of the default group.

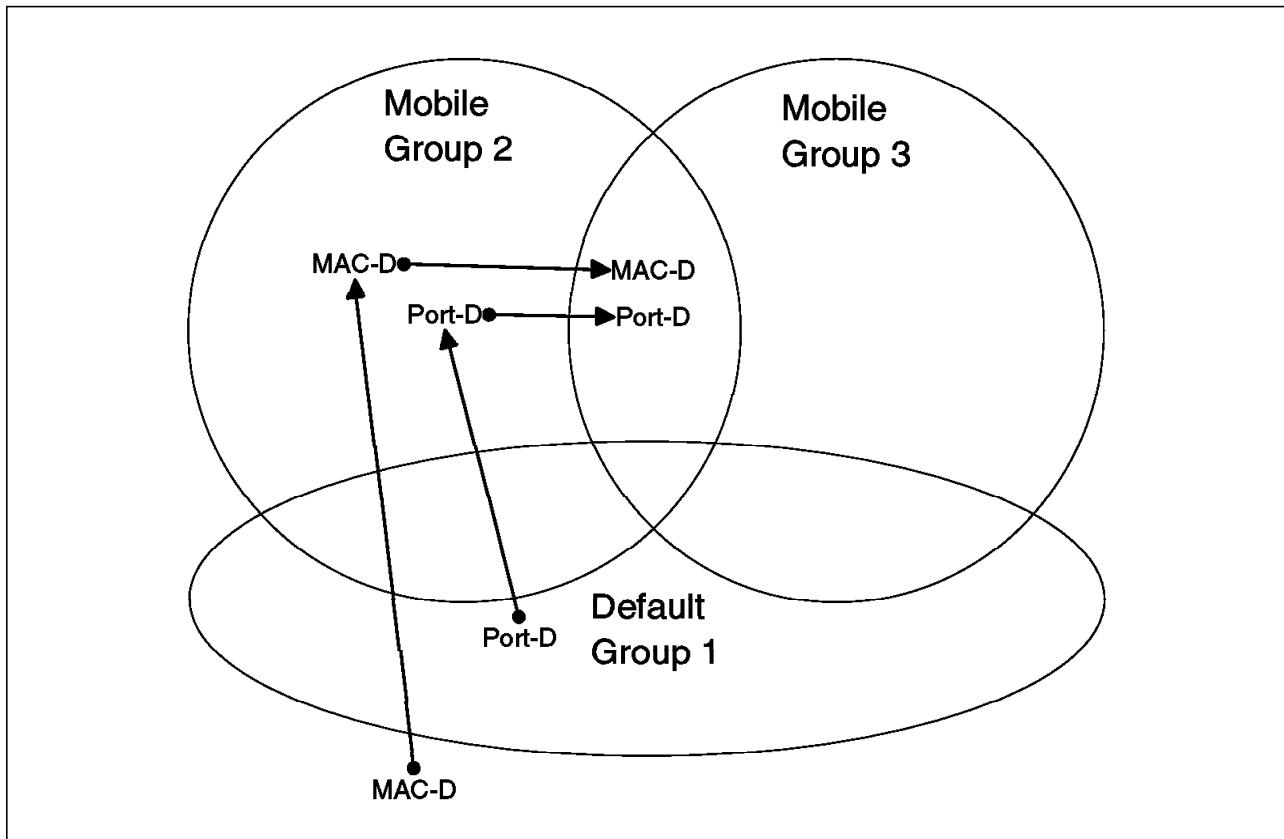


Figure 37. Broadcasts from a Default Group Matching Several Policies

- MAC-D is unknown and the first broadcast frame matches the policy for group 2.
- Port-D is in the default group.
- The first broadcast will be forwarded to all ports that are member of the default group.
- The frame is analyzed by the MPM.
- Port-D becomes a member of group 2 and MAC-D is added to the address table for group 2.
- MAC-D now sends a broadcast frame that matches the policy for group 3.
- The frame is analyzed by the MPM and dropped. Another broadcast will be needed to establish communication.
- Port-D remains a member of group 2 and also becomes a member of group 3.
- MAC-D remains in the address table for group 2 and also gets added to the address table for group 3.
- All subsequent broadcast frames matching the policy for group 2 will be forwarded to all ports that are members of group 2.
- All subsequent broadcast frames matching the policy for group 3 will be forwarded to all ports that are members of group 3.

3.6.2 Unicast Frames Processing for Mobile Groups

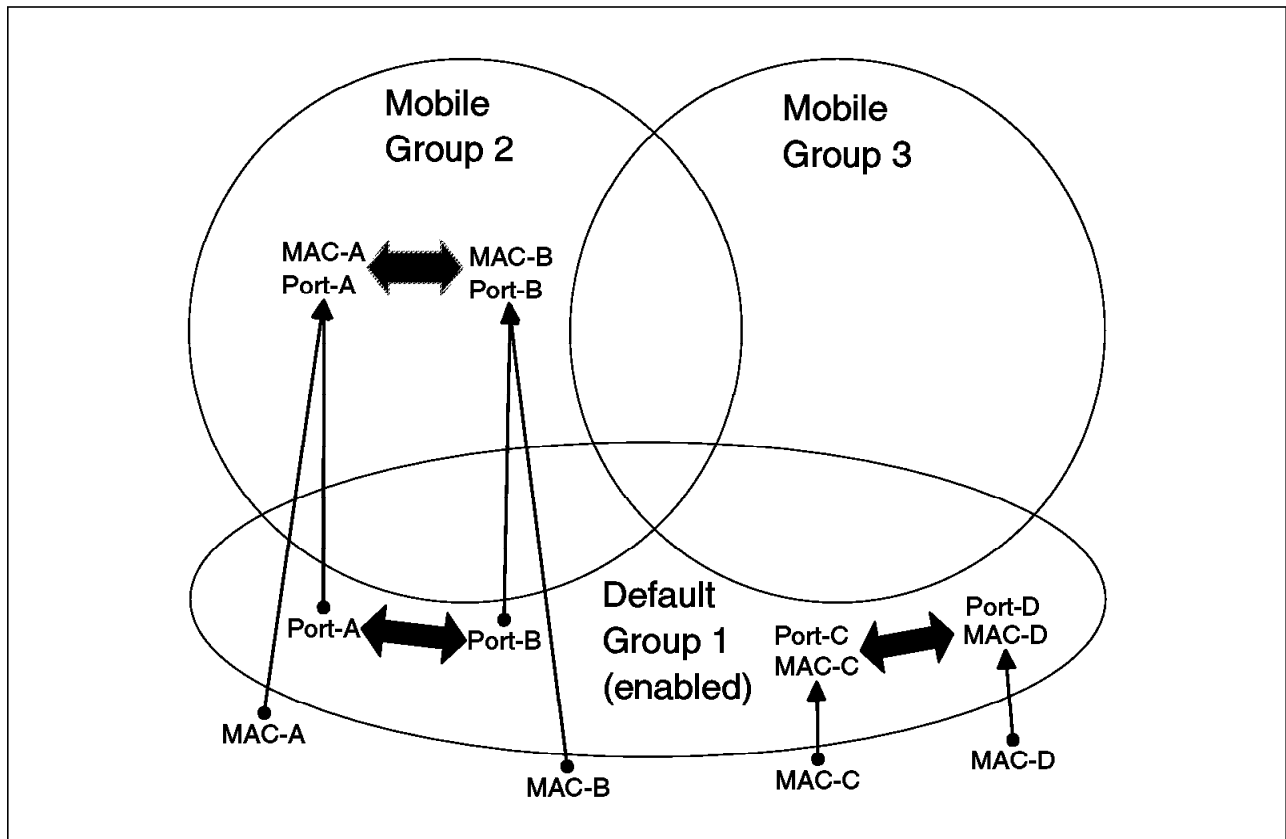


Figure 38. Unicast Processing When Source and Destination Are Unknown

Figure 38 shows two examples of unicast frame processing using mobile groups. In each case the source and destination MAC address are unknown. The default group is enabled.

1. The source unicast frame from MAC-A to MAC-B matches the policy for group 2.
 - The first unicast frame will be forwarded to all ports in the default group.
 - MAC-B will respond using the default group.
 - The MPM will analyze both the MAC-A and MAC-B frame and move both ports to group 2.
 - MAC-A and MAC-B addresses will be added to the address table for group 2.
2. The source unicast frame from MAC-C for MAC-D does not match any policy.
 - The first unicast frame will be forwarded to all ports in the default group.
 - MAC-D will respond using the default group.
 - The MPM will analyze both the MAC-C and MAC-D frame. Both Port-C and Port-D will retain membership in the default group.
 - MAC-C and MAC-D addresses will be added to the address table for the default group.

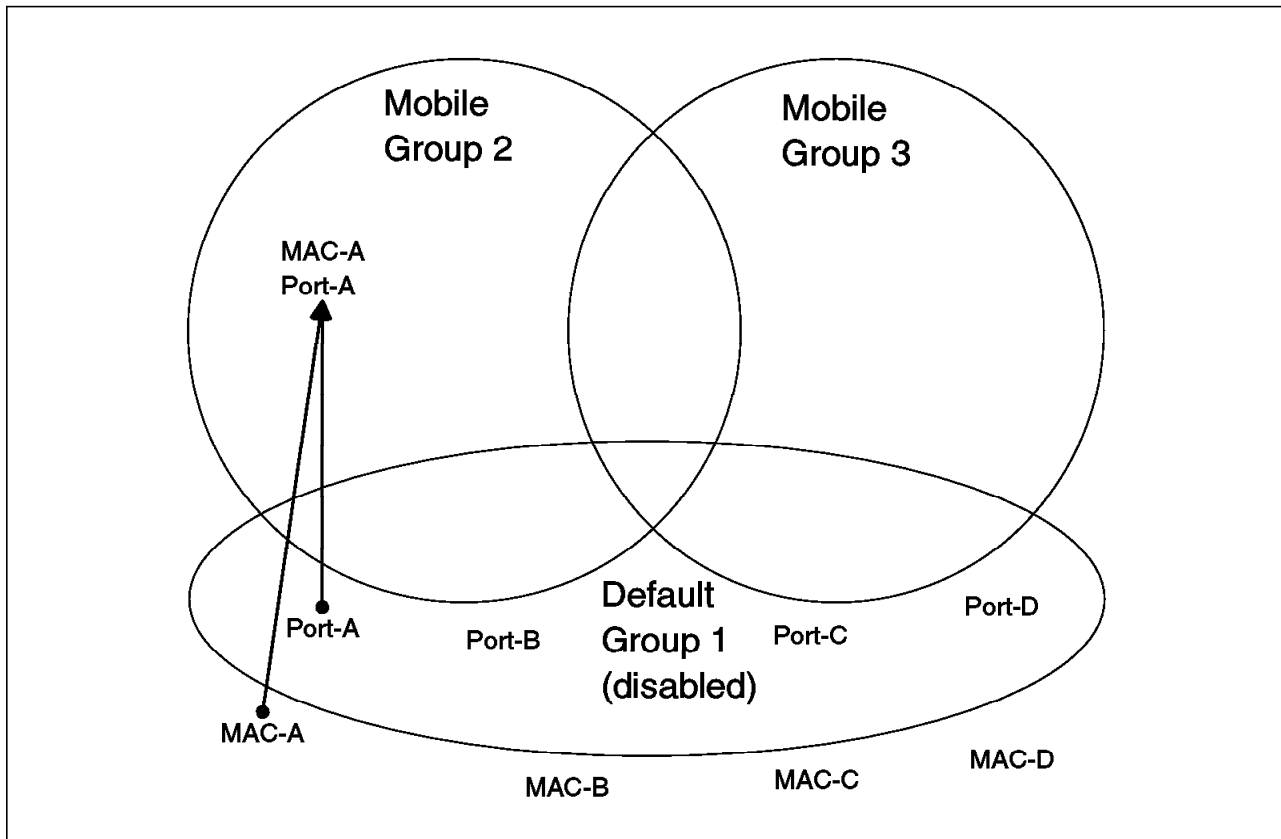


Figure 39. Unicast Processing When Source and Destination Are Unknown

Figure 39 shows two examples of unicast frame processing using mobile groups. In each case the source and destination MAC address are unknown. The default group is disabled.

1. The source unicast frame from MAC-A to MAC-B matches the policy for group 2.
 - The first unicast frame will not be forwarded to any ports that are member of group 1 (see exception note).
 - The MPM will analyze the MAC-A unicast frame.
 - Port-A will become a member of group 2.
 - MAC-A address will be added to the address table for group 2.

Exception

The first unicast frame from the first device opening a port on the RouteSwitch will be forwarded to all ports in the default group. Subsequent unicast frames from any known or unknown devices on the same port will not be forwarded to any ports in the default group.

2. The source unicast frame from MAC-C for MAC-D does not match any policy.
 - The first unicast frame will not be forwarded to any ports (see exception note).
 - The MPM will analyze both the MAC-C and MAC-D frame. Both Port-C and Port-D will retain membership in the default group.

- MAC-C address will not be added to the address table for the default group.
- MAC-C and MAC-D cannot communicate.

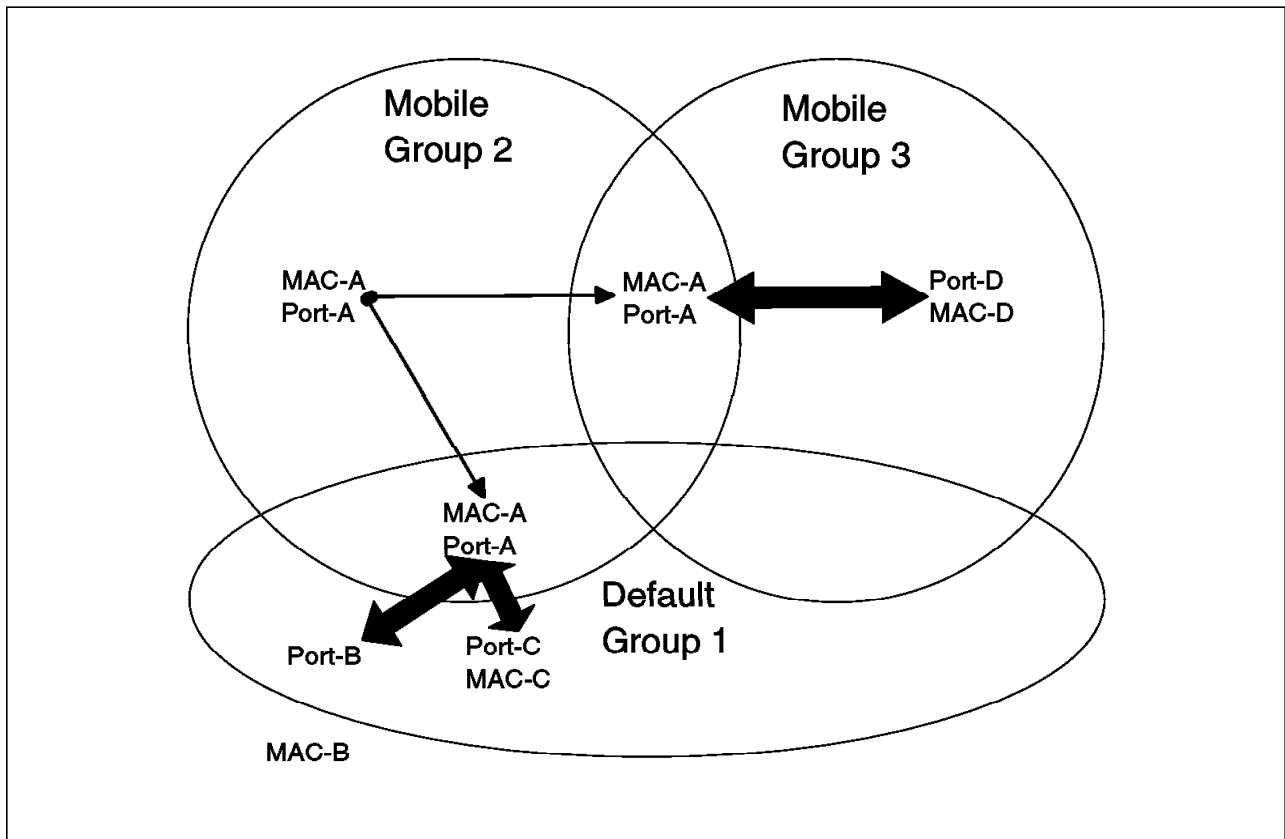


Figure 40. Unicast Frame Processing for a Known Source

Figure 40 shows examples when the source MAC address (MAC-A) is known. The default group is enabled.

1. Unicast frames from MAC-A are matching the policy for group 2.
 - Unicast frames for known destinations that are in group 2 will be switched to ports in group 2 using the address table for each port.
 - Unicast frames for the unknown address (MAC-B) or addresses in the default group (MAC-C) will not be forwarded to group 1. Connectivity between MAC-A and MAC-B or MAC-C is not possible.
2. Unicast frames from MAC-A match a policy for group 3.
 - The first unicast frame will be sent to the MPM for processing and dropped. Another unicast will be needed to establish communications.
 - Port-A will remain a member of group 2 and become a member of group 3.
 - MAC-A address will be in the address table of group 2 and 3.
 - A second unicast frame will be required to reach MAC-D.
 - Subsequent unicast frames from MAC-A will be switched to devices according to the address table for the ports that are a member of group 2 and group 3.

3. Unicast frames from MAC-A do not match any policy.

- The first unicast frame will be sent to the MPM for processing and dropped. Another unicast will be needed to establish communications.
- Port-A will retain membership in group 2 and become a member of the default group.
- MAC-A address will be added to the address table for the default group.
- For known destination, MAC-C, a second unicast frame will be switched directly to the port.
- For unknown destination, MAC-B, a second unicast frame will be forwarded to all ports in the default group.
- Connectivity between MAC-A and MAC-C or MAC-B is possible.

The following commands will help in finding out what the port's memberships are:

- `vivl` will show which group each port has joined first. There are no differences between ports assigned to non-mobile groups and ports assigned manually or by the RouteTracker to mobile groups.
- `vpl` and `vigl` will show ports assigned to a second and third group, and so on.
- `fwtvl x` (x being a group number) will show the content of the MAC addresses table for each group.

Chapter 4. Basic VLAN Configuration

This chapter is separated into two sections. The first shows the configuration of non-mobile groups and VLANs using the console configuration method. The second section shows the configuration of mobile groups, as discussed in 3.3, “RouteSwitch Terminology” on page 40.

4.1 Non-Mobile Group and VLAN Configuration

In 3.3.10, “Non-Mobile Groups” on page 47 the logical inner workings of non-mobile groups is discussed. We now go through the following setups:

- Non-mobile group configuration
- VLAN configuration in non-mobile groups

4.1.1 Non-Mobile Group Configuration Using the Console

Before a VLAN is created, it is important to know the group number it will join. If the VLAN is created for a group other than the default group (group 1), then a new group will need to be created.

Do not create an AutoTracker VLAN in group #1 if mobile groups are to be defined. Any AutoTracker VLAN defined in group #1 will be disabled when the first mobile group is defined and activated.

Figure 41 on page 68 shows the steps required to create a new group.

```

/ % crgp 1
  GROUP Number ( 2) : 2
  Description (no quotes) : Group 2 3
  Enable WAN Routing? (n): 4
  Enable ATM CIP? (n): 5
  Enable IP (y) : n 6
  Enable IPX? (y): n 7
  Enable Group Mobility on this Group ? (y/n)(n): 8
  This Group will not participate in Group Mobility

  Do you wish to configure the interface group for this Virtual LAN
    at this time? (y) y 9

  Initial Vports(Slot/Phys Intf. Range) - For example, first I/O Module
    (slot 2), second Interface would be 2/2. Specify a range of interfaces
    and/or a list as in: 2/1-3, 3/3, 3/5, 4/6-8.

    Initial Slot/Interface Assignments: 2/2-3 10
    2/2 - This interface is currently assigned to GROUP 1 -
      (Default GROUP (#1)).
      Do you wish to remove it from that GROUP and assign it (with
      new configuration values) to this GROUP
      (y|n|c to Accept defaults) (n)? y 11
    2/3 - This interface is currently assigned to GROUP 1 -
      (Default GROUP (#1)).
      Do you wish to remove it from that GROUP and assign it (with
      new configuration values) to this GROUP
      (y|n|c to Accept defaults) (n)? y

```

Figure 41. Creating a Group (1 of 2)

```

                                Modify Ether/12 Vport 2/2 Configuration 12
1) Vport                        : 5
2) Description                  :
3) Bridge Mode                  : Auto-Switched
   31) Switch Timer             : 60
4) Flood Limit                  : 192000
5) Output Format Type           : Default(IP-Eth II, IPX-802.3)
6) Ethernet 802.2 Pass Through : Yes
7) Admin, Operational Status   : Enabled, inactive
8) Mirrored Port Status        : Disabled, available

Command {Item=Value/?/Help/Quit/Redraw/Next/Previous/Save} (Previous) :

                                Modify Ether/12 Vport 2/3 Configuration 12
1) Vport                        : 5
2) Description                  :
3) Bridge Mode                  : Auto-Switched
   31) Switch Timer             : 60
4) Flood Limit                  : 192000
5) Output Format Type           : Default(IP-Eth II, IPX-802.3)
6) Ethernet 802.2 Pass Through : Yes
7) Admin, Operational Status   : Enabled, inactive
8) Mirrored Port Status        : Disabled, available

Command {Item=Value/?/Help/Quit/Redraw/Next/Previous/Save} (Previous) :

    Adding port 2/2 to GROUP 2...
    Adding port 2/3 to GROUP 2...
You may modify interfaces to this group using the addvp, modvp and rmvp
commands at a later date if you choose.
/ %
```

Figure 42. Creating a Group (2 of 2)

- 1** The command `crgp` is used to create a group.
- 2** The next available group number is assigned or a different group number can be entered.
- 3** A description for this group, up to 30 characters.
- 4** A `y` answer will enable frame relay RFC 1490 routing through this group. If only bridging or trunking using a frame relay connection is required, then frame relay RFC 1490 does not need to be enabled.
- 5** A `y` answer will allow any devices on ATM using classic IP to communicate through this group. If the answer is `y`, then only ATM CIP is supported in this group.
- 6** Enabling IP allows this group to be connected to the internal router. IP traffic will be routed to and from this group by the internal virtual router.
- 7** Enabling IPX allows this group to be connected to the internal router. IPX traffic will be routed to and from this group by the internal virtual router.
- 8** A `y` answer will enable group mobility for this group. An example of this is shown in 4.2.1, "Mobile Group Configuration Using the Console" on page 81.

9 This option allows for ports to be assigned to the newly created group immediately. If n is selected, ports can be added later using the addvp command.

10 The list of ports to be added to the newly created group. The format is slot/port interface. In this case ports 2 and 3 are added to group 2. Ports can be added later using the addvp command.

11 The RouteSwitch will ask for a confirmation to move each port to the new group.

12 The configuration of each port moved will be reviewed and can be modified if required.

The new group is now created and ports assignment can be displayed using the via command as shown in Figure 43.

GROUP Interface Attachments For All Interfaces						
GROUP: Slot/Intf	Description	Service/ Instance		Protocol	Admin Status	
=====	=====	=====	=====	=====	=====	
1.1 :*	GROUP #1.0 IP router vport	Rtr	/ 1	IP	Enabled	
1.3 :*	Net B IPX router	Rtr	/ 2	IPX	Enabled	
1.2 :*		Rtr	/ 3	IPX	Enabled	
1:2/1	ETH V2 port	Brg	/ 1	Tns	Enabled	
2:2/2	Group 2 Port 2/2	Brg	/ 1	Tns	Enabled	
2:2/3	Group 2 Port 2/3	Brg	/ 1	Tns	Enabled	
1:2/4	Virtual port (#4)	Brg	/ 1	Tns	Enabled	
1:2/5	Virtual port (#5)	Brg	/ 1	Tns	Enabled	
1:2/6	Virtual port (#6)	Brg	/ 1	Tns	Enabled	
1:2/7	Virtual port (#7)	Brg	/ 1	Tns	Enabled	
1:2/8	Virtual port (#8)	Brg	/ 1	Tns	Enabled	
1:2/9	Virtual port (#9)	Brg	/ 1	Tns	Enabled	
1:2/10	Virtual port (#10)	Brg	/ 1	Tns	Enabled	
1:2/11	Virtual port (#11)	Brg	/ 1	Tns	Enabled	
1:2/12	Virtual port (#12)	Brg	/ 1	Tns	Enabled	
1:3/1	Virtual port (#15)	Lne	/ 1	Tns	Enabled	
1:3/1	Virtual port (#14)	Lne	/ 2	Tns	Enabled	
1:4/2	Virtual port (#13)	Brg	/ 1	Tns	Enabled	
/ %						

Figure 43. Port Assignments

In Figure 43, the highlights show port 2/2 and 2/3 as part of group 2, while all other ports are part of the default group.

A group can be removed using the command rmgrp.

All ports must be moved to another group before the group can be deleted. A port must always be assigned to a defined group The command addvp is used to move ports.

4.1.2 VLAN Configuration Using the Console

The following steps are used to build VLANs with different rules (policies) using the console interface. Setting up a console interface is described in 2.3.1, “Basic Console Setup” on page 14.

4.1.2.1 Basic Switch Configuration

A list of commands to configure VLANs is available in the AutoTracker subdirectory. To get to the directory, type at on the command line.

The command `cratvl` as shown in Figure 44 is used to create VLANs.

```
/VLAN/Auto-Tracker % cratvl 1
Enter the VLAN Group id for this VLAN ( 1): 2 2
Enter the VLAN Id for this VLAN ( 2): 3 3
Enter the new VLAN's description: VLAN#2 4
Enter the Admin status for this vlan (e)nable/(d)isable (d): e 5
Select rule type: 6
1. Port Rule
2. MAC Address Rule
3. Protocol Rule
4. Network Address Rule
5. User Defined Rule
6. Binding Rule
7. DHCP PORT Rule
8. DHCP MAC Rule
Enter rule type (1):
```

Figure 44. Example of Initial VLAN Configuration Using the Console

- 1** The `cratvl` command will start the process to create a VLAN.
- 2** The group where the VLAN will be created. If a group other than the default group (group 1) is selected, the group must be defined first. VLAN numbers have to be unique within a group.
- 3** The VLAN number can be entered or the next available VLAN can be selected by pressing the Enter key.
- 4** A description of the VLAN, up to 30 characters.
- 5** This parameter enables or disables the VLAN. A VLAN can be created and disabled. Once ready to be activated, the VLAN can then be enabled.
- 6** The eight rules (policies) for VLAN creation are described starting with 4.1.2.2, “Port-Based VLAN.”

Once created, a VLAN can be modified using the `modatvl` command. When using the modified VLAN command, the group number and VLAN number must be specified. `modatvl 1:2` indicates that VLAN 2 in group 1 will be modified.

4.1.2.2 Port-Based VLAN

When rule 1 is applied, the defined port(s) will join the VLAN. The VLAN and the virtual router port will be activated. In Figure 45 on page 72, ports 5, 6, 7 and 9 of slot 2 are configured to join VLAN 2.

Important

When a port rule is defined for a VLAN, only the port joins the VLAN. The MAC addresses on that port will remain in the default VLAN and will never join a port rule only VLAN. The original port rule for the RouteSwitch was that when a port joined a VLAN, all MAC addresses seen on this port would also joined the same VLAN. The original port rule can be enabled on the RouteSwitch by entering the following line in mpm.cmd: reg_port_rule=1.

```
Select rule type:
1. Port Rule
2. MAC Address Rule
3. Protocol Rule
4. Network Address Rule
5. User Defined Rule
6. Binding Rule
7. DHCP PORT Rule
8. DHCP MAC Rule
Enter rule type (1): 1 1
Set Rule Admin Status to [(e)nable/(d)isable] (d): e 2
Enter the list of ports in Slot/Intf/Service/Instance format*colon. 2/5-7 2/9 3
Configure more rules for this vlan y/n (n): n 4
VLAN 1: 2 created successfully
Enable IP? (y): n 5
Enable IPX? (y): n
/ %
```

Figure 45. Example of Port-Based VLAN Configuration Using the Console

- 1** The port rule was selected by entering a 1.
- 2** The administration status was enabled. It is possible to define a new policy on an existing VLAN and leave the new policy disabled until such time as the policy is required.
- 3** A single port or multiple ports can be entered at the same time. For a single port the format is slot/interface. 2/5 indicates port 5 located on the module in slot 2. If consecutive ports are entered, then the format is: 2/5-7. In this case, ports 5, 6 and 7 located on the module in slot 2 will be entered in the port rule. If a service needs to be entered, then the instance also needs to be specified: 3/1 3.
- 4** It is possible to define more than one rule for a specific VLAN. The same rule can be defined more than once. In these examples, only one rule was defined at a time.
- 5** IP and IPX routing were not defined here. See 5.4, "VLAN and Group Routing" on page 125 for an explanation of the virtual router in the RouteSwitch.

4.1.2.3 MAC-Based VLAN

Rule 2 will act upon defined MAC addresses and only when a defined MAC address is seen by the MPM will the MAC address join the VLAN. The VLAN and the virtual router port will be activated when one MAC address joins the VLAN.

A list of known MAC addresses will not be displayed and thus a knowledge of the MAC addresses to be in the VLAN is required.

```

Select rule type:
1. Port Rule
2. MAC Address Rule
3. Protocol Rule
4. Network Address Rule
5. User Defined Rule
6. Binding Rule
7. DHCP PORT Rule
8. DHCP MAC Rule
Enter rule type (1): 2 1
Set Rule Admin Status to [(e)nable/(d)isable] (d): e 2
Enter the list of MAC addresses (AABBCC:DDEEFF)
in Canonical format.
(Enter save to end): 020056:5B6C01 3
08005A:BB0C59
0800F8:CCDD01
save 4
Configure more rules for this vlan [y/n] (n): n 5
VLAN 1: 2 created successfully
Enable IP? (y): n 6
Enable IPX? (y): n
/ %

```

Figure 46. Example of MAC Address-Based VLAN Configuration Using the Console

- 1** The MAC address rule was selected by entering a 2.
- 2** The administration status was enabled. It is possible to define a new policy on an existing VLAN and leave the new policy disabled until such time as the policy is required.
- 3** A single MAC address or multiple MAC addresses can be entered.
- 4** Once all MAC addresses are entered, the save command is used to save the list.
- 5** It is possible to define more than one rule for a specific VLAN. The same rule can be defined more than once. In these examples, only one rule was defined at the time.
- 6** IP and IPX routing were not defined here. See 5.4, “VLAN and Group Routing” on page 125 for an explanation of the virtual router in the RouteSwitch.

4.1.2.4 Protocol-Based VLAN

Rule 3 will act on protocol type and the MPM will assign the port and MAC address by analyzing the protocol in each frame.

```

Select rule type:
1. Port Rule
2. MAC Address Rule
3. Protocol Rule
4. Network Address Rule
5. User Defined Rule
6. Binding Rule
7. DHCP PORT Rule
8. DHCP MAC Rule
Enter rule type (1): 3 1
Set Rule Admin Status to [(e)nable/(d)isable] (d): e 2
Select Protocol:
1. IP
2. IPX
3. DECNET
4. APPLETALK
5. Protocol specified by ether-type (in hex)
6. Protocol specified by DSAP and SSAP (in hex)
7. Protocol specified by SNAP (in hex)
Enter protocol type (1): 1 3
Configure more rules for this VLAN [y/n] (n): n 4
VLAN 1: 2 created successfully
Enable IP? (y): n 5
Enable IPX? (y): n
/ %

```

Figure 47. Example of Protocol-Based VLAN Configuration Using the Console

- 1** The protocol rule was selected by entering a 3.
- 2** The administration status was enabled. It is possible to define a new policy on an existing VLAN and leave the new policy disabled until such time as the policy is required.
- 3** There are seven choices to define the protocol rule for the VLAN. Once selected, options 1, 2, 3 and 4 have no more options. Options 5, 6 and 7 have more options in order to define the protocol. The additional options for options 5, 6 and 7 are shown in the following figures:
 - 5. Protocol specified by Ether-Type (in hex); Figure 48 on page 75.
 - 6. Protocol specified by DSAP and SSAP (in hex); Figure 49 on page 75.
 - 7. Protocol specified by SNAP (in hex); Figure 50 on page 75.
- 4** It is possible to define more than one rule for a specific VLAN. The same rule can be defined more than once. In these examples, only one rule was defined at the time.
- 5** IP and IPX routing were not defined here. See 5.4, "VLAN and Group Routing" on page 125 for an explanation of the virtual router in the RouteSwitch.

```
Select Protocol:
1. IP
2. IPX
3. DECNET
4. APPLETALK
5. Protocol specified by ether-type (in hex)
6. Protocol specified by DSAP and SSAP (in hex)
7. Protocol specified by SNAP (in hex)
Enter protocol type (1): 5
Enter the Ether-type value in hex: 8137
```

Figure 48. Protocol Rule, Ether-Type

Refer to Appendix B, “Ether-Types and SAP Listings” on page 341 for a list of valid Ether-Type values.

```
Select Protocol:
1. IP
2. IPX
3. DECNET
4. APPLETALK
5. Protocol specified by ether-type (in hex)
6. Protocol specified by DSAP and SSAP (in hex)
7. Protocol specified by SNAP (in hex)
Enter protocol type (1): 6
Enter the DSAP value in hex: f0
Enter the SSAP value in hex: f0
```

Figure 49. Protocol Rule, DSAP and SSAP

Refer to the end of Appendix B, “Ether-Types and SAP Listings” on page 341 for a list of valid SAP values.

```
Select Protocol:
1. IP
2. IPX
3. DECNET
4. APPLETALK
5. Protocol specified by ether-type (in hex)
6. Protocol specified by DSAP and SSAP (in hex)
7. Protocol specified by SNAP (in hex)
Enter protocol type (1): 7
Enter the SNAP value in hex: 0000008137
```

Figure 50. Protocol Rule, SNAP

Refer to Appendix B, “Ether-Types and SAP Listings” on page 341 for a list of valid SNAP values.

4.1.2.5 Network Address-Based VLAN

A network address rule can be assigned to a VLAN. The option is only valid for IP and IPX protocols.

```

Select rule type:
1. Port Rule
2. MAC Address Rule
3. Protocol Rule
4. Network Address Rule
5. User Defined Rule
6. Binding Rule
7. DHCP PORT Rule
8. DHCP MAC Rule
Enter rule type (1): 4 1
Set Rule Admin Status to [(e)nable/(d)isable] (d): e 2
Select the Network Protocol:
1. IP
2. IPX
Enter protocol type: 1 3
Enter the IP Address: 9.24.105.0 4
Enter the IP Mask (255.0.0.0): 255.255.255.0 5
Configure more rules for this VLAN [y/n] (n): n 6
VLAN 1: 2 created successfully
Enable IP? (y): n 7
Enable IPX? (y): n
/ %

```

Figure 51. Example of Network-Based VLAN Configuration Using the Console

- 1** The network address rule was selected by entering a 4.
- 2** The administration status was enabled. It is possible to define a new policy on an existing VLAN and leave the new policy disabled until such time as the policy is required.
- 3** IP was selected by entering 1.
- 4** The IP network address is specified. In this example, any devices with a network address of 9.24.105 will join the VLAN. A single IP address can also be defined so that a unique host joins the VLAN.
- 5** The subnet mask used on this IP network.
- 6** It is possible to define more than one rule for a specific VLAN. The same rule can be defined more than once. In these examples, only one rule was defined at the time.
- 7** IP and IPX routing were not defined here. See 5.4, "VLAN and Group Routing" on page 125 for an explanation of the virtual router in the RouteSwitch. This VLAN is IP-based but cannot use the internal virtual router to communicate to other VLANs since IP routing was not enabled.

If IPX is selected as the network protocol rule for this VLAN, then a network number will have to be entered as well as an encapsulation type. An example of this is seen in Figure 52 on page 77.

```

Select the Network Protocol:
1. IP
2. IPX
Enter protocol type: 2
Enter the IPX Network Number: 0001
Select the IPX Network Encapsulation
1. Ethernet
2. IEEE 802.2
3. IEEE 802.3 SNAP
4. IPX Proprietary
Enter the IPX Network Encapsulation (1): 1 1
Configure more rules for this VLAN y/n (n): n
VLAN 1: 2 created successfully
Enable IP? (y): n
Enable IPX? (y): n
/ %

```

Figure 52. Enabling the IPX Router Arm

1 The encapsulation type used on this VLAN (see 5.7, “Mixed Media” on page 150).

The IPX client boots up with no network number because it gets it from the first server it finds. Thus, the RouteSwitch will not be able to assign the client to a VLAN based on a network number alone. When the client sends out a broadcast, it may get a response from a server on a totally different VLAN. The RouteSwitch will then assign the client to the incorrect VLAN. To insure that the IPX client joins the correct VLAN, encapsulation type is used. This enables the client to join the correct VLAN and attach to the server on that VLAN. The encapsulation method of the router port attached to this VLAN must match the encapsulation type selected during VLAN creation.

4.1.2.6 User-Defined Based VLAN

If none of the rules 1 to 4 applies, it is possible to define a rule using a specific offset in the frame.

```

Select rule type:
1. Port Rule
2. MAC Address Rule
3. Protocol Rule
4. Network Address Rule
5. User Defined Rule
6. Binding Rule
7. DHCP PORT Rule
8. DHCP MAC Rule
Enter rule type (1): 5 1
Set Rule Admin Status to [(e)nable/(d)isable] (d): e 2
Enter the Offset into the frame ( < 64 ): 20 3
Enter the value (as a hex string) of the pattern to match: ffef 4
Enter the mask (as a hex string) for the pattern to match: ffff 5
Configure more rules for this VLAN [y/n] (n): n 6

```

Figure 53. User-Defined Rule

1 The user-defined rule was selected by entering a 5.

2 The administration status was enabled. It is possible to define a new policy on an existing VLAN and leave the new policy disabled until such time as the policy is required.

3 This is the offset in the frame where the pattern to match is located. The first byte of the MAC header is byte 1 or offset 0. An offset of 20 indicates that byte 21 of the frame is where the pattern to match is located.

4 The pattern of the data to match in hex. The MPM will analyze the frame and if the pattern matches, the port and MAC address where this frame originates will join the VLAN. The pattern can be a maximum of 8 bytes.

5 The mask for the pattern to match. This allows the MPM to ignore some bits within the pattern. For example a mask of FFFF (binary 1111 1111 1111 1111) indicates that every bit of the pattern has to match for this rule. A mask of FFF7 (binary 1111 1111 1111 0111) indicates that the value of the third bit of the pattern (0 is the rightmost bit) is ignored. A mask of FFF8 (binary 1111 1111 1111 1000) indicates that the value of bit 0, 1 and 2 are ignored and the pattern has to match bit 3 to 15.

6 It is possible to define more than one rule for a specific VLAN. The same rule can be defined more than once. In these examples, only one rule was defined at the time.

4.1.2.7 Binding Rule-Based VLAN

The binding rule allows you to bind IP addresses, MAC addresses and ports together.

```
Select rule type:
1. Port Rule
2. MAC Address Rule
3. Protocol Rule
4. Network Address Rule
5. User Defined Rule
6. Binding Rule
7. DHCP PORT Rule
8. DHCP MAC Rule
Enter rule type (1): 6 1
Set Rule Admin Status to [(e)nable/(d)isable] (d): e 2
Please select one of the following bindings:
1. Bind IP Address to a Port and a MAC Address.
2. Bind MAC Address to a Protocol and a Port
Enter the type of binding (1): 1 3
Enter the port in the form of slot/interface: 2/5 4
Enter the IP Address: 9.24.105.66 5
Enter the Canonical MAC address in AABCC:DDEEFF format:
400000:aabbcc 6
```

Figure 54. Binding Rule

1 The binding rule was selected by entering a 6.

2 The administration status was enabled. It is possible to define a new policy on an existing VLAN and leave the new policy disabled until such time as the policy is required.

3 Binding an IP address to a port and a MAC address was selected by entering a 1.

4 Enter the port you want to bind the IP address to.

5 Enter the IP address. No sub-net mask is required.

6 Entering the MAC address completes the configuration for this VLAN.


```

Please select one of the following bindings:
1. Bind IP Address to a Port and a MAC Address.
2. Bind MAC Address to a Protocol and a Port
Enter the type of binding (1): 2 1
Enter the port in the form of slot/interface: 3/5 2
Enter the Canonical MAC address in AABCC:DDEEFF format: 400000:bbccdd 3
Specify the protocol as:
1. IP
2. IPX
3. DECNET
4. APPLETALK
5. Protocol specified by ether-type (in hex)
6. Protocol specified by DSAP and SSAP (in hex)
7. Protocol specified by SNAP (in hex)
Enter protocol type (1): 4 4

```

Figure 55. Binding Rule

- 1** Binding the MAC address to a protocol and port was selected by entering 2.
- 2** Enter the port that the MAC address and protocol are going to be bound to.
- 3** Enter the MAC address in canonical format.
- 4** The protocols that can be selected are the same as shown in 4.1.2.4, “Protocol-Based VLAN” on page 73.

4.1.2.8 DHCP Port-Based VLAN

The DHCP port rule allows the port to join a VLAN if a DHCP frame is seen on that port. Like the port rule, the MAC address will not join the VLAN and a second rule should be configured so that the MAC address sending the DHCP frame joins the VLAN.

```

Select rule type:
1. Port Rule
2. MAC Address Rule
3. Protocol Rule
4. Network Address Rule
5. User Defined Rule
6. Binding Rule
7. DHCP PORT Rule
8. DHCP MAC Rule
Enter rule type (1): 7 1
Set Rule Admin Status to ((e)nable/(d)isable) (d): e
Enter the list of ports in Slot/Interface format: 3/5 2

```

Figure 56. DHCP Port Rule

- 1** The DHCP port rule was selected by entering a 7.
- 2** The list of ports that are to be added to this rule are defined.

4.1.2.9 DHCP MAC Address-Based VLAN

The DHCP MAC address rule allows for a MAC address-based VLAN that will only allow DHCP clients whose MAC addresses are listed to join.

```
Select rule type:
1. Port Rule
2. MAC Address Rule
3. Protocol Rule
4. Network Address Rule
5. User Defined Rule
6. Binding Rule
7. DHCP PORT Rule
8. DHCP MAC Rule
Enter rule type (1): 8 1
Set Rule Admin Status to ((e)nable/(d)isable) (d): e
Enter the list of MAC addresses (AABBCC:DDEEFF) in Canonical format.
(Enter save to end): 400000:ccddee 2
save
```

Figure 57. DHCP MAC Address Rule

1 The DHCP MAC address rule was selected by entering a 8.

2 The list of MAC addresses is added to the VLAN definition. To end this list type save.

4.1.3 The Precedence of the Policies

When multiple policies are defined for a VLAN or mobile group the RouteSwitch will look at these policies in a specific order.

Note: It is recommended to have only one policy defined for each VLAN or mobile group.

1. MAC address policy
2. Network address policy
3. Protocol or port policy (depending on which one was defined first)
4. User-defined policy

When multiple policies are defined, the RouteSwitch will stop looking at policies once the first defined policy has been matched.

The line precedence=0 can be added to the mpm.cmd so that the RouteSwitch will now look at all of the defined policies and keep making VLAN assignment with all of the matched defined policies.

4.2 Basic Mobile Group Configuration

In 3.3.11, "Mobile Groups" on page 50 mobile groups are discussed in detail. This section shows the configuration steps required to define mobile groups and associate RouteTracker policies to them.

4.2.1 Mobile Group Configuration Using the Console

The configuration and workings of mobile groups vary from the traditional non-mobile groups.

There are two steps in creating mobile groups:

- Enable group mobility for the RouteSwitch using the `gmcfg` command
- Create each group and define the policy(ies) for each using the `crgp` command

This example shows how to configure mobile groups and helps explain the logic of how they work.

4.2.2 Configuration of the RouteSwitch

Mobile groups can be created with group mobility disabled. However the mobile groups defined will not act on any policy. Use the `gmcfg` command to enable group mobility for the RouteSwitch. Mobile groups are configured using the `crgp` command. The group mobility option also needs to be selected while creating the group. The following RouteTracker policies were defined for this example to determine membership in the mobile group.

- Disable IP from group 1. There can be no VLANs defined in group 1. If there are any VLANs defined in group 1, they will be disabled once group mobility is enabled. In this scenario IP was disabled in group 1. See the note in 2.3.2, “SNMP Setup” on page 18 when disabling IP in the default group.
- Mobile group 2 is to be configured with the NetBIOS protocol policy configured. All the network nodes running NetBIOS will become member of this group 2.
- Mobile group 3 is to be configured with the IP network 9.24.105.0 and subnet mask of 255.255.255.0 policy. All the network nodes with IP network addresses in that range will become a member group 3.

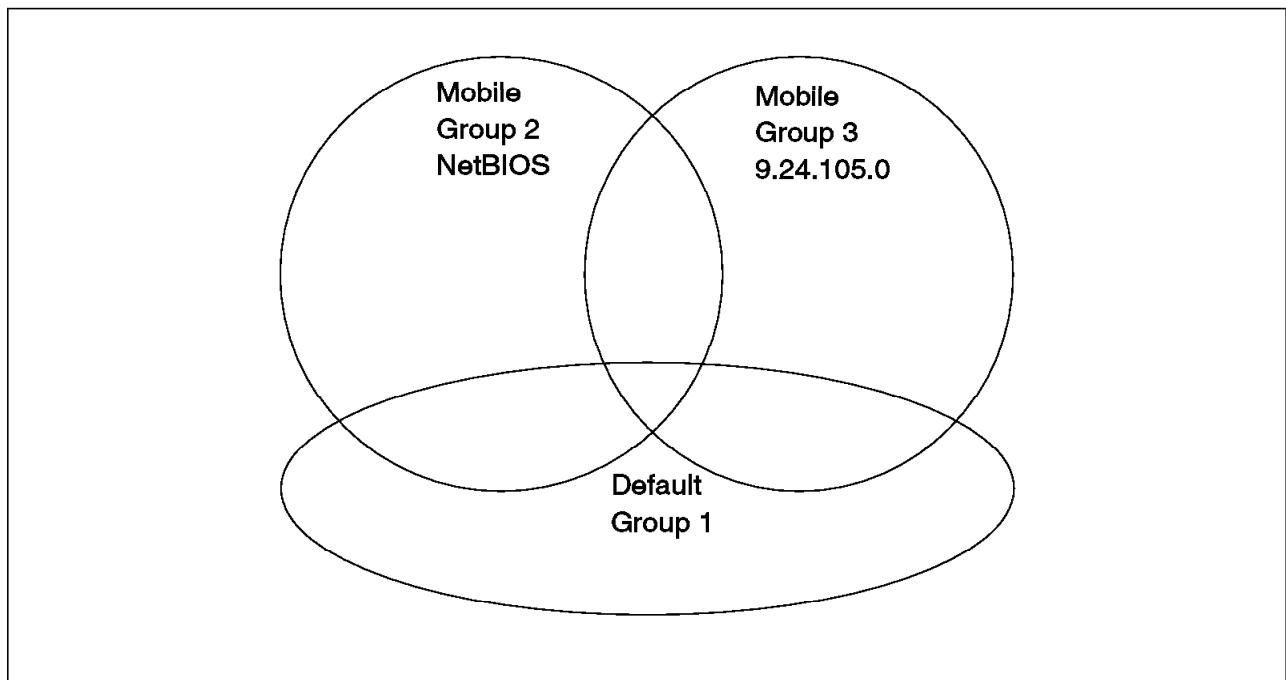


Figure 58. Logical Configuration of Mobile Groups

1. The gmcfg command is used to enable group mobility for the RouteSwitch.

```
/ % gmcfg
Group Mobility is Disabled. Enable Group Mobility ? yes/no (no): y 1
move_to_def is set to Disabled. Set to Enable ? yes/no (no): y 2
def_group is set to Enable. Set it to Disable ? yes/no (no): 3
/ %
```

Figure 59. Enabling Group Mobility

- **1** This will enable or disable group mobility for all of the groups in a RouteSwitch.
- **2** mode_to_def will decide if a port in a mobile group will be moved to the default group once the port is disabled. If move_to_def is disabled, then a port will remain in a mobile group and never rejoin the default group.
- **3** When the default group is enabled, then communication is possible within the default group. This would apply to devices that are not matching any policies. If the default group is disabled, only broadcasts and multicasts frames will be allowed in the default group.

2. Disable IP from the default group.

```
/ % modv1 1
Current values associated with GROUP 1.1 are as follows:

1) GROUP Number          - 1:1
2) Description            - Default GROUP (#1)
IP parameters:
3) IP enabled             - Y
4) IP Network Address     - 9.24.105.99
5) IP Subnet Mask         - 255.255.255.0
6) IP Broadcast Address   - 9.24.105.255
7) Router Description     - GROUP #1.0 IP router vport
8) RIP Mode               - Silent
                          {Active(a), Inactive(i), Deaf(d), Silent(s)}
9) Routing disabled       - N
10) NHRP enabled          - N
11) Default Framing       - Ethernet II
                          {Ethernet II(e), Ethernet 802.3(8), fddi(f),
                          token ring(t), source route token ring(s)}
IPX parameters:
12) IPX enabled           - N

(save/quit/cancel)
: 3=n
: ?
1) GROUP Number          - 1:1
2) Description            - Default GROUP (#1)
IP parameters:
3) IP enabled             - N
IPX parameters:
4) IPX enabled            - N

: save
```

Figure 60. Disabling IP in the Default Group

3. Create mobile group 2 and assign NetBIOS protocol policy to it. The internal virtual router port will not be configured.

```

/ % crgp
GROUP Number ( 2):
Description (no quotes) : Mobile Netbios Group 2
Enable WAN Routing? (n):
Enable ATM CIP? (n):
Enable IP (y) : n
Enable IPX? (y): n
Enable Group Mobility on this Group ? y/n(n): y
Enable User Authentication for this Group y/n(n):

Do you wish to configure the interface group for this Virtual LAN
at this time? (y) n
GROUP 2 has been added to the system.
You may add interfaces to this group using the addvp command at a later d
For now, the GROUP is inactive until you add interfaces.
Configure Auto-Activated LEC service ? y/n(y): n
Configure AutoTracker rules for this group y/n(y): y
Select rule type:
1. Port Rule
2. MAC Address Rule
3. Protocol Rule
4. Network Address Rule
5. User Defined Rule
6. Binding Rule
7. DHCP PORT Rule
8. DHCP MAC Rule
Enter rule type (1): 3
Set Rule Admin Status to (e)nable/(d)isable (d): e
Select Protocol:
1. IP
2. IPX
3. DECNET
4. APPLETALK
5. Protocol specified by ether-type (in hex)
6. Protocol specified by DSAP and SSAP (in hex)
7. Protocol specified by SNAP (in hex)
Enter protocol type (1): 6
Enter the DSAP value in hex: f0
Enter the SSAP value in hex: f0
Configure more rules for this vlan y/n (n): n
VLAN 2: 1 created successfully

```

Figure 61. Creating a NetBIOS Mobile Group

4. Create mobile group 3 and assign IP network 9.24.105.0 policy defined to it.

```

/ % crgp
  GROUP Number ( 3 ) :
  Description (no quotes) : IP Net 9.24.105.0 Group 3
  Enable WAN Routing? (n):
  Enable ATM CIP? (n):
  Enable IP (y) :
IP Address : 9.24.105.99
  IP Subnet Mask (0xff000000) : 255.255.255.0
  IP Broadcast Address (9.24.105.255 ) :
  Description (30 chars max) :
  Disable routing? (n) :
  Enable NHRP? (n) :
  IP RIP mode {Deaf(d),
    Silent(s),
    Active(a),
    Inactive(i)} (s) :
  Default framing type {Ethernet II(e),
    fddi(f),
    token ring(t),
    Ethernet 802.3 SNAP(8),
    source route token ring(s)} (e) :
  Enable IPX? (y): n
  Enable Group Mobility on this Group ? y/n(n): y
  Enable User Authentication for this Group y/n(n):

Do you wish to configure the interface group for this Virtual LAN
at this time? (y) y

Initial Vports(Slot/Phys Intf. Range) - For example, first I/O Module
(slot 2), second Interface would be 2/2. Specify a range of interfaces
and/or a list as in: 2/1-3, 3/3, 3/5, 4/6-8.

Initial Slot/Interface Assignments: 3/12
3/12 - This interface is currently assigned to GROUP 1 -
(Default GROUP (#1)).
Do you wish to remove it from that GROUP and assign it (with
new configuration values) to this GROUP y|n|c to Accept defaults (

```

Figure 62 (Part 1 of 2). Creating an IP Network Address Mobile Group

```

Modify Ether/12 Vport 3/12 Configuration

1) Vport : 12
2) Description :
3) Bridge Mode : Auto-Switched
   31) Switch Timer : 60
4) Flood Limit : 192000
5) Output Format Type : Default(IP-Eth II, IPX-802.3)
6) Ethernet 802.2 Pass Through : Yes
7) Admin, Operational Status : Enabled, active
8) Mirrored Port Status : Disabled, available

Command {Item=Value/?/Help/Quit/Redraw/Save} (Redraw) : q
Adding port 3/12 to GROUP 3...
You may modify interfaces to this group using the addvp, modvp and rmvp
commands at a later date if you choose.
Configure Auto-Activated LEC service ? y/n(y): n
Configure AutoTracker rules for this group y/n(y): y
Select rule type:
1. Port Rule
2. MAC Address Rule
3. Protocol Rule
4. Network Address Rule
5. User Defined Rule
6. Binding Rule
7. DHCP PORT Rule
8. DHCP MAC Rule
Enter rule type (1) :: 4
Set Rule Admin Status to (e)nable/(d)isable (d):: e
Select the Network Protocol::
1. IP
2. IPX
Enter protocol type:: 1
Enter the IP Address:: 9.24.105.0
Enter the IP Mask (255.0.0.0):: 255.255.255.0
Configure more rules for this vlan y/n (n):: n
VLAN 3: 1 created successfully
/ %

```

Figure 62 (Part 2 of 2). Creating an IP Network Address Mobile Group

5. The viatr1 command is used to view the VLAN configuration.

```

/ % viatr1
VLAN  VLAN  Rule Rule      Rule      Rule
Group::Id  Num  Type      Status    Definition
-----
      2      1  PROTOCOL RULE Enabled  Protocol DSAP = 0xF0, SSAP = 0xF0
      3      1  NET ADDR RULE Enabled  IP Addr = 9.24.105.0
                                   IP Mask = 255.255.255.0
/ %

```

Figure 63. Viewing the Policies for Mobile Groups

6. The gp command is used to view the groups defined in the RouteSwitch.

```

/ % gp
Group ID          Group Description          Network Address  Proto/
( :VLAN ID)                               (IP Subnet Mask) Encaps
=====
1 Default GROUP (#1)
2 Mobile Netbios Group 2          192.168.3.1      IP /
                                   (ff.ff.ff.00 )  ETH2
3 IP Net 9.24.105.0 Group 3      9.24.105.99     IP /
                                   (ff.ff.ff.00 )  ETH2
/ %

```

Figure 64. Viewing Mobile Group Policies

- The vivl command is used to view the port allocation to the first mobile group each port joins.

```

/ % vivl
          Virtual Interface VLAN Membership
Slot/Intf/Service/Instance  Group  Member of VLAN#
-----
2 /1 /Rtr /1          3      1
2 /1 /Rtr /2          2      1
3 /1 /Brg /1          1      1
3 /2 /Brg /1          1      1
3 /3 /Brg /1          1      1
3 /4 /Brg /1          1      1
3 /5 /Brg /1          1      1
3 /6 /Brg /1          1      1
3 /7 /Brg /1          1      1
3 /8 /Brg /1          1      1
3 /9 /Brg /1          1      1
3 /10 /Brg /1         2      1
3 /11 /Brg /1         1      1
3 /12 /Brg /1         3      1
4 /1 /Lne /1          1      1
4 /2 /Lne /1          1      1
5 /1 /Brg /1          1      1
5 /2 /Brg /1          1      1
5 /3 /Brg /1          1      1
5 /5 /Brg /1          1      1
5 /6 /Brg /1          1      1
/ %

```

Figure 65. Viewing Mobile Group Allocations

- The vpl command is used to view additional port allocation to mobile groups.

```

/ % vpl
=====
Group ID          Physical Port          Virtual Port
=====
Group ID: 2  NULL Port List
Group ID: 3  NULL Port List
Group ID: 1  NULL Port List
/ %

```

Figure 66. Viewing Additions Mobile Group Allocation

4.2.2.1 Group Membership for the Above Configuration

Following are the different scenarios in which a network node becomes a member of the different mobile groups created in the RouteSwitch. A network node can be a member of multiple mobile groups at a time. In these examples only NetBIOS, TCP/IP and SNA protocols are considered. NetBIOS and TCP/IP protocols are selected specifically because these protocol policies are defined for the groups in the above example, whereas SNA is considered as a representation of any protocol which is not defined for the group membership policy. Hence, if the node uses any protocol other than or in addition to the NetBIOS and/or TCP/IP protocol, then its group membership will be similar to the network node running SNA protocol in this example

Figure 67 shows the logical configuration of the mobile groups.

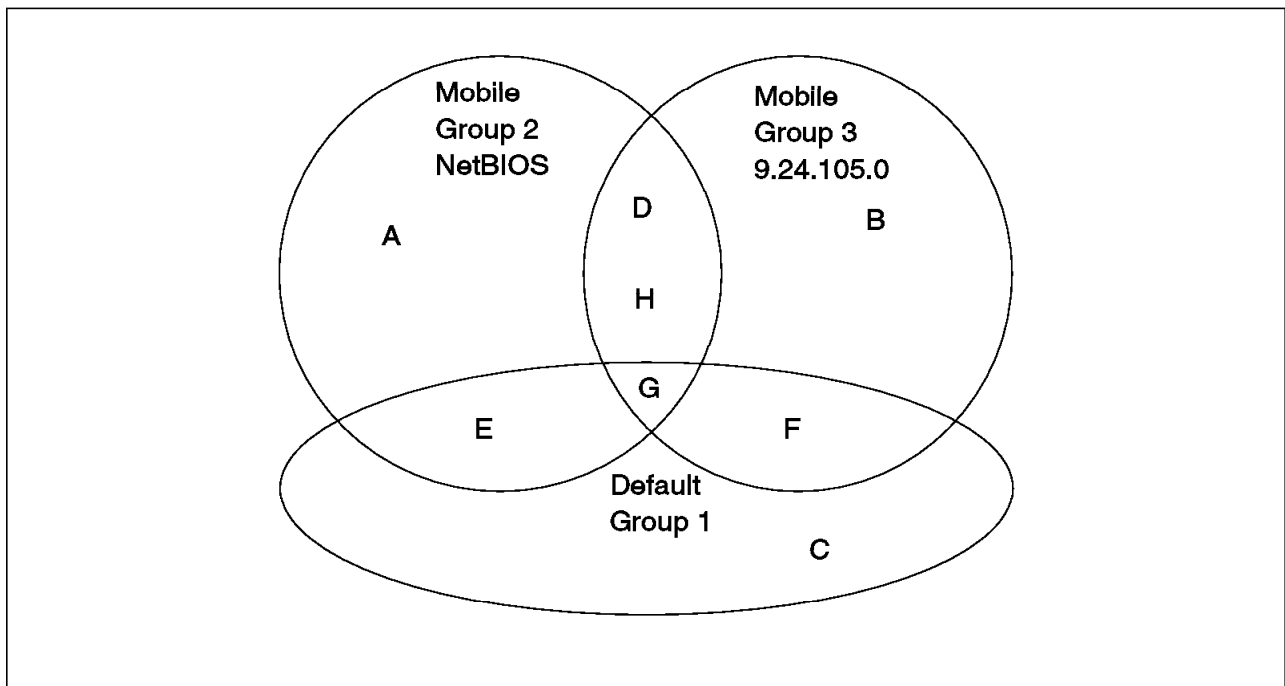


Figure 67. Ports in Mobile Groups

1. Network node **A** becomes a member of mobile group 2 as it is running only NetBIOS protocol on it.
2. Network node **B** becomes a member of mobile group 3 as it is running only TCP/IP and the IP address of the node is 9.24.105.5.
3. Network node **C** remains a member of default group, that is, group 1, as it is neither running NetBIOS protocol nor is it in the TCP/IP network 9.24.105.0. It is running only SNA protocol for which there is no policy defined to any of the mobile groups.
4. Network node **D** is a member of mobile groups 2 and 3 as it is running NetBIOS protocol as well as its TCP/IP address is 9.24.105.6.
5. Network node **E** becomes a member of group 2 as well as remaining a member of the default group, as it is running the NetBIOS and SNA protocols.
6. Network node **F** becomes a member of group 3 as well as remaining a member of the default group, as its IP address is 9.24.105.7 and it is running the SNA protocol.

7. Network node **G** becomes a member of group 2 and 3, at the same time remaining a member of the default group, as it is running NetBIOS, SNA and TCP/IP with IP address 9.24.105.8.
8. Network node **H** has been placed in group 2 using the addvp command. Network node **H** then issued a TCP/IP frame that matched the policy for group 3.

Chapter 5. Bridging, Routing, Switching and Trunking

A RouteSwitch can be compared to a transparent bridge, with the added functionality of IP and IPX routing.

In Chapter 3, "VLAN and Mobile Groups Concepts" on page 37, the concepts of VLANs and mobile groups were explained, but even in dividing our networks for any reason, there always arises the need for interconnectivity. It is important then to fully understand the concepts and workings of bridging and routing, and thus discuss these topics first as well as the topic of the spanning tree.

In the first parts of this chapter, the RouteSwitch is viewed as another bridge in the network. The concepts of VLAN bridging and VLAN switching are explained later in this chapter.

5.1 Bridging Methods

A bridge makes its primary filtering/forwarding decision based on the contents of the media-specific header of a LAN frame, that is, the MAC header.

The following types of bridges exist today:

Transparent Bridges

These bridges forward frames based on the destination MAC address of a LAN frame. The bridge builds a table for each active port based on the source MAC addresses seen. If the destination address of a frame is known to be on the same LAN as the source address, then no forwarding will take place. If the destination address is on a different port than the source MAC address, then the bridge will forward the frame. If the destination MAC address is unknown, the bridge will forward the frame to all active ports. The IEEE standard 802.1d defines the operation of transparent bridges. These bridges are predominantly used with Ethernet LANs. However, they may be used with other LAN types such as token-ring. A switch is basically a multiport transparent bridge.

Source Routing Bridges

These bridges are used to provide interconnection between token-ring and FDDI LANs. They forward frames based on a routing information field (RIF), which is part of the MAC frame header. The RIF defines a route that a frame will take to get from its source to its destination. The bridge will forward the frame only if the RIF contains the correct ring numbers and bridge number. The RIF field also contains the allowable maximum frame size to cross all of the bridges between the source and the destination.

Translational Bridges

These bridges connect LAN segments of different MAC types. Their operation has facets usually associated with routers; their operation is more protocol-specific than the previous bridge types. The translational bridge will need to convert frames from one type to the other type. In some cases, such as with the RIF field, translational bridges will have to store the RIF on behalf of the devices.

Source Route Transparent Bridges

These bridges combine both a source route and a transparent bridge. The decision to forward a frame is based on the presence of a RIF (source route) or absence of a RIF (transparent).

5.1.1 Transparent Bridging

As mentioned earlier, the transparent bridging method is primarily used to connect Ethernet LANs; however, it is used to interconnect other LANs such as token-ring. The operation of a transparent bridge is totally *transparent* to the endstations. The endstations are not aware of the presence of the bridge and view all the interconnected LANs as a single LAN.

There are two key points about transparent bridges:

1. They need a filtering database to decide whether to forward or discard a frame.
2. They allow only one active path between a pair of interconnected LANs. This is necessary to prevent frames from looping in the network and is implemented by the so-called spanning tree algorithm defined in the IEEE 802.1d standard.

Once operational, a transparent bridge builds a filtering database by listening to frames exchanged on the LAN and learning the MAC addresses of the stations attached to any LAN. It does this by recording the source addresses of the frames seen on the LANs connected to each of its ports that is in forwarding state. This results in the creation of the dynamic part of the filtering database. An aging mechanism ensures the removal of addresses that are not seen for a predetermined period of time. The timeout period is determined by the aging time parameter, which is a user-definable option.

In general, a transparent bridge applies the following rules to determine if a frame should be forwarded or discarded:

- If the destination address (DA) is associated with the receiving port, then the frame is not forwarded.
- If the DA is associated with a specific port that is in the forwarding state, the frame is forwarded.
- If the DA is not associated with a specific port, the frame is forwarded on all ports of the bridge that are in forwarding state.

Note: User-definable filters can be used to affect the forwarding/discarding of frames.

5.1.1.1 Transparent Bridging on Token-Ring

There are two issues when using transparent bridging on a token-ring: the maximum frame size that can be sent between source and destination and the address recognize indicator/frame copied indicator (ARI/FCI) issues. These two issues are now looked at in detail.

Maximum Frame Size: Unlike source route bridging, there is no process in transparent bridging to limit the size of a frame through the network. The maximum frame size that the 8274 can forward is 4472 bytes. Any frames larger than 4472 bytes will be discarded by the 8274. Each endstation must be manually configured to limit the maximum size of the frame that can be

transmitted. Endstations configured incorrectly will work fine until a request is made to transfer a large file.

Address Recognize Indicator/Frame Copied Indicator (ARI/FCI): The ARI/FCI indicators are located at the end of every token-ring frame. They are used to indicate to the source if the destination has recognized (ARI=1) the destination address in the frame and if the destination has copied (FCI=1) the frame correctly.

In 1989 the IEEE 802.5 standards committee set a recommendation on the use of the Address Recognized Indicator/Frame Copied Indicator. The new recommendation is that the setting of the bits should not be used to confirm the receipt or delivery of frames. The 8274, as other IBM switches, was developed with the new recommendation. Issues can arise in a network with devices drivers that were developed to the old recommendation, or with devices/drivers developed since 1989 that do not adhere to the current recommendations of the IEEE 802.5 standards committee.

The 8274 will set the ARI/FCI bits to 1 for every LLC frame that passes by the 8274 (default). The console command `tpcfg` allows the behavior of any of the 8274 token-ring ports to be altered so that the 8274 will no longer modify the ARI/FCI bits.

```
/ % tpcfg 2/1

ACTIVE STATUS: port mode = Station, ring speed = 16 Mbps

New Ring Speed (16 or 4 Mbps) (current configured speed is 16) :
Active Monitor Participation (currently configured No) :
Frame-Copied Bit Always Set (currently configured yes): n
The new value is saved into configuration and will be
activated on next Reset command or after reboot.
New Frame-Copied Mode is saved in configuration.
Immediate Command [Open, Reset, Close, <ret>No action]: reset
Reset the token-ring port 2/1 may cause disruption to the ring.
Are you sure you want to do this ? y/n [n]: y

Enter Slot/Interface [<ret> to exit]:

/ %
```

Figure 68. Setting of the Frame-Copied Bit

Note: Duplicate address tests sent from a device inserting to a token-ring segment rely on the ARI/FCI bits to insure that there is not a duplicate MAC address on the segment. The 8274 RouteSwitch does not alter the ARI/FCI bits of MAC frames so the duplicate address test is not affected.

Two issues to be discussed here are soft errors (frame-copied errors) and failure to establish connections along with the loss of connections.

1. The first of the two issues occurs when the 8274 is left at the default value (setting of the ARI/FCI bits to 1). All frames that pass by the 8274 before getting to the destination will have the ARI/FCI bits set to 1 by the 8274. This includes frames that do not have to pass through the 8274 to get to the destination address.




Figure 69. Address Recognized/Frame Copied Indicators

Figure 69 illustrates a ring with two endstations, S and D. The 8274 port is located between S and D. Endstation S sends a frame to endstation D. As the frame passes by the 8274, the 8274 sets the ARI/FCI bits to 1. D recognizes its MAC address in the frame and copies the frame; when D reaches the end of the frame it notices that the ARI/FCI bits are set to 1. This indicates to D that some other endstation has recognized D's MAC address (ARI=1) and copied the frame (FCI=1). D will then start a 2 second timer. During that time it will count how many more frames it receives with the ARI/FCI set to 1. Any other soft errors will also be included in their respective counters. When the 2 second timer expires, D will send a frame to the Ring Error Monitor (REM) indicating how many frame-copied error(s) and other soft error(s) it has accumulated. The process will start again when D detects another soft error. There is no problem when D sends a frame to S since endstation S will set the ARI/FCI to 1. In this case, the 8274 will not be able to alter the setting of the ARI/FCI bits when the frame goes by its port.

If the 8274 is configured to *not* modify the ARI/FCI bits, then no frame copied errors will be seen by the endstations.

2. The second issue is connection loss and/or the ability to establish a connection. In most cases, this can be attributed to the device and/or driver needing the ARI/FCI bits to be set to 1 (pre-1989 standard). This time, the source and destination are on two different ports of the 8274, and the 8274 has been configured to not modify the ARI/FCI bits.




Figure 70. Loss of Connection

Endstation S in Figure 70 sends a frame and the 8274 copies the frame to the appropriate port where endstation D is located. The ARI/FCI bits are left to 0. S removes the frame and notices that the destination has not recognized its address (ARI=0) and that the frame was not copied (FCI=0). The source will normally keep resending the frame. In the mean time D receives the frame sent by S and replies. S receives the reply frame from D and may ignore it or proceed to the next frame depending on the drivers/software installed.

If the 8274 is left at the default value, setting the ARI/FCI bits to 1, then the transactions between S and D will proceed normally.

Refer to Appendix A, “Correcting ARI/FCI Issues” on page 335 on the procedures to get drivers that are compliant to the new IEEE’s recommendations.

5.1.1.2 Transparent Bridge Port States

The ports on a transparent bridge can be in one of the following five states:

Disabled	Not participating in spanning tree protocol, not learning, not forwarding frames.
Blocking	Participating in spanning tree protocol, not learning, not forwarding frames.
Listening	Participating in spanning tree protocol, not learning, not forwarding frames, in transition from blocking to learning. The purpose of this state is to prevent the bridge from building its filtering database, based on incorrect station information, before the spanning tree becomes stable.
Learning	Participating in spanning tree protocol, learning, not forwarding frames. The purpose of this stage is to minimize the unnecessary forwarding of frames by ensuring that the bridge has built up its filtering database before beginning to forward frames.
Forwarding	Participating in spanning tree protocol, learning and forwarding frames.

Note

All of the above states, except the disabled state, are determined by the spanning tree protocol.

At V3.2.x and later, the RouteSwitch will open its ports in optimized mode. If more than one MAC address is seen on a port, then the RouteSwitch will go through the spanning tree algorithm.

When a port joins a mobile group, the port will participate in the spanning tree of the first group that the port joins. The `vivl` command will show which group the port join first.

5.1.2 Source Route Bridging

Source route bridging is the scheme used by IBM token-ring LANs to control the route a frame will traverse in a multisegment LAN. Source route bridges are the result of placing the responsibility for *navigating* through a multisegment LAN on the endstations. This is in contrast to transparent bridging in which the endstations have no knowledge of the route a frame will travel to reach its destination.

In a source route bridging environment, the route through the network is described by the sequence of *rings* and *bridges* that the frame should traverse to reach its destination. This information is stored in the *routing information field* (RIF) of a token-ring frame.

RIF is an optional part of the MAC header of the token-ring frame. The presence of this optional field in the data frame is indicated by the *routing information indicator* (RII) bit, which is the high-order bit (individual/group bit) of the source MAC address. If set to 1, it signals the existence of the RIF in the frame.

If present, the RIF contains at least a two-byte routing control field and optionally may contain up to a maximum of eight (or 14) route designator fields. IEEE's new recommendation has expanded the route designator to 14. It is not recommended to exceed the seven hops limit since not all token-ring cards will be able to understand more than eight route designators.

Each route designator field is two bytes in length and consists of a *ring number* (12 bits) and a *bridge number* (4 bits).

The routing control field specifies if the frame is non-broadcast, single-route broadcast or all-routes broadcast. It also specifies the length of the RIF and the largest frame size that can be sent over this path. In addition, the *direction bit* of this field indicates whether the route designators are to be interpreted from left to right or from right to left.

A source route bridge definition will contain a ring number associated with one side of the bridge, a bridge number and a second ring number associated with the ring on the other side of the bridge.

The route designators map out the route through a multisegment LAN. The bridges will examine the RIF and forward the frame if they recognize the two ring numbers and their bridge number.




Figure 71. Source Route Bridge Network

In Figure 71, Bridge 5 will only forward a frame if it sees 001-5-002 in the RIF. Bridge A will have to see 002-A-003 in order to forward a frame.

The broadcast indicators in the routing control field also control the way a bridge treats the frame. The types of source routed frames are:

Non-Broadcast Also known as *routed* frames. The frame will travel a specific route as defined in the RIF.

The term *specifically routed frame* (SRF) is used by the IEEE to describe these frames.

All-Routes Broadcast Also known as *general broadcast* frames. The frame will be forwarded across the bridge provided certain conditions are met. These conditions are described later in this chapter.

In IEEE terminology, these frames are known as *all routes explorers* (ARE).

Single-Route Broadcast Also known as *limited broadcast* frames. The frame will be forwarded across a spanning tree, reaching every LAN segment once.

In IEEE terminology, these frames are known as *spanning tree explorers* (STE).

Typically, all-routes broadcast and single-route broadcast frames are used to discover a route during session setup. Once the route is established, non-broadcast frames are generally used.

5.1.3 Translational Bridges

Translational bridges interconnect two different LAN types using bridging techniques common to each network bridged. In an OSI standard-based world this would not cause much difficulty, but unfortunately:

- Not all protocols use IEEE 802.2 LLC headers. One prime example is Ethernet, which often uses Ethernet V2 frame formats rather than IEEE frame formats.

- Some protocols use layer 2 addresses within some of their data frames.
- Some LAN protocols use source routing fields in their MAC headers (token-ring) and some do not (Ethernet).

Generally, the preferred method of connecting dissimilar LAN types is with routers. However, in some specific combinations, MAC layer bridges can be used. The IBM 8209/8229 is a good example of a MAC layer bridge that can connect two dissimilar LAN types. These products interconnect source route token-rings to transparent bridged Ethernet networks for a common set of protocols. These bridges are known as source route to transparent bridges (SR-TB). One side of the bridge (token-ring) will act as a source route bridge while the other side (Ethernet) will act as a transparent bridge.

5.1.4 Source Route Transparent Bridges

An SRT bridge is a bridge that combines both a source route bridge and a transparent bridge. The presence or absence of an RII/RIF will trigger the bridge to function like a source route bridge or a transparent bridge.

1. If a frame arrives at an SRT bridge port without an RII/RIF in the MAC header, then the SRT bridge will function like a transparent bridge as described in 5.1.1, "Transparent Bridging" on page 90.
2. If a frame arrives at an SRT bridge port with an RII/RIF in the MAC header, then the SRT bridge will function like a source route bridge as described in 5.1.2, "Source Route Bridging" on page 94.

5.2 The Spanning Tree

The spanning tree enables transparent bridges to dynamically discover a loop-free network tree and provide a single physical path between any two stations attached to the network. The RouteSwitch has one spanning tree per group. The spanning tree can be enabled, disabled or change from the 802.1d to the IBM spanning tree for each group defined in the RouteSwitch.

Transparent bridges and source route bridges use a different spanning tree algorithm:

- Source route bridges use the IBM spanning tree algorithm.
- Transparent bridges use the 802.1d spanning tree algorithm.

5.2.1 The Spanning Tree in Source Route Bridges

IBM source route bridges use the spanning tree algorithm to determine the route that will be taken by the *single-route broadcast* frames through a multisegment token-ring LAN.

The spanning tree algorithm used in source route bridges is identical to the spanning tree algorithm used in transparent bridging with the following exceptions:

1. The Hello BPDU is sent to the bridge functional address X'C0000000100'.
2. The port ID for a source route bridge consists of a *ring identifier* and *bridge number* while the port ID for a transparent bridge consists of a *port priority* and *port number*.

3. The spanning tree in source route bridges is used only by single-route broadcast frames. This means that bridges that are in blocking state will only block single-route broadcast frames. They will forward any all-routes broadcast frames, as well as frames that carry the appropriate routing information.

This means that unlike transparent bridges, source route bridges support active parallel paths that can be used for load-balancing across bridges as well as providing a backup path in case of bridge failures.

4. As there is no learning process in source route bridges, they can be in one of three states:
 - Blocking
 - Listening
 - Forwarding
5. Source route bridges do not support the topology change notification protocol, as it is needed only to update the transparent bridge filtering database.

5.2.2 The Spanning Tree in Transparent Bridging

The spanning tree algorithm enables transparent bridges to dynamically discover a loop-free network (*tree*) and provide a single physical path between any two stations attached to the network (*spanning*).

The IEEE 802.1d standard defines how the spanning tree will work in a transparent bridging environment.

If there were loops in the network topology, there would be network configurations where:

- Frames endlessly circulate within the network.
- Communication is prevented because a bridge may incorrectly assert that a source and destination address are on the same side of the bridge.

If there is more than one bridge between any two interconnected LANs, the spanning tree algorithm will ensure that only one of them will be included in the spanning tree. This bridge will be in a *forwarding* state and is the only one that will be passing frames between the two interconnected LANs. All the other bridges that are parallel to the forwarding bridge will be excluded from the spanning tree and will not perform any frame forwarding. These bridges are said to be in *blocking* state.

Note: In the case of multiport bridges, a single bridge may be in forwarding state on some of its ports while it is in blocking state on the others.




Figure 72. A Network with Spanning Tree

Note: In the following explanation the words *bridge* and *switch* are interchangeable.

The switches in Figure 72 have many connections to various segments. A switch works similar to a multiport bridge. The dark line represents physical connections that are forwarding and the broken line represents physical connections that have been blocked by the spanning tree to insure that there are no loops in the network.

The spanning tree algorithm will always insure that there is a path to the root switch. In this example switch 1 became the root switch.

Many configurable parameters will affect which ports will be forwarding and which ports will be blocking. See 5.2.5, "Setting the Parameters That Control the Spanning Tree" on page 103 for the various spanning tree parameters.

5.2.3 The 802.1d Spanning Tree Protocol

To participate in the spanning tree protocol, each bridge will initially assume it is the root bridge and will transmit a Hello BPDU on each of its ports.

See 5.2.3.1, "Hello BPDU" on page 100 for details about the contents of the Hello BPDU. This message will be sent every *Hello time*. Hello time is one of the spanning tree configuration parameters that can be specified for each bridge during the bridge configuration. This Hello BPDU will have the following characteristics:

1. The source address will be the address of the transmitting bridge.
2. The destination address will be X'800143000000'.
3. The source and destination SAPs will be X'42'.
4. The Root ID field will contain the ID of the transmitting bridge.
5. The Bridge ID field will contain the ID of the transmitting bridge.

6. The Path Cost field will contain 0.
7. It will be sent out by the bridge on all its ports.

Each Hello BPDU sent out on a bridge port will be received by all the other bridges that are connected to the LAN attached to that port.

BPDU's are not broadcast frames. They are sent to bridges, and selectively updated and forwarded to other bridges.

Each bridge uses the information received in the Hello BPDU's to determine the root bridge, the designated bridges, and the designated ports within each designated bridge. To do this, each bridge will continue transmitting a Hello BPDU on each of its ports until it receives a *better* Hello BPDU than the one it is transmitting on that port.

The better Hello BPDU will be determined based on the following information contained in the Hello BPDU (listed in order of their significance):

1. The lowest root ID
2. The lowest path cost to the bridge
3. The lowest transmitting bridge ID
4. The lowest port ID

As soon as a bridge receives such a Hello BPDU on a port, it will stop transmitting any further Hello BPDU's on that port and will use the information received in the better Hello BPDU to transmit a new Hello BPDU on all its other ports. The new Hello BPDU will have the following characteristics:

1. The source address will be the address of this bridge.
2. The destination address will be X'800143000000'.
3. The source and destination SAP will be X'42'.
4. The Root ID field will contain the root ID received in the better Hello BPDU.
5. The Bridge ID field will contain the ID of this bridge.
6. The Path Cost field will be the sum of the path cost received in the better Hello BPDU plus the path cost defined for the bridge port on which the better Hello BPDU was received.
7. The new Hello BPDU will be sent out by the bridge on all its ports except the port on which the better Hello BPDU was received.

This process will be repeated by all the bridges until:

1. There is one bridge (root bridge) remaining that is still transmitting its original Hello BPDU.
2. One bridge (designated bridge) on each LAN is transmitting the Hello BPDU based on the Hello BPDU received from the root bridge.

On the designated bridge, the port on which the best Hello BPDU is received is the root port and all the ports to which the Hello BPDU's are transmitted are the designated ports.

Note

There may be some ports on the designated bridge, over which the bridge will not be transmitting Hello BPDU's due to the fact that the received BPDU's on those ports are better than the one this bridge would be able to transmit (but they are not better than the Hello BPDU received on its root port).

Once the root and designated bridges have been elected, the root ports and the designated ports will be put in forwarding state and all the other ports will be put in blocking state.

5.2.3.1 Hello BPDU

The format of the Hello BPDU is shown in Table 6.

Table 6. Configuration BPDU	
Field	Size (bytes)
Protocol Identifier	2
Protocol Version Identifier	1
BPDU Type	1
Flags	1
Root Identifier	8
Root Path Cost	4
Bridge Identifier	8
Port Identifier	2
Message Age	2
Max Age	2
Hello Time	2
Forward Delay	2

The meaning of the various fields in the Hello BPDU are as follows:

- Protocol Identifier
Identifies the spanning tree protocol. This field contains 0.
- Protocol Version Identifier
Identifies the version number of the spanning tree protocol used. This field contains 0.
Note: The IEEE 802.1d committee has proposed Version 1 for *remote* bridges only.
- BPDU Type
Denotes the type of BPDU.
 - 0 for a configuration (Hello) BPDU.
 - 128 for Topology Change Notification (TCN) BPDU (see 5.2.4, “Transparent Bridges and Network Topology Changes” on page 102).
- Flags
 - Topology Change
This is the least significant bit of the Flag field, and if set to 1, denotes that the receiving bridge should use the *forward delay timer* rather than the *aging timer* for aging out the entries in the filtering database. See 5.2.4.2, “Topology Change Notification” on page 103 for more details about topology change notification.

- Topology Change Acknowledgment (TCA)

This is the most significant bit of the Flag field and if set to 1, it indicates that the bridge no longer needs to send TCN BPDUs. See 5.2.4.2, “Topology Change Notification” on page 103 for more details about TCA.

- Root ID

Specifies the bridge identifier of the root bridge. This field consists of the *priority* of the root bridge (2 bytes) and *bridge address* of the root bridge. Priority is assigned to a bridge during the configuration, and bridge address is the MAC address of the port with the lowest port identifier.

- Root Path Cost

Total cost from the bridge that transmitted this BPDU to the root bridge. BPDUs transmitted by the root bridge will have 0 in this field.

- Bridge ID

Bridge ID of the bridge transmitting the Hello BPDU. This field consists of 2 bytes of the bridge priority followed by the MAC address of the port with the lowest port identifier. The bridge transmitting a Hello BPDU is either the root or a designated bridge.

- Port Identifier

Port priority (1 byte) plus the port number (1 byte). Port priority is a user-configurable option.

- Message Age

Indicates the approximate age of the BPDU since it originated at the root bridge.

When a bridge receives a Hello BPDU, it starts a timer that is incremented every second. The initial value of this timer is the value contained in the Message Age field of the received BPDU.

When the designated bridge transmits its own Hello BPDU, it puts the value of this timer in the Message Age field.

- Max Age

This is the time after which the Hello BPDU stored in the bridge is deleted. Once the Message Age timer has reached this value, the bridge will assume the root bridge is not active and it will begin to establish itself as the root bridge.

- Hello Time

Denotes the frequency with which the root bridge should send the Hello BPDUs. This is a user-configurable option.

- Forward Delay Time

Specifies the length of the time that the bridge should stay in each of the *listening* and *learning* states before moving from blocking to forwarding state. As discussed in 5.2.4.1, “Filtering Database Update” on page 102, this timer can also be used for aging out the entries in the filtering database.

5.2.4 Transparent Bridges and Network Topology Changes

Bridges using the spanning tree algorithm automatically adjust to changes in network topology to ensure that a loop-free network is maintained. A change in the network topology can occur in the following circumstances:

1. When bridges enter or leave the network
2. When spanning tree parameters change, causing bridge ports to change state or causing a change in the choice of the root bridge

The result of any of the above changes is that:

1. The spanning tree has to be reconfigured using a Topology Change Notification protocol.
2. The filtering database must be updated.

5.2.4.1 Filtering Database Update

A transparent bridge builds a filtering database for each of its ports by listening to the frames exchanged on the LAN attached to that port. This database contains the addresses of stations attached to that LAN segment and are used to forward/discard frames across the bridge.

When the network topology changes due to the bridge addition, removal or reconfiguration, it is important that the bridges can update their filtering database quickly enough to cope with these changes in a manner that:

1. Ensures that the stations can continue to communicate with each other through the bridges.
2. Ensures the performance of the network is not affected due to the bridges forwarding the frames incorrectly and flooding the network.

To ensure the above, an *aging timer* is used by the bridges to delete entries within the filtering database that have not been used recently.

This timer should be able to cope with changes that happen as a result of stations physically moving from one LAN to another, as well as changes happening as a result of a bridge addition/removal (spanning tree reconfiguration). The latter will normally result in a group of stations logically moving from one LAN to another.

In general, to cope with the changes occurring due to the station moves, a longer aging timer is required than the one required to cope with the spanning tree reconfiguration. Therefore, the standard defines two timer values for the aging timer:

1. A longer timer value is to be used in coping with normal changes due to station additions, removals or timeouts. This is a user-configurable parameter and is referred to as *aging time*.
2. A shorter timer value is to be used when the bridge is in a state of topology change (see 5.2.4.2, "Topology Change Notification" on page 103). The *forward delay timer* of the root bridge is used for this purpose.

Note: The forward delay timer is specified for each bridge during its configuration, but all the bridges will use the value defined in the current root bridge.

5.2.4.2 Topology Change Notification

Spanning tree topology change is detected by a bridge whenever:

- A port enters forwarding state
- A port leaves forwarding state
- A new bridge becomes the root bridge

When a topology change happens, the following actions will be performed:

1. The bridge detecting the change issues a Topology Change Notification (TCN) BPDU. This frame will be sent on the root port to the destination address X'800143000000'.
2. The designated bridge on this port will acknowledge this frame by sending back a Hello BPDU with Topology Change Acknowledgment (TCA) set to 1.
3. The designated bridge will issue, on its root port, its own TCN BPDU.
4. This process repeats until a TCN BPDU reaches the root bridge.
5. The root bridge will start transmitting a Hello BPDU with the TCN set to 1 for a period equal to the sum of forward delay time and maximum age time.
6. The bridges that receive the Hello BPDU with TCN set to 1, will start using the shorter aging timer (forward delay) to age out filtering database entries. The forward delay timer will be used as the aging timer until a Hello BPDU with TCN set to 0 is received.

5.2.5 Setting the Parameters That Control the Spanning Tree

Table 7 shows the configurable spanning tree parameters that are defined as part of the standard for transparent bridging.

Table 7. Spanning Tree Parameters		
Parameter	Meaning	Default
Bridge Max Age	Maximum age of received BPDU	20 seconds
Bridge Hello Time	Time interval between configuration BPDUs	2 seconds
Bridge Forward Delay	Time spent in Listening state, time spent in Learning state, short aging timer	15 seconds
Bridge Priority	Priority portion of bridge identifier	32768
Path Cost	Cost for entering this port	1000/LAN_speed (Mbps)
Port Priority	Priority portion of port identifier	128

The spanning tree configuration parameters for the RouteSwitch can be verified by entering the command `sts` at the console.

```

/ % sts 1
  Spanning Tree Parameters for Group 1 (Default GROUP (#1))
Spanning Tree Status :          ON 1
Bridge Protocol Used :          IEEE 802.1D 2
Priority :          32768 (0x8000) 3
Bridge ID :          8000-0020DA:71C2C1 4
Designated Root :          000F-08005A:910371 5
Cost to Root Bridge :          100 6
Root Port :          Slot 2 Interface 1 Brg/1 7
Hold Time :          1 8
Topology Changes :          1 9
Last Topology Change :          2 minutes, 8 seconds ago 10
Bridge Aging Timer :          300 11

      Current Parameters 12
-----
Max Age          20 secs
Forward Delay    15 secs
Hello Time       2 secs

      Parameters system uses when 13
      attempting to become root
-----
System Max Age   20 secs
System Forward Delay 15 secs
System Hello Time 2 secs

```

Figure 73. IEEE 802.1d Spanning Tree

```

/ % sts 1
  Spanning Tree Parameters for Group 1 (Default GROUP (#1))
Spanning Tree Status :          ON 1
Bridge Protocol Used :          IBM Spanning Tree 2
Priority :          32768 (0x8000) 3
Bridge ID :          8000-0020DA:71C2C1 4
Designated Root :          0000-000000:000000 5
Cost to Root Bridge :          0 6
Root Port :          None 7
Hold Time :          1 8
Topology Changes :          0 9
Last Topology Change :          5 seconds ago 10
Bridge Aging Timer :          300 11

      Current Parameters 12
-----
Max Age          6 secs
Forward Delay    4 secs
Hello Time       2 secs

      Parameters system uses when 13
      attempting to become root
-----
System Max Age   6 secs
System Forward Delay 4 secs
System Hello Time 2 secs

```

Figure 74. IBM Spanning Tree

Figure 73 shows the output generated by the sts command when the RouteSwitch spanning tree is configured for the 802.1d spanning tree and Figure 74 shows the output when configured for the IBM spanning tree. The 1 after the sts command indicates that the bridge for group 1 is being viewed.

1 Spanning Tree Status: Indicates that the spanning tree is currently active. Care should be taken if the spanning tree is turned off since all ports on the switch will start forwarding. If the spanning tree needs to be disabled, first verify that there are no blocked port on the switch. If any ports are blocked, disable them first.

2 Bridge Protocol Used: The spanning tree is using the IEEE 802.1D version for transparent bridging. The IBM spanning tree protocol can be selected for source route bridging.

3 Priority: The bridge priority is utilized by the spanning tree algorithm to decide which bridge will become the root bridge. The priority can be set by entering a decimal number from 0 to 65,265. Zero being the highest priority.

4 Bridge ID: This is the current bridge ID of this switch.

5 Designated Root: This is the bridge ID of the root bridge. In this example, the root bridge is a different switch/bridge than the RouteSwitch.

6 Cost to the Root Bridge: This is the cost of the path from the RouteSwitch to the root bridge.

7 Root Port: The port on the RouteSwitch that has a path to the root bridge.

8 Hold Time: The interval length, in seconds, during which no more than two BPDUs can be transmitted.

9 Topology Change: The total number of topology changes detected by the RouteSwitch since the MPM was last reset or initialized.

10 Last Topology Change: Indicates the length of time elapsed since the last topology change.

11 Bridge Aging Timer: A timeout period in seconds for aging out dynamically learned forwarding information.

12 Current Parameters: The maximum age, forward delay, and hello time parameters defined in the root bridge.

13 Parameters System Uses when Attempting to Become Root: The maximum age, forward delay, and hello time defined in the RouteSwitch.

Note: The maximum age, forward delay and hello time are global parameters. To avoid confusion, they should be identically set on all bridges. Once the topology stabilizes, the values used throughout the network are those defined in the root bridge.

The stc command is used to modify the spanning tree configuration parameters.

```
/ % stc 1
Spanning Tree Parameters for Group 1 (Default GROUP (#1))

Spanning Tree is ON for this Group, set to OFF ?      (y/n) : n
IEEE spanning Tree for this Group, set to IBM ?      (y/n) : n
New Priority (0..65535) (current value is 32768·0x8000') :
New Bridge Hello Time (1..10 secs) (current value is 2) :
New Bridge Max Age (6..40 secs) (current value is 20) :
New Bridge Forward Delay (4..30 secs) (current value is 15) :
Aging Time (10..1000000 sec) (current value is 300) :
Auto-Tracker VLAN Aging Time (10..1000000 sec)(current value is 1200)

/ %
```

Figure 75. Spanning Tree Parameter Modifications Using the stc Command

The spanning tree parameters can also be viewed and modified from the IBM RouteSwitch Manager.



Figure 76. Selecting Spanning Tree Parameters from RouteSwitch Manager

The screenshot shows the 'IBM8274A:VLAN 1 Switching Spanning Tree Parameters' window. It displays various configuration parameters for the Spanning Tree Protocol (IEEE 802.1).

Parameter	Current Value	New Value
Priority:	32768	
Maximum Age (sec):	20	
Maximum Age when root (sec):	20	
Hello Time (sec):	2	
Hello Time when root (sec):	2	
Forward Delay (sec):	15	
Forward Delay when root (sec):	15	
Time Since Topology Change:	03:11:47.0	
Topology Changes:	1	
Designated Root:	000F08005A910371	
Root Port (Slot/Port/Service):	4.1.Bridge	
Root Cost:	107	
Hold Time (secs):	1	

At the bottom of the window, there is a status bar that reads: "Data received successfully, operation completed." Below this, there are four buttons: "Close", "Apply", "Update", and "Help".

Figure 77. Spanning Tree Parameters from the RouteSwitch Manager

The aging time can be viewed and modified from the IBM RouteSwitch Manager as seen in Figure 78 on page 107 and Figure 79 on page 107.



Figure 78. Selecting Aging Time Parameters from the RouteSwitch Manager

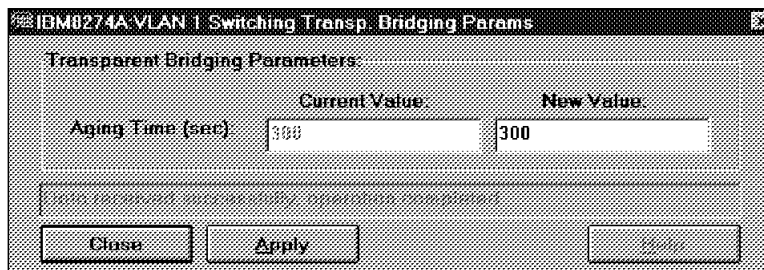


Figure 79. Aging Time Parameters

The spanning tree configuration parameters for all ports of the RouteSwitch can be verified by entering the command stps at the console.

```

/ % stps 1
Spanning Tree Port Summary for Group 1 (Default GROUP (#1))

```

Slot	Service					Path	Desig	Des	Rt	Swt	Fw	Root	Bridg
Intf	Inst	Pri	State	MAC	Cost	Cost	Pt	Pt	Pt	Pt	Tx	Design	Bridg
4/ 1	Brg/	1	128	FORWD	6FA8E0	7	100	No	Yes	No	1	000F-08005A:	
												8000-0020DA:	
5/ 1	Brg/	1	128	BLOCK	76B210	10	100	No	No	No	0	000F-08005A:	
												8000-0020DA:	

Figure 80. Port Spanning Tree

Figure 80 shows the output from the stps command. Only active ports are displayed.

- Slot/Intf: The slot number/port number.
- Service Inst: The type of service, in this case a bridge, and also indicates if there is more than one service defined on a port.
- Pri: The priority of the port (from 0 to 256).
- State: The current state of the port as defined by the spanning tree algorithm. The state values are: disabled, blocking, listening, learning and forwarding.
- MAC: The last three bytes of the MAC address for the port.
- Path Cost: The cost of this port. The spanning tree algorithm will use this value to make a decision on the lowest cost paths to reach the root bridge. The RouteSwitch assigns the path cost based on the formula: $1000/\text{media speed}$. Port 4/1 is an ATM port with a media speed of 155 Mbps; $1000/155=6.45$, rounded to 7. Port 5/1 is a 100 Mbps Ethernet port, $1000/100=10$. Port 4/1 having the lowest cost is forwarding frames while port 5/1 is blocked.
- Desig Cost: The path cost of the designated port of the segment connected to this port; this value is always 0 on the root bridge.
- Rt Pt: The root port, if yes, has the lowest cost path to the root bridge; no indicates that another port has the lowest path cost.
- Swt pt: Switch port; if yes, indicates that this port is in optimize switch mode. The RouteSwitch will not transmit BPDU on this port. If a BPDU is receive on this port, then the RouteSwitch will automatically revert the port back to normal operation and process BPDU frames.
- FWD Transition: The number of times this port has changed from a learning state to a forwarding state.
- Root Bridge ID: The bridge identification number of the root bridge.
- Desig Bridge ID: The unique bridge identifier of the designated bridge for this port.

The spanning tree port parameters can be modified using the stpc command, as shown in Figure 81 on page 109.

```

/ % stpc 1
Spanning Tree Port Configuration for Group 1 (Default GROUP (#1))

```

Index	Slot/Intf/Service/Inst	Port Priority (a)	Path Cost (b)	Enable Spanning Tree (c)	tx FA (d)	Manual Mode (e)
1	3/ 1/ Brg/ 1	128	100	y	NA	n
2	3/ 2/ Brg/ 1	128	100	y	NA	n
3	3/ 3/ Brg/ 1	128	100	y	NA	n
4	3/ 4/ Brg/ 1	128	100	y	NA	n
5	3/ 5/ Brg/ 1	128	100	y	NA	n
6	3/ 6/ Brg/ 1	128	100	y	NA	n
7	4/ 1/ Brg/ 1	128	7	y	NA	n
8	4/ 2/ Brg/ 1	128	7	y	NA	n
9	5/ 1/ Brg/ 1	128	10	y	NA	n

```

save|cancel|next|prev :

```

Figure 81. Spanning Tree Port Configuration Menu

The parameters for each port can be modified by entering <index><column>=<new parameter> at the command line.

```

/ % stpc 1
Spanning Tree Port Configuration for Group 1 (Default GROUP (#1))

```

Index	Slot/Intf/Service/Inst	Port Priority (a)	Path Cost (b)	Enable Spanning Tree (c)	tx FA (d)	Manual Mode (e)
1	3/ 1/ Brg/ 1	128	100	y	NA	n
2	3/ 2/ Brg/ 1	128	100	y	NA	n
3	3/ 3/ Brg/ 1	128	100	y	NA	n
4	3/ 4/ Brg/ 1	128	100	y	NA	n
5	3/ 5/ Brg/ 1	128	100	y	NA	n
6	3/ 6/ Brg/ 1	128	100	y	NA	n
7	4/ 1/ Brg/ 1	128	7	y	NA	n
8	4/ 2/ Brg/ 1	128	7	y	NA	n
9	5/ 1/ Brg/ 1	128	10	y	NA	n

```

save|cancel|next|prev : 1c=n
save|cancel|next|prev : save
All items saved
save|cancel|next|prev : cancel
/ %

```

Figure 82. Modifying the Spanning Tree Port Configuration

Typing 1c=n means that:

1. Modify the entry at index 1.
2. Change the parameter for column c; enable the spanning tree.
3. Set the new parameter to no.

The commands save and cancel were entered to save the new parameters and return to the prompt of the RouteSwitch console.

Figure 83 on page 110 shows the result of the 1c=n command the next time the command stpc is entered.

```

/ % stpc 1
Spanning Tree Port Configuration for Group 1 (Default GROUP (#1))

Index Slot/Intf/Service/Inst  Port Priority (a)  Path Cost (b)  Enable Spanning Tree (c)  tx FA (d)  Manual Mode (e)
-----
1      3/ 1/ Brg/ 1      128      100      n      NA      n
2      3/ 2/ Brg/ 1      128      100      y      NA      n
3      3/ 3/ Brg/ 1      128      100      y      NA      n
4      3/ 4/ Brg/ 1      128      100      y      NA      n
5      3/ 5/ Brg/ 1      128      100      y      NA      n
6      3/ 6/ Brg/ 1      128      100      y      NA      n
7      4/ 1/ Brg/ 1      128       7      y      NA      n
8      4/ 2/ Brg/ 1      128       7      y      NA      n
9      5/ 1/ Brg/ 1      128      10      y      NA      n
save|cancel|next|prev :
/ %

```

Figure 83. Modified Spanning Tree Port Configuration

The following parameters can be modified in the spanning tree port configuration menu:

- Port priority (a): The priority of the port. The default is 128. The port priority has little effect on the spanning tree algorithm.
- Path cost (b): The contribution of this port to the path cost of paths towards the spanning tree root which includes this port. 802.1d-1990 recommends that the default value of this parameter be inversely proportional to the speed of the attached LAN. The RouteSwitch uses the formula 1000/media speed.
- Enable spanning tree (c): To enable or disable the spanning tree on a port. The values are y or n.
- tx FA: Transmit functional address.

The values are:

- NA - Not applicable.
- y - Yes transmit functional address instead of the normal spanning tree multicast address.
- n - (Default) transmit normal spanning tree multicast address.

The spanning tree port configuration can also be viewed and modified with the IBM RouteSwitch Manager.



Figure 84. Selecting Port Spanning Tree Parameters

IBM8274A Switching Spanning Tree Port Table

Switching Spanning Tree Port Table

	Slot	Port	Service	Instance	Priority	Port State	Admin State	Path Cost	Fwd. Trans.
1	3	1	Bridge	1	128	Disabled	Disabled	100	0
2	3	2	Bridge	1	128	Disabled	Enabled	100	0
3	3	3	Bridge	1	128	Disabled	Enabled	100	0
4	3	4	Bridge	1	128	Disabled	Enabled	100	0
5	3	5	Bridge	1	128	Disabled	Enabled	100	0
6	3	6	Bridge	1	128	Disabled	Enabled	100	0
7	3	7	Bridge	1	128	Disabled	Enabled	100	0
8	3	8	Bridge	1	128	Disabled	Enabled	100	0
9	3	9	Bridge	1	128	Disabled	Enabled	100	0
10	3	10	Bridge	1	128	Disabled	Enabled	100	0
11	3	11	Bridge	1	128	Disabled	Enabled	100	0
12	3	12	Bridge	1	128	Disabled	Enabled	100	0
13	4	1	Bridge	1	128	Forwarding	Enabled	7	0
14	4	2	Bridge	1	128	Disabled	Enabled	7	0
15	5	1	Bridge	1	128	Blocking	Enabled	10	0

Close Update Modify Help

Figure 85. Port Spanning Tree Parameters

Slot	Port	Service	Instance	Priority	Port State	Admin State	Path Cost	End Trans
3	1	Bridge	1	128	Disabled	Disabled	100	0
3	2	Bridge	1	128	Disabled	Enabled	100	0
3	3	Bridge	1	128	Disabled	Enabled	100	0
3	4	Bridge	1	128	Disabled	Enabled	100	0
3	5	Bridge	1	128	Disabled	Enabled	100	0
3	6	Bridge	1	128	Disabled	Enabled	100	0
3	7	Bridge	1	128	Disabled	Enabled	100	0
3	8	Bridge	1	128	Disabled	Enabled	100	0
3	9	Bridge	1	128	Disabled	Enabled	100	0
3	10	Bridge	1	128	Disabled	Enabled	100	0
3	11	Bridge	1	128	Disabled	Enabled	100	0
3	12	Bridge	1	128	Disabled	Enabled	100	0
4	1	Bridge	1	128	Forwarding	Enabled	7	0
4	2	Bridge	1	128	Disabled	Enabled	7	0
5	1	Bridge	1	128	Blocking	Enabled	10	0

Close Update Modify Help

Figure 86. Modifying the Port Spanning Tree Parameters

5.2.6 Summary of the IEEE 802.1d Spanning Tree Algorithm

From the previous discussion, it is clear that the key parameters governing the topology of the spanning tree are the bridge priority and the path cost. The port priority is unlikely to have any effect on the spanning topology.

In most circumstances, adjusting the bridge priority and using the IEEE defaults for all other parameters should provide acceptable control over the active topology. One approach could be to:

1. Choose three values of bridge priority: a low value, a medium value and the IEEE default.
2. Assign the low value to the bridge to be the normal root bridge. This bridge becomes the center of the network.
3. Assign the medium value to any other bridge that may become the root. This allows for failure of the root bridge.
4. Allow all other bridges in the network to use the IEEE or IBM default.

The result of the spanning tree algorithm for transparent bridges is a loop-free network in which the endstations require no knowledge of the network topology to be able to communicate with the other stations through one or more bridges. However, another result is a network in which there is no load balancing over bridges and in case of parallel bridges all but one of the bridges will be idle (blocking state).

5.3 Configuring Source Route Bridging

By default, all ports of the RouteSwitch will forward frames using transparent bridging.

In some cases, source route bridging might be needed in a network. For example, to have load balancing paths (used commonly with SNA host access), or to preserve ring and bridge numbers for network management purpose.

It is possible to define the ports on the token-ring and FDDI modules to be a part of a source route bridge (SR) or source route transparent bridge (SRT).

As mentioned in 5.2, “The Spanning Tree” on page 96, only one type of spanning tree can be defined for a group. When there is a need to use both the 802.1d and the IBM spanning tree for SR or SRT bridging, then the token-ring ports will have to be assigned to a new group in order to enable the IBM spanning tree for this new group, leaving the 802.1d spanning tree for the original group.

Important

The RouteSwitch does not store RIF information. Therefore, translational bridging is not supported. It is possible to enable RIF stripping in the RouteSwitch so that if a frame contains a broadcast RIF with no RIF information, the RouteSwitch will remove the RIF frame and forward the frame.

RIF stripping can be enabled on the entire RouteSwitch if the following conditions are met:

- The RIF is 2 bytes.
- The frame is switched between token-ring and Ethernet.
- The token-ring port is set to SRT.

To enable RIF stripping, add the following line to the mpm.cmd:
rifStripping=1

Token-ring frames with an RII/RIF can only be forwarded using the SR or SRT bridging function.

Token-ring frames without an RII/RIF will be translated to other frame types using the any-to-any switching function of the RouteSwitch. See 5.7, “Mixed Media” on page 150, for a further explanation of any-to-any switching.

SR and SRT bridges require a ring number and bridge number.




Figure 87. RouteSwitch As a Source Route Bridge

Figure 87 shows a simple source route bridge network using the 8274 as the bridge. This network will now be defined using the RouteSwitch console. The token-ring module is installed in slot 2 and six ports are assigned to group #2.

```

/ % src 2

Source Routing Parameters for Group 2 (Test GROUP #2)

Slot  Type/  Ring  Bridge  Largest  HopCnt Port  Block
Intf  Inst/Srv  Number Number  frame   In  Out Type  ARE
-----
1. 2/ 1 Brg/  1/ na  ***** not configured *****
2. 2/ 2 Brg/  1/ na  ***** not configured *****
3. 2/ 3 Brg/  1/ na  ***** not configured *****
4. 2/ 4 Brg/  1/ na  ***** not configured *****
5. 2/ 5 Brg/  1/ na  ***** not configured *****
6. 2/ 6 Brg/  1/ na  ***** not configured *****

Enter index of the entry to configure (e.g. 1) <RETURN> to exit :
  
```

Figure 88. Defining Source Route Bridging

Figure 88 shows the output of entering the command src 2. The 2 indicates that ports assigned to group 2 will be modified.

```

/ % src 2

Source Routing Parameters for Group 2 (Test GROUP #2)

  Slot  Type/      Ring      Bridge      Largest      HopCnt  Port  Block
  Intf  Inst/Srvc   Number    Number     frame      In  Out  Type  ARE
-----
1. 2/ 1 Brg/  1/ na  ***** not configured *****
2. 2/ 2 Brg/  1/ na  ***** not configured *****
3. 2/ 3 Brg/  1/ na  ***** not configured *****
4. 2/ 4 Brg/  1/ na  ***** not configured *****
5. 2/ 5 Brg/  1/ na  ***** not configured *****
6. 2/ 6 Brg/  1/ na  ***** not configured *****

Enter index of the entry to configure (e.g. 1) <RETURN> to exit : 1
Ring Number (1 - 4095, 0 to disable) (-----) : 1 1
Virtual Ring (y/n) (n) : n 2
Bridge Number (1 - 15, 0 to disable) (---) : 6 3
Max Outbound Hop Count (7) : 7 4
Max Inbound Hop Count (7) : 7 4
Largest Frame size (4472) : 5 5
Turn Transparent Bridging ON (y/n) (y) : n 6
Block ARE on non-forward state (y/n) (n) : 7 7
Save the new configuration? (y/n) (n) : y 8
Enter index of the entry to configure (e.g. 1) <RETURN> to exit :
/ %

```

Figure 89. Defining Port 2/1 for Source Route Bridging

The following explains the parameters used in Figure 89.

- 1** Ring number 1 defined on port 2/1. If the ring number is entered with no parameters, then the ring number will be converted from decimal to hexadecimal. The format to specify a hexadecimal ring number is 0xrrr where rrr is the ring number. For example entering 588 will convert to a ring number of 0x24c. Entering 0x588 will provide a ring number of 0x588.
- 2** For virtual ring definition, see 5.3.1, “Virtual Token-Rings” on page 120.
- 3** Bridge number 6. The bridge number must always be the same within a group.
- 4** The Maximum Outbound and Inbound hop count were left at the default of 7.
- 5** The maximum frame size was left at the default of 4472.
- 6** Port 2/1 is an SR bridge only. Only frames with an RII and the correct RIF will be forwarded by this port.
- 7** All route explorer frames will be sent on this port even if the spanning tree has blocked the port.
- 8** The configuration is saved. Entering n or pressing the Return key would discard the configuration just defined.

```

Enter index of the entry to configure (e.g. 1) <RETURN> to exit : 2
Ring Number (1 - 4095, 0 to disable) (-----) : 5 1
Virtual Ring (y/n) (n) : n 2
Bridge Number (1 - 15, 0 to disable) (---) : 6 3
Max Outbound Hop Count ( 7) : 7 4
Max Inbound Hop Count ( 7) : 7 4
Largest Frame size (4472) : 5 5
Turn Transparent Bridging ON (y/n) (y) : y 6
Block ARE on non-forward state (y/n) (n) : 7
Save the new configuration? (y/n) (n) : y 8
Enter index of the entry to configure (e.g. 1) <RETURN> to exit :
/ %

```

Figure 90. Defining Port 2/2 for Source Route Transparent Bridging

The following explains the parameters used in Figure 90.

- 1** Ring number 5 defined on port 2/2.
- 2** For virtual ring definition, see 5.3.1, “Virtual Token-Rings” on page 120.
- 3** Bridge number 6. The bridge number must always be the same within a group.
- 4** The Maximum Outbound and Inbound Hop Counts were left at their defaults of 7.
- 5** The maximum frame size was left at the default of 4472.
- 6** Port 2/2 is an SRT bridge. Frames with an RII and the correct RIF will be forwarded using source route bridging and frames with no RIF will be forwarded using transparent bridging.
- 7** All route explorer frames will be sent on this port even if the spanning tree has blocked the port.
- 8** The configuration is saved. Typing n or pressing the Return key would discard the configuration just defined.

```

Enter index of the entry to configure (e.g. 1) <RETURN> to exit : 3
Ring Number (1 - 4095, 0 to disable) (-----) : 8 1
Virtual Ring (y/n) (n) : n 2
Bridge Number (1 - 15, 0 to disable) (---) : 6 3
Max Outbound Hop Count ( 7) : 7 4
Max Inbound Hop Count ( 7) : 7 4
Largest Frame size (4472) : 1542 5
Turn Transparent Bridging ON (y/n) (y) : y 6
Block ARE on non-forward state (y/n) (n) : 7
Save the new configuration? (y/n) (n) : y 8
Enter index of the entry to configure (e.g. 1) <RETURN> to exit :
/ %

```

Figure 91. Defining Port 2/3 for Source Route Transparent Bridging

The following explains the parameters used in Figure 91.

- 1** Ring number 8 defined on port 2/3.
- 2** For virtual ring definition, see 5.3.1, “Virtual Token-Rings” on page 120.
- 3** Bridge number 6. The bridge number must always be the same within a group.

4 The Maximum Outbound and Inbound Hop Counts were left at their defaults of 7.

5 The maximum frame size was reduced to 1542.

6 Port 2/3 is an SRT bridge. Frames with an RII and the correct RIF will be forwarded using source route bridging and frames with no RII/RIF will be forwarded using transparent bridging.

7 All route explorer frames will be sent on this port even if the spanning tree has blocked the port.

8 The configuration is saved. Typing n or pressing the Return key would discard the configuration just defined.

The same configuration can be done using the RouteSwitch Manager as shown in Figure 92 through Figure 97 on page 120.



Figure 92. Selecting Source Route Bridging Function

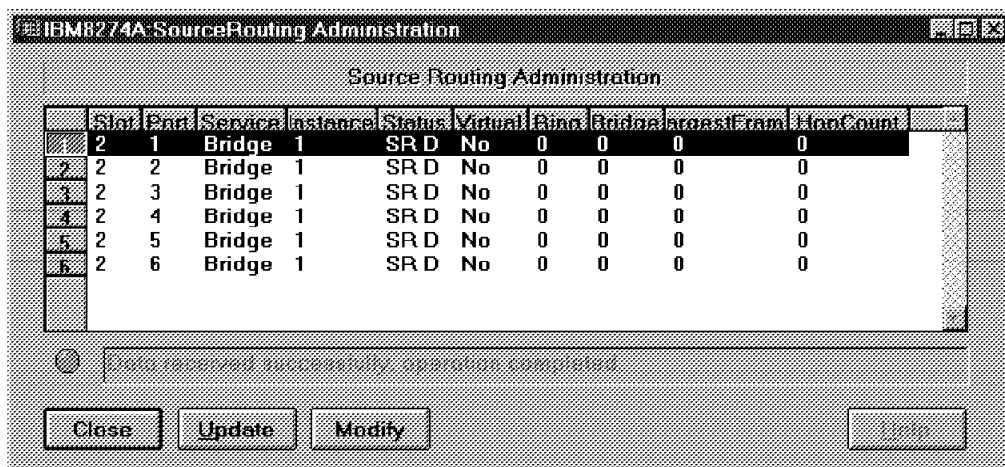


Figure 93. Selecting a Port to Be Configured

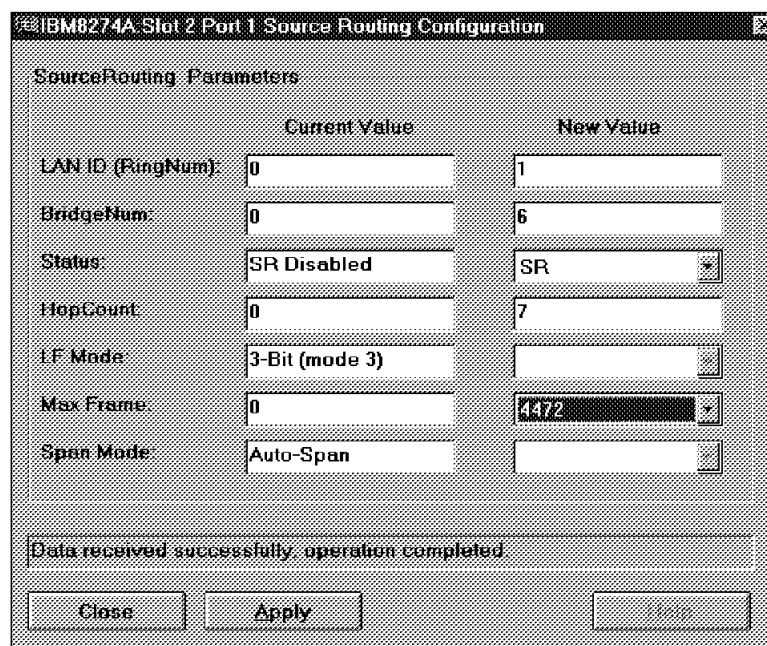


Figure 94. Configuring Port 2/1

IBM8274A Slot 2 Port 2 Source Routing Configuration

SourceRouting Parameters

	Current Value	New Value
LAN ID (RingNum):	0	5
BridgeNum:	0	6
Status:	SR Disabled	SRT
HopCount:	0	7
LF Mode:	3-Bit (mode 3)	
Max Frame:	0	4472
Span Mode:	Auto-Span	

Data received successfully, operation completed.

Close Apply Help

Figure 95. Configuring Port 2/2

IBM8274A Slot 2 Port 3 Source Routing Configuration

SourceRouting Parameters

	Current Value	New Value
LAN ID (RingNum):	0	0
BridgeNum:	0	6
Status:	SR Disabled	SRT
HopCount:	0	7
LF Mode:	3-Bit (mode 3)	
Max Frame:	0	1542
Span Mode:	Auto-Span	

Data received successfully, operation completed.

Close Apply Help

Figure 96. Configuring Port 2/3

Source Routing Administration										
	Slot	Port	Service	Instance	Status	Virtual	Ring	Bridge	Transit	Frame
1	2	1	Bridge	1	SR	No	1	6	4472	7
2	2	2	Bridge	1	SRT	No	5	6	4472	7
3	2	3	Bridge	1	SRT	No	8	6	1542	7
4	2	4	Bridge	1	SR D	No	0	0	0	0
5	2	5	Bridge	1	SR D	No	0	0	0	0
6	2	6	Bridge	1	SR D	No	0	0	0	0

Info received successfully, operation completed

Close Update Modify Help

Figure 97. Viewing New Source Route Bridge Configuration

5.3.1 Virtual Token-Rings

The 8274 RouteSwitch allows one ring number to be assigned to multiple token-ring ports; this is known as virtual rings. The traffic between the physical rings will be switched. The traffic going to a different ring number will be bridged by the SRT bridge. When using virtual rings, only an SRT bridge can be configured for these ports.




Figure 98. Network with Virtual Rings

Figure 98 shows a network with virtual rings, both 001 rings. When the 8274 RouteSwitch sees a frame destined for ring 001, it will switch the frame to port 2/1 or 2/3 according to the destination MAC address of the frame.

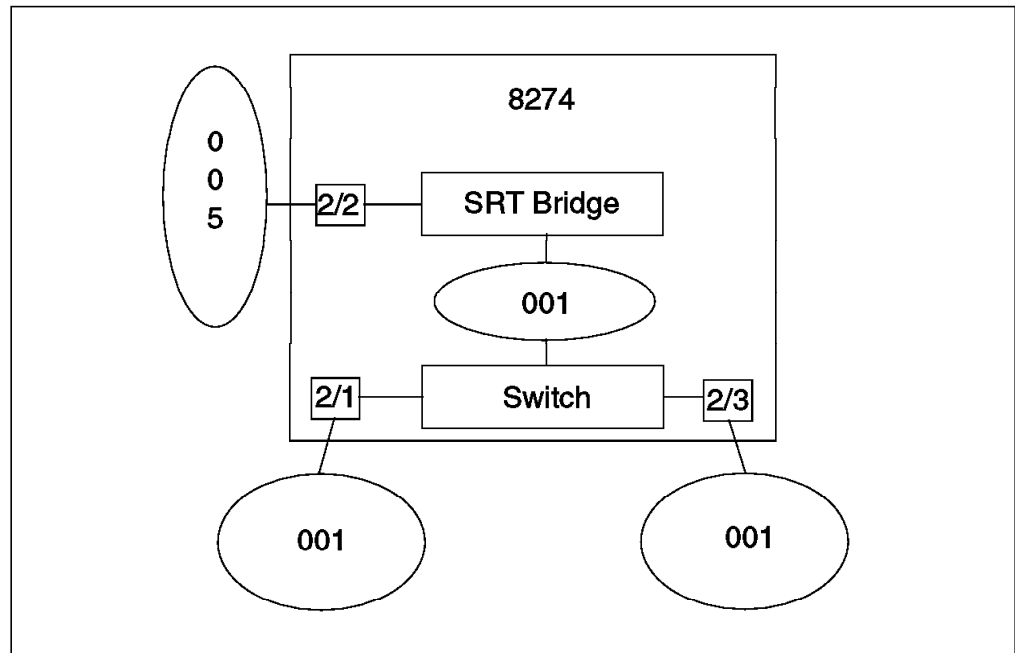


Figure 99. Logical Representation of Virtual Rings

Figure 99 represents the logical view of having virtual rings mixed with physical rings. Ring 001 logically exists in the 8274 RouteSwitch. Frames to and from ring 005 will use the SRT bridge to exchange frames with ring 001. Once the frames are in ring 001, the 8274 RouteSwitch will switch the frames to the appropriate physical port, 2/1 or 2/3 following the transparent bridging rules.

The users on the ring attached to port 2/1 and 2/3 will appear to be on ring 001.

```

Enter index of the entry to configure (e.g. 1) <RETURN> to exit : 1
Ring Number (1 - 4095, 0 to disable) (-----) : 2
Virtual Ring (y/n) (n) : y
Bridge Number (1 - 15, 0 to disable) (---) : 6
Max Outbound Hop Count ( 7) : 7
Max Inbound Hop Count ( 7) : 7
Largest Frame size (4472) :
Block ARE on non-forward state (y/n) (n) :
Save the new configuration? (y/n) (n) : y
Enter index of the entry to configure (e.g. 1) <RETURN> to exit :

```

Figure 100. Defining Port 2/1 for Virtual Ring

Figure 100 shows how to define a virtual ring.

Since only SRT bridging is possible when selecting a virtual ring, the question on transparent bridging is not asked.

5.3.2 Invalid Configuration

It is possible to enter an invalid source route bridge configuration in the 8274 RouteSwitch.

Do not:

- Configure more than one bridge number per group.

- Configure two rings with the same ring number in the same group without having virtual rings enabled for these rings.
- Create a loopback to a ring using an external SR or SRT bridge.

5.3.3 Source Route Bridging between RouteSwitches

Virtual rings cannot be extended between two or more RouteSwitches when the backbone is token-ring. In this case, the 8274 RouteSwitch will act as an SR or SRT bridge.

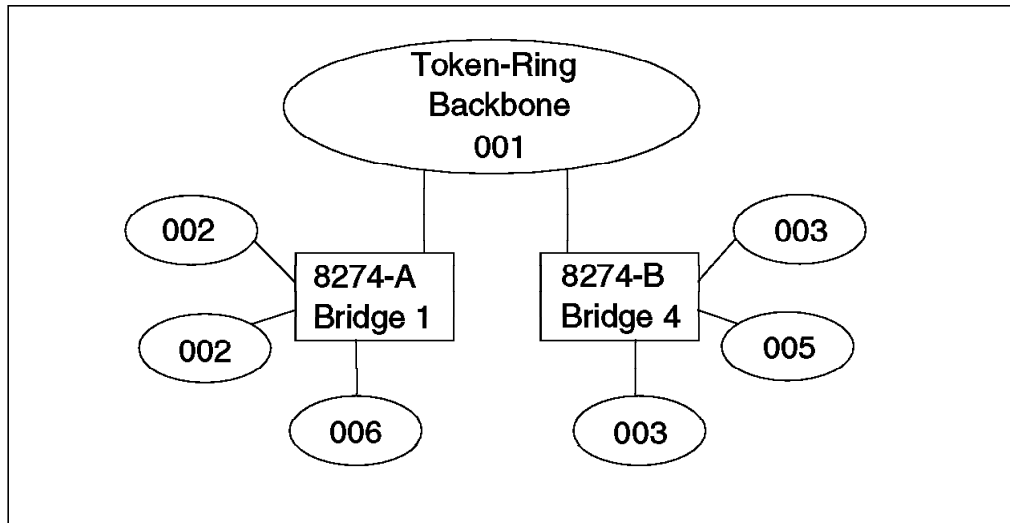


Figure 101. Source Route Bridging between RouteSwitch

Figure 101 shows two RouteSwitches, each sharing the backbone ring 001. Rings 002 on 8274-A and rings 003 on 8274-B must be defined as virtual rings as described in 5.3.1, "Virtual Token-Rings" on page 120.

Notice that the ring numbers are different for each ring or each virtual ring across the network. The bridge numbers are also different.

When the link between two or more RouteSwitches is an FDDI or ATM link and the link is a trunk, it is then possible to extend the virtual rings across the RouteSwitches.

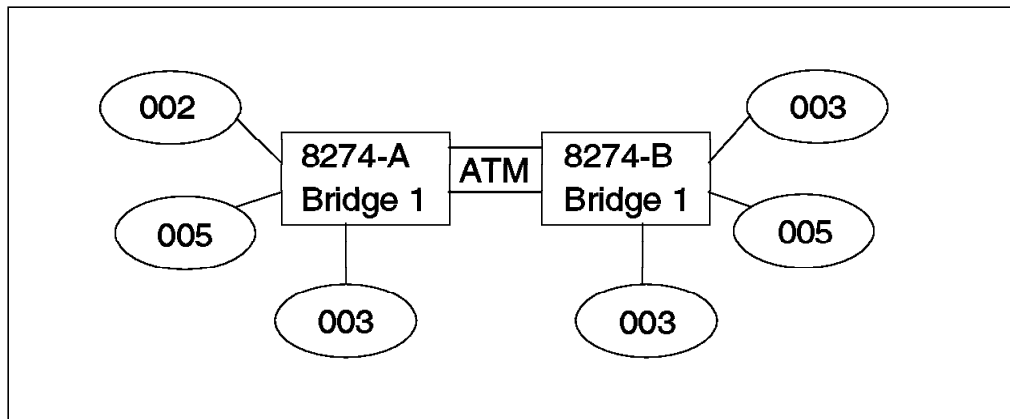


Figure 102. Source Route Bridging with Trunking

Figure 102 shows that ring 005 on 8274-A and ring 005 on 8274-B are virtual rings. A frame with an RIF for ring 005 will go to both RouteSwitches and then the frame will be switched to the appropriate port as per the destination MAC address of the frame. Virtual ring 003 has three ports associated with it.

Notice that in this case, both RouteSwitches bridge numbers are the same and will act as a single source route bridge number 1.

The extension of virtual rings between RouteSwitches are only supported when using FDDI or ATM trunking. ATM LEC will not work to extend virtual rings between RouteSwitches.

5.3.4 Summary of Bridge Types

Each token-ring port of the RouteSwitch can be configured with a different bridge type: TB, SR or SRT. Mixing bridge types within a group may lead to a failure to establish connections between devices on different ports.

Figure 103 along with Table 8 on page 124 shows what devices are able to communicate together in an environment where multiple bridge types are used.

This diagram does not take frame type and protocol type into consideration when frames are moved to/from the Ethernet segment.

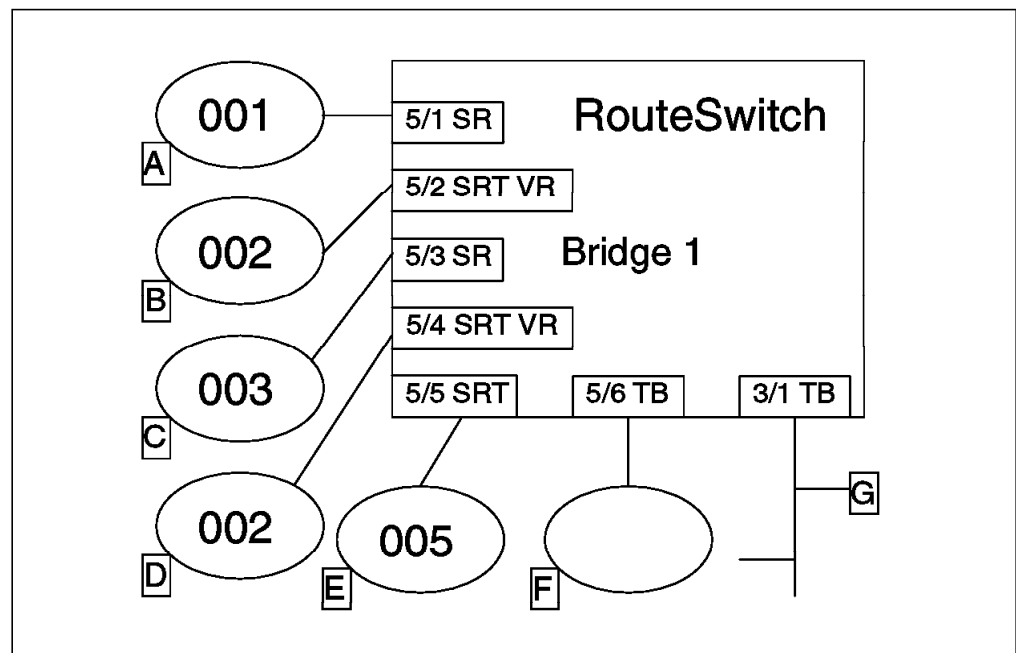


Figure 103. Bridging within the RouteSwitch

The following was configured:

- One RouteSwitch, one group using bridge 1
- No VLANs configured
- Port 5/1: Ring 001, token-ring, source route bridge only
- Port 5/2: Ring 002, token-ring, source route transparent bridge with virtual ring enabled, associated to port 5/4
- Port 5/3: Ring 003, token-ring, source route bridge only

- Port 5/4: Ring 002, token-ring, source route transparent bridge with virtual ring enabled, associated to port 5/2
- Port 5/5: Ring 005, token-ring, source route transparent bridge and virtual ring disabled
- Port 5/6: Token-ring, transparent bridge (default)
- Port 3/1: Ethernet, transparent bridge

```

/ % srs 2

Source Routing Parameters for Group 1 (Default GROUP (#1))

Slot Type/      Ring      Bridge      Largest      HopCnt      Port      Block
Intf  Inst/Srv  Number      Number      frame      In   Out  Type  ARE
-----
1. 5/ 1 Brg/  1/ na      1 (0x001)  1 (0x1)    4472    14  14  SR    n
2. 5/ 2 Brg/  1/ na (V)    2 (0x002)  1 (0x1)    4472    14  14  SRT   n
3. 5/ 3 Brg/  1/ na      3 (0x003)  1 (0x1)    4472    14  14  SR    n
4. 5/ 4 Brg/  1/ na (V)    2 (0x002)  1 (0x1)    4472    14  14  SRT   n
5. 5/ 5 Brg/  1/ na      5 (0x005)  1 (0x1)    4472    14  14  SR    n
6. 5/ 6 Brg/  1/ na      *****  not configured *****

Enter index of the entry to see details (e.g. 1) <RETURN> to exit :

```

Figure 104. Token-Ring Bridging Configuration

The srs command will show the bridge configuration for each token-ring port.

In Figure 104 port 5/6 shows not configured. This means that port 5/6 is a transparent bridge port.

Table 8 shows which devices can communicate together and exchange data.

Table 8. Bridging within the RouteSwitch							
	A	B	C	D	E	F	G
A	-	SR	SR	SR	SR	N/A	N/A
B	SR	-	SR	T	SR-T	T	T
C	SR	SR	-	SR	SR	N/A	N/A
D	SR	T	SR	-	SR-T	T	T
E	SR	SR-T	SR	SR-T	-	T	T
F	N/A	T	N/A	T	T	-	T
G	N/A	T	N/A	T	T	T	-

The following explains the acronyms used in Table 8:

- N/A: Not Available.
- SR: The traffic will flow using the source route bridge function.
- T: The traffic will flow using the transparent bridge function.
- SR-T: The traffic will flow using either the source route bridge function or the transparent bridge function.

Important

The transparent bridge (TB) function in the 8274 will not forward frames that contain a RIF (unlike the 8272). This also includes frames sent with a broadcast RIF. RIF stripping can be enabled on the entire RouteSwitch if the following conditions are met:

- The RIF is 2 bytes.
- The frame is switched between token-ring and Ethernet.
- The token-ring port is set to SRT.

To enable RIF stripping, add the following line to the mpm.cmd:
rifStripping=1. Some protocols, in their implementation will:

- Always include an RIF broadcast within their broadcast frames.
- Include an RIF broadcast in their frame when replying to a frame with no RIF.

5.4 VLAN and Group Routing

The RouteSwitch contains an internal virtual router that can route IP and IPX frames between VLAN and/or groups. Please take into consideration the introduction of mobile groups in 3.3.11, "Mobile Groups" on page 50. This section describes the routing between VLANs in non-mobile groups. Routing ports can also be defined in mobile groups with the same policies.

Each RouteSwitch can have up to thirty-two active virtual router ports.

The router can forward thirty to forty thousands packets per seconds.

Important

A VLAN or a group can only have one router port assigned to it.

Only one IP and/or IPX network address can be assigned to each router port.




Figure 105. Logical View of the RouteSwitch Virtual Router

In Figure 105, a second group was defined within a RouteSwitch. VLAN 2 was added to the default group. VLAN 2 and VLAN 3 were added to group 2.

Virtual router ports are defined when creating (cratvl) or modifying (modvl) a VLAN.

See Chapter 4, “Basic VLAN Configuration” on page 67 for the procedure to define groups and VLANs.

VLAN 2 of group 2 does not have a connection to the virtual router. An endstation that is a member of a VLAN that is not linked to the virtual router cannot send or receive any IP or IPX frame from the other VLANs. As an example, endstation X in Figure 105 cannot send or receive any IP or IPX frames via the virtual router.

However, an endstation on an isolated VLAN, (VLAN 2, Group 2) which is also a member of another VLAN (VLAN 3 Group 2) that is attached to the virtual router, may be able to communicate to endstations on other VLANs.


```

/ % cratv1
Enter the VLAN Group id for this VLAN ( 1) :
Enter the VLAN Id for this VLAN ( 2) :
Enter the new VLAN's description: IP based VLAN
Enter the Admin status for this vlan (e)nable/(d)isable (d): e
Select rule type:
  1. Port Rule
  2. MAC Address Rule
  3. Protocol Rule
  4. Network Address Rule
  5. User Defined Rule
  6. Binding Rule
  7. DHCP PORT Rule
  8. DHCP MAC Rule
Enter rule type (1): 3
Set Rule Admin Status to (e)nable/(d)isable (d): e
Select Protocol:
  1. IP
  2. IPX
  3. DECNET
  4. APPLETALK
  5. Protocol specified by ether-type (in hex)
  6. Protocol specified by DSAP and SSAP (in hex)
  7. Protocol specified by SNAP (in hex)
Enter protocol type (1): 1
Configure more rules for this vlan y/n (n): n
VLAN 1: 2 created successfully

```

Figure 106. Configuration of an IP Router Port (1 of 2)

```

Enable IP? (y): y 1
  IP Address : 9.24.106.1 2
  IP Subnet Mask (0xff000000) : 255.255.255.0 3
  IP Broadcast Address (9.24.106.255 ) : 4
  Description (30 chars max) : Router port for VLAN 2 5
  Disable routing? (n) :
  Enable NHRP? (n) :
  IP RIP mode {Deaf(d),
    Silent(s),
    Active(a),
    Inactive(i)} (s) : a 6
  Default framing type {Ethernet II(e),
    fddi(f),
    token-ring(t),
    Ethernet 802.3(8),
    source route token-ring(s)} (e) : 7
Created router port for vlan 1: 2
Enable IPX? (y): y 8
  IPX Network : 01234567 9
  Description (30 chars max) : IPX router on VLAN 2 10
  IPX RIP and SAP mode {RIP and SAP active(a),
    RIP only active(r),
    RIP and SAP inactive(i)} (a) : a 11
  Default router framing type for: {
    Ethernet Media:
      Ethernet II(0),
      Ethernet 802.3 LLC(1),
      Ethernet 802.3 SNAP(2),
      Novell Ethernet 802.3 raw(3),
    FDDI Media:
      fddi SNAP(4),
      source route fddi SNAP(5),
      fddi LLC(6),
      source route fddi LLC(7),
    Token-Ring Media:
      token-ring SNAP(8),
      source route token-ring SNAP(9),
      token-ring LLC(a),
      source route token-ring LLC(b)} (0) : 0 12
Created router port for vlan 1: 2

```

Figure 107. Configuration of an IP Router Port (2 of 2)

Figure 106 on page 127 and Figure 107 show how to define a router port using the `cratvl` command.

- 1** Enabling IP will activate the IP protocol on the router port for the VLAN just defined. Entering `n` will bypass all IP configurations.
- 2** The IP address associated with this router port.
- 3** The subnet mask of the IP address defined in **2**.
- 4** The broadcast address.
- 5** A brief description of this router port, up to 30 characters.
- 6** The status of the RIP protocol on this router port.
 - Deaf: The router port will not listen to RIP updates from the network, but will transmit RIP updates.
 - Silent: The router port will listen to RIP updates, but will not transmit any RIP updates.

- Active: The router port will listen and transmit RIP updates.
- Inactive: The router port will not be listening or transmitting RIP updates.

7 Select the default frame type for the frames that will be generated by this router port and propagated over this VLAN to the outbound ports. Set the framing type to the encapsulation type that is most prevalent in this VLAN. If this VLAN contains devices using encapsulation types other than those defined here, the MPM must translate those frames, which slows down throughput:

- Ethernet Version II
- FDDI
- Token-ring
- Ethernet 802.3
- Source route token-ring

8 Enabling IPX will activate the IPX protocol on the router port for the VLAN just defined. Typing n will bypass all IPX configurations.

9 IPX Network: An IPX network address contains eight hex characters. If less than eight characters are entered, the RouteSwitch will prefix your entry with 0s.

10 A brief description of this router port, up to 30 characters.

11 Router Internet Protocol (RIP) or Service Access Protocol (SAP) for this router port can be enabled.

The choices are:

- RIP and SAP active: The router will transmit and receive RIP and SAP.
- RIP only: The router will only transmit and receive RIP.
- RIP and SAP inactive: The IPX router port is active but does not participate in the RIP and SAP protocol.

12 Select the default frame type for the frames that will be generated by this router port and propagated over this VLAN to the outbound ports. Set the framing type to the encapsulation type that is most prevalent in this VLAN. If this VLAN contains devices using encapsulation types other than those defined here, the MPM must translate those frames, which slows down throughput.

Important

The frame type must match the frame type defined in the VLAN policy: network address rule, IPX.

The following IP protocols are supported by the virtual router:

- Routing Information Protocol (RIP)
- Internet Control Message Protocol (ICMP)
- User Datagram Protocol (UDP)
- Simple Network Management Protocol (SNMP)
- Transmission Control Protocol (TCP)
- File Transfer Protocol (FTP)

- TELNET
- Serial Line Internet Protocol (SLIP)
- Rlogin

5.4.1 RouteSwitch Handling of Router Ports

RouteTracker functions on the assumption that the data in a frame is associated with the source MAC address contained in the frame.

Frames coming from a router port may contain several different IP or IPX network numbers. SAP and RIP sent by the router port will be associated with the IP and IPX network defined on that router port.

Following this logic, RouteTracker would then assign the MAC address of the virtual router port to all VLANs that have a policy matching the data in the frames sent by the router.

RouteTracker analyzes the frame further and if one of the following types of frames is seen coming out of a MAC address, then RouteTracker flags this MAC address as a router port and removes all VLAN assignments for the MAC address of the virtual router port.

- IP protocol: RIP, OSPF, BGP4, DVRMP and IGRP
- IPX protocol: IPX RIP, SAP

5.5 Advanced Routing

In Release 3.x of the RouteSwitch software, more routing protocols are introduced. No additional hardware is required, but a separately orderable module of code called *Gated.img* must be loaded into the flash.

Gate Daemon (GateD) originated from a UNIX environment. It is developed, maintained, and licensed by Merit Network Inc.'s GateD Consortium and is a dynamically loadable program.

GateD supports multiple routing protocols, such as:

- OSPF
- RIP I and II
- EGP
- BGP
- SNMP

The advantages of GateD on the RouteSwitch, allow the integration of OSPF and RIP and RIP II routing. It is also an extensively tested, industry standard routing platform, and being modular new routing protocols can be added in the future.

For a comprehensive guide on the advanced routing function please see the *Nways RouteSwitch Advanced Routing User's Guide*, GC30-3965.

For GateD to run on the 8274, it requires both *gated.img* and *gated.conf* files in the 8274 flash at bootup. If *gated.conf* is not in the flash or does not compile at bootup, the advanced routing function will not be started.

Gated.conf (Configuration file) must be written in a text editor and downloaded as ASCII. Changes may be made by editing the file in the 8274's buffer. To run GateD it requires 16 MB DRAM and 4 MB Flash, it also requires V3.X.X of the RouteSwitch software.

5.5.1 Advanced Routing Example 1

In this example an 8274 and an MSS were configured with various IP subnets and the OSPF protocol was used to exchange routing information. A single backbone (0.0.0.0) OSPF area was created for all the routing updates between the 8274 and the MSS, as seen in Figure 108.

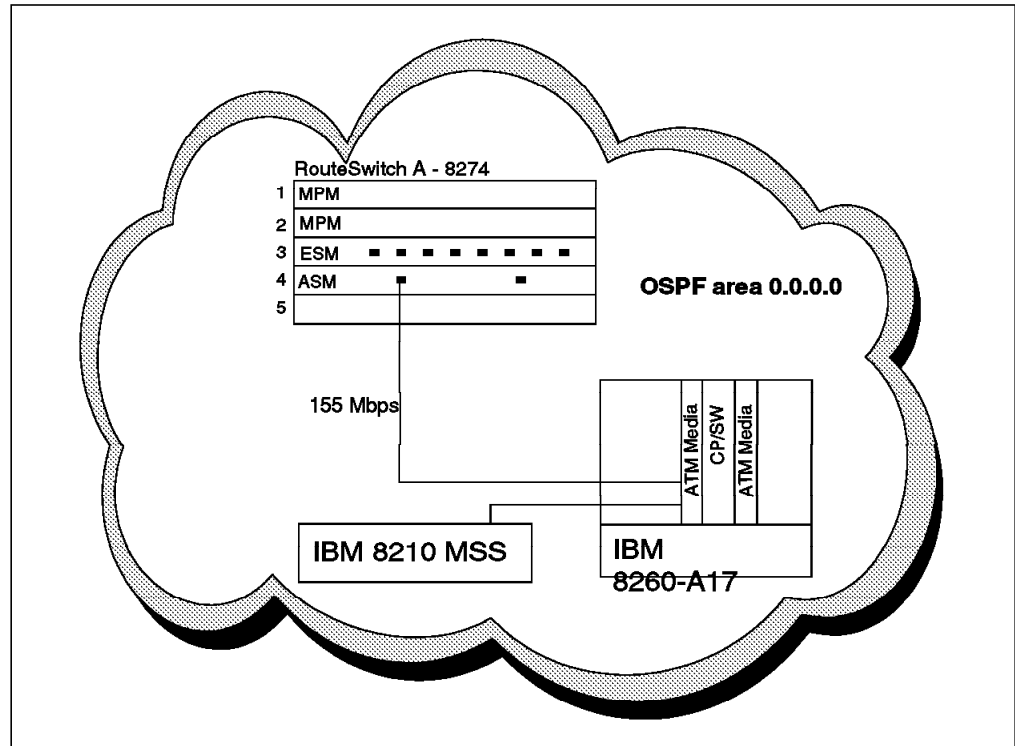


Figure 108. Physical View of Autonomous OSPF Area

An Ethernet LEC was set up on the 8274 to join an Ethernet ELAN configured on the MSS. In Figure 109 on page 132, a service on the ASM port 4/1 was modified to join this Ethernet ELAN. If you are running token-ring ELANs, then a token-ring LEC should be configured instead of an Ethernet LEC.

```

8274/Networking/Gated % mas 4/1 2

Slot 4 Port 1 Service 2 Configuration

1) Description (30 chars max)           : LAN Emulation Service 2
2) LAN Emulated Group                   : 1
   21) LAN type { 802.3 (1),
                  802.5 (2) }           : 802.3
   22) Change LANE Cfg { NO (1),
                        YES (2) }       : NO
3) LECS Address (40-char-hex)           :
   4700790000000000000000000000A03E00000100
4) Admin Status { disable(1),
                  enable(2) }           : Enable
6) Connection Type { PVC(1),
                    SVC(2) }            : SVC
   60) SEL for the ATM address           : 02

Enter (option=value/save/cancel) : 22=2

Slot 4 Port 1 Service 2 LANE Configuration Parameters

1) Proxy { NO (1), YES (2) }            : YES
2) Max Frame Size { 1516 (1), 4544 (2),
                    9234 (3), 18190 (4) } : 1516
3) Use translation options{NO (1), YES (2) : Yes
   (use Swch menu to set)
4) Use Fwd Delay time { NO (1), YES (2) } : NO
5) Use LE Cfg Server (LECS){ NO (1), YES (2)}: YES
6) Use Default LECS address { NO(1), YES (2)}: YES
7) Control Time-out (in seconds)         : 10
8) Max Unknown Frame Count               : 10
9) Max Unknown Frame Time (in seconds)   : 1
10) VCC Time-out Period (in minutes)     : 20
11) Max Retry Count                     : 2
12) Aging Time (in seconds)              : 300
13) Expectd LE_ARP Resp Time (in seconds) : 1
14) Flush Time-out (in seconds)          : 4
15) Path Switching Delay (in seconds)    : 6
16) ELAN name (32 chars max)             : eth1

Enter (option=value/save/cancel) : save

```

Figure 109. Modify the Service on the ASM Port to Join the eth1 ELAN

Viewing the ASM ports in slot 4 and the services configured on the ports is shown in Figure 110 on page 133.


```

8274/ % viatr1
VLAN  VLAN  Rule Rule      Rule      Rule
Group: Id   Num Type      Status    Definition
-----
      1: 2      1 NET ADDR RULE Enabled  IP Addr = 192.168.7.0
                                IP Mask = 255.255.255.0
                                2 PORT RULE    Enabled  4/1/Lne/1

      1: 3      1 NET ADDR RULE Enabled  IP Addr = 192.168.22.0
                                IP Mask = 255.255.255.0
                                2 PORT RULE    Enabled  4/1/Lne/1

      1: 4      1 PROTOCOL RULE Enabled  Protocol = IPX
      1: 5      1 NET ADDR RULE Enabled  IP Addr = 192.168.5.0
                                IP Mask = 255.255.255.0
                                2 PORT RULE    Enabled  4/1/Lne/1

      1: 6      1 NET ADDR RULE Enabled  IP Addr = 10.0.1.0
                                IP Mask = 255.255.240.0
                                2 PORT RULE    Enabled  4/1/Lne/1

8274/ %

```

Figure 111. Viewing VLANs Defined

The gated.conf file that was used is shown in Figure 112 on page 135. Loading the gated.conf file into the buffer on the 8274 allows for it to be edited and then saved back from the buffer to the flash. The switch will then have to be rebooted for changes to take effect.


```

8274/File % edit
8274/File/Edit % ?

Command      Edit Menu
-----
ab           Append line(s) to the buffer
cb           Clear the buffer
db           Delete line from the buffer
eb           Edit a buffer line
ib           Insert buffer line
lb           List contents of the buffer
nb           Name file for buffer
rb           Read file into buffer
wb           Write buffer to file

Main      File      Summary  VLAN      Networking
Interface Security System  Services Help
8274/File/Edit % cb

8274/File/Edit % rb gated.conf

8274/File/Edit % lb
00: traceoptions state ;
01: routerid 9.24.105.99 ; 1
02: ospf yes { 2
03:     backbone { 3
04:         authtype none ;
05:         interface all { 4
06:             priority 1 ;
07:         } ;
08:     } ;
09: } ;
10: export proto ospfase { 5
11:     proto ospfase {
12:         all
13:         metric 1 ; } ;
14:     proto static {
15:         all
16:         metric 1 ; } ;
17:     proto direct {
18:         all
19:         metric 1 ; } ;
20:     proto rip {
21:         all
22:         metric 1 ; } ;
23: } ;

Work buffer is unnamed
8274/File/Edit % cb

```

Figure 112. Viewing the GateD.conf File

- 1** The router ID is used to identify this router in routing tables throughout the network. It is advisable to make it unique, and can consist of numeric and alphanumeric characters.
- 2** This enables the OSPF protocol.
- 3** The OSPF area defined is 0.0.0.0 (backbone). Depending on individual requirement multiple OSPF areas can exist in a single 8274. Further configuration will be required in order to match the OSPF areas to interfaces.
- 4** All interfaces have been included in the backbone area.

5 In this section of the `gated.conf` file we specify all the protocols and learned routes that must be advertised and the interfaces that they will be advertised on.

In Figure 112 on page 135, the `gated.conf` file has been indented to make it more legible. Although using indentation makes the file easier to read it is not a requirement in writing the `gated.conf` file. Writing the file without indentation would look something like this:

```
traceoptions state ;
routerid 9.24.105.99 ;
ospf yes {
backbone {
authtype none ;
interface all {
priority 1 ;
} ;
} ;
} ;
.
.
.
```

Figure 113 on page 137 shows how the protocols defined in the advanced routing setup can be viewed.

```

8274/ % gated
parameters for gated root command are:

          stat          GateD status display
          routes        Displays GateD internal routing table
          reconfig       Reloads configuration file

8274/Networking/Gated % gated stat

Gated is running

RIP:    not configured
OSPF:   enabled
ICMP Router Discovery: not configured

8274/Networking/Gated % gated routes

          ***** Gated Routing Table *****
Destination  Subnet Mask      Gateway      Metric    Tag    Pref  Protocol
=====
9.24.105.0   255.255.255.0    9.24.105.99    0         0        0   Direct
10.0.0.0     255.255.240.0    10.0.1.1       0         0       120   Direct
127.0.0.0    255.0.0.0        127.0.0.1      0         0        0   Static
127.0.0.1    255.255.255.255  127.0.0.1      0         0        0   Direct
192.168.5.0  255.255.255.0    192.168.5.200  0         0        0   Direct
192.168.7.0  255.255.255.0    192.168.7.200  0         0        0   Direct
192.168.20.10 255.255.255.255  192.168.5.10   1         0       10   OSPF
192.168.21.10 255.255.255.255  192.168.5.10   1         0       10   OSPF
192.168.22.0  255.255.255.0    192.168.22.200 0         0       120   Direct

8274/Networking/Gated % ospf stat
          ***** OSPF Status *****

OSPF running with 1 neighbor(s), Router ID = 9.24.105.99

Hello Packets          TX:1865 Hello Packets          RX:351
Database Description packets TX: 10 Database Description packets RX: 6
LSR packets            TX: 22 LSR packets            RX: 2
LSU packets            TX: 0  LSU packets            RX: 7
LSA packets            TX: 6  LSA packets            RX: 19

8274/Networking/Gated % ospf nl
          ***** Neighbours List *****
Area      IP Address      Rtr Id      Priority    State    (M/S)
=====
0.0.0.0   192.168.5.10    192.168.20.10 1          Full    (S)

8274/Networking/Gated % ospf al
          ***** Area List *****

ASBDRs = Autonomous System Border Routers
ABDRs  = Area Border Routes

Area      Transit/Stub SPF runs ASBDRs  ABDRs  LSA count  Full neighbors
=====
0.0.0.0   Neither      9       1         0         3         1

8274/Networking/Gated %

```

Figure 113. Viewing OSPF Statistics

The configuration steps that were required on the MSS are also shown. Figure 114 on page 138 illustrates the IP addresses for each interface, which are not automatically added with an OSPF interface.

```

MSS_A *t 6
MSS_A Config>p ip
Internet protocol user configuration

MSS_A IP config>li add

IP addresses for each interface:
  intf 0 192.168.20.10 255.255.255.0 Local wire broadcast, fill 1
          192.168.21.10 255.255.255.0 Local wire broadcast, fill 1
  intf 1 192.168.3.10 255.255.255.0 Local wire broadcast, fill 1
  intf 2 192.168.2.10 255.255.255.0 Local wire broadcast, fill 1
  intf 3                                     IP disabled on this interface
  intf 4                                     IP disabled on this interface
  intf 5                                     IP disabled on this interface
  intf 6                                     IP disabled on this interface
  intf 7 192.168.5.10 255.255.255.0 Local wire broadcast, fill 1
  intf 8 9.24.104.115 255.255.255.0 Local wire broadcast, fill 1
          192.168.4.10 255.255.255.0 Local wire broadcast, fill 1
  intf 9 192.168.207.10 255.255.255.0 Local wire broadcast, fill 1

MSS_A IP config>ex
MSS_A Config>p ospf
Open SPF-Based Routing Protocol configuration console

MSS_A OSPF Config>li int

--Interface configuration--
IP address      Area          Cost  Rtrns  TrnsDly  Pri  Hello  Dead
192.168.20.10   0.0.0.0          1      5       1        1    10     40
192.168.21.10   0.0.0.0          1      5       1        1    10     40
192.168.3.10    0.0.0.0          1      5       1        1    10     40
192.168.2.10    0.0.0.0          1      5       1        1    10     40
192.168.5.10    0.0.0.0          1      5       1        1    10     40

```

Figure 114. Viewing IP and OSPF Interfaces on the MSS

Figure 115 on page 139 shows how to add an IP interface so that it will run the OSPF protocol. This is confirmed when we list the OSPF interface configuration table.

```

MSS_A *t 6

MSS_A OSPF Config>set int
Interface IP address (0.0.0.0)? 192.168.207.10
Attaches to area (0.0.0.0)?
Retransmission Interval (in seconds) (5)?
Transmission Delay (in seconds) (1)?
Router Priority (1)?
Hello Interval (in seconds) (10)?
Dead Router Interval (in seconds) (40)?
Type Of Service 0 cost (1)?
Authentication Key ()?
Retype Auth. Key ()?
MSS_A OSPF Config>

MSS_A OSPF Config>li int

--Interface configuration--
IP address      Area          Cost  Rtrns  TrnsDly  Pri  Hello  Dead
192.168.20.10   0.0.0.0       1     5      1        1    10     40
192.168.21.10   0.0.0.0       1     5      1        1    10     40
192.168.3.10    0.0.0.0       1     5      1        1    10     40
192.168.2.10    0.0.0.0       1     5      1        1    10     40
192.168.5.10    0.0.0.0       1     5      1        1    10     40
192.168.207.10  0.0.0.0       1     5      1        1    10     40

MSS_A OSPF Config>

```

Figure 115. Configuring the OSPF Interface

Figure 116 on page 140 shows a dump of OSPF routes that are defined on the MSS and also learned from elsewhere in the network.

```

MSS_A *t 5

CGW Operator Console

MSS_A +
MSS_A +p os
Open SPF-Based Routing Protocol console
MSS_A OSPF>dump
Type  Dest net      Mask      Cost      Age      Next hop(s)

Stat* 0.0.0.0        00000000  1         3525     192.168.20.60
Sbnt  9.0.0.0        FF000000  1         1366     None
Dir*  9.24.104.0      FFFFFFFF  1         1427     TKR/3
SPF   9.24.105.0      FFFFFFFF  2         1386     192.168.5.200
Sbnt  10.0.0.0         FF000000  1         1366     None
SPF   10.0.0.0         FFFFFFFF  2         1386     192.168.5.200
Dir*  192.168.4.0      FFFFFFFF  1         1427     TKR/3
SPF*  192.168.5.0      FFFFFFFF  1         1427     Eth/3
SPF   192.168.7.0      FFFFFFFF  2         1143     192.168.5.200
Dir*  192.168.20.0     FFFFFFFF  1         3496     ATM/0
SPF   192.168.20.10   FFFFFFFF  0         3525     ATM/0
Dir*  192.168.21.0     FFFFFFFF  1         3496     ATM/0
SPF   192.168.21.10   FFFFFFFF  0         3525     ATM/0
SPF   192.168.22.0     FFFFFFFF  2         1386     192.168.5.200
Dir*  192.168.207.0    FFFFFFFF  1         1427     Eth/3

Default gateway in use.
Type Cost      Age      Next hop
Stat 1         3525     192.168.20.60

Routing table size: 768 nets (49152 bytes), 14 nets known

```

Figure 116. Viewing the OSPF Routing Table

5.5.2 Advanced Routing Example 2

This example shows two RouteSwitches that are connected together in an OSPF autonomous area. There are multiple OSPF areas as well as RIP II running on the 8274s, as is shown in Figure 117 on page 141. Extracts from the gated.conf files from both switches are listed to show the required parameters that were used.

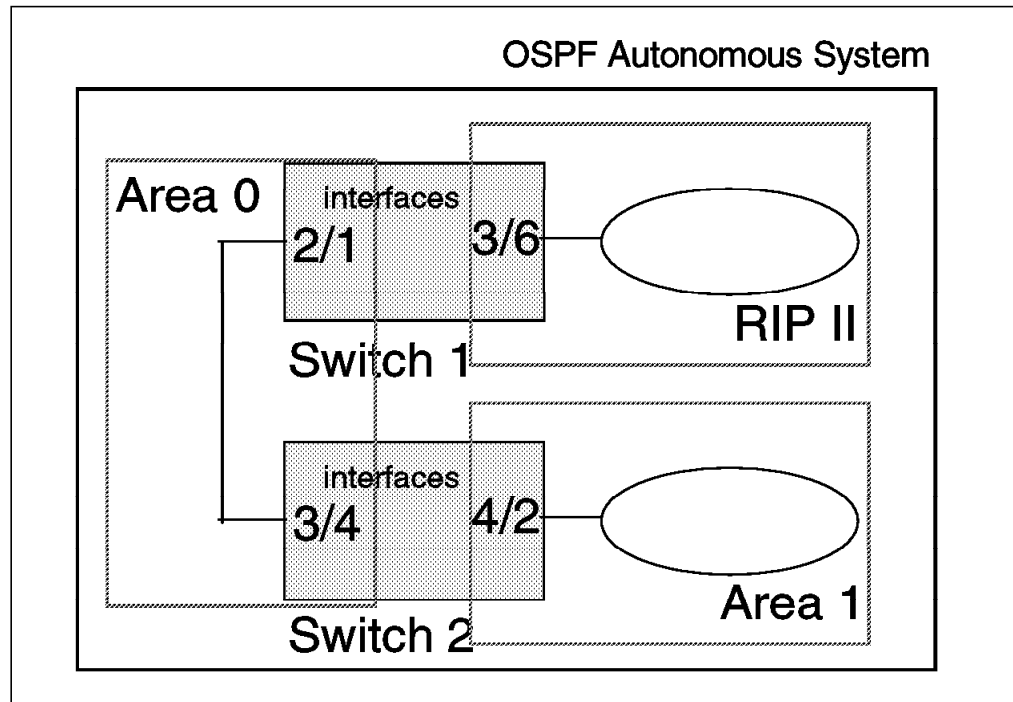


Figure 117. Physical View of Autonomous OSPF Area

In Appendix C, “Sample GATED Configuration File” on page 345, a detailed sample gated configuration file lists all the parameters that can be used when writing the gated.conf file. In Figure 118, the protocol section in the gated.conf file for switch 1 shows RIP II and OSPF running in the same RouteSwitch.

```
.
.
.
rip yes {
    interface 3/6 version 2 ; 1
};
ospf yes {
    backbone { interface 2/1; 2
    };
};
export proto ospfase interface 2/1 { 3
    proto rip interface 3/6 {
        9.24.105.0 mask 255.255.255.0 ;
    };
};
export proto rip interface 3/6 { 4
    proto ospfase interface 2/1 ;
};
.
.
.
```

Figure 118. Extract of gated.conf File on Switch 1

- 1** RIP II is enabled on port 3/6.
- 2** OSPF area 0 (backbone) is enabled on port 2/1.
- 3** The RIP II routing information that is learned on port 3/6 about the IP subnet 9.24.105.0 is exported as OSPF routing information on port 2/1.

4 The OSPF routing information that is learned on port 2/1 is exported in RIP II format on port 3/6.

An extract from the gated.conf file for switch 2 is shown in Figure 119.

```
.  
. .  
rip no ;  
ospf yes {  
    backbone { interface 3/4 ; 1  
    } ;  
    area 1 { interface 4/2 ; 2  
    } ;  
} ;  
. .  
.
```

Figure 119. Extract of gated.conf File on Switch 2

1 OSPF is enabled on port 3/4 with OSPF area 0.

2 OSPF is enabled on port 4/2 with OSPF area 1.

5.6 Next Hop Resolution Protocol (NHRP)

Next Hop Resolution Protocol (NHRP) is an IETF protocol that can improve network performance by eliminating router hops on Non-Broadcast Multi-Access (NBMA) subnetworks such as ATM. When internetworking protocols such as IP are run over NBMA subnetworks, the routed path through the network can include multiple hops across the same subnetwork.




Figure 120. Conventional Routing over ATM

As an example, consider the network shown in Figure 120, where the ATM subnetwork has been partitioned into IP subnets, and ATM is used primarily as a high-speed data link technology (that is, a fat pipe) that provides connectivity between hosts and routers. Packets from Host x to Host y encounter three router hops and traverse the ATM subnetwork four times: (1) Host x to Router 1, (2) Router 1 to Router 2, (3) Router 2 to Router 3, and (4) Router 3 to Host y. NHRP was designed to better utilize the capabilities of the underlying switched infrastructure in configurations such as this by enabling the establishment of short-cut routes across ATM subnetworks. The short-cut may be directly to the destination or, if that is not possible, to the egress router nearest to the destination.

NHRP is a client/server protocol. NHRP clients (NHCs) issue requests to NHRP servers (NHSs), and NHSs either respond to the requests or forward the requests along the routed path.

There are also cases where an NHS will respond to NHRP resolution requests with an ATM address associated with its co-located router. One such case is when the NHS is serving as the egress router to destinations that are not resident on the ATM subnetwork. Figure 121 on page 144 provides an example where the destination (IP Host y) is behind a router (Router 4) that is connected to the egress NHS (NHS 3) via an FDDI LAN.




Figure 121. Shortcut to an Egress Router

Protocol address-to-ATM address mappings provided in NHRP resolution responses include a holding time that indicates how long the mapping information may be considered valid. When the holding time expires, the associated short-cut can no longer be used. Thus, the MSS server's NHC attempts to refresh mappings, provided that they are still being used, before the holding time expires in order to avoid lapses in short-cut service. NHCs include the holding time when registering their (protocol and ATM) addresses with an NHS.

Since NHRP resolution requests follow the routed path, NHSs are coupled with routers, and routers without an NHS can discontinue NHRP reachability. In general, short-cuts are not possible when there is an intermediate router that does not support NHRP.

5.6.1 ELAN Support

NHRP packets that are to follow the routed path are forwarded over ELANs as well as Classical IP subnets. While this seems like the natural thing to do, and certainly is efficient, since existing VCCs are utilized, other implementations may only send NHRP packets over VCCs employing RFC 1483 encapsulation. The NHRP specification is unclear in this regard, specifying that NHRP packets must follow the routed path, but referring only to RFC 1483 when defining ATM encapsulations. (The ambiguity arises because the routed path may be over an ELAN, but LANE Version 1.0 does not employ RFC 1483 encapsulation.) In addition to the superfluous connections, implementations that do not utilize ELAN

VCCs may impose the unnecessary configuration burden of requiring users to specify the ATM addresses to be used when establishing these RFC 1483 VCCs.

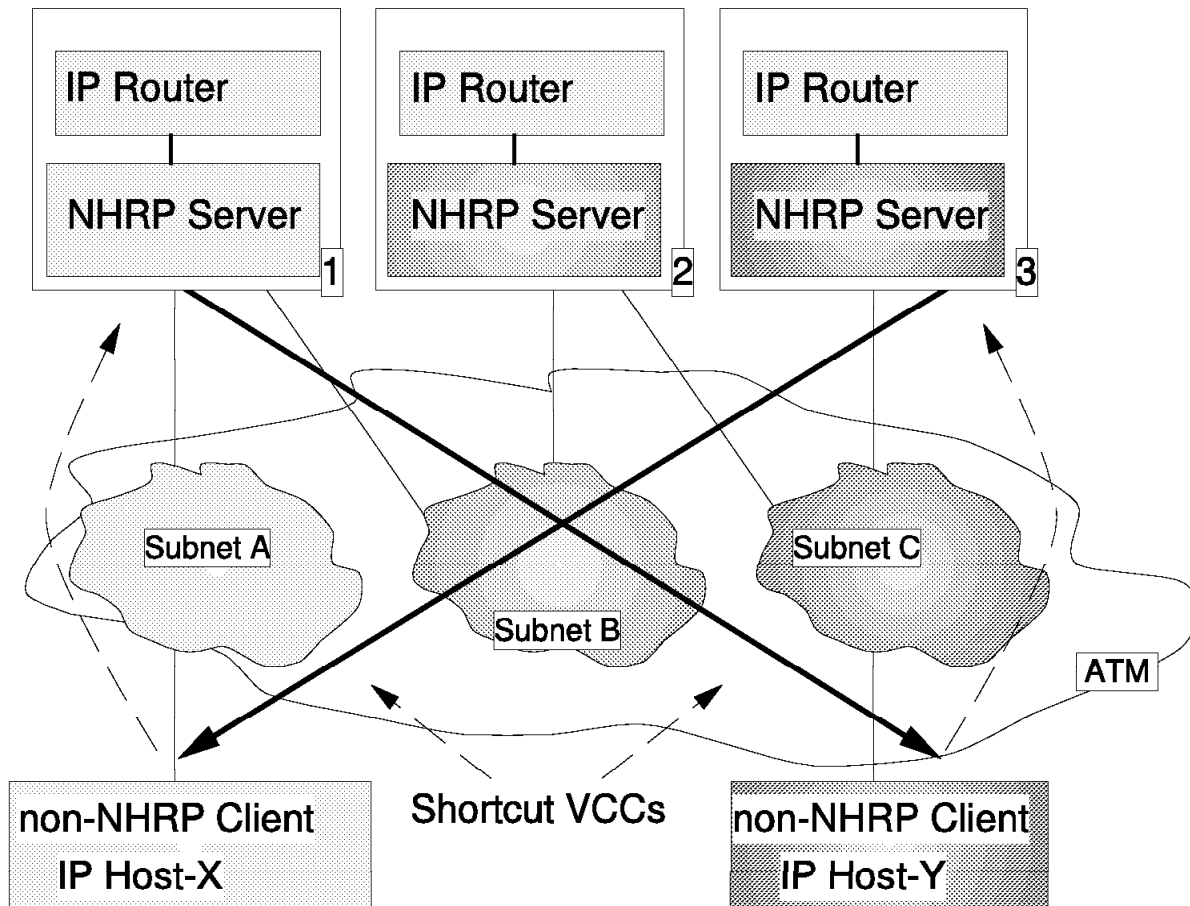


Figure 122. One-Hop Routing

One-hop routing with LANE short-cuts is illustrated in Figure 123 on page 146. In this example, both source and destination IP stations reside on legacy LANs behind ATM-attached LAN switches. MSS Server 1 initiates a short-cut to IP Host y on behalf of IP Host x, and includes the LANE short-cut extensions in the NHRP resolution request. The request is sent to MSS Server 2 over IP Subnet 9.1.2, which may be an ELAN or a LIS. MSS Server 2 recognizes that the destination is accessible via a local ELAN subnet (that is, 9.1.3), and fills in the extensions with the MAC address of IP Host y and the ATM address of the LEC representing IP Host y in LAN Switch 2. MSS Server 2 learns the MAC address of IP Host y through IP ARP procedures, and the associated ATM address is obtained via LE_ARP procedures. When the NHRP resolution reply is returned, MSS Server 1 establishes a short-cut LANE data direct VCC to LEC 2, and subsequent traffic for IP Host y is transmitted directly over this VCC. Shortcut traffic destined for other hosts represented by LEC 2, such as IP Host z, will also be transmitted over this VCC.

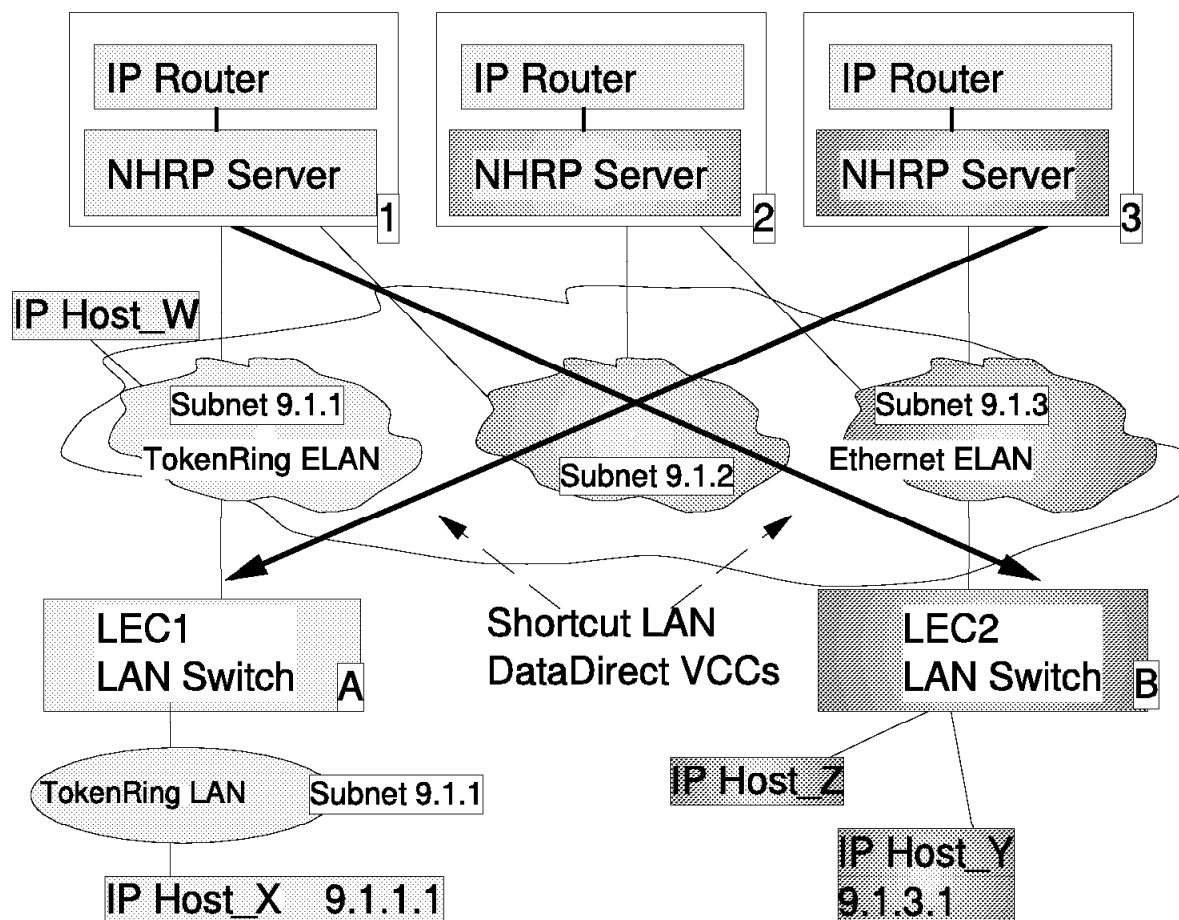


Figure 123. One-Hop Routing with LANE Shortcuts

NHS 2 establishes a short-cut for traffic to IP Host x in a similar manner. However, in this case, the LANE short-cut extensions also include the token-ring routing information field (RIF) necessary for sending traffic to IP Host x, which resides behind a downstream source route bridge. Although both hosts reside on legacy LANs, the short-cut routing mechanisms are equivalent for ATM-attached hosts with LANE interfaces, such as IP Host w.

The 8274 NHRP support is for IP traffic over ATM only. This allows for direct connections to be established to LAN Emulation clients and to Classical IP (CIP) (RFC 1577) clients.

The 8274 is also only an NHRP client and thus requires an NHRP server in the network in order for NHRP to run. The IBM MSS is used for the NHRP server function.

5.6.2 NHRP Example 1

A simple configuration, shown in Figure 124 on page 147, was used to test if the 8274 NHRP client would work with the MSS NHRP server, which according to all documentation it should. A logical representation of the network is shown in Figure 125 on page 147.

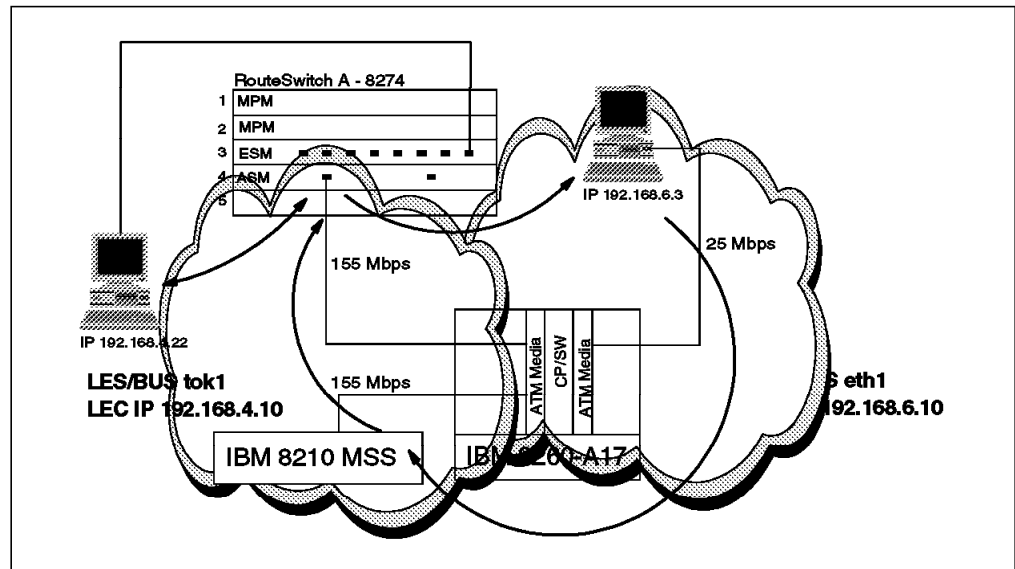


Figure 124. Physical View of Network - NHRP Example

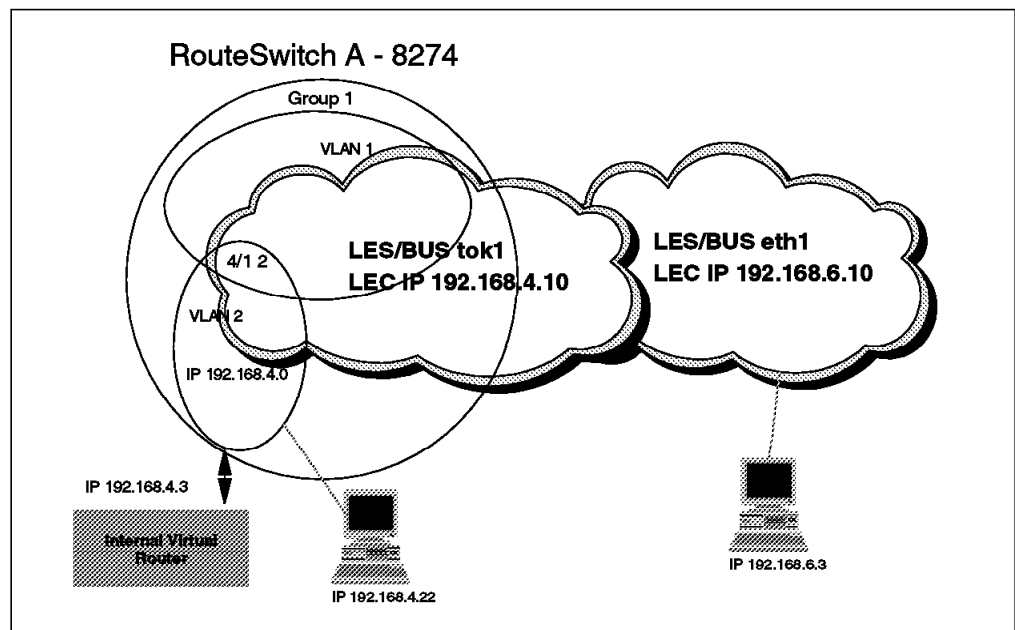


Figure 125. Logical View of Network - NHRP Example

5.6.2.1 NHRP Example 1 Configuration

A token-ring LEC was defined on port 4/1 service number 3, and the default Ethernet LEC that is defined by factory defaults on all ASM ports was deleted. Viewing the services for the ASM ports in Figure 126 on page 148, shows the result of this configuration.

```
8274/ >vas
```

ATM Services

Slot	Port	Serv Num	Service Description	Service Type
4	1	2	VCM Service 2	VCM
4	1	3	LAN Emulation Service 3	802.5 LEC
4	2	1	LAN Emulation Service 1	802.3 LEC
4	2	2	VCM Service 2	VCM

ATM Services

Slot	Port	Serv Num	VC Typ	Oper Status	SEL Groups	Conn VCI's/Addresses
4	1	2	SVC	Enabled	02 1	
4	1	3	SVC	LANE Op.	03 1	568 569 570 571
4	2	1	SVC	Initial	01 1	
4	2	2	SVC	Disabled	02 1	

FDDI Services do not exist!

Mammoth Ethernet Services do not exist!

```
8274/ >
```

Figure 126. Token-Ring LEC Is Configured on Port 4/1 3 to Join tok1

The RouteTracker VLANs that we defined in Group 1 are shown in Figure 127. It is important to note that although the service that was created on port 4/1 to join the token-ring ELAN was service 3 on port 4/1, when adding that port and service to VLAN 2 it shows that we have added 4/1/Lne/2. This is because it is the second sequential service that is running on port 4/1 and when it is added to the VLAN we can only choose a service that is defined and running on an ASM port. In this case service 3 on port 4/1 was the second sequential service running on ASM port 4/1.

```
8274/ >viatr1
```

VLAN Group:	VLAN Id	Rule Num	Rule Type	Rule Status	Rule Definition
1: 2		1	NET ADDR RULE	Enabled	IP Addr = 192.168.4.0 IP Mask = 255.255.255.0
		2	PORT RULE	Enabled	4/1/Lne/2
		3	PORT RULE	Enabled	3/12/Brg/1
1: 3		1	NET ADDR RULE	Enabled	IP Addr = 192.168.3.0 IP Mask = 255.255.255.0

```
8274/ >
```

Figure 127. VLAN Configuration

During RouteTracker VLAN configuration we are prompted to enable IP, it is during this IP configuration that we are asked if NHRP should be enabled. If yes

is selected to enable NHRP, then NHRP is enabled for that VLAN. This is shown in Figure 128 on page 149.

```

Enable IP? (y):
  IP Address                : 192.168.4.3
  IP Subnet Mask             (0xffffffff) : 255.255.255.0
  IP Broadcast Address (192.168.4.255) :
  Description (30 chars max) :
  Disable routing?          (n) :
  Enable NHRP?              (n) : y
  IP RIP mode {Deaf(d),
    Silent(s),
    Active(a),
    Inactive(i)}            (s) : a
  Default framing type {Ethernet II(e),
    fddi(f),
    token ring(t),
    Ethernet 802.3(8),
    source route token ring(s)} (e) : t
Created router port for vlan 1: 2
Enable IPX? (y): n

```

Figure 128. VLAN Configuration

The Ethernet workstation with IP address 192.168.4.22 on ESM port 3/12 then sent a ping to the ATM-attached workstation with IP address 192.168.6.3 and member of the eth1 ELAN. The statistics shown in Figure 129, Figure 130, and Figure 131 on page 150, were then seen.

```

8274/Networking/NHRP >vinvl
VLANs with NHRP enabled
=====
gp:vl   Router IP @      State   Rtr  CIP?
-----
  1:3   192.168.3.3      Inactive 252  No
  1:2   192.168.4.3      Active  253  No
=====
8274/Networking/NHRP >

```

Figure 129. NHRP Enabled VLANs

```

8274/Networking/NHRP >vinrc

NHRP Resolution Cache Entries
=====
Dest Address NextHop Address State HTime MTU  VCM Rtr VCI  Type
-----
192.168.6.3  0.0.0.0      Samp  253  0   0   253  0   unresolved
=====
8274/Networking/NHRP >

```

Figure 130. NHRP Cache Entries

```

8274/Networking/NHRP >vinrce

Enter destination address (0.0.0.0) : 192.168.6.3
=====
Destination:    192.168.6.3
PrefixLength:   0xFF
NextHop:        unknown
NbmaAddress:    unknown
VCI:            0
Shortcut Type:  unresolved
EntryType:      7 (other)
EntryState:     4 (Sampling)
MTU size:       0
IfIndex:        0
Holding time:   227
VCM port:       0
Router port:    253
Flags:          0x01000000
=====

8274/Networking/NHRP >

```

Figure 131. NHRP Cache Entry Detail

5.7 Mixed Media

One way for data to get from one media type to another is via routing. Another way is using an SR-TB bridge. The router removes the media-specific headers of a received frame and appends the new media-specific header for the network type of the destination port. In this process the frame itself is not transmitted from one media type to another, only the information within it.

Routing involves heavy computation, requiring table lookups to guide this header deletion/creation and additional router-to-router protocols to set up and maintain these tables.

Routing is not restricted, nor even primarily intended for moving data between different types of media but instead seeks to break networks down into a number of smaller networks, each of which is a broadcast domain. Historically, networks based on different technologies and media naturally form distinct broadcast domains.

However, the advent of LAN switching is rewriting these rules. Today, the formation of broadcast domains and the allocation of devices to them is driven by logical requirements such as virtual LANs and LAN switches. They seek to break free of topology and network constraints imposed by mere media differences.

5.7.1 Any-to-Any Switching

Because the RouteSwitch is a LAN switch that carries frames from multiple media types on its backplane fabric, it offers the facility to switch frames from any media to any other media for example, an Ethernet frame onto a token-ring frame. This feature is referred to as any-to-any switching.

Within this new paradigm there is still a place for routing. The installed base of clients and servers must communicate by established routing protocols, but the broadcast domains handled by a router are no longer a restraint to a single media type.

To support this paradigm, a LAN switch must “transform” a frame of one media type into a frame of another media type in such a way that the frame is still acceptable to the routing protocols.

5.7.1.1 Any-to-Any Switching Is Protocol-Specific

Unfortunately, the requirements for this transformation algorithm are specific to the various protocols that currently exist.

There is no single, simple algorithm that will allow the frames to be switched transparently between media types to the higher level protocols and frame formats.

This leads to a fairly complex set of configuration options and limitations on the applicability of the any-to-any switching features.

In order to understand why these options and limitations arise and to better understand the configuration options available, it is advisable to understand, as background, the theory of operation for any-to-any switching. This material also may help to determine the applicability of any-to-any switching for a protocol not described in the reference material.

5.7.2 Theory of Operation

Every frame, of any protocol, on any media type, consists of the following parts:

1. **MAC Header:** This consists of a source and destination address (MAC addresses) specifying the transmitting station and the intended recipient(s), as well as other media-specific fields for example, Access Control (AC) and Field Control (FC) fields in token-ring, FC in FDDI, etc.
2. **RIF:** Optional field defined by the source routing standard and is only found on token-ring and FDDI media.
3. **Encapsulation:** This is defined by the various standards for the media, many of which reference common standards. For example, on Ethernet media, as defined by Ethernet II, this is a 16-bit type field. On Ethernet media, as defined by the IEEE 802.3 committee, this is a length field together with any encapsulation defined by the IEEE 802.2 logical link control (LLC) committee. On token-ring and FDDI, it is any encapsulation defined by the IEEE 802.2 LLC committee.
4. **Network Header:** This is defined by the organization responsible for the particular routing protocol whose data is being carried within the frame. The values of fields defined in the encapsulation area allow the recipient to identify which protocol standard to use to decode the network header part of the frame.
5. **Data:** The payload being carried between the endstations.

In a routing implementation, the first three fields are the ones stripped and rebuilt when the frame is forwarded. These are the three areas that have to be manipulated.

Each of these areas and their interactions will be examined to see the media and protocol dependencies. As these are understood, the reasons why any-to-any switching requires protocol-specific switch support will become apparent.

5.7.2.1 The MAC Header

The format and values defined for this area are covered in the media standards, but here we show the choices that are dictated by the upper layer protocol:

- Canonical versus non-canonical

The first requirement of the switch is to transform the bit ordering of the address fields. Token-ring and FDDI use the non-canonical format or most significant bit (MSB) first. Ethernet uses the canonical format or least significant bit (LSB) first. Thus when a frame is moved between these media types, the addresses must be bit-swapped.

- Abbreviated addresses

Normally, a MAC address is 48 bits long. FDDI and 802.5 token-ring media allows for the use of small 16-bit MAC addresses.

The RouteSwitch only supports a 48-bit MAC address. Abbreviated address-based protocols are not supported.

- Functional addresses and multicasts

The IEEE 802.5 media also have different rules for the formation of multicast addresses or group addresses. In Ethernet, a single bit defines the address as a multicast. In 802.5, a single bit also indicates a multicast, but the remaining bits are structured into functional address groups with pre-assigned meanings and functions.

The RouteSwitch does not map multicasts and functional addresses, thus protocols dependent on these features may not be switchable any-to-any.

5.7.2.2 The RIF Field

The same source routing standard is supported by FDDI and token-ring so the RIF fields can be switched without problems between these media.

Ethernet does not support source routing, thus token-ring frames with RIF fields cannot be switched onto the Ethernet media. RIF stripping can be enabled by entering the following command in the mpm.cmd: `rifStripping=1`. RIF stripping only works when going to and from token-ring and Ethernet. The token-ring ports must be set to SRT and the RIF field must be 2 bytes long.

The alternative is stripping the RIF fields, remembering them for each MAC address and reinserting them on replies to terminate a source routed connection and act as a proxy to a transparent device (such as SR-TB on 8209/8229). This is not well standardized and is difficult to execute and manage. SR-TB bridging is not supported by the RouteSwitch.

Ethernet frames are allowed onto a token-ring if the token-ring supports transparent bridging; for example, the port is configured as either transparent (TB) or source route transparent (SRT).

Token-ring frames with an RIF will not be switched to Ethernet when the token-ring port is configured for SRT.

Token-ring frames without RIF will be switched to Ethernet when the token-ring port is configured for SRT. Also all communication between SR configured token-ring ports and transparent Ethernet ports is barred.

5.7.2.3 Encapsulation

Encapsulation is the biggest problem for implementing a transformation algorithm in support of any-to-any switching. All media types provide a choice of more than one encapsulation and not all encapsulations are available on all media. Additionally, the methodology of these encapsulations vary from protocol to protocol.

An ideal protocol would dictate a single encapsulation which would be the same on all media.

Most protocols make use of more than one encapsulation. For example, IP uses Ether-Type most of the time on Ethernet and SNAP on FDDI and token-ring. In this case, there may be clearly established rules for transforming frames from one encapsulation to another as media types are traversed.

Some protocols may allow more than one encapsulation even on a single media type. Some might use the encapsulation to separate functional parts of the protocols, for example, routing table updating protocols from user data forwarding protocols. Others, such as IPX may simply allow the user to arbitrarily choose the encapsulation method.

Some, most notably IPX, may entangle the notion of encapsulation in with the notion of the network level broadcast domain to create multiple logical networks over a single physical broadcast domain.

Clearly, there is no single algorithmic rule by which the any-to-any transformation function can switch arbitrary protocols.

There are two choices available to address this situation:

1. The switch must be configurable, per device, per protocol, per media to select the transformation of encapsulations.
2. The switch performs a single transformation and the user must configure all endstations and routers to use this single choice made by the switch.

The RouteSwitch uses the first approach for IP and IPX as they are the most dominant protocols in the market. It uses the second approach for all other protocols.

Protocols Other Than IP and IPX: On Ethernet media, three encapsulations are possible:

- Ether-Type
- IEEE 802.2 LLC
- IEEE 802.2 SNAP

Note: SNAP is an instance of an LLC encapsulation defined by the IEEE 802.2 committee to support the transformation of Ether-Type Ethernet frames to media that do not support that encapsulation.

On token-ring and FDDI, two encapsulations are permitted by the standards:

- IEEE 802.2 LLC

- IEEE 802.2 SNAP

The SNAP Conversion: The intent of the 802.2 committee is that Ether-Type frames are transformed to SNAP on crossing from Ethernet media to 802 media and restored to Ether-Type in the reverse direction.

The RouteSwitch could follow this rule for all protocols including IP; however, this would prevent AppleTalk from inter-working between Ethernet and FDDI. The RouteSwitch explicitly checks for the AppleTalk protocol. If found, the rule is not applied. In addition, the RouteSwitch checks for the Banyan Vines protocol and translates according to the media type.

As there may be other protocols with this problem, the SNAP to Ether-Type transformation is configurable for all protocols other than AppleTalk.

Other Conversions: There are no equivalent algorithmic approaches that the transformation function can adopt for dealing with protocols that require Ether-Type on Ethernet and some form of LLC encapsulation on FDDI and/or token-ring. The mapping between Ether-Type values and LLC values is arbitrary requiring tables indexed by protocol.

The approach followed in the RouteSwitch is therefore to simply pass LLC encoding between Ethernet, FDDI and token-ring with no changes other than to insert/strip the length field required by IEEE 802.3 on Ethernet. This leaves protocols that require transformations between Ether-Type and LLC encapsulations as unswitchable unless the clients and servers can be configured to use SNAP.

Summary of Non-IPX Encapsulation Transformation Rules: Ether-Type/SNAP transformations are configurable on an all protocols, except for AppleTalk and Banyan Vines. Ether-Type frames going to FDDI or token-ring are translated to SNAP unconditionally.

SNAP frames going to Ethernet are translated to Ether-Type or left as SNAP as per the RouteSwitch configuration, unless the protocol is AppleTalk, in which case they are left as SNAP.

LLC frames are passed unchanged in value but with the length field required on Ethernet media stripped/inserted.

IPX Encapsulation Transformation Rules: For IPX, the encapsulation problems described above are compounded by the introduction of a fourth encapsulation on Ethernet media.

Novell introduced a frame format, when the IEEE 802.3 standards committee produced its version of Ethernet, which was incompatible with Ethernet.

Novell places its network header and data within a raw IEEE 802.3 Ethernet frame with no intervening IEEE 802.2 LLC header. This is in direct contravention of the standards but has become a de facto standard encapsulation.

Novell refers to this as Ethernet 802.3. It is also widely known as Novell Proprietary, Novell Raw, Raw 802.3, etc. Such frames are identifiable only by the fact that the Novell network header starts with a two-byte field called the checksum, which is never used and assumes the value 0xFFFF.

Therefore, routers, bridges and switches check for this value after an 802.3 length field. In effect, Novell has usurped the value 0xFF for the destination and source SAP addresses (DSAP/SSAP) of an LLC header.

Thus, on Ethernet media, there are four encapsulations for IPX as represented by Figure 132.

- Ether-Type - Value 0x8137
- Novell Proprietary (raw 802.3)
- LLC - SAP value 0XE0
- SNAP - Protocol Identifier 0x0000008137

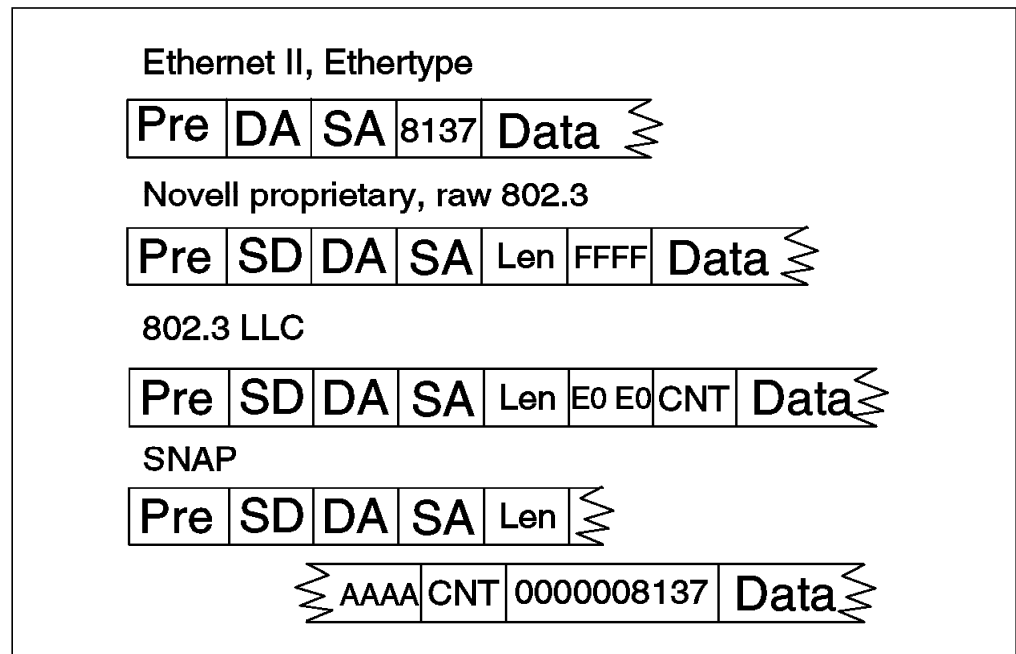


Figure 132. IPX Frame Types on Ethernet

On token-ring and FDDI, the same LLC and SNAP encapsulations are found as on Ethernet (without the length field).

Fortunately, this leaves an aggregate of only four encapsulations across all media with only two being universal (LLC and SNAP).

Unfortunately, the SNAP conversion rule isn't applicable and there is no algorithmic determination for the use of particular encapsulations on any media; it is purely the choice of the network administrator.

Worse, multiple encapsulations can be found on a single media to create multiple logical networks over a single physical broadcast domain.

The RouteSwitch, therefore, allows configuration of the encapsulation transformations of IPX frames. Before transmission of a frame occurs, the switch first determines the current encapsulation of the frame. It then consults the configuration information to determine which of the permitted encapsulations for the media the frame is to be transmitted on is required. Thus, the administrator can choose not only a single output option but an option per possible received encapsulation.

For example, over FDDI media, LLC and SNAP are permissible so the administrator might configure one of the following:

1. LLC and SNAP encapsulations received from other FDDI, token-ring or Ethernet media are translated to SNAP.
2. Ether-Type and proprietary encapsulations from Ethernet are translated to LLC.

Essentially, for each encapsulation, transformation to each of the other three encapsulations is available, but may simply be left as is. This choice may be further constrained by the output media type. For example, Ether-Type is not a valid option on FDDI or token-ring.

5.7.2.4 The Network Header

There are essentially two requirements for the any-to-any switching transformation function to address the network header fields:

1. In every protocol, there is a mechanism for mapping global network-wide addresses required in the local broadcast domain.
2. Different media have different minimum and maximum frame sizes leading to the issues of padding insertion/stripping and fragmentation/reassembly or maximum frame size negotiation protocols at the network level.

Address Mapping: There are almost as many ways to map a global network level address to a local subnetwork MAC address as there are routing protocols. These may or may not be affected by any-to-any switching.

Some may construct MAC addresses algorithmically, for example, DECNET model. Some may involve table lookups with an additional protocol to build and maintain these tables, for example, the IP/ARP model. Others may involve some form of building the network address around the MAC address as in the IPX model.

In all cases, these mechanisms are susceptible, without good design and forethought, to the problem of canonical versus non-canonical representation of MAC addresses in the network header area.

Address Mapping in IP - ARP: To map a 32-bit IP network address into the MAC address of a locally connected station, a router uses the Address Resolution Protocol (ARP) to build an ARP table. The router broadcasts a request containing the IP address in the body of the frame. The station with that IP address responds with its MAC address in the body of an ARP reply frame. The router inserts these two addresses in its ARP table and can now use the MAC address received to transmit any frames addressed to that IP address.

Since a router can have interfaces to Ethernet ports (canonical MAC addresses), FDDI, and token-ring (non-canonical MAC addresses), it is crucial that the router keep track of what media type it receives on each port.

If an IP ARP were defined such that all MAC addresses when conveyed in the body of an ARP frame were in canonical format, switching would be easy. A router, when taking an address from the ARP table and using it as the destination MAC address on an Ethernet port would use the address as is. If sending to FDDI or token-ring, it would bit swap the address to non-canonical format as required by the media.

Given this model of implementation, a station responding with an ARP on Ethernet that was switched to FDDI would result in the same representation of the MAC address in the ARP table of the router. The router would then use the bit swapped form in the MAC address of subsequent frames to the FDDI ring and the switch would bit swap these MAC header addresses as it transformed the frame onto Ethernet, resulting in the correct representation to be received by the original station.

Unfortunately, this model has only been defined in IP for Ethernet and FDDI. Token-ring stations place MAC addresses into the body of ARP frames in their native, non-canonical format and routers use addresses from the ARP table as is when sending to token-ring ports.

To achieve any-to-any switching with IP, it is necessary for the RouteSwitch to be sensitive to ARP frames and to bit swap the MAC addresses in the body of the ARP when switching a frame between token-ring and FDDI or Ethernet.

Because IP is well designed, the issue of address mapping being confined to the ARP protocol, this is sufficient to isolate the problem allowing all subsequent IP frames to be switched any-to-any.

Address Mapping in IPX: A network address in IPX consists of three parts:

1. Network Number: A globally unique identifier of a particular broadcast domain.
Strictly, because of the formation of logical networks using encapsulations, this is not equivalent to a physical broadcast domain but the distinction can be put aside for the purposes of this particular discussion.
2. Node Address: The MAC address of a station on that domain.
3. Socket Number: The task (process) within that station which should process the message.

Just as in IP, routers move a frame along hop by hop on the basis of the network number portion of the destination address. To do this, IPX needs the MAC address of the next hop router. This is obtained from the RIP table which is built up from the RIP updates sent out by all routers. When a router receives an RIP update frame, it uses the source node address in the frame as the MAC address for the next hop router.

Although there is not an explicit ARP-like protocol for mapping addresses as in IP, this same function is being achieved by the use of source node addresses in RIP frames.

In IPX, as in IP, the canonical versus non-canonical representation of addresses still applies. In switching, this needs to be considered for the source node address in IPX frames.

The node address in IPX Ethernet and FDDI observe a convention of using MAC addresses in the IPX header in canonical format. Therefore no translation takes place in the IPX network node address. For token-ring, the node addresses are non-canonical, therefore the node address will be translated between canonical and non-canonical when transversing the RouteSwitch.

The RouteSwitch will not modify any addresses found in the data portion of a SAP frame. If the client and the service are in the same IPX network and across a RouteSwitch using different media, the client will use the node address found

in the data portion of the SAP frame. Since that address has not been translated when transversing the RouteSwitch the client will never be able to contact the service. If the client and the service are on different IPX networks, then the client uses a RIP frame to contact the service. Novell has an internal IPX network which will cause the client to send a RIP frame. Network printers and other applications do not have an internal IPX network. These will fail if the service and the client are on different media across a RouteSwitch.

Proprietary Token-Ring IPX Switching: The RouteSwitch therefore offers the facility to modify IPX frames switching between token-ring and FDDI or Ethernet.

ARP bit swapping for IP is a de facto standard widely implemented in the industry. This is not the case with IPX. The switch must be able to coexist with bridges that do not support any-to-any switching or applications where this feature is not required. Therefore, this feature can be configured on or off.

Frame Size Requirements: The frame size requirements for the different media causes two problem areas which have to be addressed by the any-to-any switching transformation function.

Ethernet has a minimum frame size requirement. This requires that padding is inserted on frames switched to Ethernet which are below the minimum size and stripped from frames switched from Ethernet.

All media have different maximum frame size requirements. This gives rise to the problems of fragmenting large frames and/or negotiating maximum frame sizes.

Insertion of Frame Padding: Ethernet has a minimum frame size of 64 bytes. For frames smaller than 64 bytes it is a simple task for the RouteSwitch to perform padding.

Stripping such padding from Ethernet frames when switching to FDDI or token-ring is not so easy.

In most implementations of IP that have been tested, the presence of padding on FDDI or token-ring frames appears not to cause any problems. However, IPX implementations are adversely affected by its presence. Therefore, the RouteSwitch takes a conservative approach for all frames, regardless of protocol type, and strips padding where it can be detected.

Stripping of Padding for all IEEE 802.3 Frames: Ethernet frames in IEEE 802.3 format can be stripped of padding because of the presence of the length field. This includes all LLC and hence SNAP encapsulated protocols as well as Novell Proprietary format.

No Stripping of Non-IPX Ether-Type Frames: Padding can only be detected for Ether-Type encapsulated frames if the protocol is known and the protocol has some length information that can allow the valid data size to be inferred. This is protocol-specific and is currently only performed for IPX frames. Therefore, the RouteSwitch does not strip padding from non-IPX Ether-Type encapsulated frames including IP.

IPX-Specific Stripping: For IPX, the RouteSwitch performs pad stripping for all frame types including Ether-Type. This is possible because all IPX frames have a common header that includes the data length and allows the frame size to be

inferred. In fact, for IPX, the length in the IPX header is used to strip padding in all frame encapsulations, including the 802.3-based formats. This is because as well as padding frames for minimum length, many IPX Ethernet implementations pad frames to an even byte length. This single-byte pad, when performed on 802.3-based frames, is included in the 802.3 length field. Thus, the generic 802.3-based stripping technique is not sufficient to strip this odd-byte padding. When performing any-to-any switching FDDI implementations of IPX were found to be tolerant of this extra byte, whereas token-ring implementations would not work with it present. By adopting the single IPX stripping strategy of using the IPX header length, these problems are avoided, thus the RouteSwitch unconditionally strips all padding from IPX frames.

Also, it does not support odd-byte pad insertion when switching to Ethernet. This was a feature added to overcome limitations of some NIC cards, which is now of only historical importance and in fact, NetWare 4.1 servers provide this insertion as a port configuration option.

MTU Handling: Routers address the problem of maximum frame size limitations with the notion found in many protocols of a maximum transmission unit (MTU) size. Protocols use this notion in two possible ways.

- PDU fragmentation/reassembly

The router is configured with the MTU of each port. If a frame that is too large is required to be sent on a port, the protocol data unit (PDU) within the frame is fragmented into many smaller PDUs, each of which is re-encapsulated and sent as a frame that fits within the MTU.

- Connection-oriented end-to-end MTU negotiation

When an endstation enters into a protocol to communicate with another station, the initial PDU exchanges are guaranteed to fit all possible MTUs. In the handshaking between endstations to establish the connection, a phase is entered where large frames are sent. If an intervening link has an MTU too small for these frames, it will be dropped and the handshaking will time out. The endstations send progressively smaller frames until the handshaking succeeds and hence, establish the MTU to be used between the two stations for the remainder of their connection use.

IP supports the former mechanism and IPX the latter.

IP Fragmentation: The RouteSwitch Ethernet interfaces will use IP fragmentation if they are allowed to (if the don't fragment bit is not set). Fragmentation by FDDI and token-ring is not supported, though technically the token-ring could send frames larger than those supported by FDDI, and LAN Emulation could generate frames larger than both.

ICMP-Based MTU Discovery: IP uses the don't fragment bit to support an MTU discovery protocol, which superficially resembles the negotiation of IPX. The difference is that when IP stations attempt to discover an MTU size for their use, which doesn't require fragmentation by intermediate routers, the protocol expects a protocol response by the intermediate router that is an ICMP reporting that a frame was dropped because it couldn't be fragmented.

The RouteSwitch transformation function of any-to-any switching does not support this ICMP generation but just silently drops IP frames that can't be fragmented. The IP virtual router in the RouteSwitch does honor this protocol and supports ICMP. It is only the any-to-any switching that doesn't support this

protocol because it is not a router and may not even have an IP address with which to respond.

IPX Packet Size Negotiation: For IPX, the requirement of intervening devices is simply to drop frames that are too large to be forwarded. This is what the RouteSwitch does.

Other Protocols: Dropping oversized frames is the approach for all protocols other than IP. If the protocol in question is modeled like IPX, this will be the correct thing to do and will not cause problems. If the protocol is modeled like IP and expects fragmentation to occur or requires explicit response from the RouteSwitch, then the protocol will not succeed in any-to-any switching.

Banyan Vines: Banyan Vines supports Ethernet, FDDI, and token-ring networks. Each type of network generates a different frame format so the RouteSwitch performs translations for frames moving from one network type to another. The Banyan Vines protocol only uses one frame format per network type; no user configuration of translations is necessary. This protocol uses Ethernet II frames on Ethernet, SNAP frames on FDDI, and IEEE 802.2 (LLC) frames on token-ring. The RouteSwitch uses these frame formats when translating Banyan Vines frames.

Note: Checksums for Banyan Vines frames are automatically set to the null checksum, 0xFFFF, so that the checksum header does not require recalculation. Receiving stations will ignore this field and assume the sender is not using checksums.

5.7.2.5 Implementation Details

The previous material has covered the requirements for any-to-any switching in terms of what an abstract transformation function needs to do. The following sections cover the actual implementation of this function on the RouteSwitch.

The any-to-any switching function is done by each module when the frame is received from the RouteSwitch backplane.

Forwarding versus Flooding: Such frames will be handled in two ways:

- Forwarded frames - If the frame has a unicast destination address that has been learned on a particular port, the encapsulation translation choices are driven by options associated with the destination MAC address.
- Flooded/multicast frames - If the frame has a unicast destination address that has not been learned on a particular port, or if the destination address is a multicast address, then the frame has to be transmitted potentially on many ports. In this case, the encapsulation translation choices are driven by options associated with each destination port.

Port-Based Translation Options: The translation options for ports allow configuration of IP and IPX protocols on a per encapsulation basis.

MAC Address-Based Translation Options: The translation options for MACs arise from two possible sources:

1. Inheritance from port options during source address learning

When a source MAC address is learned, the translation options of the port on which it is learned are copied into the MAC-based database.

2. Automatic determination by RouteTracker

When a frame is processed by RouteTracker as part of determining the VLAN to be associated with the MAC, the frames protocol type and encapsulation are also determined. This information is used to update/set the translation options in the MAC-based database.

The AutoEncaps Option: Autoencapsulation is a technique employed by the RouteTracker software to learn the protocol and encapsulation type used by a source MAC address and automatically translate frames bound to that MAC address to the proper encapsulation type.

Normally, all devices attached to a switch port receive frames translated according to the translation options defined on that port. However, some devices attached to the same port may require a different frame format.

For example, one endstation may support IPX 802.3 frames and another endstation may support IPX SNAP frames. The switch port may be configured to translate incoming IPX 802.3 frames to LLC frames, which would satisfy only one of the endstations. If autoencapsulation is turned on, then the switch will translate frames for the first IPX endstation to IPX 802.3 and for the second endstation to IPX SNAP. Autoencapsulation operates only on learned unicast frames. It does not work on broadcast, multicast or unlearned unicast frames. For this reason, it is recommended that the *autocaps* option be on for ports with client endstations only. The server ports would be sending a high volume of broadcast and the server ports should be defined to use the correct encapsulation type for the server.

Which of these options is used is determined by setting the AutoEncaps option.

Native versus Non-Native on Ethernet: For Ethernet, one further distinction is made. If the frame received from the backplane is an Ethernet media type frame from another Ethernet switching module in the same chassis, then no encapsulation translations are applied. Such frames are referred to as *native* frames.

If a frame of an Ethernet media type was put onto the backplane by some other type of switching module, for example, the frame came from an FDDI card via a trunk port, or from the MPM via routing, then encapsulation translations are applied. Such frames are referred to as *non-native* frames.

Native versus Non-Native on FDDI and Token-Ring: For FDDI, token-ring, and LAN Emulation on ATM, a native/non-native distinction is not made. Instead, no encapsulation translations are applied by these switching modules to frames that are of their own media type.

No Translation on Trunk or PTP Ports: Switching modules that support encapsulation mechanisms, such as trunking ports on FDDI and token-ring, and point-to-point ports on ATM, do not apply translation to frames destined to such ports.

All other aspects of the transformation process are driven by the media type of the frame, the media type of the port on which the frame is to be transmitted and the protocol type determined for the frame. Thus frame padding insertion/stripping, IP fragmentation, IP ARP bit swapping, etc., are all automatic.

The Proprietary Token-Ring IPX Option: The one area that remains configurable is the bit swapping of source addresses for IPX in order to allow token-ring to work with FDDI and Ethernet. This is the equivalent function to IP ARP bit swapping.

This option is configurable and by default is off.

Earlier releases of the RouteSwitch did not provide extensive support for any-to-any switching and only a limited set of options were provided.

In more recent releases, full support is provided and all aspects are configurable.

In order to not encumber users upgrading to later releases who do not require these new features, much of the previous interface has been preserved allowing configuration and management of a subset of the full features. A new interface has been provided for the convenience of those users requiring these new features and these two interfaces are able to coexist. The factory defaults for all options are chosen such that a later release's factory defaults are compatible with previous releases.

In earlier releases, the user interface is the vi command to observe encapsulation configuration for ports, and the modvp, addvp and crgp commands where encapsulation transformation options can be set.

The commands of the latest release's user interface are found in a new menu, the Switch menu. This menu is covered later in this chapter.

In SNMP, the old interface is the vportEncapsulation item of the vportTable. The new interface is the vportSwitchTable and the vportSwitchDefaultTable.

5.7.2.6 The User Interface

Essentially, the forwarding code is now capable of applying the transformation function per:

- Protocol
- Encapsulation
- Port for flooded/mcast traffic
- Protocol
- Encapsulation
- Destination MAC address for forwarded unicast traffic

The old interface provides a small subset of these possible port translation options.

The addvp, modvp and crgp Commands: All of these commands include in their dialog an output format question for ports and a subsidiary IEEE 802.2 passthrough option. The options offered are:

- Default
- Ether-Type
- SNAP
- LLC

Each of these represents a set of translation options for the IP and IPX protocols. The names chosen for these sets basically represent the translations for IPX with the translation for IP being implied.

For example, LLC represents a translation set where all IPX encapsulations are configured to translate to IEEE 802.2. This is not a valid encapsulation for IP which is therefore configured to a default appropriate to the media, Ether-Type for Ethernet ports and SNAP for FDDI and token-ring ports. The translation of all other protocol types and encapsulations is fixed by the RouteSwitch. Thus, AppleTalk is never translated and Ether-Type/SNAP-based protocols follow the IP option.

For those options that imply a translation of IEEE 802.2 IPX frames to something else, a subsidiary question is asked, IEEE 802.2 IPX Pass Through(y/n):. An IEEE 802.2 passthrough option is provided because Novell 4.1 servers use this encapsulation by default and it is becoming Novell's encapsulation of choice.

The Default Translation Option: The meaning of the default is determined separately for each media type and is fully configurable. The factory defaults are chosen so that the latest release is fully compliant with earlier ones. The default translation option is provided to allow a single point of configuration of all ports capability. When the default option for a media is changed, all ports of that media type whose encapsulation is configured as default will inherit the new translation setting. All MAC address-based translation options that were inherited from those ports, as opposed to those set by RouteTracker, will also be updated. Ports that have an encapsulation setting other than default will be unaffected.

5.7.3 Default Translation Tables

The following tables detail the default types of translations for the different media types. The factory defaults are chosen so that the latest release is fully compliant with earlier ones.

<i>Table 9. Ethernet Media - Default Mode</i>
Ethernet Media - Default Mode
No translation is performed on outbound Ethernet frames where the inbound interface is Ethernet.
IP frames of any encapsulation are transmitted as Ethernet II frames.
IPX frames are transmitted as IEEE 802.3 Proprietary as the default setting. The only exception is when LLC passthrough mode is enabled, then the IEEE 802.2 (LLC) frames are forwarded as is.
No translation is performed on AppleTalk frames, and only AppleTalk Phase II (SNAP format) is supported.
Other than IP and IPX, all other Ethernet II and SNAP encapsulated protocols are sent as Ethernet II frames.
All other IEEE 802.3 with LLC encapsulated protocols are not translated.

<i>Table 10. FDDI Media - Default Mode</i>
FDDI Media - Default Mode
IP of any encapsulation is encapsulated SNAP.
IPX encapsulations are encapsulated SNAP except for IEEE 802.2 which is forwarded as is.
All other Ether-Type and SNAP encapsulated protocols are sent as for IP.
All other LLC encapsulated protocols are forwarded as is.

<i>Table 11. Token-Ring Media - Default Mode</i>
Token-Ring Media - Default Mode
IP of any encapsulation is encapsulated SNAP.
IPX encapsulations are encapsulated SNAP except for IEEE 802.2 which is forwarded as is.
All other Ether-Type and SNAP encapsulated protocols are sent as for IP.
All other LLC encapsulated protocols are forwarded as is.

<i>Table 12. Ethernet Media - Ethernet II Mode</i>
Ethernet Media - Ethernet II Mode
No translation is performed on outbound Ethernet frames where the inbound interface is Ethernet.
IP frames are transmitted as Ethernet II frames.
All IPX frames are transmitted as Ethernet II frames. The only exception is when LLC passthrough mode is enabled, then the IEEE 802.2 (LLC) frames are forwarded as is.
No translation is performed on AppleTalk frames, and only AppleTalk Phase II (SNAP format) is supported.
Other than IP and IPX, all other Ethernet II or SNAP frames are transmitted as Ethernet II frames.
Other IEEE 802.3 with LLC are not translated.

<i>Table 13. Ethernet Media - SNAP Mode</i>
Ethernet Media - SNAP Mode
No translation is performed on outbound Ethernet frames where the inbound interface is Ethernet.
IP frames are transmitted as SNAP frames.
All IPX frames are transmitted as IEEE 802.2 (LLC) frames.
No translation is performed on AppleTalk frames, and only AppleTalk Phase II (SNAP format) is supported.
Other than IP and IPX, all other Ethernet II or SNAP frames are transmitted as Ethernet II frames.
Other IEEE 802.3 with LLC are not translated.

<i>Table 14. FDDI/Token-Ring Media - SNAP Option</i>
FDDI/Token-Ring Media - SNAP Option
No translation is performed on outbound FDDI or token-ring frames where the inbound interface is the same media type.
IP frames of any encapsulation type are transmitted as IEEE 802.5 with SNAP frames.
IPX frames received that do not have an IEEE 802.2 encapsulation type, are transmitted as IEEE 802.5 with SNAP.
IPX frames received that are of IEEE 802.2 encapsulation type are transmitted as IEEE 802.5 with SNAP if the LLC passthrough is disabled. If the LLC passthrough is enabled, these frames will not be translated.
No translation is performed on AppleTalk frames, only AppleTalk Phase II is supported.
All other LLC encapsulated protocols are left as is.

<i>Table 15. Ethernet Media - LLC Mode</i>
Ethernet Media - LLC Mode
No translation is performed on outbound Ethernet frames where the inbound interface is Ethernet.
IP frames are transmitted as IEEE 802.2 (LLC) frames.
All IPX frames are transmitted as IEEE 802.2 (LLC) frames.
No translation is performed on AppleTalk frames, and only AppleTalk Phase II (SNAP format) is supported.
Other than IP and IPX, all other Ethernet II or SNAP frames are transmitted as Ethernet II frames.
Other IEEE 802.3 with LLC are not translated.

<i>Table 16. FDDI/Token-Ring Media - LLC Mode</i>
FDDI/Token-Ring Media - LLC Mode
IP frames are transmitted as IEEE 802.5 with SNAP frames.
All IPX frames are transmitted as IEEE 802.2 (LLC) frames.
No translation is performed on AppleTalk frames, and only AppleTalk Phase II (SNAP format) is supported.
Other than IP and IPX, all other Ethernet II or SNAP frames are transmitted as IEEE 802.5 with SNAP frames.
No translation is performed on outbound FDDI or token-ring frames where the inbound interface was the same media type.
Other IEEE 802.3 with LCC are not translated.

5.8 Example of Bridging, Routing and Switching

Bridging, routing and switching in the RouteSwitch are best summarized using an example.

The example described in this section has been tested using an 8274 RouteSwitch.

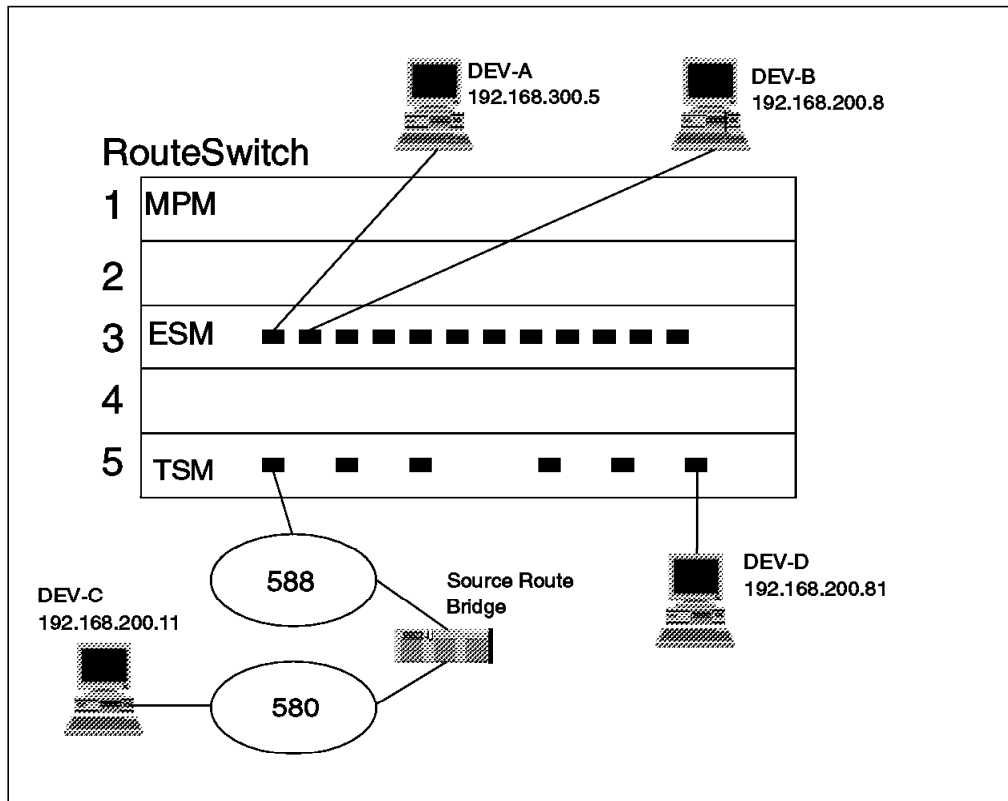


Figure 133. Physical View

Figure 133 shows the physical view of the network.

There are four endstations:

- DEV-A and DEV-B are directly connected to the ESM (Ethernet) module of the RouteSwitch. DEV-A is on port 3/1 and DEV-B is on port 3/2.
- DEV-C is attached to token-ring segment 580 via a source route bridge. Port 5/1 of the RouteSwitch is attached to token-ring segment 588.
- DEV-D is attached directly to TSM (token-ring) port 5/6.

Ports 5/1 and 5/6 are configured for SRT bridging with virtual ring enabled. The segment number is 588 and the bridge number in group 2 is 1.

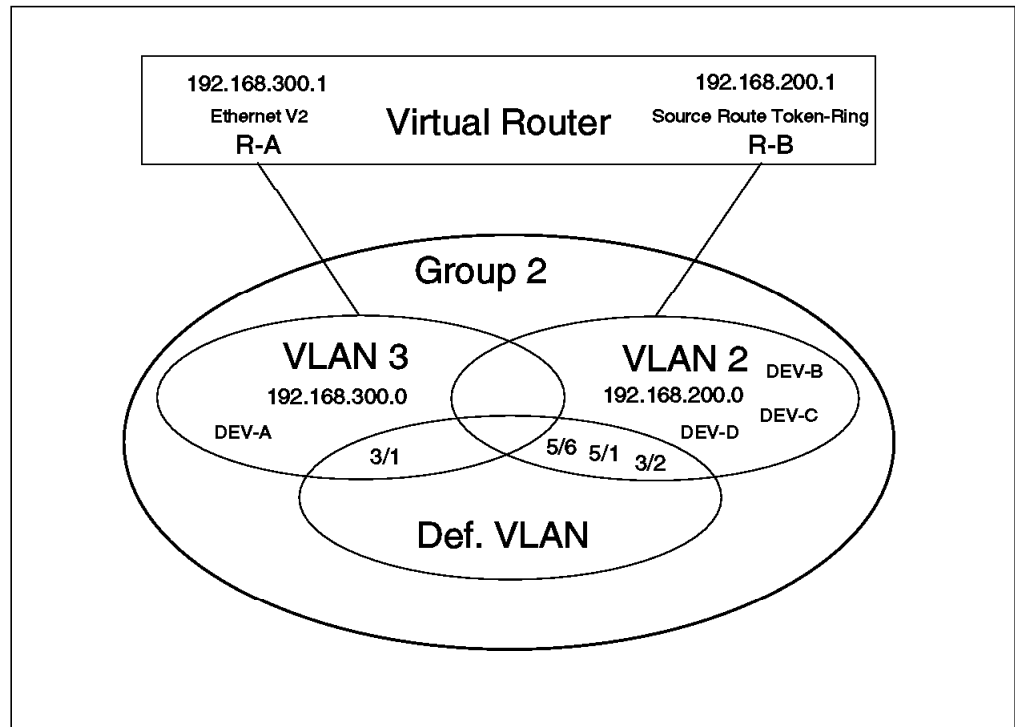


Figure 134. Logical View

Figure 134 is a logical representation of the physical network shown in Figure 133 on page 166.

- Group 2 is defined, but IP is not enabled in group 2.
- The default VLAN is not modified.
- VLAN 2 is created and assigned a network rule of 9.168.200.0. IP is enabled and the router arm R-B is assigned address 9.168.200.1. The frame type for the virtual router arm R-B was changed to source route token-ring.
- VLAN 3 is created and assigned a network rule of 9.168.300.0. IP is enabled and the router arm is assigned address 9.168.300.1. The frame type for the virtual router arm R-A was left at the default: Ethernet V2 frame type for IP protocol.

As described in Chapter 2, “Basic RouteSwitch Setup” on page 11, the MPM will have to see an IP frame with the correct IP network in order to assign ports to the appropriate VLAN.

One of the problems encountered during the testing was that once a MAC is assigned to a VLAN, the RouteSwitch will only forward broadcast to ports that belong to the same VLAN as the MAC address that sent the broadcast. For example, in Figure 134 if DEV-D had just powered up and had not sent an IP frame yet, then DEV-D would never see an IP ping from DEV-B. The solution is to either send an IP frame from DEV-D so that port 5/6 joins VLAN 2, or force port 5/6 to join VLAN 2 using a port rule policy.

The following is a summary of the endstations that can communicate together using the network represented by Figure 133 on page 166 and Figure 134.

<i>Table 17. Bridging within the RouteSwitch</i>					
	DEV-A	DEV-B	DEV-C	DEV-D	R-A
DEV-A	-	YES	YES	YES	YES
DEV-B	YES	-	NO	YES	YES
DEV-C	YES	NO	-	YES	NO
DEV-D	YES	YES	YES	-	YES

- DEV-A to DEV-B: The IP traffic will flow through the virtual router. The frame will be converted from Ethernet V2 to source route token-ring by the router. The frame will convert again to Ethernet V2 by module 3 because port 3/2 is left at a default of Ethernet V2.
- DEV-B to DEV-A: The IP traffic will flow through the virtual router. No frame conversion will take place since the frame type is the same for port 3/1 and router arm R-A.
- DEV-A to DEV-C: The IP traffic will flow through the virtual router. The frame will convert from Ethernet V2 to source route token-ring by the router. No further frame conversion will take place since router arm R-B and port 5/1 have the same frame type.
- DEV-C to DEV-A: The IP traffic will flow through the virtual router. The frame will convert from source route token-ring to Ethernet V2 by the router. The router will act as a source route termination and keep the RIF on behalf of DEV-A.
- DEV-B to DEV-C or DEV-C to DEV-B: No communication possible. The frame from DEV-C contains an RIF. The RouteSwitch any-to-any switching will not forward a frame with an RIF when the port is Ethernet or token-ring TB.
- DEV-B to DEV-D or DEV-D to DEV-B: Since ports 5/1 and 5/6 are on the same virtual ring, the frame from DEV-B will be received using the source route bridge configured in the RouteSwitch. The RouteSwitch will then switch the frame to port 5/6. See 5.3.1, "Virtual Token-Rings" on page 120 for an explanation of virtual rings.
- DEV-C to DEV-D: Port 5/1 and 5/6 are configured to be on the same virtual ring. Therefore the frame will be switched by the MAC address between the ports.
- DEV-C to R-A: It appears that the router port cannot handle a frame with a RIF field. The router however is able to provide RIF route termination for devices located on the Ethernet side.

Further examples include:

- Port 5/6 was an SR only port: Then no communication would be possible between DEV-C and DEV-D. A token-ring port must be TB or SRT in order to communicate with Ethernet ports or token-ring TB ports.
- Port 5/6 was left at the default of transparent bridging: Then communication would be possible between DEV-A and DEV-D, DEV-B and DEV-D. However, there will be no communication between DEV-C and DEV-D because of the RIF field in the frames sent by DEV-C.
- Router arm R-B was left at the default of Ethernet V2: Everything would still work, except for communication to DEV-C. Port R-B would no longer be able

to handle RIF route termination. Module 5 would convert the frame from Ethernet V2 to token-ring for port 5/6 using any-to-any switching.

5.9 Trunking

VLAN trunking provides the capability to transport multiple VLAN groups through high-speed RouteSwitch to RouteSwitch links. Only ATM, FDDI and WAN links support trunking. The extended VLAN group (one that would span multiple switches) must be connected through contiguous VLAN groups or the VLAN group will be segmented.

Groups can be connected together with their own single connection as in Figure 135. However, this would be a waste of resources. Each connection between extended groups requires its own virtual circuit Trunking provides the ability to multiplex VLAN groups across a single virtual connection as seen in Figure 136 on page 170.

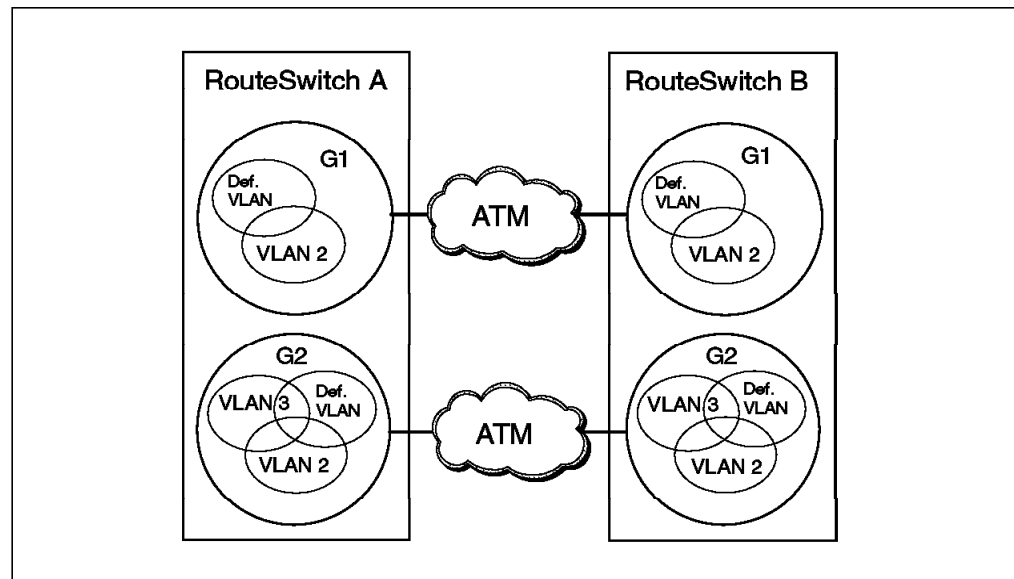


Figure 135. Multiple Groups without Trunking




Figure 136. Trunking Allows Multiple Groups to Use a Single Physical Connection

When created between RouteSwitches, trunk ports maintain data separation, keep spanning trees distinct to a VLAN group and allow multiple VLAN groups to span the trunking media.

When using trunking, the frames are encapsulated within a proprietary frame, implicitly identifying them with their VLAN group number to provide the data separation and to prevent interaction with non-trunked stations.

5.9.1 ATM Trunking

An ATM trunk can be configured as a PVC or an SVC connection.

5.9.1.1 Configuring an ATM Trunk over a PVC Connection

The following network example shows how to set up a RouteSwitch with an ATM trunk over a PVC connection. Two RouteSwitches, ibm8274a and ibm8274b, with an ASM module, were used in this example. The ATM ports of both the RouteSwitches were interconnected for trunking. In both the RouteSwitches a non-mobile group 2 was created with a default VLAN. VLAN 2 was defined in group 2. Mobile group 3 was created in both the RouteSwitches with a IP protocol policy defined to it.

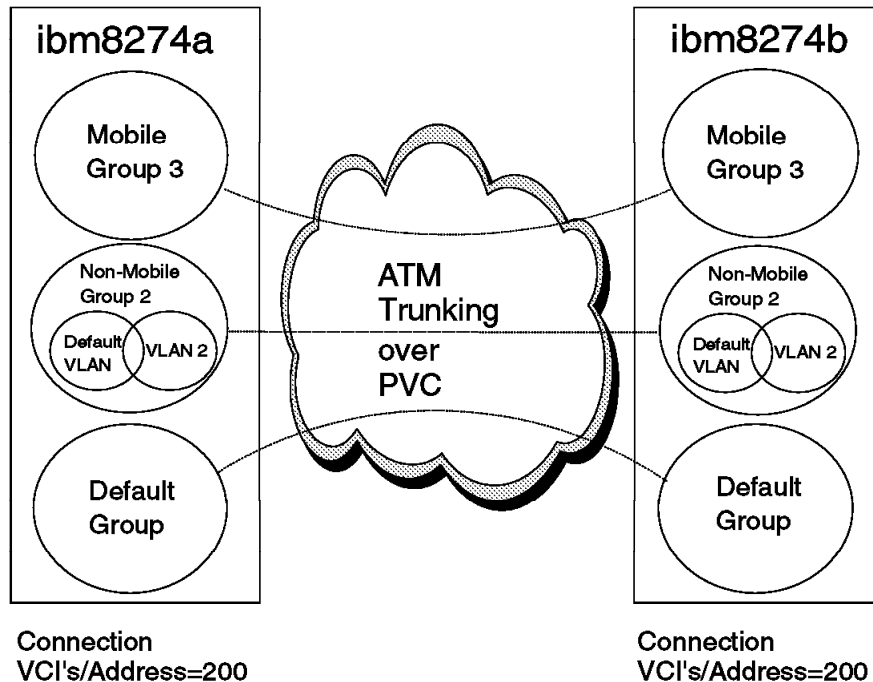


Figure 137. Logical View of Using ATM Trunking over PVC

Trunking can only be configured from the console or by using a telnet session.

To create an ATM trunking service, enter the command `cas` followed by the slot and port on which the service will be created (see Figure 138 through Figure 143 on page 174). The same configuration needs to be done on both the RouteSwitches: Switch ibm8274a and ibm8274b. Here it is assumed that the configuration of both RouteSwitches is identical. If the configuration of the switches is not identical, then the correct slot number for the ASM module needs to be considered while doing the configuration.

```
IBM8274A- />>cas 4/1

Slot 4 Port 1 Service 2 Configuration

1) Description (30 chars max)      : PTOP Bridging Service 2
2) Service type { LAN Emulation (1),
                  Trunking (4),
                  Classical IP(5),
                  PTOP Bridging(6),
                  VLAN cluster(7) } : PTOP Bridging
10) Encaps Type { Private(1),
                 RFC1483(2) }      : Private
3) Connection Type { PVC(1),
                   SVC(2) }        : PVC
4) PTOP Group                      : 1
5) PTOP connection                  : none
6) Admin Status { disable(1),
                 enable(2) }        : Enable

Enter (option=value/save/cancel) :
```

Figure 138. Configuring an ATM Trunk over a PVC Connection (1 of 5)

Define the service on port 4/1 to be trunking by typing 2=4.

```
Enter (option=value/save/cancel) :2=4

Slot 4 Port 1 Service 2 Configuration

1) Description (30 chars max)      : Trunking Service 2
2) Service type { LAN Emulation (1),
                  Trunking (4),
                  Classical IP(5),
                  PTOP Bridging(6),
                  VLAN cluster(7) } : Trunking
3) Connection Type { PVC(1),
                    SVC(2) }       : PVC
4) Trunked Groups                : 1
5) Connection                    : none
6) Admin Status { disable(1),
                  enable(2) }      : Enable
```

Figure 139. Configuring an ATM Trunk over a PVC Connection (2 of 5)

Define which group will be multiplexed on this trunk by entering 4=(group number).

In this example group 2 is added to the trunk.

```
Enter (option=value/save/cancel) : 4=2

Slot 4 Port 1 Service 2 Configuration

1) Description (30 chars max)      : Trunking Service 2
2) Service type { LAN Emulation (1),
                  Trunking (4),
                  Classical IP(5),
                  PTOP Bridging(6),
                  VLAN cluster(7) } : Trunking
3) Connection Type { PVC(1),
                    SVC(2) }       : PVC
4) Trunked Groups                : 1 2
5) Connection                    : none
6) Admin Status { disable(1),
                  enable(2) }      : Enable

Enter (option=value/save/cancel) :
```

Figure 140. Configuring an ATM Trunk over a PVC Connection (3 of 5)

Similarly group 3 is added to the trunk.

```

Enter (option=value/save/cancel) : 4=3

Slot 4 Port 1 Service 2 Configuration
1) Description (30 chars max)      : Trunking Service 2
2) Service type { LAN Emulation (1),
                  Trunking (4),
                  Classical IP(5),
                  PTOP Bridging(6),
                  VLAN cluster(7) } : Trunking
3) Connection Type { PVC(1),
                   SVC(2) }       : PVC
4) Trunked Groups      : 1 2 3
5) Connection          : none
6) Admin Status { disable(1),
                 enable(2) }     : Enable

Enter (option=value/save/cancel) :

```

Figure 141. Configuring an ATM Trunk over a PVC Connection (4 of 5)

Define Connection VCI for this connection by entering 5=(VCI number, which should be the same as that configured on the other switch on the trunk).

In this example Connection is configured to VCI as 200.

```

Enter (option=value/save/cancel) : 5=200

Slot 4 Port 1 Service 2 Configuration
1) Description (30 chars max)      : Trunking Service 2
2) Service type { LAN Emulation (1),
                  Trunking (4),
                  Classical IP(5),
                  PTOP Bridging(6),
                  VLAN cluster(7) } : Trunking
3) Connection Type { PVC(1),
                   SVC(2) }       : PVC
4) Trunked Groups      : 1 2 3
5) Connection          : 200
6) Admin Status { disable(1),
                 enable(2) }     : Enable

Enter (option=value/save/cancel) : save
Creating service, please wait...

Enabling service...
IBM8274A-/>>

```

Figure 142. Configuring Connection Parameters (5 of 5)

The command vas displays all the services as shown in Figure 143 on page 174.

IBM8274A- />>vas

ATM Services

Slot	Port	Serv Num	Service Description	Service Type
4	1	1	PTOP Bridging Service 1	PTOP Priv
4	1	2	Trunking Service 2	Trunking
4	2	1	PTOP Bridging Service 1	PTOP Priv

ATM Services

Slot	Port	Serv Num	VC Typ	Oper Status	SEL Groups	Conn VCI's/Addresses
4	1	1	PVC	Disabled	N/A 1	100
4	1	2	PVC	Enabled	N/A 1 2 3	200
4	2	1	PVC	Disabled	N/A 1	100

More? [<SPACE> for next page, <RETURN> for next line, Quit]

Figure 143. Displaying Services

Note: If there are ATM switches between the two RouteSwitches when using an ATM PVC connection, the PVC will also need to be configured on the ATM switches to have RouteSwitch-to-RouteSwitch connectivity.

5.9.1.2 Configuring an ATM Trunk over an SVC Connection

Configuration of the two RouteSwitches using an ATM trunk over an SVC can be done as given in the following example.




Figure 144. Logical View of Using ATM Trunking over SVC

Logical view of the network is as shown in Figure 144. Two RouteSwitches, ibm8274a and ibm8274b, were connected using the ATM port of the ASM module installed in both the RouteSwitches. In both the RouteSwitches, the physical port as well as the service, need to be defined as an SVC. Steps required to configure both the RouteSwitches are as follows:

1. Configure the ATM port of the RouteSwitch ibm8274a as an SVC port.

```

ibm8274a->>map 4/1

Slot 4 Port 1 Configuration

1) Description (30 chars max)           : ATM PORT
2) Conn Type { PVC(1), SVC(2) }         : PVC
3) Max VCCs (1-1023)                   : 1023
4) Max VCI bits (1..10)                 : 10
5) UNI Type { Pub(1), Priv(2) }         : Private
6) Tx Segment Size (2048-131072)        : 16384
7) Rx Segment Size (2048-131072)        : 16384
8) Tx Buffer Size (1800-8192)            : 4600
9) Rx Buffer Size (1800-8192)            : 4600
10) Pl Scramble {(False(1),True(2))}    : True
11) Timing Mode {(Loop(1),Local(2))}    : Local
12) Loopback Config { NoLoop(1), DiagLoop(2),
    LineLoop(3) }                       : NoLoop
13) Phy media { SONET(1),SDH(2)}        : SONET

Enter (option=value/save/cancel) :

```

Figure 145. Configuring an ATM Port for SVC on ibm8274a (1 of 11)

2. Change connection type to SVC.

```

Enter (option=value/save/cancel) : 2=2

1) Description (30 chars max)           : ATM PORT
2) Conn Type { PVC(1), SVC(2) }         : SVC

    30) Sig version { 3.0(1) 3.1(2) }    : 3.0
    31) Signaling VCI (0..1023)          : 5
    32) ILMI Enable {(False(1),True(2))} : True
    33) ESI (12 hex-chars)               : 0020da6fa8e1
    34) ILMI VCI (0..1023)               : 16

3) Max VCCs (1-1023)                   : 1023
4) Max VCI bits (1..10)                 : 10
5) UNI Type { Pub(1), Priv(2) }         : Private
6) Tx Segment Size (4096-131072)        : 16384
7) Rx Segment Size (4096-131072)        : 16384
8) Tx Buffer Size (1800-8192)            : 4600
9) Rx Buffer Size (1800-8192)            : 4600
10) Pl Scramble {(False(1),True(2))}     : True
11) Timing Mode {(Loop(1),Local(2))}     : Local
12) Loopback Config { NoLoop(1), DiagLoop(2),
    LineLoop(3) }                       : NoLoop
13) Phy media { SONET(1),SDH(2)}         : SONET

Enter (option=value/save/cancel) :

```

Figure 146. Configuring an ATM Port for SVC on ibm8274a (2 of 11)

3. Change ILMI Enable status to False.

```

Enter (option=value/save/cancel) : 32=1

Slot 4 Port 1 Configuration

1) Description (30 chars max)           : ATM PORT
2) Conn Type { PVC(1), SVC(2) }         : SVC

    30) Sig version { 3.0(1) 3.1(2) }    : 3.0
    31) Signaling VCI (0..1023)          : 5
    32) ILMI Enable {(False(1),True(2))} : False
    33) Net Prefix & ESI (38 hex-chars)  : 3911223344556677889900

3) Max VCCs (1-1023)                   : 1023
4) Max VCI bits (1..10)                 : 10
5) UNI Type                             : Private
6) Tx SAR Buffer Size (2048-131072)      : 16384
7) Rx SAR Buffer Size (2048-131072)      : 16384
8) Tx Frame Buffer Size (1800-16384)     : 4600
9) Rx Frame Buffer Size (1800-16384)     : 4600
10) Pl Scramble {(False(1),True(2))}     : True
11) Timing Mode {(Loop(1),Local(2))}     : Local
12) Loopback Config { NoLoop(1), DiagLoop(2),
    LineLoop(3) }                       : NoLoop
13) Phy media { SONET(1),SDH(2)}         : SONET

Enter (option=value/save/cancel) :

```

Figure 147. Configuring an ATM Port for SVC on ibm8274a (3 of 11)

4. Enter appropriate Net Prefix at option 33.

```

Enter (option=value/save/cancel):
33=39112233445566778899001101400000082 74a00

Slot 4 Port 1 Configuration

1) Description (30 chars max) : ATM PORT
2) Conn Type { PVC(1), SVC(2) } : SVC

30) Sig version { 3.0(1) 3.1(2) } : 3.0
31) Signaling VCI (0..1023) : 5
32) ILMI Enable {(False(1),True(2))} : False
33) Net Prefix & ESI (38 hex-chars)
    : 3911223344556677889900110240000008274a00

3) Max VCCs (1-1023) : 1023
4) Max VCI bits (1..10) : 10
5) UNI Type : Private
6) Tx SAR Buffer Size (2048-131072) : 16384
7) Rx SAR Buffer Size (2048-131072) : 16384
8) Tx Frame Buffer Size (1800-16384) : 4600
9) Rx Frame Buffer Size (1800-16384) : 4600
10) Pl Scramble {(False(1),True(2))} : True
11) Timing Mode {(Loop(1),Local(2))} : Local
12) Loopback Config { NoLoop(1), DiagLoop(2),
    LineLoop(3) } : NoLoop
13) Phy media { SONET(1),SDH(2)} : SONET

Enter (option=value/save/cancel) :

```

Figure 148. Configuring an ATM Port for SVC on ibm8274a (4 of 11)

5. Save the configuration on the ibm8274a RouteSwitch.

```

Enter (option=value/save/cancel) : save

Reset all services on slot 4 port 1 (n)? : y
Resetting port, please wait...

ibm8274a-/>>

```

Figure 149. Configuring an ATM Port for SVC on ibm8274a (5 of 11)

6. Define service on port 4/1 of the RouteSwitch ibm8274a.

```

ibm8274a- />>cas 4/1

Slot 4 Port 1 Service 2 Configuration

1) Description (30 chars max)      : PTOP Bridging Service 2
2) Service type { LAN Emulation (1),
                  Trunking (4),
                  Classical IP(5),
                  PTOP Bridging(6),
                  VLAN cluster(7) } : PTOP Bridging
10) Encaps Type { Private(1),
                 RFC1483(2) }      : Private
3) Connection Type { PVC(1),
                   SVC(2) }        : PVC
4) PTOP Group                      : 1
5) PTOP connection                  : none
6) Admin Status { disable(1),
                 enable(2) }       : Enable

Enter (option=value/save/cancel):

```

Figure 150. Configuring an ATM Port for SVC on ibm8274a (6 of 11)

7. Configure Service type as Trunking.

```

Enter (option=value/save/cancel): 2=4

Slot 4 Port 1 Service 2 Configuration

1) Description (30 chars max)      : Trunking Service 2
2) Service type { LAN Emulation (1),
                  Trunking (4),
                  Classical IP(5),
                  PTOP Bridging(6),
                  VLAN cluster(7) } : Trunking
3) Connection Type { PVC(1),
                   SVC(2) }        : PVC
4) Trunked Groups                  : 1
5) Connection                      : none
6) Admin Status { disable(1),
                 enable(2) }       : Enable

Enter (option=value/save/cancel):

```

Figure 151. Configuring an ATM Port for SVC on ibm8274a (7 of 11)

8. Change the connection type from PVC to SVC by typing 3=2.

```

Enter (option=value/save/cancel): 3=2

Slot 4 Port 1 Service 2 Configuration

1) Description (30 chars max)           : Trunking Service 2
2) Service type { LAN Emulation (1),
                  Trunking (4),
                  Classical IP(5),
                  PTOp Bridging(6),
                  VLAN cluster(7) } : Trunking
3) Connection Type { PVC(1),
                    SVC(2) }           : SVC
30) SEL for the ATM address             : 02
4) Trunked Groups                       : 1
5) Address (40-char-hex)                : none
6) Admin Status { disable(1),
                  enable(2) }           : Enable

Enter (option=value/save/cancel):

```

Figure 152. Configuring an ATM Port for SVC on ibm8274a (8 of 11)

9. Add which group(s) will be using this trunking service by typing 4=(group number). In our example group 2 was added. The ATM address of the ATM switch is also defined.

```

Enter (option=value/save/cancel): 4=2

Enter (option=value/save/cancel) : 5=3911223344556677889900110240000008274b00

Address 3911223344556677889900110240000008274b00':
        doesn't exist, this address will be created with default values

Connection Address 3911223344556677889900110240000008274b00 Configuration

1) Description (30 chars max) : Address 1

2) Requested Tx QoS Class { Unspecified(0) } : Unspecified
3) Requested TX Best Effort { False (1), True (2) } : True
4) Requested Tx Traffic Descriptor { NoCLPNoSCR(2) } : NoCLP NoSCR
20) Peak Cell Rate (cells/sec) for CLP=0+1 : 353208

5) Requested Rx QoS Class { Unspecified(0) } : Unspecified
6) Requested RX Best Effort { False (1), True (2) } : True
7) Requested Rx Traffic Descriptor { NoCLPNoSCR(2) } : NoCLP NoSCR
30) Peak Cell Rate (cells/sec) for CLP=0+1 : 353208

14) Tx Maximum Frame Size : 4520
15) Rx Maximum Frame Size : 4520
Enter (option=value/save/cancel) :

```

Figure 153. Configuring an ATM Port for SVC on ibm8274a (9 of 11)

10. Save the port configuration.

```

Enter (option=value/save/cancel) : save
Creating address connection, please wait...
  Slot 4 Port 1 Service 2 Configuration

1) Description (30 chars max)           : Trunking Service 2
2) Service type { LANE client(1),
                  Trunking (4),
                  Classical IP(5),
                  PTOP Bridging(6),
                  VLAN cluster(7) } : Trunking
3) Connection Type { PVC(1),
                    SVC(2) }           : SVC
   30) SEL for the ATM address          : 02
4) Trunked Groups                       : 1 2
5) Address (40-char-hex)                : 3911223344556677889900110240000008274b00
6) Admin Status { disable(1),
                  enable(2) }           : Enable

Enter (option=value/save/cancel) :

```

Figure 154. Configuring an ATM Port for SVC on ibm8274a (10 of 11)

11. Save the service configuration.

```

Enter (option=value/save/cancel): save
Creating service, please wait...

Enabling service...
ibm8274a- />>

```

Figure 155. Configuring an ATM Port for SVC on ibm8274a (11 of 11)

12. Configure ATM port of the RouteSwitch ibm8274b as an SVC port.

```

ibm8274b- />> map 4/1

  Slot 4 Port 1 Configuration

1) Description (30 chars max)           : ATM PORT
2) Conn Type { PVC(1), SVC(2) }          : PVC
3) Max VCCs (1-1023)                   : 1023
4) Max VCI bits (1..10)                 : 10
5) UNI Type { Pub(1), Priv(2) }         : Private
6) Tx Segment Size (2048-131072)        : 16384
7) Rx Segment Size (2048-131072)        : 16384
8) Tx Buffer Size (1800-8192)            : 4600
9) Rx Buffer Size (1800-8192)            : 4600
10) Pl Scramble {(False(1),True(2))}    : True
11) Timing Mode {(Loop(1),Local(2))}    : Local
12) Loopback Config { NoLoop(1), DiagLoop(2),
                    LineLoop(3) }        : NoLoop
13) Phy media { SONET(1),SDH(2)}        : SONET

Enter (option=value/save/cancel) :

```

Figure 156. Configuring an ATM Port for SVC on ibm8274b (1 of 11)

13. Change connection type to SVC.

```

Enter (option=value/save/cancel) : 2=2

1) Description (30 chars max)           : ATM PORT
2) Conn Type { PVC(1), SVC(2) }         : SVC

    30) Sig version { 3.0(1) 3.1(2) }    : 3.0
    31) Signaling VCI (0..1023)          : 5
    32) ILMI Enable {(False(1),True(2))} : True
    33) ESI (12 hex-chars)               : 0020da6fa8e1
    34) ILMI VCI (0..1023)               : 16

3) Max VCCs (1-1023)                   : 1023
4) Max VCI bits (1..10)                 : 10
5) UNI Type { Pub(1), Priv(2) }         : Private
6) Tx Segment Size (4096-131072)        : 16384
7) Rx Segment Size (4096-131072)        : 16384
8) Tx Buffer Size (1800-8192)            : 4600
9) Rx Buffer Size (1800-8192)            : 4600
10) Pl Scramble {(False(1),True(2))}     : True
11) Timing Mode {(Loop(1),Local(2))}     : Local
12) Loopback Config { NoLoop(1), DiagLoop(2),
    LineLoop(3) }                       : NoLoop
13) Phy media { SONET(1),SDH(2)}         : SONET

Enter (option=value/save/cancel) :

```

Figure 157. Configuring an ATM Port for SVC on ibm8274b (2 of 11)

14. Change ILMI Enable status to False.

```

Enter (option=value/save/cancel) : 32=1

Slot 4 Port 1 Configuration

1) Description (30 chars max)           : ATM PORT
2) Conn Type { PVC(1), SVC(2) }         : SVC

    30) Sig version { 3.0(1) 3.1(2) }    : 3.0
    31) Signaling VCI (0..1023)          : 5
    32) ILMI Enable {(False(1),True(2))} : False
    33) Net Prefix & ESI (38 hex-chars)  : 3911223344556677889900

3) Max VCCs (1-1023)                   : 1023
4) Max VCI bits (1..10)                 : 10
5) UNI Type                             : Private
6) Tx SAR Buffer Size (2048-131072)      : 16384
7) Rx SAR Buffer Size (2048-131072)      : 16384
8) Tx Frame Buffer Size (1800-16384)     : 4600
9) Rx Frame Buffer Size (1800-16384)     : 4600
10) Pl Scramble {(False(1),True(2))}     : True
11) Timing Mode {(Loop(1),Local(2))}     : Local
12) Loopback Config { NoLoop(1), DiagLoop(2),
    LineLoop(3) }                       : NoLoop
13) Phy media { SONET(1),SDH(2)}         : SONET

Enter (option=value/save/cancel) :

```

Figure 158. Configuring an ATM Port for SVC on ibm8274b (3 of 11)

15. Enter the appropriate Net Prefix at option 33.

```

Enter (option=value/save/cancel) : 33=3911223344556677889900110240000008274b00

Slot 4 Port 1 Configuration

1) Description (30 chars max) : ATM PORT
2) Conn Type { PVC(1), SVC(2) } : SVC

30) Sig version { 3.0(1) 3.1(2) } : 3.0
31) Signaling VCI (0..1023) : 5
32) ILMI Enable {(False(1),True(2))} : False
33) Net Prefix & ESI (38 hex-chars) : 3911223344556677889900110240000008274b00

3) Max VCCs (1-1023) : 1023
4) Max VCI bits (1..10) : 10
5) UNI Type : Private
6) Tx SAR Buffer Size (2048-131072) : 16384
7) Rx SAR Buffer Size (2048-131072) : 16384
8) Tx Frame Buffer Size (1800-16384) : 4600
9) Rx Frame Buffer Size (1800-16384) : 4600
10) Pl Scramble {(False(1),True(2))} : True
11) Timing Mode {(Loop(1),Local(2))} : Local
12) Loopback Config { NoLoop(1), DiagLoop(2),
LineLoop(3) } : NoLoop
13) Phy media { SONET(1),SDH(2)} : SONET

Enter (option=value/save/cancel) :

```

Figure 159. Configuring an ATM Port for SVC on ibm8274b (4 of 11)

16. Save the configuration on the ibm8274b RouteSwitch.

```

Enter (option=value/save/cancel) : save

Reset all services on slot 4 port 1 (n)? : y
Resetting port, please wait...

ibm8274b-/>>

```

Figure 160. Configuring an ATM Port for SVC on ibm8274b (5 of 11)

17. Define the service on port 4/1 of the RouteSwitch ibm8274b.


```
ibm8274b- />>cas 4/1
```

Slot 4 Port 1 Service 2 Configuration

- 1) Description (30 chars max) : PTOP Bridging Service 2
- 2) Service type { LAN Emulation (1),
Trunking (4),
Classical IP(5),
PTOP Bridging(6),
VLAN cluster(7) } : PTOP Bridging
- 10) Encaps Type { Private(1),
RFC1483(2) } : Private
- 3) Connection Type { PVC(1),
SVC(2) } : PVC
- 4) PTOP Group : 1
- 5) PTOP connection : none
- 6) Admin Status { disable(1),
enable(2) } : Enable

Enter (option=value/save/cancel):

Figure 161. Configuring an ATM Port for SVC on ibm8274b (6 of 11)

18. Configure Service type as Trunking.

```
Enter (option=value/save/cancel): 2=4
```

Slot 4 Port 1 Service 2 Configuration

- 1) Description (30 chars max) : Trunking Service 2
- 2) Service type { LAN Emulation (1),
Trunking (4),
Classical IP(5),
PTOP Bridging(6),
VLAN cluster(7) } : Trunking
- 3) Connection Type { PVC(1),
SVC(2) } : PVC
- 4) Trunked Groups : 1
- 5) Connection : none
- 6) Admin Status { disable(1),
enable(2) } : Enable

Enter (option=value/save/cancel):

Figure 162. Configuring an ATM Port for SVC on ibm8274b (7 of 11)

19. Change the connection type from PVC to SVC by typing 3=2.

```

Enter (option=value/save/cancel): 3=2

Slot 4 Port 1 Service 2 Configuration

1) Description (30 chars max)           : Trunking Service 2
2) Service type { LAN Emulation (1),
                  Trunking (4),
                  Classical IP(5),
                  PTOP Bridging(6),
                  VLAN cluster(7) } : Trunking
3) Connection Type { PVC(1),
                    SVC(2) }           : SVC
30) SEL for the ATM address             : 02
4) Trunked Groups                       : 1
5) Address (40-char-hex)                : none
6) Admin Status { disable(1),
                  enable(2) }          : Enable

Enter (option=value/save/cancel):

```

Figure 163. Configuring an ATM Port for SVC on ibm8274b (8 of 11)

20. Add which group(s) will be using this trunking service by typing 4=(group number). In our example group 2 was added. The ATM address of the ATM switch is also defined.

```

Enter (option=value/save/cancel): 4=2

Enter (option=value/save/cancel) : 5=3911223344556677889900110140000008274a00
Address 3911223344556677889900110140000008274a00':
        doesn't exist, this address will be created with default values

Connection Address 3911223344556677889900110140000008274a00 Configuration

1) Description (30 chars max) : Address 1

2) Requested Tx QoS Class { Unspecified(0) } : Unspecified
3) Requested TX Best Effort { False (1), True (2) } : True
4) Requested Tx Traffic Descriptor { NoCLPNoSCR(2) } : NoCLP NoSCR
20) Peak Cell Rate (cells/sec) for CLP=0+1 : 353208

5) Requested Rx QoS Class { Unspecified(0) } : Unspecified
6) Requested RX Best Effort { False (1), True (2) } : True
7) Requested Rx Traffic Descriptor { NoCLPNoSCR(2) } : NoCLP NoSCR
30) Peak Cell Rate (cells/sec) for CLP=0+1 : 353208

14) Tx Maximum Frame Size : 4520
15) Rx Maximum Frame Size : 4520
Enter (option=value/save/cancel) :

```

Figure 164. Configuring an ATM Port for SVC on ibm8274b (9 of 11)

21. Save the port configuration.

```

Enter (option=value/save/cancel) : save
Creating address connection, please wait...
  Slot 4 Port 1 Service 2 Configuration

1) Description (30 chars max)      : Trunking Service 2
2) Service type { LANE client(1),
                  Trunking (4),
                  Classical IP(5),
                  PTOP Bridging(6),
                  VLAN cluster(7) } : Trunking
3) Connection Type { PVC(1),
                    SVC(2) }       : SVC
   30) SEL for the ATM address     : 02
4) Trunked Groups                  : 1 2
5) Address (40-char-hex)          : 3911223344556677889900110140000008274a00
6) Admin Status { disable(1),
                  enable(2) }      : Enable

Enter (option=value/save/cancel) :

```

Figure 165. Configuring an ATM Port for SVC on ibm8274b (10 of 11)

22. Save the service configuration.

```

Enter (option=value/save/cancel): save
Creating service, please wait...

Enabling service...
ibm8274b- />>

```

Figure 166. Configuring an ATM Port for SVC on ibm8274b (11 of 11)

5.9.2 Group Multiplexing on FDDI

Using FDDI, it is possible to multiplex several groups via a single FDDI connection. Two methods can be used to multiplex groups:

- Trunking
- 802.10

Trunking is the recommended method since it encapsulates rather than translates frames. When using trunking, the media type and source routing information are retained since the original frame is not modified. The original frame is appended with a header containing the group information for the next RouteSwitch. Once the frame comes out of the trunk, the header is analyzed and removed. The original frame is then delivered to the correct group.

802.10 is the method used by Cisco and other vendors for multiplexing trunks over an FDDI backbone. When using 802.10, the frame is first translated, then tagged with a header containing the group information. 802.10 does not support source routing.

The commands `vi` and `fwt` will display the type of trunk used on FDDI ports. The two types will be displayed as `Trk` for trunking and `T10` for 802.10.

By default, every LAN port of the RouteSwitch has a virtual transparent bridge port associated with it. The FDDI virtual bridge port cannot be deleted, but can be moved between groups.

The FDDI virtual bridge port is, by default, attached to the default group (group 1).

When defining a group multiplexing service (trunking or 802.10) on an FDDI port, a virtual trunk port will be created on the FDDI physical port.

A physical FDDI port can have up to 31 virtual trunk ports and only one bridge port.

An FDDI virtual bridge port and an FDDI trunk port cannot coexist within the same group.

The virtual bridge port will have to be moved to a group that will not be using group multiplexing. If such a group is not available, then one must be created.




Figure 167. Group Multiplexing on FDDI

Figure 167 shows that groups 2 and 3 are using trunking and the frame tagging will keep the data separated from other traffic on the physical FDDI port. Group 1 is using the virtual bridge, therefore the frame will be translated to an FDDI frame.

The spanning tree is maintained and separated for all three groups. The spanning tree BPDU frames for the multiplexed groups are encapsulated, tagged with the group information and sent across the trunk port.

When trunking (not 802.10) is selected, the RouteSwitch sends a trunking hello message every 30 seconds. The trunking hello message contains the groups supported by this trunking service.

When a RouteSwitch receives a trunking hello message, it learns which groups are supported by the sender of the hello message.

The receiving RouteSwitch will ignore any groups in the trunking hello message that are not supported by the trunking service at the receiving RouteSwitch.

802.10 trunking does not use trunking hello messages. The group information is contained in the translated frame, and that information is used to forward the frame to the appropriate group. If the receiving RouteSwitch does not support the group contained in the frame, then the frame is discarded.

<i>Table 18. Receiving a Frame on an FDDI Trunk</i>	
Trunking	802.10
Learns the source MAC address in the encapsulated frame and associates it with the source MAC address of the remote trunk service. The header is stripped and frame forwarded to the RouteSwitch back plane for processing.	Checks the group ID in the header and forwards the frame to the appropriate group after stripping the header.
The frame will be dropped if the group is not active on this trunk service.	The frame will be dropped if the group does not exist.

<i>Table 19. Transmitting a Frame on an FDDI Trunk</i>	
Trunking	802.10
Known destination MAC address: The frame is encapsulated and forwarded on the appropriate trunk. Unknown destination MAC address: The frame is encapsulated and forwarded to all trunk services.	The RouteSwitch first verifies that the source group can be multiplexed. The frame is then translated to FDDI, the 802.10 header is added, and the frame is forwarded to the FDDI port.

It is possible to view the MAC addresses learned from the trunking service via the rts command.

Trunking over ATM is done in almost the same way; the only difference is that a virtual circuit identifier (VCI) has to be defined for each group so that a virtual circuit is defined in the permanent connection. When configuring the trunking, a connection name will need to be entered, which is the VCI. This name must be the same when configuring the trunking on the other switch.

5.9.2.1 FDDI Trunking Commands

The same commands are used to configure and manage FDDI trunking:

- Type cas followed by the slot and port number of the interface where the service will be created.
- Select the type of service (1=Trunking, 2=802.10).
- A screen displays the settings selected so far and a description can now be entered. The groups participating in this trunk service will need to be entered.

An example of this is seen in Figure 168 on page 188.

```

/ % cas 7/1
Service Type (1=Trunk, 2=802.10) (1) : 1
Slot 7 Station 1 FDDI Trunk Service
1) Description (30 chars max)      :
2) Group IDs                       : none
   : 1=FDDI Trunking service
Slot 7 Station 1 FDDI Trunk Service
1) Description (30 chars max)      : FDDI Trunking service
2) Group IDs                       : none
   : 2=2
Slot 7 Station 1 FDDI Trunk Service
1) Description (30 chars max)      : FDDI Trunking service
2) Group IDs                       : 2
   : save
Created trunk service for Group 2 on 7/1 (slot/station)
/ %

```

Figure 168. Example of the Configuration of Trunking over FDDI

Chapter 6. RouteSwitch ATM LAN Emulation

This chapter is intended to introduce the basic concepts of ATM Forum-compliant (FC) LAN Emulation and the implementation on the IBM Nways RouteSwitches.

6.1 LAN Emulation Benefits

Today's networking applications are running primarily on Ethernet and token-ring networks that interface with LAN adapters via standard interfaces such as ODI and NDIS. ATM application programming interfaces (APIs) are under development, which will allow applications to interface directly with the ATM layer and take advantage of all of ATM's advanced features (such as quality of service). In the meantime, a service is required that will allow existing applications written for Ethernet and token-ring networks to take advantage of at least some of ATM's benefits today, such as high-speed switched connections and scalability. This service is called LAN Emulation (LANE).

LANE allows applications running on ATM-attached stations, as well as token-ring and Ethernet stations, to communicate over an ATM network without any changes. For directly attached ATM stations (via 25,100 or 155Mbps), a driver is installed at the data link layer which presents a standard token-ring or Ethernet interface to the upper layers, while at the same time converting LAN frames to ATM cells to present to the ATM network. Endstations that do not have an ATM interface, but are on token-ring or Ethernet segments, can access the ATM network by means of a bridge or switch that has an ATM uplink. This type of device is known as a proxy as it does the frame-to-cell conversion, and connection management, on behalf of its LAN-attached endstations.

6.2 Emulated LANs (ELANs)

There is much confusion in the industry over what is meant by the terms ELAN and VLAN, as these terms are often used interchangeably. To avoid further confusion, these terms are reviewed as they apply in this publication:

VLAN: A VLAN (Virtual LAN) is a logical grouping of hosts, independent of physical location in the network, which defines what stations can communicate to each other. A VLAN can be based on different policies, such as protocol or network address. Each endstation joining a VLAN will share a broadcast domain with other endstations in the same VLAN.

ELAN: The term ELAN (Emulated LAN) is a specific implementation of a Virtual LAN, as it relates to LAN Emulation in ATM networks. An ELAN consists of one or more LAN Emulation clients (LEC), which share the same LAN Emulation Server and Broadcast and Unknown Server (LES/BUS). Broadcasts by any member of the ELAN are contained within the boundaries of that ELAN, and, ELAN membership can be assigned based on configurable policies.

The LES, BUS and LECS are collectively known as the LE Service components. The LE clients may obtain the address of their LES/BUS from a LAN Emulation Configuration Server (LECS), or alternatively, may be pre-configured with their LES/BUS address. It is preferable to allow the LECS to assign the LES/BUS address to the LE client, as the LECS can act as the central administration point for the creation of ELANs based on policies. For instance, an ELAN based on an

ELAN_NAME policy, could assign all LE clients with the name ACCOUNTS to one ELAN, and all the LE clients with the name ENGINEERING to another ELAN. For ATM-attached endstations, this would allow the physical re-location of the endstation without requiring reconfiguration of that endstation; that is, an ENGINEERING workstation would still belong to the ENGINEERING ELAN, regardless of where it is located.

Figure 169 shows a physical and logical view of a simple LAN Emulation network consisting of two ELANs.

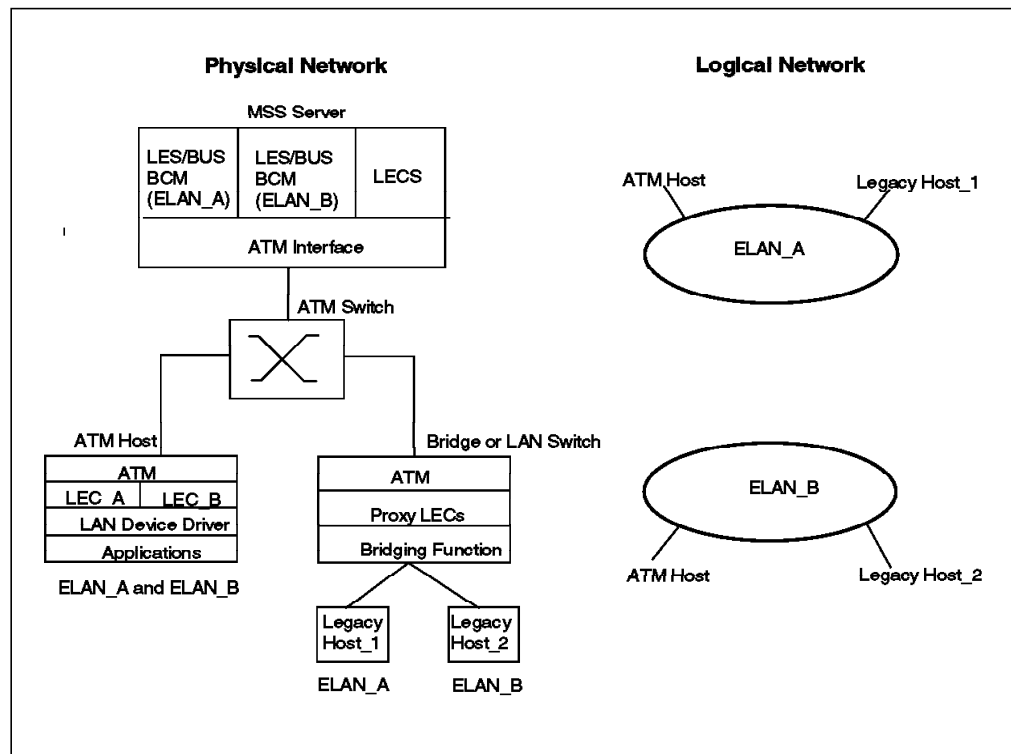


Figure 169. Physical and Logical Views of a Simple LAN Emulation Network

The LAN Emulation 1.0 specification does not specify how ELANs should be interconnected. To enable ELANs to be interconnected, the MSS server has extensive bridging and routing functions. These bridging and routing functions are accessed by creating internal LE clients (LECs), which are assigned to the ELANs to be bridged or routed. These internal LECs convert the ATM cells to token-ring or Ethernet frames which are then processed by the MSS's internal bridge or router, before being converted back to ATM cells to be forwarded to the destination ELAN. All members of an ELAN must be of the same type, that is, Ethernet or token-ring. ELANs of different types can be bridged or routed by the MSS server.

The MSS also contains the broadcast manager (BCM) function. When the BCM function is enabled, it will learn the network address associated to the ATM address. If the destination address has been learned, then a broadcast from an endstation will be changed to a unicast and forwarded to the appropriate ATM address. If the destination address is unknown to the BCM, then the broadcast is sent to the LES/BUS for processing.

A LEC can only belong to a single ELAN; however, an endstation that has more than one ATM adapter (or an adapter that supports more than one LEC), can have LECs belonging to different ELANs.

Before going into LAN Emulation in more detail, the basics of ATM addressing and the functions of ILMI will be reviewed as they relate to LAN Emulation.

6.3 ATM Addresses

ATM uses 20-byte hierarchical addressing. The first 13 bytes of an ATM address are called the network prefix. End systems obtain the network prefix component of their addresses from their adjacent ATM switch. The next six bytes of the address are called the end system identifier (ESI). The final byte is called the selector. End systems form their addresses by appending their ESI and selector to the network prefix provided by the ATM switch. The selector is only significant within the end system; it is not used to route calls within the ATM network, but is used within end systems to uniquely identify called/calling parties.

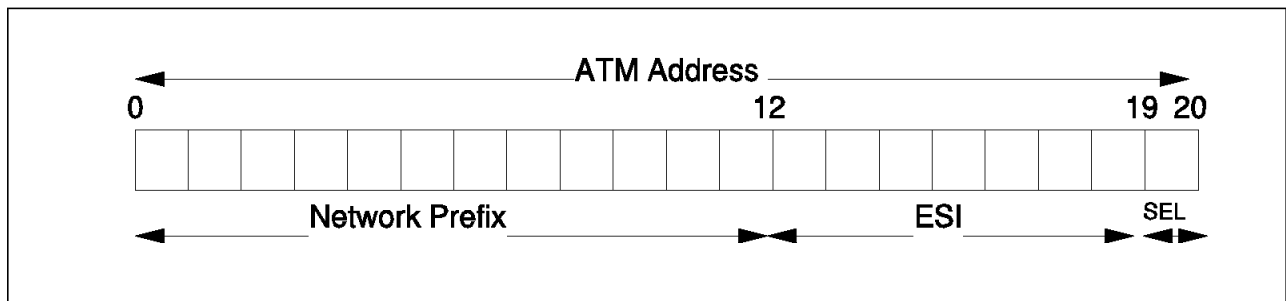


Figure 170. ATM Addresses

The network prefix and ESI components of ATM addresses must register with an ATM switch before calls can be placed or received. If the address is not unique (that is, if it duplicates an address already registered with the switch), the switch will reject the registration. One way to guarantee a unique ATM address is to use the burned-in (universally administered) IEEE MAC address as the ESI.

6.3.1 ATM Addresses of LAN Emulation Components

In general, ATM addresses must be unique among LAN Emulation components. The only exception is a LES and BUS serving the same ELAN may share an ATM address. (This is the case on the MSS server.) LAN Emulation components are configured for a particular ATM interface. The user may decide to use the burned-in MAC address as the ESI portion of the component's ATM address or select one of the locally administered ESIs defined for the ATM interface. Multiple LE components may share the same ESI if they have unique selector bytes. By default, the configuration interface assigns each LE component a unique selector byte value for the configured ESI; however, the user may override this assignment and explicitly configure a particular selector byte value.

An ATM interface parameter determines the number of selectors per ESI reserved for explicit assignment. (The remainder are available for dynamic assignment by the ATM interface at run time.) LE components only use the selectors reserved for explicit assignment; by default, 200 of the 256 possible selectors per ESI are reserved for explicit assignment. Run-time selector

assignment is beneficial when the user does not need to control the assigned selector. (Classical IP clients are an example.)

While ATM addresses must be unique among LE components, LE components may use the same ATM addresses as non-LE components such as Classical IP clients/servers.

6.4 Overview of ILMI Functions

The Interim Local Management Interface (ILMI) defines a set of SNMP-based procedures used to manage the user-to-network interface (UNI) between an ATM end system and an ATM switch. The following three ILMI functions are particularly relevant to LAN Emulation:

1. ATM address registration
2. Dynamic determination of UNI version being run on the switch
3. Acquisition of the LECS ATM address(es)

ILMI is the method of choice for locating the LECS. The ILMI MIB at the ATM switch includes a list of LECS ATM addresses that may be retrieved by the LECs. This is useful having the LECS ATM address(es) configured at the ATM switches only, not at LECs. There are fewer switches than LECs.

6.4.1 LAN Emulation Components

The individual components that make up an ELAN are now reviewed. An emulated LAN is comprised of the following components:

- One LAN Emulation server (LES)
- One broadcast and unknown server (BUS)
- One LAN Emulation configuration server (LECS)
- LAN Emulation clients (LECs), such as user workstations, bridges, routers, etc.

Users connect to the ELAN via the LEC, which requests services through the LAN Emulation User-to-Network Interface (LUNI). The three components (LES, LECS, and BUS) may be distributed over different physical systems or may be grouped together in one system, but logically they are distinct functions. The LAN Emulation services may be implemented in ATM intermediate systems (for example, switches such as the 8260 and 8285) as part of the ATM network, or in one or more ATM end systems, such as the MSS.

As illustrated in Figure 171 on page 193, each LEC has to support a variety of VCCs across the LUNI for transport of control and data traffic.

6.4.1.1 LAN Emulation Server (LES)

The basic function of the LE server is to provide directory and address resolution services to the LECs of the emulated LAN. Each emulated LAN must have an LE server. An LE client registers the LAN address(es) it represents with the LE server. When an LE client wants to establish a direct connection with a destination LEC, it gets the destination's MAC address from the higher layer protocols and has to ask the LE server for the destination's ATM address. The LES will either respond directly (if the destination client has registered that address) or forward the request to other clients to find the destination.

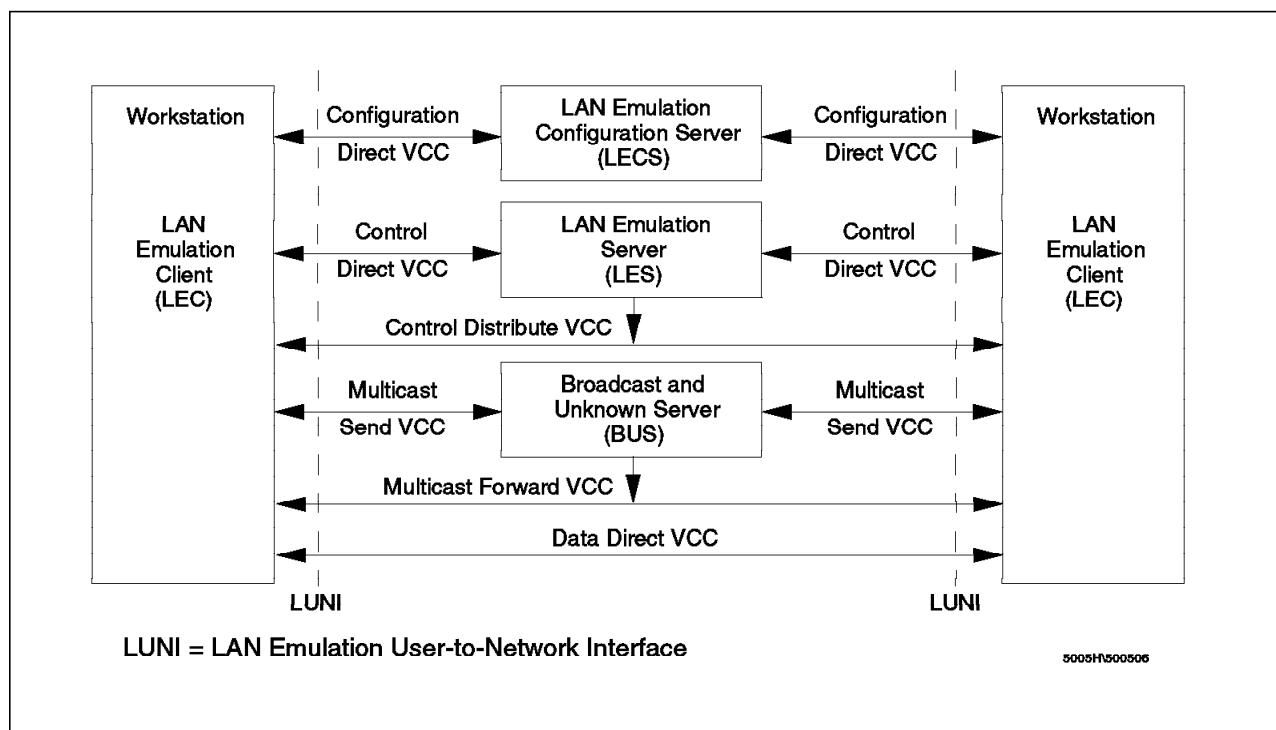


Figure 171. LAN Emulation Components

An emulated token-ring LAN cannot have members that are emulating an Ethernet LAN (and vice versa). Thus, an instance of a LES is dedicated to a single type of LAN emulation. The problems of translational bridging between different LAN types is not addressed in the ATM Forum's LAN Emulation. The MSS will do translational bridging between Ethernet and token-ring for NetBIOS and SNA protocols.

The LES may be physically internal to the ATM network or provided in an external device, but logically it is always an external function that simply uses the services provided by ATM to do its job.

6.4.1.2 LE Configuration Server (LECS)

The LECS assigns the individual LECs to the different ELANs that can exist in the ATM network. During initialization, a LEC requests the ATM address of the LES for the ELAN to which it should be connected. A LEC is not required to request this information from the LECS; a LES's ATM address may be configured (system-defined) in the LEC. In this case a LECS is not required.

Using a LECS to assign LECs to the different ELANs allows for central configuration and administration of multiple ELANs in an ATM network. The LECS could make its decision to assign a LES, for example, based on a client's ATM or MAC address according to a defined policy, or simply based on a system-defined database.

To take advantage of an MSS server's flexibility in creating policy-based ELANs, it is recommended that a LEC be configured to use the LECS to obtain the ATM address of their LES, where possible.

6.4.1.3 Broadcast and Unknown Server (BUS)

The BUS has two main functions:

- Distribute multicast and broadcast frames to all LECs in the ELAN
- Forward unicast frames to the appropriate destination

The BUS is required because legacy application on Ethernet and token-ring relies on broadcast to find their partners. Since ATM is a point-to-point connection, BUS service is required to be able to resolve the broadcasts on the ATM network. To avoid creating a bottleneck at the BUS, the rate at which a LEC can send unicast frames to the BUS is limited.

6.4.1.4 LAN Emulation Client (LEC)

Each workstation connecting to the ELAN has to implement the LE layer (also called LE entity), which performs data forwarding and control functions such as address resolution, establishment of the various VCCs, etc. The LE layer functions could be implemented completely in software, in hardware on a specialized LAN Emulation ATM adapter, or in a combination of both. The layered structure of the LEC is shown in Figure 172.

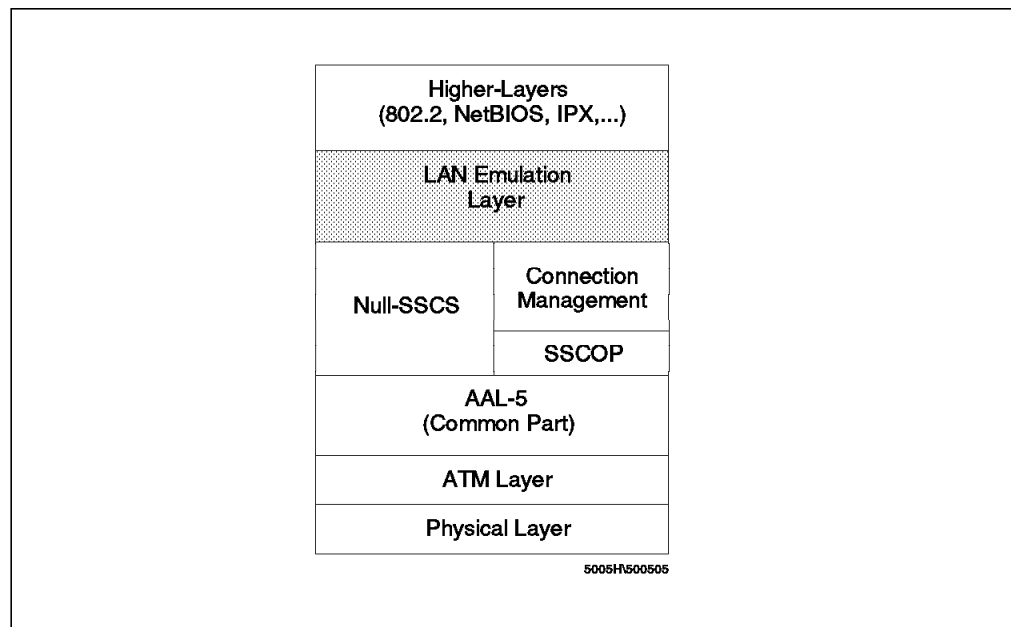


Figure 172. LAN Emulation Client Functional Layers

The LE layer provides the interface to existing higher-layer protocol support (such as IPX, IEEE 802.2 LLC, NetBIOS, etc.) and emulates the MAC-level interface of a real shared-media LAN (802.3/Ethernet or token-ring). This means that no changes are needed in existing LAN application software to use ATM services. The LE layer implements the LUNI interface when communicating with other entities in the emulated LAN.

The primary function of the LE layer is to transfer LAN frames (arriving from higher layers) to their destination.

A separate instance of the LE layer is needed in each workstation for each different LAN or type of LAN to be supported. For example, if both token-ring and Ethernet LAN types are to be emulated within a single station, then two LE layers are required. In fact, they will probably just be different threads within the

same copy of the same code but they are logically separate LE layers. Separate LE layers would also be used if one workstation needed to be part of two different ELANs both emulating the same LAN type (for example, token-ring). Each separate LE layer needs to have a different MAC address and must be attached to its own LE server, but it can share the same physical ATM connection (adapter).

6.4.2 LAN Emulation VC Connections

Data transfer in the LE system (consisting of control messages and encapsulated LAN frames) uses a number of different ATM VCCs as illustrated in Figure 171 on page 193.

6.4.2.1 Configuration and Control Connections

Control VCCs connect a LEC to the LECS and the LES, but they are never used for user data traffic. These connections may be permanent or switched and are established when a LEC connects to the ELAN.

Configuration Direct VCC

A bidirectional, point-to-point configuration direct VCC may be established between a LEC and the LECS to obtain configuration information (for example, the LES's ATM address).

Control Direct VCC

A bidirectional, point-to-point control direct VCC must be established (and kept active) between each LEC and the LES. This is used for the exchange of control traffic (for example, address resolution) between the LEC and the LES.

Control Distribute VCC

The LES may optionally establish a unidirectional control distribute VCC to distribute control information (for example, query for an unregistered MAC address) to all LECs connected to the ELAN. This can be a point-to-point VCC to each LEC. If the ATM supports point-to-multipoint connections, then the LES might instead establish one point-to-multipoint VCC to all LECs. (The clients will be added or deleted as leaves on this point-to-multipoint tree as they enter or leave the ELAN.)

6.4.2.2 Data Connections

Data connections are direct VCCs from a LEC to another LEC and to the BUS. They are used to carry user data traffic and never carry control traffic (except for a flush message for cleanup).

Data Direct VCC

For unicast data transfer between end systems, data direct VCCs are set up through ATM signaling as bidirectional, point-to-point connections once the LEC has received the destination's ATM address from the LES.

The LEC will send its first frame to the LES/BUS. The LES/BUS will resolve the destination LEC's ATM address and send the destination address to the originating LEC. The originating LEC will now set up a data direct VCC to the destination LEC and send the data on that newly created VCC.

Data direct VCCs stay in place until one of the partner LECs decides to end the connection based on installation options defining relevant timeouts, etc.

Multicast Send VCC

During initialization, a LEC has to establish a bidirectional, point-to-point multicast send VCC to the BUS (the BUS's ATM address is provided to the LEC by the LES) and must keep this VCC alive while being connected to the ELAN. This VCC is used by the LEC to send broadcast and multicast data frames.

Multicast Forward VCC

When a LEC establishes its multicast send VCC to the BUS, the BUS learns about the new member of the ELAN. The BUS then will initiate signaling for the unidirectional multicast forward VCC to the LEC. This VCC can be either point-to-point or point-to-multipoint (A point-to-multipoint VCC is more effective for multicast operations.)

6.4.3 LE Service Operation

In operation, the LAN Emulation service performs the following functions:

Initialization

During initialization, the LEC discovers its own ATM address from the ATM switch, which is needed for the client to set up direct VCCs. It obtains the LES's ATM address from the LECS and establishes the control VCCs with the LES and the BUS. The BUS address is provided to the LEC by the LES.

Address Registration

Clients use this function to provide address information to the LES. A client must either register all LAN destinations for which it is responsible, or join as a proxy. The LAN destinations may also be unregistered as the state of the clients changes. A LES can respond to address resolution requests if LECs register their LAN destinations (MAC addresses, or for source routing IEEE 802.5 LANs only, route descriptors) with the LES.

Address Resolution

This is the method used by an ATM client to associate a LAN destination with the ATM address of another client or the BUS. Address resolution allows clients to set up data direct VCCs to carry cells. This function includes mechanisms for learning the ATM address of a target station, mapping the MAC address to an ATM address, storing the mapping in a table, and managing the table.

For the LES, this function provides the means for supporting the use of direct VCCs by endstations. This includes a mechanism for mapping the MAC address of an end system to its ATM address, storing the information, and providing it to a requesting endstation.

Connection Management

In SVC environments, the LEC, the LES, and the BUS set up connections between each other using UNI signaling. This function is beyond the scope of this book.

Data Transfer

To transmit a frame, the sending LE layer must do the following:

- Decide on which of its VCCs (to destination LEC or BUS) a frame is to be transmitted

- Encapsulate the frame (AAL-5 is used.)

It must also decide when to establish and release data direct VCCs. To do this, it may need to access the LES for address resolution purposes.

6.4.3.1 Operation in Real Systems

In many practical LAN networks, this system is going to work very well. While LANs allow any-to-any data transport, typical LAN users connect to very few servers (such as communication servers, file servers, and print servers) at one time.

In this situation, the LE architecture described above can be extremely effective. After a short time, workstations will have established VCCs with all of the servers that they usually communicate with. Data transfer is then direct and very efficient. Timeouts can be set so that the switched VCCs are maintained during normal session usage.

Further details and an explanation of ATM Forum LAN Emulation is available in the ATM Forum Technical Committee's *LAN Emulation Over ATM* technical specification.

6.4.4 LAN Emulation Summary

LAN Emulation provides a relatively efficient means of transporting existing LAN-based applications across an ATM network, providing them with the benefits of high-speed switched connections, and scalability. A LAN Emulation system consists of the following basic components:

1. LAN Emulation Clients (LECs), which reside on endstations or proxy-bridges/switches. These provide the interface between traditional LAN traffic based on frames, and the ATM cell-based network.
2. A LAN Emulation Server (LES), which provides ATM-to-MAC address resolution.
3. A Broadcast and Unknown Server (BUS), which forwards all broadcast traffic to the LE clients.
4. A LAN Emulation Configuration Server (LECS), which assigns LES/BUS addresses to attaching LE clients. The LECS is optional, but if available, provides a central administration point for the assignment of policy-based ELANs.

6.5 LAN Emulation on the RouteSwitch

The RouteSwitch supports both Ethernet ELANs and token-ring ELANs.

Each ATM port can support up to 16 services. These services could be a mix of LEC, PTOP bridging, ATM trunking and VLAN cluster.

Because of the conversion to/from Ethernet, the default frame type might cause a connectivity problem.

The RouteSwitch without an FCSM is not an ATM switch, but can use the services of an 8210 (MSS) or an 8260 for example. A RouteSwitch with the FCSM and CSM modules can support LECs attached to it. With Version 3 Release 2 the separately orderable LSM code can be used to create ELANs that other LECs

could join as well as the frame-based side of the 8274. See Chapter 8, “LECS and LES/BUS Functionality” on page 213 for details on this functionality.

Within the RouteSwitch, a LEC is associated to a group and only one group. A LEC can only be connected to one LES/BUS. A group can have more than one LEC defined if each LEC is connected to a different LES/BUS. See 10.4, “IP Routing at the MSS” on page 291 on the configuration and utilization LECs with the RouteSwitch.

An example of how to configure a LEC on a ASM port can be found in 10.3, “RouteSwitch Interconnection Using MSS and 8260” on page 284.

6.5.1 SAR Option

All ATM switching modules support ATM Forum LAN Emulation 1.0. Two options are available on all modules (except for DS-3 modules) to specify SRAM size use for SAR buffer space. Generally speaking, large LANE environments will require the larger SRAM option.

Figure 173 illustrates the variables affected by SRAM size, namely logical connections (SVCs/PVCs) and SAR buffer size per connection. Increasing the size of the SAR buffer commensurately reduces the maximum number of connections that can be supported. Determining the maximum required size of the SAR buffer for a particular application will indicate the maximum number of connections that can be supported. If the 500 KB SAR option cannot support the required number, it will be necessary to utilize the 2 Mb SAR option. If the 2 Mb option cannot support the required number of connections, it will be necessary to segment the emulated LAN.




Figure 173. SAR Option Chart

If, for example, 24 VCs are required and the ATM module has a 500 KB buffer, then the maximum memory available for each of these 24 VCs is 16000 bytes. If Ethernet is the only media used, then the SAR for each VC will be able to contain ten Ethernet frames ($16000/1518=10.5$). On the other hand, the SAR

buffer per VC will increase to 64000 bytes if the 2 MB option is installed on the ATM module. This will allow to store 42 Ethernet frames in the SAR of each VC.

The recommended option for SAR values are:

- 8 KB - Ethernet-only connections
- 16 KB - Minimum for token-ring, FDDI and ATM LANE
- 32 KB - Best performance for token-ring, FDDI and ATM LANE

Chapter 7. ATM Cell Switching

This chapter describes the 8274 as an ATM switch. As described in the last chapter most ATM networks utilize LANE or Classical IP. These services can be provided in an 8210 (MSS) or with the new LSM (LANE Services Module) described in Chapter 8, "LECS and LES/BUS Functionality" on page 213.

7.1 Cell Switching Modules

The cell switching modules are known as CSMs and are required in order to do ATM cell switching.

Each (CSM) supports:

- ATM UNI 3.0/3.1
- NNI (PNNI 1.0 or IISP)
- 4 Input Output Processor (IOP) ASICs
- 1 ATM Fabric ASIC
- All ports support point-to-point and point-to-multipoint connections
- Virtual Path Connections (VPCs)
- Virtual Channel Connections (VCCs)
- Permanent Virtual Circuits (PVCs)
- Switched Virtual Circuits (SVCs)
- 1 OC-3c port supports 4096 connections in hardware
- 1 OC-12c port supports 65536 connections in hardware

7.1.1 Frame to Cell Switching Module

The frame-to-cell switching Module (FCSM) provides the internal link between the backplane's frame bus and cell matrix. The FCSM is required when the 8274 is configured with both frame and cell (LAN-to-ATM) switching, or just for ATM cell switching. It is not required for a pure frame switch configuration.

The FCSM does not contain any physical ports, as can be seen in Figure 174 on page 202, but it does contain an internal backplane port that can be configured and viewed with the user interface (UI). This internal port is functionally the same as an ASM port (see 6.5, "LAN Emulation on the RouteSwitch" on page 197) and is hardwired into the ATM cell matrix. ATM services such as LAN Emulation, Classical IP and trunking can be configured on this port.

This internal ATM port is also functionally the same as a CSM port, thus logically these seemingly two ports are actually halves of the same virtual port. This is seen in Figure 174 on page 202, where there is one IOP supporting two SARs, one supporting the ASM function, the other the CSM function.

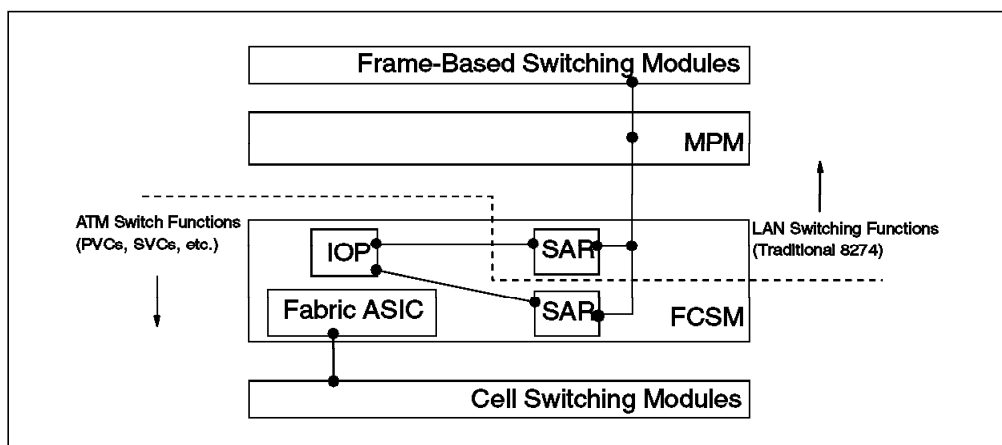


Figure 174. FCSM Conversion between Frame and Cell Switching

7.2 Cell Switching Matrix

The ATM cell matrix is distributed across the CSM modules with a total aggregate bandwidth of 13.2 Gbps. As an ATM switch the 8274 can be scaled from 2.0 Gbps up to 13.2 Gbps in 1.6 Gbps increments.

Each CSM module that is added to the 8274 provides an additional 1.6 Gbps of backplane capacity, due to the design, where each module has processors that allow for distributed cell switching.

Figure 175 shows the aggregation of bandwidth as modules are added.




Figure 175. CSM Aggregate Bandwidth

7.3 Traffic Management

There are various methods used to manage traffic in the 8274, all with the goal of preventing congestion. When there is congestion it can cause traffic queuing, degradation of performance and even cell discard.

The following methods are used in the 8274 to manage the traffic and prevent congestion:

- Buffer management
- Traffic contract parameters
- Class of service
- Virtual circuit priority

7.3.1 Buffer Management

Cells that are being switched from input ports to output ports can have collisions with other cells destined for the same output port. These collisions on the cell fabric create the need for queueing/buffering.

The 8274 uses a dynamic buffering system where the buffers are physically located on the input ports but are controlled by the output ports.

There are 8,192 cells expandable to 32K cell buffers per OC-3c/STM-1 interface, 128K cell buffers per OC-12c/STM-4c and 1K cells expandable to 4,096 cell buffers per ATM25.

These buffers are allocated to the various connections (VCs) based on connection's class of service. Each output port can see all traffic that is destined for it. This allows for the traffic to be released based on Class of Service and fairness algorithms.

7.3.2 Traffic Contract Parameters

The traffic contract parameters only affect the virtual circuit on which they are defined. When the congestion levels reach a threshold that can no longer be managed by buffers, the Class of Service and user priority variables are used to determine the traffic that will receive priority on the physical link.

The following methods are used by the 8274 to police, tag, and discard traffic:

1. Cell Loss Priority (CIP)

The cell loss priority is a 1-bit field in the ATM header and is used to indicate the priority of the cell. It is generally used during times of congestion to discard cells of a low priority and keep cells of a high priority.

CIP=0 High priority

CIP=1 Low priority

Due to traffic contract parameters, congestion can cause the status to be changed from high priority to low priority.

2. Traffic Contract Descriptors

Each virtual circuit has a defined traffic contract which includes:

- A description of traffic type.

- A class of service expected on that circuit.
- A descriptor that quantifies the cell rate allowed.
- Traffic descriptors are defined for each direction of the connection.
- These values can differ for each direction.
- Values are specified through software.

The traffic contract descriptors supported are:

- Peak Cell Rate (PCR), which is the maximum number of cells per second allowed on a virtual circuit, defined for all ATM traffic.
- Sustainable Cell Rate (SCR), which is the maximum average cell rate per second allowed for traffic. It is always less than PCR, and is not specified for constant bit rate traffic, as PCR will be equal to SCR.
- Maximum Cell Rate (MBS), is the maximum number of cells that can be sent in a burst at the peak cell rate. MBS is not specified for CBR traffic.

3. Traffic Contract Enforcement Methods

The traffic descriptors indicate the cell flow to monitor and the parameters to check cell flow against. When traffic violates any of the specified parameters, some form of traffic enforcement will take place. Generally, traffic that violates the specified contract will be tagged for discard eligibility or will be discarded.

The 8274 can do both congestion-based traffic enforcement and standard static enforcement. The difference being:

- Congestion-based traffic enforcement only discards cells when there is congestion. If there is available bandwidth on the connection, then no cells are discarded.
- Static traffic enforcement discards any cell that violates the traffic contract parameters regardless of the availability of bandwidth.

7.3.3 Class of Service

The Classes of Service that the 8274 supports as defined by the ATM Forum are:

- Class 1 - Constant Bit Rate (CBR)
This is the highest priority and guarantees zero cell loss. It is usually used for real-time traffic, for example compressed voice or video, as the loss of a cell can affect the decompression at the destination.
- Class 2 - Real-Time Variable Bit Rate
For real-time applications (for example, uncompressed voice and video). Losing a cell here will not impact on the application.
- Class 3 - Non-Real-Time Variable Bit Rate
This is used for non-real-time connection-oriented traffic.
- Class 4 - Available Bit Rate and Unspecified Bit Rate
This supports connectionless data protocols, for example, IP.
- Unspecified Class - Unspecified Bit Rate.
Traffic is transmitted on a "best effort" basis.

7.4 Private Network-to-Network Interface (PNNI)

Private Network-to-Network Interface (PNNI) is the ATM routing protocol that is used to implement ATM networks. PNNI determines the path a cell/cells should use to traverse the ATM network so as to meet bandwidth and quality of service requirements. PNNI is a hierarchical structure that is built dynamically by ATM switches once it is enabled.

7.4.1 Understanding PNNI Networks

PNNI's hierarchical structure requires nodes within the network to have information about other nodes and paths within the network. This information is then used to compute routes through the network.

Terminology used in PNNI networks:

- Peer Group - Made up of nodes with the same network identifier.
- Designated Transit List (DSL) - Route between nodes within a peer group.
- Horizontal Link - Link between two nodes in a single peer group.
- Border Node - Lies at the edge of the peer group and has a physical link to another border node in a different peer group.
- Exterior Link - A link between border nodes.

Each node can be a member of one or more peer group; nodes within a peer group share information on paths and other nodes within that peer group.




Figure 176. PNNI Peer Groups

7.4.2 PNNI Network Initialization

In order for connections to be set up between ATM endstations the ATM switches need to initialize and learn about the network topology.

The process that they follow is:

1. Discover directly attached endstations.

This includes ATM workstations and LAN switches with ATM uplinks. The 8274 uses Interim Local Management Interface (ILMI) to discover directly attached endstations.

2. Discover neighbor nodes.

All neighboring nodes are discovered within the peer group.

3. Discover topology information.

Each node exchanges topology information about the attached devices that it has discovered.

4. Computation of topology.

A shortest path first (SPF) algorithm is run at user-defined intervals. This computes the fastest path through the network.

7.4.3 PNNI Data Path

The following steps explain how a point-to-point connection is established, so that two endstations can communicate.

1. Call request received by switch.

An ATM endstation requests a call to another endstation. In the case of a SVC this includes Class of Service and VPI/VCI identifiers.

2. Determine called parties.

The ATM switch searches its topology database to determine if the destination device is reachable. If it is not reachable, the call is terminated.

3. Select path through network.

Now that the destination is determined, a path is selected through the network that will satisfy the Class of Service and bandwidth requirements.

4. Setup message sent.

A setup message is then sent to the next node, this node will process the request and then pass it on to the following node in the path.

5. Setup message is processed.

At each node the setup message is processed to determine if the requested bandwidth and class of service can be supported for the duration of the call. This is known as call admission control. If the requirements are met, then a VPI/VCI is set up and bandwidth is allocated for the connection.

6. Call proceeding message is sent.

Once the connection request has been accepted, the node forwards a call proceeding message to the previous node and a setup message to the next node in the path.

7. Connect message is sent.

After all nodes in the path have accepted the setup message, the destination endstation sends a connect message. This message is sent back to the endstation that originated the call. Once this message is received, data can be sent.

7.4.4 PNNI and IISP

With the current implementation of PNNI on the 8274, there is only support for single peer groups using PNNI. If you are implementing a network that requires multiple peer groups, you will have to run Interim Inter-Switch Signalling Protocol (IISP) between the peer groups.

IISP is a static ATM routing protocol and routes between peer groups have to be configured manually.

Note: Because the 8274 supports just one peer group in this release, there is just one level in the PNNI hierarchy. The default number for this level is 99.

7.4.5 PNNI Implementation in the 8274

To best illustrate the implementation of PNNI in the 8274 the following example is used:




Figure 177. Example of PNNI between 8260 and 8274

The hardware configuration used for this example was:

1. 8274 with:
 - FCSM
 - CSM 155 Mb Multimode Fiber (MM) Module
2. 8260 with:
 - A-CPSW (ATM Cell Switch)
 - 155 Mb Multimode Fiber (MM) Module

The following screens show the setup of the 8260 for PNNI.

```

8260ATM> set module 13 connected enable 1
Slot 13:Module set

8260ATM> set port 13.1 enable pnni ilmi_vpi_vci:none 2
13.01:Port set

8260ATM> show port 13.1 verbose 3

```

Type	Mode	Status

13.01:PNNI enabled	UP	4

```

No ILMI 5
NNI Bandwidth      : 155000 kbps
RB Bandwidth       : unlimited
Signaling vci      : 0.5
Routing vci        : 0.18
Administrative weight: 5040
VPI.VCI range      : 15.1023 (4.10 bits)
Connector          : SC DUPLEX
Media              : multimode fiber
Port speed         : 155000 kbps
Remote device is active

Frame format       : SONET STS-3c
Scrambling mode    : frame and cell
Clock mode         : internal

8260ATM>

```

Figure 178. Example of PNNI between 8274 and 8260 - Setup on 8260

- 1 Enable the OC-3 module so that the ports can be configured and connected to the backplane.
- 2 Enable the port on the module and also enable PNNI for the port. ILMI must be disabled for PNNI to work and this is disabled in the same command.
- 3 Display the settings of the port in verbose mode.
- 4 This line shows that PNNI is enabled and is currently up and running.
- 5 No ILMI: This shows ILMI is disabled.

```

8260ATM> set pnni node_0 atm_address: 1
Enter ATM address :
39.99.99.99.99.00.00.99.99.01.10.00.00.00.82.60.02.00

To activate issue COMMIT after your last 'set pnni...' entry.

8260ATM> save all 2

8260ATM> commit pnni 3
COMMIT execution will first SAVE pnni configuration
updates then RESET Hub.
Are you sure ? (Y/N) Y

8260ATM>

```

Figure 179. Example of PNNI between 8274 and 8260 - Setup on 8260

- 1** The ATM address for the 8260 must also be set. The 13th byte of the address must be unique for each switch in the peer group.
- 2** Save the settings before resetting the 8260.
- 3** Commit PNNI: This command will save the PNNI settings and then reset the 8260.

That completes the setup for PNNI on the 8260. The following screens show the setup for PNNI on the 8274.

```

8274/ >pnni 1
8274/Networking/PNNI >?
Command      ATM PNNI Menu
-----
Pconfig      Enter the PNNI configuration submenu
Proute       Enter the PNNI route management submenu
Pinfo        Enter the PNNI information submenu
Pstats       Enter the PNNI statistics submenu
Padmin       Enter the PNNI administration submenu

Main         File      Summary  VLAN      Networking
Interface    Security System    Services  Help
8274/Networking/PNNI >pconfig 2
8274/Networking/PNNI/Config >?
Command      ATM PNNI Configuration Menu
-----
pgcfg        Configure PNNI general parameters
pncfg        Configure PNNI node-specific operation parameters
ppcfg        Configure PNNI port (interface) operating parameters

Related Menus:
Pconfig  Proute  Pinfo    Pstats  Padmin
8274/Networking/PNNI/Config >

```

Figure 180. Example of PNNI between 8274 and 8260 - Setup on 8274

- 1** PNNI menu of commands on the 8274.
- 2** The PNNI configuration menu shows three levels for PNNI configuration, these are general, node and port.

No changes need to be made for the general parameters.


```

8274/Networking/PNNI/Config >map 5/7

Slot 5 Port 7 Configuration

1) Description (30 chars max)      : CSM PORT
2) Atm Address (40 hex-chars)      :
   00000000000000000000000000000000
3) Max VPI bits (1..12)           : 2
4) Max VCI bits (1..12)           : 10
5) I/F Type {Pub UNI(1), Pri UNI(2),
   PNNI(3), IISP(4)}               : PNNI 1
6) Phy Protocol {SONET(1), SDH(2)} : SONET
7) Signaling Ver {3.0(1), 3.1(2)} : 3.0
8) ILMI Enable {False(1), True(2)} : Disable 2

Enter (option=value/save/cancel) : save

Reset all connections on slot 5 port 7 (n)? : y
Resetting port, please wait...
send port change event
send port change event

8274/Networking/PNNI/Config >

```

Figure 182. Example of PNNI between 8274 and 8260 - Setup on 8274

1 The CSM port must be modified for PNNI.

2 ILMI must be disabled on the port doing PNNI.

Chapter 8. LECS and LES/BUS Functionality

Version 3.2.x of the Nways RouteSwitch Software Program added support to the RouteSwitch for basic ATM functions. These functions include:

- LECS
- LES
- BUS

These additional functions are a chargeable extension to the Nways RouteSwitch Software Program. See Chapter 6, "RouteSwitch ATM LAN Emulation" on page 189 for a detailed explanation of these ATM functions.

These new functions allow the RouteSwitch to operate in a manner similar to the integrated functions of the 8260 and 8285. The addition of the LECS takes the RouteSwitch one step further by making it easier to configure clients into a RouteSwitch ATM environment by being able to specify just an ELAN name and not the entire ATM address of the LES/BUS. Also it makes it possible to provide backup LES/BUS functions on multiple switches or ATM interfaces. These new functions are based on IBM's 8210 MSS Server.

Additionally this function is able to run on any RouteSwitch in the network, not just switches that have the ATM cell switching capability. This provides flexibility in placing and backing up the functions throughout the network. For example, if your network presently includes an ATM-capable 8260 running the internal LES/BUS function, these services could be distributed throughout the network by running these extensions on a number of RouteSwitches including the backup LES/BUS capability that the 8260 and 8285 do not include.

8.1 LANE Services Configuration

In order for the LANE services to be enabled, the lsm.img file must be present on the RouteSwitch. You can verify that the lsm.img file is loaded on the RouteSwitch by entering the ls command to see if it is listed.

The configuration options are available from the LANE Service Menu. By entering lsm at the RouteSwitch command line the screen in Figure 183 on page 214 is displayed.

```

/ % 1sm
Command      LANE Service Menu
-----
autolesbus   Automatic Configuration of LES/BUS, LECS and/or LECS
              database
lsmcfg       Configuration of LES/BUS, LECS and/or LECS database
ls1b         Show status of a LES/BUS pair
vlb          Show status of a LES/BUS pair
vlbs         Show statistics of a LES/BUS pair
vlbc         Show configuration of a LES/BUS pair
vlec         List all LE clients per LES/BUS pair
vmac         List registered MAC address of a given LES/BUS pair
vrd          List registered route descriptor of a given LES/BUS pair
vlecd        List detail LE client information by LEC id
vlecs        Show status of a LECS
vlecss       Show statistics of a LECS
vlecs        Show configuration of a LECS
velan        List elan(s) configured in the LECS database
vpolicy      List policy value assigned to an elan in the LECS
              database

Main      File      Summary  VLAN      Networking
Interface Security System  Services  Help
/Interface/LSM %

```

Figure 183. LANE Services Menu

All of the commands available on the LSM menu, as well as the submenus, use the command syntax of command <slot>/<port>, where the slot is where the ASM or FCSM is located and the appropriate port on the blade.

Note

One important consideration in implementing these functions on the RouteSwitch is that the port setup for these services must be changed from the default of PVC to SVC since LANE functions are SVC only. This can be changed using the map command. See Figure 267 on page 287 for an example of this procedure.

8.2 LSM Configuration Examples

The next set of examples begins by very simply creating a single ELAN environment and then follows up by manually creating additional ELANs and builds in redundancy.

8.2.1 Auto-Creating an ATM Environment

Most ATM networks contain at least one LECS and a LES/BUS pair. If you have a simple ATM environment and need only the LECS and a single LES/BUS pair (one ELAN), then the autolesbus command may be used. This command will allow you to quickly set up your RouteSwitch to support one ELAN and to configure LECS services. Figure 186 on page 216 shows the command autolesbus entered and the resulting output. This command creates a LECS with an entry for an ELAN to match an ELAN name policy in the LECS database. In addition an ELAN with the name default is created.

To demonstrate the use of the `autolesbus` command, a simple network has been created using an 8273, 8274 and an 8260. Please see Figure 184 on page 215, which is a physical diagram of the network setup.




Figure 184. Physical View of the Network for Auto-Creating an ATM Environment

Figure 185 shows the logical configuration of the LECS, LES/BUS and LECs.




Figure 185. Logical View of Network for Auto-Creating an ATM Environment

The configuration screens that follow describe the necessary steps that were followed in order to make the network functional.

In Figure 186 on page 216, the 8274 is configured with a LES/BUS pair that creates an ELAN with the name `default` and a LECS is also configured. The ELAN is added to the LECS database, and a policy is defined for the ELAN in the LECS database.

In this example all this is done using the `autolesbus` command where 4/1 is the ASM blade in slot four, port one.

```
8274/Interface/LSM > autolesbus 4/1

Creating LSM service ... please wait

LSM service created for slot 4, port 1
Creating LECS on slot 4, port 1, please wait...

LECS created on slot 4, port 1

Creating LES/BUS pair for elan 'default' on slot 4, port 1, please wait..

LES/BUS pair for elan 'default' created on slot 4, port 1

Creating default ELAN 'default' for LES/BUS pair on slot 4, port 1,
please wait..

default elan default added to the LECS database
default policy (ELAN_NAME) added to LECS database for elan 'default'
8274/Interface/LSM >
```

Figure 186. Auto Create the LES/BUS Pair and LECS

In Figure 187 on page 217, the status of the LECS and the LES/BUS pair are shown. This also shows us the ATM addresses that are being used. The ESI was changed by the `map 4/1` command. An example of use of the `map` command can be seen in 10.3, “RouteSwitch Interconnection Using MSS and 8260” on page 284.

```
8274/Interface/LSM > vlec 4/1

LECS status at slot 4, port 1

State: Operating normally (88)
Time of last state change: 00.01.11.04
Elapsed time since last change: 00.02.44.59
Error Log: No error (0)
Local ATM address: 399999999999999999990000999990101400082740000c1
Well-known address: 470079000000000000000000000000a03e00000100

8274/Interface/LSM > vlb 4/1 default
ELAN Name: default
ELAN Type: Ethernet
# of Proxy LEC's: 1
# of Non-Proxy LEC's: 0
LES ATM Address: 399999999999999999990000999990101400082740000c2

-Status-
LES-BUS State: OPERATIONAL
Major Reason LES-BUS was last Down: none
Minor Reason LES-BUS was last Down: none
LES-BUS State last changed at: 00.01.11.08
(System Up Time)
LES-LEC Status Table changed at: 00.01.14.54
(System Up Time)
BUS-LEC Status Table changed at: 00.01.15.46
(System Up Time)

-Current Configuration-
LES-BUS Enabled/Disabled: Enabled
ELAN Type: (S2) Ethernet
Max Frame Size: (S3) 1516
Control Timeout: (S4) 120
Max Frame Age: (S5) 1
Redundancy: Disabled

8274/Interface/LSM >
```

Figure 187. View Status of LES/BUS and LECS

Figure 188 on page 218, shows the reachable addresses on the 8260. These are ATM addresses of LEC devices that are connected to the 8260 ATM switch in our network. Note that on the 8260, blade 13 port 2 shows the ATM address of the 8274, along with the well-known address associated with the LECS on the 8274.

```

8260_HUB1> show reachable_address all
Port  Len Address                                     Active Idx
-----
12.01  96 39.99.99.99.99.99.99.00.00.99.99.03. . . . . Y 1
12.03 152 39.99.99.99.99.99.99.00.00.99.99.01.01.08.00.5A.99.0A.B3 Y
                                           Dyn 0
13.02 152 39.99.99.99.99.99.99.00.00.99.99.01.01.40.00.82.74.00.00 Y
                                           Dyn 0
13.02 152 47.00.79.00.00.00.00.00.00.00.00.00.00.00.00.A0.3E.00.00.01 Y
                                           Dyn 0
13.03 152 39.99.99.99.99.99.99.00.00.99.99.01.01.40.00.82.73.00.00 Y
                                           Dyn 0
17.03 152 39.99.99.99.99.99.99.00.00.99.99.01.01.00.04.AC.2C.35.94 Y

8260_HUB1>

```

Figure 188. Reachable Addresses on the 8260

Figure 189 on page 219, shows a LAN Emulation service being set up on an ASM port in the 8273. The LECS well-known address is used, and the ELAN name is added.

```

8273/ > mas 3/1 4

Slot 3 Port 1 Service 4 Configuration

1) Description (30 chars max)           : LAN Emulation Service 4
2) LAN Emulated Group                   : 1
   21) LAN type { 802.3 (1),
                802.5 (2) }             : 802.3
   22) Change LANE Cfg { NO (1),
                        YES (2) }        : NO
3) LECS Address (40-char-hex)           :
   4700790000000000000000000000A03E00000100
4) Admin Status { disable(1),
                  enable(2) }           : Enable
6) Connection Type { PVC(1),
                    SVC(2) }            : SVC
60) SEL for the ATM address              : 04

Enter (option=value/save/cancel) : 22=2

Slot 3 Port 1 Service 4 LANE Configuration Parameters

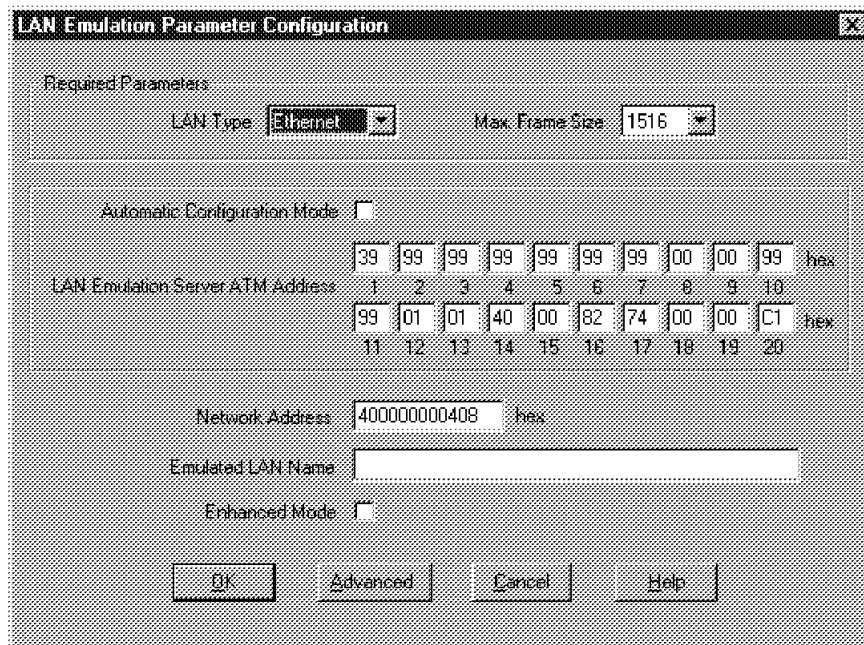
1) Proxy { NO (1), YES (2) }            : YES
2) Max Frame Size { 1516 (1), 4544 (2),
                    9234 (3), 18190 (4) } : 4544
3) Use translation options{NO (1), YES (2) : Yes
   (use Swch menu to set)
4) Use Fwd Delay time { NO (1), YES (2) } : NO
5) Use LE Cfg Server (LECS){ NO (1), YES (2)}: YES
6) Use Default LECS address { NO(1), YES (2)}: YES
7) Control Time-out (in seconds)         : 120
8) Max Unknown Frame Count               : 1
9) Max Unknown Frame Time (in seconds)   : 1
10) VCC Time-out Period (in minutes)     : 20
11) Max Retry Count                     : 1
12) Aging Time (in seconds)              : 300
13) Expectd LE_ARP Resp Time (in seconds) : 1
14) Flush Time-out (in seconds)          : 4
15) Path Switching Delay (in seconds)    : 6
16) ELAN name (32 chars max)             : default

Enter (option=value/save/cancel) :

```

Figure 189. Modify the Service on the ASM Port in the 8273

A Windows NT workstation is also configured to join the ELAN on the 8274 by configuring the LES-BUS address into the driver screen. The configuration of the adapter can be seen in Figure 190 on page 220. It would also have been possible to check the Automatic Configuration Mode box and input an Emulated LAN Name (default in our case).



LAN Emulation Parameter Configuration

Required Parameters

LAN Type: **Ethernet** Max. Frame Size: **1516**

Automatic Configuration Mode: ☐

LAN Emulation Server ATM Address:

1	2	3	4	5	6	7	8	9	10	hex
39	99	99	99	99	99	99	00	00	99	
11	12	13	14	15	16	17	18	19	20	hex
99	01	01	40	00	82	74	00	00	C1	

Network Address: **400000000408** hex

Emulated LAN Name:

Enhanced Mode: ☐

OK Advanced Cancel Help

Figure 190. Configuring the Windows NT Driver

Viewing the status of the default ELAN, as seen in Figure 191 on page 221, we see that there are two proxy LECs and one non-proxy LEC that have joined the default ELAN. The service that was created on the 8273 and on the 8274 are the proxy LECs shown here, and the Windows NT workstation is the non-proxy LEC.

Also in Figure 191 on page 221, a list of all the registered MAC addresses for the default ELAN are shown.

```
8274/Interface/LSM > vlb 4/1 default
ELAN Name:                        default
ELAN Type:                        Ethernet
# of Proxy LEC's:                 2
# of Non-Proxy LEC's:             1
LES ATM Address:
3999999999999000099990101400082740000c2

-Status-
LES-BUS State:                    OPERATIONAL
Major Reason LES-BUS was last Down: none
Minor Reason LES-BUS was last Down: none
LES-BUS State last changed at:    00.01.11.08
                                   (System Up Time)
LES-LEC Status Table changed at:  00.13.14.43
                                   (System Up Time)
BUS-LEC Status Table changed at:  00.13.14.59
                                   (System Up Time)

-Current Configuration-
LES-BUS Enabled/Disabled:         Enabled
ELAN Type:                        (S2)      Ethernet
Max Frame Size:                   (S3)      1516
Control Timeout:                  (S4)      120
Max Frame Age:                    (S5)      1
Redundancy:                       Disabled

8274/Interface/LSM > vmac 4/1 default

Number of Registered MAC's to display: 3

Registered                                     LEC
MAC Address   Registering ATM Address           Type   ID
-----
400082730000  3999999999999000099990101400082730000004 R  0003
020000002010  39999999999990000999901010004ac2c359481 R  0002
0020da6fa8e0  399999999999900009999010140008274000001 R  0001

8274/Interface/LSM >
```

Figure 191. Viewing the Status of Members in the LES/BUS

8.2.2 LSM Example: Routing between ELANs

Following on from 8.2.1, “Auto-Creating an ATM Environment” on page 214, we configured another LES/BUS pair (ELAN) manually and then defined AutoTracker VLANs, so as to route between the two ELANs.

The physical layout and equipment used was unchanged, as can be seen in Figure 192 on page 222, but with the added complexity of multiple ELANs the diagram differs slightly from Figure 184 on page 215.




Figure 192. Physical View of Network for ELAN Routing

Figure 193 shows the extended VLAN configuration, as well as the new ELAN, which was not configured in Figure 185 on page 215.

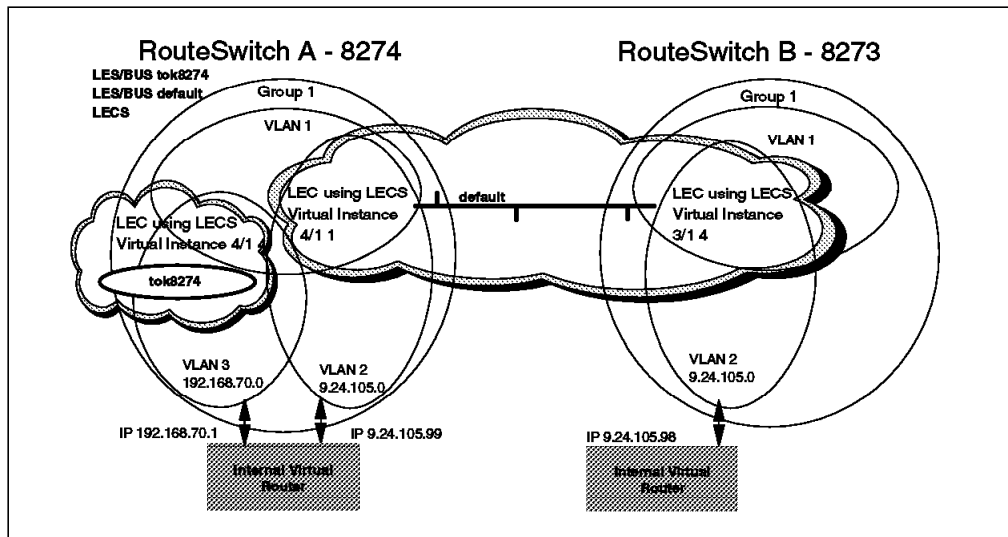


Figure 193. Logical View of Network for ELAN Routing

We now look in detail at the configuration screens. In Figure 194 on page 223, we enter the LSM configuration menu and select the option to create a LES/BUS. The settings that need to be defined are:

- ELAN name: We set to tok8274 for the new ELAN.
- ELAN type: Because the ELAN type is token-ring we must set it to 802.5 which is option 2.
- Max Data Frame Size. The frame size for token-ring is 4544.


```
8274/Interface/LSM > lsmcfg 4/1
```

```
LSM configuration for slot 4, port 1
```

- 1) Specify Global elan name (currently none specified)
- 2) Create LES/BUS
- 3) Modify LES/BUS
- 4) Delete LES/BUS
- 5) Create LECS
- 6) Modify LECS
- 7) Delete LECS
- 8) Add elan to LECS database
- 9) Delete elan from LECS database
- 10) Add policy to elan in LECS database
- 11) Delete policy from elan in LECS database
- 12) Exit configuration menu

```
Enter option : 2
```

```
LES/BUS for Slot 4 Port 1
```

- 1) ELAN name (32 chars max) :
- 2) ELAN type { 802.3 (1), 802.5 (2) } : 802.3
- 3) Max Data Frame Size { 1516 (1), 4544 (2),
9234 (3), 18190 (4) } : 1516
- 4) Control time-out { 10 - 300 seconds } : 120
- 5) Max. Frame age { 1 - 4 seconds } : 1
- 6) Enable redundancy { No (1), Yes (2) } : NO
- 7) Admin Status { Disable (1), Enable (2) } : Enable
- 8) LES/BUS Security { Disable (1), Enable (2) } : Disable
- 9) Create default ELAN { No (1), Yes (2) } : No

```
Enter (option=value/save/cancel) : 1=tok8274
```

```
Enter (option=value/save/cancel) : 2=2
```

```
Enter (option=value/save/cancel) : 3=2
```

```
LES/BUS for Slot 4 Port 1
```

- 1) ELAN name (32 chars max) : tok8274
- 2) ELAN type { 802.3 (1), 802.5 (2) } : 802.5
- 3) Max Data Frame Size { 1516 (1), 4544 (2),
9234 (3), 18190 (4) } : 4544
- 4) Control time-out { 10 - 300 seconds } : 120
- 5) Max. Frame age { 1 - 4 seconds } : 1
- 6) Enable redundancy { No (1), Yes (2) } : NO
- 7) Admin Status { Disable (1), Enable (2) } : Enable
- 8) LES/BUS Security { Disable (1), Enable (2) } : Disable
- 9) Create default ELAN { No (1), Yes (2) } : No

```
Enter (option=value/save/cancel) : save
```

```
Creating LES/BUS pair for elan 'tok8274' on slot 4, port 1, please wait..
```

```
LES/BUS pair for elan 'tok8274' created on slot 4, port 1
```

Figure 194. Create a LES/BUS Pair - Token-Ring ELAN

After creating the ELAN called tok8274, it must be added to the LECS database. In Figure 195 on page 224, the LSM configuration menu and the options that must be configured to add the ELAN to the LECS database are shown.

Figure 195. Add ELAN to LECS Database

Now that the ELAN is added to the LECS database, we must define policies that allow entry into the ELAN. In Figure 196 on page 225, from the LSM configuration menu we select the option that allows us to add a policy to an ELAN in the LECS database.

In this example the ELAN name was used to allow entry into the ELAN. After completing the configuration, the settings defined are saved automatically when exiting the menu.

```

LSM configuration for slot 4, port 1

1) Specify Global elan name (currently none specified)
2) Create LES/BUS
3) Modify LES/BUS
4) Delete LES/BUS
5) Create LECS
6) Modify LECS
7) Delete LECS
8) Add elan to LECS database
9) Delete elan from LECS database
10) Add policy to elan in LECS database
11) Delete policy from elan in LECS database
12) Exit configuration menu

Enter option : 10

Enter (elan name) : tok8274

Add policy value to elan 'tok8274'

1) By Elan name
2) By Elan type
3) By ATM address prefix
4) By MAC address
5) By Max. Frame size
6) By Route Descriptor
7) Exit

Enter option : 1

Enter elan name : tok8274
policy added to LECS database for elan 'tok8274'

Add policy value to elan 'tok8274'

1) By Elan name
2) By Elan type
3) By ATM address prefix
4) By MAC address
5) By Max. Frame size
6) By Route Descriptor
7) Exit

Enter option : 7

```

Figure 196. Add a Policy for the Token-Ring ELAN

Creating a service on the 8274 to join the tok8274 ELAN is required for routing between the ELANs. Each service that is created on port 4/1, for both the token-ring ELAN and the Ethernet ELAN, will be added to separate AutoTracker VLANs using the port rule. These AutoTracker VLANs also have network address rules defined, based on IP subnets. Each AutoTracker VLAN is a different IP subnet, as can be seen in Figure 193 on page 222.

```

8274/Interface/LSM > cas 4/1 4

Slot 4 Port 1 Service 4 Configuration

1) Description (30 chars max)           : LAN Emulation Service 4
2) Service type { LANE client(1),
                  Trunking (4),
                  Classical IP(5),
                  PTOB Bridging(6),    1
                  VLAN cluster(7) } : LAN Emulation
21) LAN type { 802.3 (1),
               802.5 (2) }           : 802.5 2
22) Change LANE Cfg { NO (1),
                     YES (2) }       : NO
3) Connection Type { PVC(1),
                    SVC(2) }         : SVC
30) SEL for the ATM address           : 04
4) LAN Emulated Group                 : 1
5) LECS Address (40-char-hex)         :
   4700790000000000000000000000A03E00000100
6) Admin Status { disable(1),
                 enable(2) }         : Enable

Enter (option=value/save/cancel) : 22=2 3

Slot 4 Port 1 Service 4 LANE Configuration Parameters

1) Proxy { NO (1), YES (2) }          : YES
2) Max Frame Size { 1516 (1), 4544 (2),
                   9234 (3), 18190 (4) } : 4544 4
3) Use translation options{NO (1), YES (2) } : Yes (use Switch menu to set
4) Use Fwd Delay time { NO (1), YES (2) } : NO
5) Use LE Cfg Server (LECS){ NO (1), YES (2)}: YES
6) Use Default LECS address { NO(1), YES (2)}: YES
7) Control Time-out (in seconds)       : 10
8) Max Unknown Frame Count             : 10
9) Max Unknown Frame Time (in seconds) : 1
10) VCC Time-out Period (in minutes)   : 20
11) Max Retry Count                    : 2
12) Aging Time (in seconds)            : 300
13) Expectd LE_ARP Resp Time (in seconds) : 1
14) Flush Time-out (in seconds)        : 4
15) Path Switching Delay (in seconds)  : 6
16) ELAN name (32 chars max)           : tok8274 5
17) Ring Number (1 - 4095)            : 0
18) Bridge Number (1 - 15)            : 0

Enter (option=value/save/cancel) : save

```

Figure 197. Create a Service on the ASM Port in the 8274

For a full explanation of creating a service please see 10.3, "RouteSwitch Interconnection Using MSS and 8260" on page 284. A summary of creating the service for this example is shown in Figure 197

- 1** Select the LANE service.
- 2** The ELAN that the service is going to join is a token-ring ELAN, so we select 802.5.
- 3** We are required to make changes to the LANE configuration.
- 4** Set the frame size for token-ring.

5 Enter the ELAN name the service must join.

A summary of the AutoTracker VLANs that are configured is shown in Figure 198.

```
8274/ >viatr1
```

VLAN Group:	VLAN Id	Rule Num	Rule Type	Rule Status	Rule Definition
1: 2		1	PORT RULE	Enabled	3/11/Brg/1
		2	NET ADDR RULE	Enabled	IP Addr = 9.24.105.0 IP Mask = 255.255.255.0
		3	PORT RULE	Enabled	4/1/Lne/1
1: 3		1	NET ADDR RULE	Enabled	IP Addr = 192.168.70.0 IP Mask = 255.255.255.0
		2	PORT RULE	Enabled	4/1/Lne/2

```
8274/ >
```

Figure 198. VLANs Configured on the 8274

Viewing the status of the tok8274 LES/BUS pair in Figure 199 on page 228 shows one proxy LEC and one non-proxy LEC. The proxy LEC is the service that has been configured on the ASM port of the 8274, and the non-proxy LEC is the NT workstation whose configuration has been changed to join the tok8274 ELAN.

A list of the registered MAC addresses for the tok8274 ELAN is also shown.

Figure 199. View Status of LES/BUS - Token-Ring ELAN

A list of the registered MAC addresses for the default ELAN is also shown.

```

8274/ >v1b 4/1 default
ELAN Name:                default
ELAN Type:                Ethernet
# of Proxy LEC's:        2
# of Non-Proxy LEC's:    0
LES ATM Address:          399999999999000099990101400082740000c2

-Status-
LES-BUS State:            OPERATIONAL
Major Reason LES-BUS was last Down: none
Minor Reason LES-BUS was last Down: none
LES-BUS State last changed at: 00.01.11.08 (System Up Time)
LES-LEC Status Table changed at: 02.05.24.08 (System Up Time)
BUS-LEC Status Table changed at: 02.05.24.08 (System Up Time)

-Current Configuration-
LES-BUS Enabled/Disabled: Enabled
ELAN Type: (S2)           Ethernet
Max Frame Size: (S3)      1516
Control Timeout: (S4)     120
Max Frame Age: (S5)       1
Redundancy:              Disabled

8274/ >vmac 4/1 default

Number of Registered MAC's to display: 2

Registered
MAC Address   Registering ATM Address   Type   LEC ID
-----
400082730000  399999999999000099990101400082730000004  R  0003
0020da6fa8e0  39999999999900009999010140008274000001  R  0001
8274/ >

```

A view of the IP routing table in Figure 201 shows that the internal virtual router has learned the IP subnets that were defined in the AutoTracker VLANs and their respective gateways. Routing will now take place between the AutoTracker VLANs and thus routing can take place between the ELANs.

Figure 201. IP Routing Table on 8274

8.2.3 Example of Redundant LES/BUS

The LSM code of the 8274 supports redundant LES/BUS pairs. In this example we show a very simple redundant LES/BUS configuration.

The primary LES/BUS is configured on the MSS and the backup LES/BUS is configured on the 8274. Both the 8274 and the MSS have the configuration of the ELAN eth1, and the definition of primary and redundant LES/BUS ATM addresses are defined on the devices.

If the primary LES/BUS fails for any reason, the redundant LES/BUS takes over. We also configured a LECS on both the MSS and the 8274. The ATM switch uses the LECS it finds first. During the test we found that it used the LECS on the MSS. However, when we switched off the power to the MSS the ATM switch used the LECS on the 8274. When the MSS was powered back on, it resumed the operation of the LES/BUS, but the LECS continued to be done by the 8274. Only when the 8274 was rebooted did the ATM switch use the LECS on the MSS.

A physical representation of the example is shown in Figure 202.




Figure 202. Physical View of the Network

The logical diagram is shown Figure 203 on page 231.




Figure 203. Physical View of the Network

The screens that follow show the configuration steps we used. A Qconfig was done on the MSS and the following was defined:

- Ethernet ELAN.
- LECS.
- LEC.
- An IP address was added to the LEC of 192.168.5.10.
- No bridging or IPX was configured.

After the Qconfig, the configuring of the LES/BUS redundancy on the MSS was done, as seen in Figure 204.

```
*t 6

Config>net 0
ATM user configuration
ATM Config>le-s

LE Services config>les-bus
LES-BUS config for ELAN 'eth1'>enable red
Redundancy protocol role
    (1) Primary LES-BUS
    (2) Backup LES-BUS

Enter Selection: (1) 1 1

ATM address of backup les-bus ()?
39.99.99.99.99.99.00.00.99.99.01.01.40.00.82.74.00.00.c1 2
Selection "Enable Redundancy" Complete
```

Figure 204. Configuration of Redundant LES/BUS on MSS

- 1** The LES/BUS is primary on the MSS.
- 2** This is the ATM address on the redundant LES/BUS on the 8274.


```

*t 6

Config>net 0
ATM user configuration
ATM Config>list
Command error
ATM Config>le-s
LAN Emulation Services user configuration
LE Services config>list
List of Configured LES-BUS(s)

ELAN Type (E=Ethernet/802.3, T=Token Ring/802.5)
| Interface #
| Enabled
| ELAN Name          LES ESI      Sel  Max.  Redundancy
|                   |             |     Frame Size  Role
|-----|-----|-----|-----|-----|
T 0 Y tok1          400082100000 x02   4544 (disabled)
E 0 Y eth1          400082100000 x03   1516 Primary
LE Services config>

```

Figure 206. View LES/BUS Active on MSS

Detailed in the following screen captures are the configuration steps on the 8273 for configuring a LEC to join the eth1 ELAN

The LAN Emulation service on the 8273 is modified to join the eth1 ELAN. These configuration steps can be seen in Figure 207 on page 234, and are the same as if the LEC was configured to join any other ELAN.

```

8273/ >mas 3/1 2

Slot 3 Port 1 Service 2 Configuration

1) Description (30 chars max)           : LAN Emulation Service 2
2) LAN Emulated Group                   : 1
   21) LAN type { 802.3 (1),
                  802.5 (2) }           : 802.3
   22) Change LANE Cfg { NO (1),
                          YES (2) }      : NO
3) LECS Address (40-char-hex)           :
   4700790000000000000000000000A03E00000100
4) Admin Status { disable(1),
                  enable(2) }           : Enable
6) Connection Type { PVC(1),
                    SVC(2) }            : SVC
60) SEL for the ATM address              : 02

Enter (option=value/save/cancel) : 22=2

Slot 3 Port 1 Service 2 LANE Configuration Parameters

1) Proxy { NO (1), YES (2) }             : YES
2) Max Frame Size { 1516 (1), 4544 (2),
                    9234 (3), 18190 (4) } : 1516
3) Use translation options{NO (1), YES (2) : Yes
   (use Swch menu to set)
4) Use Fwd Delay time { NO (1), YES (2) } : NO
5) Use LE Cfg Server (LECS){ NO (1), YES (2)}: YES
6) Use Default LECS address { NO(1), YES (2)}: YES
7) Control Time-out (in seconds)          : 120
8) Max Unknown Frame Count                 : 1
9) Max Unknown Frame Time (in seconds)     : 1
10) VCC Time-out Period (in minutes)       : 20
11) Max Retry Count                       : 1
12) Aging Time (in seconds)                : 300
13) Expectd LE_ARP Resp Time (in seconds) : 1
14) Flush Time-out (in seconds)            : 4
15) Path Switching Delay (in seconds)      : 6
16) ELAN name (32 chars max)               : eth2

Enter (option=value/save/cancel) : 16=eth1

```

Figure 207 (Part 1 of 2). Creating a Service on 8273

```

Slot 3 Port 1 Service 2 LANE Configuration Parameters

1) Proxy { NO (1), YES (2) } : YES
2) Max Frame Size { 1516 (1), 4544 (2)
   9234 (3), 18190 (4) } : 1516
3) Use translation options{NO (1), YES (2) : Yes
   (use Swch menu to set)
4) Use Fwd Delay time { NO (1), YES (2) } : NO
5) Use LE Cfg Server (LECS){ NO (1), YES (2)}: YES
6) Use Default LECS address { NO(1), YES (2)}: YES
7) Control Time-out (in seconds) : 120
8) Max Unknown Frame Count : 1
9) Max Unknown Frame Time (in seconds) : 1
10) VCC Time-out Period (in minutes) : 20
11) Max Retry Count : 1
12) Aging Time (in seconds) : 300
13) Expectd LE_ARP Resp Time (in seconds) : 1
14) Flush Time-out (in seconds) : 4
15) Path Switching Delay (in seconds) : 6
16) ELAN name (32 chars max) : eth1

Enter (option=value/save/cancel) : save

Saving new LANE Configuration values

Slot 3 Port 1 Service 2 Configuration

1) Description (30 chars max) : LAN Emulation Service 2
2) LAN Emulated Group : 1
21) LAN type { 802.3 (1),
   802.5 (2) } : 802.3
22) Change LANE Cfg { NO (1),
   YES (2) } : NO
3) LECS Address (40-char-hex) :
   4700790000000000000000000000A03E00000100
4) Admin Status { disable(1),
   enable(2) } : Enable
6) Connection Type { PVC(1),
   SVC(2) } : SVC
60) SEL for the ATM address : 02

Enter (option=value/save/cancel) : save
Modifying service, please wait...

Resetting service, please wait...
Enabling service...
8273/ >

```

Figure 207 (Part 2 of 2). Creating a Service on 8273

Viewing the statistics on the port 3/1 and also the services on the port is shown in Figure 208 on page 236.

```

8273/ >vap

                        ATM Port Table

Slot Port      ATM Port Description      Conn Tran  Media UNI Max  VCI
=====
3    1    ATM PORT                        SVC  STS3c  Multi Pri 1023 10

Slot Port Loopback Cfg Tx Clk Source
=====
3    1    NoLoop      LocalTiming

Slot Port      ATM Network Prefix      End System  Sig Sig  ILMI  ILMI
=====
3    1    3999999999999999000099990101  400082730000 3.0 5    True  16

                        Status

Slot Port Tx Seg Sz Rx Seg Sz Tx Buff Sz Rx Buff Sz Oper SSCOP ILMI
=====
3    1      8192      8192      4600      4600 Enabled Up Up

8273/ >vas

                        ATM Services

Slot Port Serv Service Service
      Num Description Type
=====
3    1    1    PTOP Bridging Service 1    PTOP Priv
3    1    2    LAN Emulation Service 2    LANE

                        ATM Services

Slot Port Serv VC Oper
      Num Typ Status SEL Groups Conn VCI's/Addresses
=====
3    1    1    PVC Enabled N/A 1    100
3    1    2    SVC LANE Op. 02 1    39 40 41 42

FDDI Services do not exist!

```

Figure 208. Viewing Services on 8273

The AutoTracker VLAN configuration on the 8273 is shown in Figure 209 on page 237.

```

8273/ >viatr1
VLAN  VLAN  Rule Rule      Rule      Rule
Group: Id   Num Type      Status    Definition
-----
      1: 2      1  NET ADDR RULE Enabled  IP Addr = 192.168.5.0
                                   IP Mask = 255.255.255.0
                                   2  PORT RULE    Enabled  3/1/Lne/2

8273/ >

```

Figure 209. VLAN Configuration on 8273

The redundant LES/BUS on the 8274 is configured next.

In Figure 210 and Figure 211 on page 238 the LES/BUS is created.

```

8274/Interface/LSM % lsmcfg 4/1
Creating LSM service ... please wait

LSM service created for slot 4, port 1

LSM configuration for slot 4, port 1

1) Specify Global elan name (currently none specified)
2) Create LES/BUS
3) Modify LES/BUS
4) Delete LES/BUS
5) Create LECS
6) Modify LECS
7) Delete LECS
8) Add elan to LECS database
9) Delete elan from LECS database
10) Add policy to elan in LECS database
11) Delete policy from elan in LECS database
12) Exit configuration menu

Enter option : 2 1

LES/BUS for Slot 4 Port 1

1) ELAN name (32 chars max) :
2) ELAN type { 802.3 (1), 802.5 (2) } : 802.3
3) Max Data Frame Size { 1516 (1), 4544 (2),
                        9234 (3), 18190 (4) } : 1516
4) Control time-out { 10 - 300 seconds } : 120
5) Max. Frame age { 1 - 4 seconds } : 1
6) Enable redundancy { No (1), Yes (2) } : NO
7) Admin Status { Disable (1), Enable (2) } : Enable
8) LES/BUS Security { Disable (1), Enable (2) } : Disable
9) Create default ELAN { No (1), Yes (2) } : No

Enter (option=value/save/cancel) :

```

Figure 210. Configure Redundant LES/BUS on 8274

```
Enter (option=value/save/cancel) : 1=eth1 2
Enter (option=value/save/cancel) : 6=2 3
Enter (option=value/save/cancel) : 61=2 4
Enter (option=value/save/cancel) : 
62=3999999999999900009999010140008210000003 5
Enter (option=value/save/cancel) : 63=1 6

LES/BUS for Slot 4 Port 1

1) ELAN name (32 chars max) : eth1
2) ELAN type { 802.3 (1), 802.5 (2) } : 802.3
3) Max Data Frame Size { 1516 (1), 4544 (2),
    9234 (3), 18190 (4) } : 1516
4) Control time-out { 10 - 300 seconds } : 120
5) Max. Frame age { 1 - 4 seconds } : 1
6) Enable redundancy { No (1), Yes (2) } : YES
61) Redundancy Role { Pri (1), Sec(2) } : Secondary
62) Primary LES's ATM addr or addr index :

      index   ATM address
      ----   -
          1   3999999999999900009999010140008210000003

63) Primary LES local { No (1), Yes (2) } : No
7) Admin Status { Disable (1), Enable (2) } : Enable
8) LES/BUS Security { Disable (1), Enable (2) } : Disable
9) Create default ELAN { No (1), Yes (2) } : No

Enter (option=value/save/cancel) : save

Creating LES/BUS pair for elan 'eth1' on slot 4, port 1, please wait...

LES/BUS pair for elan 'eth1' created on slot 4, port 1
```

Figure 211. Configuring LES/BUS on 8274

- 1 From the LSM configuration menu, select create LES/BUS.
- 2 Set the ELAN name to eth1.
- 3 Enable LES/BUS redundancy.
- 4 Enter the role of the LES/BUS we are configuring on the 8274.
- 5 Enter the primary LES/BUS ATM address. In this example the ATM address of the ELAN eth1 on the MSS is entered.
- 6 Specify whether the primary LES/BUS is local or remote.

Next an LECS was configured on the 8274. This configuration is shown in Figure 212 on page 239.

LSM configuration for slot 4, port 1

- 1) Specify Global elan name (currently none specified)
- 2) Create LES/BUS
- 3) Modify LES/BUS
- 4) Delete LES/BUS
- 5) Create LECS
- 6) Modify LECS
- 7) Delete LECS
- 8) Add elan to LECS database
- 9) Delete elan from LECS database
- 10) Add policy to elan in LECS database
- 11) Delete policy from elan in LECS database
- 12) Exit configuration menu

Enter option : 5

Configuration for LECS at Slot 4 Port 1

- 1) Max Config Direct VCCs to LECS {1 - 128}: 128
- 2) Seconds before VCC declare idle {1 - 43200}: 60
- 3) Priority for ELAN name policies {0 - 65535}: 1
- 4) Priority for ELAN type policies {0 - 65535}: 0 - not used
- 5) Priority for ATM addr prefix policies {0 - 65535}: 0 - not used
- 6) Priority for MAC address policies {0 - 65535}: 0 - not used
- 7) Priority for Max. Frame Size policies {0 - 65535}: 0 - not used
- 8) Priority for Route Descriptor policies {0 - 65535}: 0 - not used
- 9) Admin Status { Disable (1), Enable (2) } : Enable

Enter (option=value/save/cancel) : **save**

Creating LECS on slot 4, port 1, please wait...

LECS created on slot 4, port 1

Figure 212. Configuring LECS on 8274

After configuring the LECS, the ELAN must now be added to the LECS database. This configuration is listed in Figure 213 on page 240 and Figure 214 on page 241.

- 1) Specify Global elan name (currently none specified)
- 2) Create LES/BUS
- 3) Modify LES/BUS
- 4) Delete LES/BUS
- 5) Create LECS
- 6) Modify LECS
- 7) Delete LECS
- 8) Add elan to LECS database
- 9) Delete elan from LECS database
- 10) Add policy to elan in LECS database
- 11) Delete policy from elan in LECS database
- 12) Exit configuration menu

```

1) ELAN name (32 chars max) :
2) Elan type {802.3 (1), 802.5 (2) } : 802.3
3) Max Frame Size { 1516 (1), 4544 (2)
                      9234 (3), 18190 (4) } : 1516
4) Primary LES's ATM address :
    index ATM address

```

Enter (option=value/save/cancel) :

Enter (option=value/save/cancel) : 5=2 **3**

```
Enter (option=value/save/cancel) :  
51=39999999999999000099990101400082740000c1
```

```

5) Backup LES { No (1), Yes (2) } : Yes
51) Backup LES' s ATM address      :
    index   ATM address
-----
         2  39999999999999000099990101400082740000c1

```

```
elan eth1 added to the LECS database
```

1 Select option 8 from LSM configuration menu to add the ELAN to the LECS database.

2 Enter the ELAN name.

3 Specify that there is a backup LES/BUS.

4 Enter the primary LES/BUS ATM address. The ATM address of the LES/BUS in the MSS is used in this example.

5 Enter the backup LES/BUS ATM address. In this example the ATM address of the backup LES/BUS in the 8274 is used.

The last LSM configuration needed is to add a policy to the LECS database for the eth1 ELAN. This is shown in Figure 214.

LSM configuration for slot 4, port 1

- 1) Specify Global elan name (currently none specified)
- 2) Create LES/BUS
- 3) Modify LES/BUS
- 4) Delete LES/BUS
- 5) Create LECS
- 6) Modify LECS
- 7) Delete LECS
- 8) Add elan to LECS database
- 9) Delete elan from LECS database
- 10) Add policy to elan in LECS database
- 11) Delete policy from elan in LECS database
- 12) Exit configuration menu

Enter option : 10 **1**

Enter (elan name) : eth1 **2**

Add policy value to elan 'eth1'

- 1) By Elan name
- 2) By Elan type
- 3) By ATM address prefix
- 4) By MAC address
- 5) By Max. Frame size
- 6) By Route Descriptor
- 7) Exit

Enter option : 1 **3**

Enter elan name : eth1 **4**

policy added to LECS database for elan 'eth1'

Add policy value to elan 'eth1'

Enter option : 7 **5**

Figure 214. Add ELAN Policy to LECS Database

1 Select option 10 from the LSM configuration menu to add a policy to the LECS database for an ELAN.

2 Enter the ELAN you wish to add the policy for. In this example it is eth1.

3 The policy we selected was ELAN name.

4 The name for the policy we set to eth1.

5 Enter 7 to exit the policy menu and save settings.

The LAN emulation service was then modified on the ASM port 4/1, as shown in Figure 215.

```
8274/ % mas 4/1 1

Slot 4 Port 1 Service 1 Configuration

1) Description (30 chars max)      : LAN Emulation Service 1
2) LAN Emulated Group              : 1
   21) LAN type { 802.3 (1),
                  802.5 (2) }      : 802.3
   22) Change LANE Cfg { NO (1),
                          YES (2) } : NO
3) LECS Address (40-char-hex)      :
   4700790000000000000000000000A03E00000100
4) Admin Status { disable(1),
                  enable(2) }      : Enable
6) Connection Type { PVC(1),
                     SVC(2) }      : SVC
60) SEL for the ATM address         : 01

Enter (option=value/save/cancel) : 22=2

Slot 4 Port 1 Service 1 LANE Configuration Parameters

1) Proxy { NO (1), YES (2) }       : YES
2) Max Frame Size { 1516 (1), 4544 (2),
                    9234 (3), 18190 (4) } : 1516
3) Use translation options{NO (1), YES (2) } : Yes (use Swch menu to set)
4) Use Fwd Delay time { NO (1), YES (2) }   : NO
5) Use LE Cfg Server (LECS){ NO (1), YES (2)} : YES
6) Use Default LECS address { NO(1), YES (2)} : YES
7) Control Time-out (in seconds)           : 10
8) Max Unknown Frame Count                  : 10
9) Max Unknown Frame Time (in seconds)      : 1
10) VCC Time-out Period (in minutes)        : 20
11) Max Retry Count                         : 2
12) Aging Time (in seconds)                 : 300
13) Expectd LE_ARP Resp Time (in seconds)   : 1
14) Flush Time-out (in seconds)             : 4
15) Path Switching Delay (in seconds)       : 6
16) ELAN name (32 chars max)                : eth1

Enter (option=value/save/cancel) : 16=eth1

Enter (option=value/save/cancel) : save

Saving new LANE Configuration values

Enter (option=value/save/cancel) : save
Modifying service, please wait...

Resetting service, please wait...
Enabling service...
8274/ %
```

Figure 215. Modifying a Service on the ASM Port of the 8274

```

/ % vas

ATM Services

Slot Port Serv Service Service
      Num Description Type
=====
4 1 1 LAN Emulation Service 1 802.3 LEC
4 1 2 VCM Service 2 VCM
4 1 3 LANE Service Module Service 3 LSM
4 2 1 LAN Emulation Service 1 802.3 LEC
4 2 2 VCM Service 2 VCM

ATM Services

Slot Port Serv VC Oper
      Num Type Status SEL Groups Conn VCI's/Addresses
=====
4 1 1 SVC LANE Op. 01 1 784 785 786 787
4 1 2 SVC Enabled 02 1
4 1 3 SVC Enabled 03 N/A N/A
4 2 1 SVC Initial 01 1
4 2 2 SVC Disabled 02 1

FDDI Services do not exist!

Mammoth Ethernet Services do not exist!

8274/ %

```

With the MSS still running the primary LES/BUS, the backup LES/BUS state was viewed, as shown in Figure 217.

[illegible]

In Figure 218 on page 244 the status of the ELAN eth1 is shown.

```

/Interface/LSM % vlec 4/1 eth1

Number of LEC's to display: 0 1

/Interface/LSM % vlb 4/1 eth1
ELAN Name:                eth1
ELAN Type:                 Ethernet
# of Proxy LEC's:         0 2
# of Non-Proxy LEC's:     0 3
LES ATM Address:          3999999999999999000099990101400082740000c1

-Status-
LES-BUS State:             OPERATIONAL 4
Redundancy VCC State:     ESTABLISHED 5
Major Reason LES-BUS was last Down: none
Minor Reason LES-BUS was last Down: none
LES-BUS State last changed at: 00.08.02.29 (System Up Time)
LES-LEC Status Table changed at: 00.00.00.00 (System Up Time)
BUS-LEC Status Table changed at: 00.00.00.00 (System Up Time)

-Current Configuration-
LES-BUS Enabled/Disabled: Enabled
ELAN Type: (S2)           Ethernet
Max Frame Size: (S3)      1516
Control Timeout: (S4)     120
Max Frame Age: (S5)       1
Redundancy:               Enabled

```

Figure 218. Viewing Status of the ELAN eth1 on the 8274

- 1** There are no LECs that have joined the ELAN eth1 of the backup LES/BUS on the 8274 to display, because the MSS is still running and all the LECs have joined the ELAN on the primary LES/BUS.
- 2** Shows no proxy LECs have joined the backup ELAN on the 8274, as the MSS is still running.
- 3** For the same reason no non-proxy LECs have joined the backup ELAN.
- 4** Shows the backup LES/BUS is up and available.
- 5** Shows the primary and backup LES/BUSes are communicating.

The power to the MSS was then switched off and the output in Figure 219 on page 245, Figure 220 on page 245, and Figure 221 on page 246 were observed.

```

/Interface/LSM % v1b 4/1 eth1
ELAN Name:                eth1
ELAN Type:                 Ethernet
# of Proxy LEC's:         2 1
# of Non-Proxy LEC's:     1 2
LES ATM Address:          39999999999999990000999990101400082740000c1

-Status-
LES-BUS State:             OPERATIONAL 3
Redundancy VCC State:     IDLE 4
Major Reason LES-BUS was last Down: none
Minor Reason LES-BUS was last Down: none
LES-BUS State last changed at: 00.01.02.71 (System Up Time)
LES-LEC Status Table changed at: 00.10.00.53 (System Up Time)
BUS-LEC Status Table changed at: 00.10.00.69 (System Up Time)

-Current Configuration-
LES-BUS Enabled/Disabled: Enabled
ELAN Type: (S2)           Ethernet
Max Frame Size: (S3)      1516
Control Timeout: (S4)     120
Max Frame Age: (S5)       1
Redundancy:               Enabled
Redundancy Role:          Backup LES-BUS

/Interface/LSM %

```

Figure 219. Viewing Stats on LES/BUS on 8274 - MSS Off

- 1** The proxy LECs now join the backup ELAN. The LECS gives the LEC the ATM address of the backup LES/BUS because the primary LES/BUS in the MSS is no longer available.
- 2** For the same reason the non-proxy LECs join the backup ELAN.
- 3** Shows the LES/BUS is up and running.
- 4** Shows no connection can be made to the primary LES/BUS.

Figure 220 shows the LECs that have registered in the backup ELAN in more detail.

```
/Interface/LSM % vlec 4/1 eth1
```

Number of LEC's to display: 3

LEC-LES and LEC-BUS State (UP=Up, ID=Idle, --. --.
**=Other; Show specific LEC to see actual)

LEC Primary ATM Address	Proxy	LEC ID	State LES BUS	#ATM Adrs	#Reg MACs
39999999999900009999010140008274000001	Y	0001	UP UP	1	1
39999999999900009999010140008273000002	Y	0002	UP UP	1	1
3999999999990000999901010004ac2c359481	N	0003	UP UP	1	1

```
/Interface/LSM %
```

Figure 220. LEC Registration

```

8260_HUB1> show reachable_address 13.a11
Port      Len Address                                     Active Idx      VPI
-----
13.02 152 39.99.99.99.99.99.00.00.99.99.01.01.40.00.82.74.00.00
                                Y Dyn 0
13.02 152 47.00.79.00.00.00.00.00.00.00.00.00.00.00.A0.3E.00.00.01
                                Y Dyn 0
13.03 152 39.99.99.99.99.99.00.00.99.99.01.01.40.00.82.73.00.00
                                Y Dyn 0
8260_HUB1>

```

The MSS was then powered back up and in Figure 222, there are no longer any LECs joining the ELAN on the backup LES/BUS. These LECs are moved to the ELAN in the primary LES/BUS on the MSS.

```

/Interface/LSM % vlec 4/1 eth1

Number of LEC's to display: 0 1

/Interface/LSM % vlb 4/1 eth1
ELAN Name:                eth1
  ELAN Type:                Ethernet
  # of Proxy LEC's:        0
  # of Non-Proxy LEC's:    0
  LES ATM Address:         3999999999999000099990101400082740000c1

-Status-
  LES-BUS State:            OPERATIONAL
  Redundancy VCC State:    ESTABLISHED 2
  Major Reason LES-BUS was last Down: none
  Minor Reason LES-BUS was last Down: none
  LES-BUS State last changed at: 00.01.02.71 (System Up Time)
  LES-LEC Status Table changed at: 00.22.03.43 (System Up Time)
  BUS-LEC Status Table changed at: 00.22.03.43 (System Up Time)

-Current Configuration-
  LES-BUS Enabled/Disabled: Enabled
  ELAN Type: (S2) Ethernet
  Max Frame Size: (S3) 1516
  Control Timeout: (S4) 120
  Max Frame Age: (S5) 1
  Redundancy: Enabled
  Redundancy Role: Backup LES-BUS

/Interface/LSM %

```

1 Shows no LECs on the ELAN in backup LES/BUS.

2 Shows the connection between the primary and backup LES/BUSes has been established.

Even though the MSS has resumed control of the LES/BUS, the LECS is still running on the 8274. In Figure 223 the well-known address is still shown on port 13.2 on the 8260.

```
8260_HUB1> show reachable_address 13.all
```

Port	Len	Address	Active Idx	VPI
13.01	152	39.99.99.99.99.99.00.00.99.99.01.01.00.04.13.47.39.36		
		Y Dyn 0		
13.01	152	39.99.99.99.99.99.00.00.99.99.01.01.40.00.82.10.00.00		
		Y Dyn 0		
13.02	152	39.99.99.99.99.99.00.00.99.99.01.01.40.00.82.74.00.00		
		Y Dyn 0		
13.02	152	47.00.79.00.00.00.00.00.00.00.00.00.00.00.00.A0.3E.00.00.01		
		Y Dyn 0		
13.03	152	39.99.99.99.99.99.00.00.99.99.01.01.40.00.82.73.00.00		
		Y Dyn 0		

```
8260_HUB1>
```

Figure 223. Reachable Addresses on 8260 after Powering On MSS

8.2.4 Example of LES/BUS Redundancy with IP Routing

In 8.2.3, “Example of Redundant LES/BUS” on page 230, the 8274 was the backup LES/BUS for the MSS. In MSS configurations on production networks, however, there are different protocols running. These also need to be taken into consideration when configuring the 8274 to back up the MSS.

It is important to note that the MSS is more functional than the 8274 in its protocol support. The 8274 only supports the routable protocols IP and IPX for ELANs. This means that all other protocols running on the MSS can not be supported by the 8274 when the 8274 is the backup LES/BUS.

In Release 1.1 and later of the MSS code, the same IP address of a gateway can be configured on both the primary and backup MSS. In the configuration it is specified that the IP address is primary on the primary MSS and backup on the backup MSS.

In Release 1.0 of the MSS code, the same IP address for a gateway on the primary MSS could also be configured on the backup MSS, but there were no parameters that could specify that the IP address was primary or backup. The way backup IP addresses were configured on the MSS in Release 1.0 is similar to the way we are required to configure it for the 8274.

This example shows the configuration required to back up the LES/BUS of the MSS on the 8274, and also have backup IP gateways so that no IP connectivity is lost should the MSS fail.

The configuration follows on from that in 8.2.3, “Example of Redundant LES/BUS” on page 230.

A physical representation of the example is shown in Figure 224 on page 248.




Figure 224. Physical View of the Network

The logical diagram is shown Figure 225.

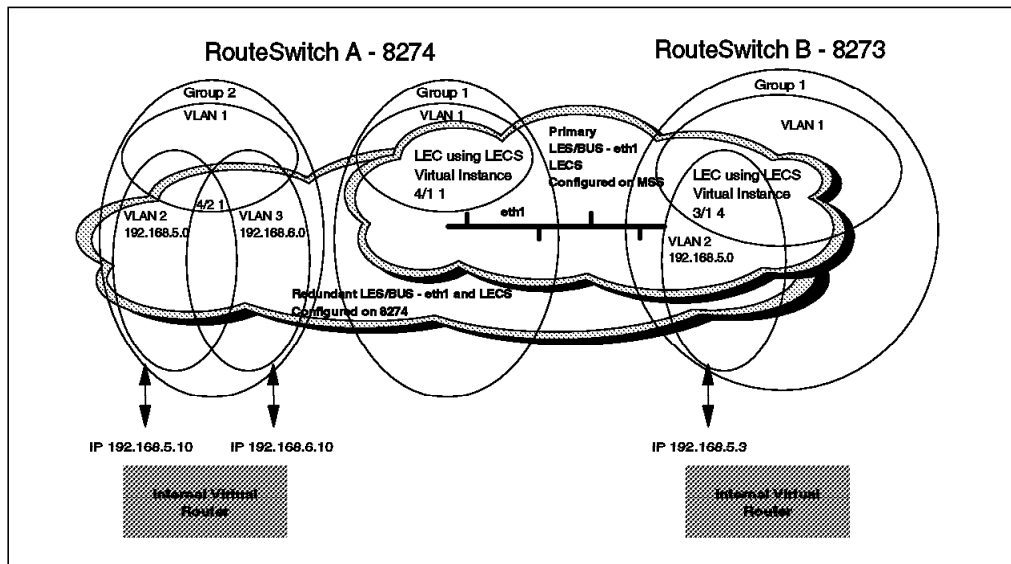


Figure 225. Physical View of the Network

In Figure 226 on page 249 an IP instance is added to the ELAN eth1 on the MSS. The list of instances is also shown.

```

IP config>add address 2
New address (0.0.0.0)? 192.168.6.10
Address mask (255.255.255.0)?

IP config>list all
Interface addresses
IP addresses for each interface:
  intf 0  192.168.21.10  255.255.255.0  Local wire broadcast, fill 1
  intf 1  192.168.4.10   255.255.255.0  Local wire broadcast, fill 1
  intf 2  192.168.5.10   255.255.255.0  Local wire broadcast, fill 1
              192.168.6.10  255.255.255.0  Local wire broadcast, fill 1

Routing

Protocols
BOOTP forwarding: disabled
IP Time-to-live: 64
Source Routing: enabled
Echo Reply: enabled
Directed broadcasts: enabled
ARP subnet routing: disabled
ARP network routing: disabled
Per-packet-multipath: disabled
OSPF: enabled
BGP: disabled
RIP: enabled
RIP default origination: disabled
  Per-interface address flags:
    intf 0  192.168.21.10  RIP disabled for this interface
    intf 1  192.168.4.10   Send net, subnet, static and default routes
                          Received RIP packets are ignored.
    intf 2  192.168.5.10   Send net, subnet, static and default routes
                          Received RIP packets are ignored.
                          192.168.6.10  Send net and subnet routes
                          Receive No Dynamic host routes

Accept RIP updates always for:
(NONE)

IP config>exit
Config>write
Config Save: Using bank A and config number 3
*reload
Are you sure you want to reload the gateway? ((yes) or (No)): y

```

Figure 226. Configuring Second IP Instance on MSS

A second group is defined on the 8274 that will be used only for the configuration of VLANs for IP routing. The IP addresses used will be the same as the ones used on the MSS, and a LEC will be defined for group 2 that will only be able to join the backup ELAN when the MSS fails. The backup IP addresses will thus only become active in the network when the MSS fails.

In Figure 227 on page 250 group 2 is created, and the port 4/2 is added to the new group.

```

/ % crgp
  GROUP Number ( 2) :
  Description (no quotes) : IP backup
  Enable WAN Routing? (n):
  Enable ATM CIP? (n):
  Enable IP (y) : n
  Enable IPX? (y): n
  Enable Group Mobility on this Group ? :[y/n:]: n
  This Group will not participate in Group Mobility

  Do you wish to configure the interface group for this Virtual LAN
  at this time? (y) y

  Initial Vports(Slot/Phys Intf. Range) - For example, first I/O Module
  (slot 2), second Interface would be 2/2. Specify a range of interfaces
  and/or a list as in: 2/1-3, 3/3, 3/5, 4/6-8.

  Initial Slot/Interface Assignments: 4/2
  Please use the 'cas' command to add ATM services.
  You may modify interfaces to this group using the addvp, modvp and rmvp
  commands at a later date if you choose.
/ %

```

Figure 227. Creating a New Group on 8274

We now create a LAN emulation service on ASM port 4/2 to join the backup ELAN on the 8274.

/VLAN % **cas 4/2**

Slot 4 Port 2 Service 1 Configuration

- 1) Description (30 chars max) : LAN Emulation Service 1
- 2) Service type { LANE client(1),
Trunking (4),
Classical IP(5),
PTOP Bridging(6),
VLAN cluster(7) } : LAN Emulation
- 21) LAN type { 802.3 (1),
802.5 (2) } : 802.3
- 22) Change LANE Cfg { NO (1),
YES (2) } : NO
- 3) Connection Type { PVC(1),
SVC(2) } : SVC
- 30) SEL for the ATM address : 03
- 4) LAN Emulated Group : 2 **1**
- 5) LECS Address (40-char-hex) :
470079000000000000000000000000A03E00000100
- 6) Admin Status { disable(1),
enable(2) } : Enable

Enter (option=value/save/cancel) : **22=2 2**

Enter (option=value/save/cancel) : **5=1 3**

Slot 4 Port 2 Service 1 LANE Configuration Parameters

- 1) Proxy { NO (1), YES (2) } : Yes
- 2) Max Frame Size { 1516 (1), 4544 (2),
9234 (3), 18190 (4) } : 1516
- 3) Use translation options{NO (1), YES (2) } : Yes
(use Swch menu to set)
- 4) Use Fwd Delay time { NO (1), YES (2) } : NO
- 5) Use LE Cfg Server (LECS){ NO (1), YES (2)}: NO
- 6) Use Default LECS address { NO(1), YES (2)}: YES
- 7) Control Time-out (in seconds) : 10
- 8) Max Unknown Frame Count : 10
- 9) Max Unknown Frame Time (in seconds) : 1
- 10) VCC Time-out Period (in minutes) : 20
- 11) Max Retry Count : 2
- 12) Aging Time (in seconds) : 300
- 13) Expectd LE_ARP Resp Time (in seconds) : 1
- 14) Flush Time-out (in seconds) : 4
- 15) Path Switching Delay (in seconds) : 6
- 16) ELAN name (32 chars max) :

Enter (option=value/save/cancel) : **save 4**

Saving new LANE Configuration values

Figure 228 (Part 1 of 2). Creating a Service on the 8274 for Group 2 on ASM Port 4/2

Figure 228 (Part 2 of 2). Creating a Service on the 8274 for Group 2 on ASM Port 4/2

7 Save the service configuration to create and enable the service.

8274/VLAN % viatr1					
VLAN Group:	VLAN Id	Rule Num	Rule Type	Rule Status	Rule Definition
2: 2		1	NET ADDR RULE	Enabled	IP Addr = 192.168.5.0 IP Mask = 255.255.255.0
		2	PORT RULE	Enabled	4/2/Lne/1
2: 3		1	NET ADDR RULE	Enabled	IP Addr = 192.168.6.0 IP Mask = 255.255.255.0
		2	PORT RULE	Enabled	4/2/Lne/1

8274/VLAN %

In Figure 230 the status of the backup LES/BUS is shown. The MSS is still up and running so no LECs can join this ELAN, including the LANE service on port 4/2.

Figure 230. Viewing Stats - 8274 LES/BUS Backup to MSS

In Figure 231 on page 254 the operational status of service 1 on port 4/2 shows that it is trying to join the backup ELAN on the 8274. It will not be able to join the backup LES/BUS until the primary LES/BUS fails.

/ % vas

		ATM Services				
		Serv	Service	Service		
Slot	Port	Num	Description	Type		
====	====	====	=====	=====		
4	1	1	LAN Emulation Service 1	802.3	LEC	
4	1	1	VCM Service 1	VCM		
4	1	3	LANE Service Module Service 3	LSM		
4	2	1	LAN Emulation Service 1	802.3	LEC	
4	2	2	VCM Service 2	VCM		

ATM Services

		Serv	VC	Oper		
Slot	Port	Num	Type	Status	SEL Groups	Conn VCI's/Addresses
====	====	====	====	=====	=====	=====
4	1	1	SVC LANE Op.	01 1		173 174 175 176
4	1	1	SVC Enabled	01 1		
4	1	3	SVC Enabled	03 N/A		N/A
4	2	1	SVC Join	01 2		
4	2	2	SVC Enabled	02 1		

Figure 231. Viewing a Service on the 8274 with MSS Still Operational

The MSS is now powered off. In Figure 232, service 1 on port 4/2 now joins the backup ELAN on the 8274. This means that devices in group 1 can use the IP gateway addresses defined in group 2 in the RouteTracker VLANs.

/ % vas

ATM Services

Service

Service

Type

Slot

Port

Serv Num

Description

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

=====

Figure 232. Viewing a Service on the 8274 with MSS Not Operational

Figure 233 on page 255 shows the ATM network prefix and ESI on each ASM port.


```

/ % vap

ATM Port Table

Slot Port      ATM Port Description      Conn Tran  Media UNI Max  VCI
====  ====  =====  =====  =====  ==  ==  ==
4      2      ATM PORT      SVC  STS3c  Multi Pri  1023  10

Slot Port Loopback Cfg Tx Clk Source
====  ====  =====  =====
4      1      NoLoop      LocalTiming
4      2      NoLoop      LocalTiming

Slot Port      ATM Network Prefix      End System  Sig Sig  ILMI  ILMI
====  ====  =====  =====  Ver VCI  Enable VCI
4      1      399999999999999999990000999990101  400082740000  3.0 5  True  16
4      2      399999999999999999990000999990101  400082740000b  3.0 5  True  16

Status
-----
Slot Port Tx Seg Sz Rx Seg Sz Tx Buff Sz Rx Buff Sz Oper SSCOP ILMI
====  ====  =====  =====  =====  =====  =====  =====
4      1      16384      16384      4600      4600  Enb(SVC)  Up  Up
4      2      16384      16384      4600      4600  Enb(SVC)  Up  Up

```

Figure 233. Viewing a Port

Viewing the status of the backup ELAN eth1 in Figure 234 on page 256, proxy and non-proxy LECs have joined the backup ELAN. On closer analysis of the LECs we see the ATM addresses shown in Figure 233 for ports 4/1 and 4/2.

Figure 234. Viewing Stats of LES/BUS on 8274

Chapter 9. Wide Area Networking

The WAN switching module (WSM) consists of high-speed high-performance serial interfaces for connecting switches across a frame relay network. The WSM functions such as management, data handling, compression, and multi-protocol encapsulation are compatible with the current frame relay standards such as RFC 1490 and FRF.9.

Switches with the WSM can also be connected "back-to-back" over leased T1/E1 lines. Interoperability exists for connecting to routers, including other vendors products.

Some of the advantages of using the wide area switching module in your 8274 network include:

- Large geographical distances
- VLAN preserved across the WAN
- Trunking across the WAN
- Compatible with any-to-any switching
- Congestion controls
- Extensive statistics
- Self configuration
- Automatic cable detection
- Supports common frame relay standards

9.1 Interfaces

Table 20 shows the cable interface specifications and the line speeds that the WSM currently supports.

<i>Table 20. Interface Specifications</i>		
Interface Cable	Minimum Speed	Maximum Speed
RS-232	9600	64 Kbps
RS-449	9600	2 Mbps
RS-530	9600	2 Mbps
V.35	9600	2 Mbps
X.21 (Europe)	9600	2 Mbps

The WSM is designed to require as little configuration as possible. It will automatically sense the attached cable type installed and map virtual circuits to virtual ports. When the WSM is installed it is automatically configured to the default settings. In default mode the WSM uses the frame relay protocol and for the serial ports (RS-232) the speed is set to 64 kbps. For other cable types the default is 2 Mbps. The configuration menu for the WSM module can be reached by entering wan on the command line. Figure 235 on page 258 shows the menu used to configure the WSM modules.

Command	Wide Area Networking Menu
wpmodify	Modify a given WAN port's parameters
wpdelete	Delete a given port's parameters, and restore defaults
wpview	View WAN port parameters for a given slot and port
wpstatus	View WAN port status of entire chassis, slot, or individual port
fr	Enter the Frame Relay submenu
ppp	Enter the PPP submenu
isdn	Enter the ISDN-specific submenu
link	Enter the link-specific submenu
Main	File Summary VLAN Networking
Interface	Security System Services Help

Figure 235. WSM Configuration Screen

9.2 Compression

Data compression allows you to get more data through the frame relay pipeline, further enhancing cost benefits. The typical data compression rate on the WSM board at the hardware level is 4:1. In addition, the compression processor (STAC 9705) has its own memory that can store up to 100 virtual circuits (on a 4-port WSM) and 200 virtual circuits on an 8-port WSM without performance degradation.

In order to use compression on the WSM you must enable compression negotiation through software and the bridge/router at the other end of the frame relay virtual circuit must support standard FRF.9 compression. (A RouteSwitch-to-RouteSwitch connection would support compression.) The compression negotiation is specified on the frmodify command. Figure 236 on page 259 shows the modify screen where option 5 would be set to enabled to allow for negotiation.

```

Modifying Frame Relay DLCI for Slot 3, Port 1, DLCI 17.
1) Administrative State .....= U
{(U)p, (D)own}
2) Committed Information Rate (CIR) in BPS ..... = 0
{0 through line speed in BPS}
3) Committed Burst Rate(Bc) ..... = 0
{0 through positive number in bits}
4) Excess Burst Rate(Be) ..... = 0
{0 through positive number in bits}
5) Compression Administrative Status ..... = Enabled
{(E)nabled, (D)isabled}
6) Compression PRetry Time ..... =3
{1..10}
7) Compression PRetry Count ..... =10
{3..255}
To change a value, enter the corresponding number, an '=', and the new
value. For example to set a new DLCI Active/Inactive Traps, use
: 5=d
When complete enter "save" to save all changes, or cancel or Ctrl-C to
cancel all changes. Enter ? to view the new configuration.

```

Figure 236. Compression Negotiation

Negotiation is necessary because if compressed data is sent to a bridge/router that does not support compression, then the remote bridge/router will not recognize the data and will automatically drop the unrecognizable frames. If compression negotiation is enabled, the WSM will query the frame relay device on the other end of the circuit (according to FRF.9 specifications) to see if it supports compression. If it does, then the WSM compresses all data except DLCMI (management) data. If it doesn't, then data on that virtual circuit is sent uncompressed.

Note

Compression is not supported on the 2-port WSM modules.

9.3 Virtual Circuits and DLCIs

The WSM can support up to 256 permanent virtual circuits (PVCs). The WSM does not support switched virtual circuits (SVCs). PVCs may either be configured or dynamically learned. User-configured PVCs consist of management, control, and any configured data PVCs. Management VCs are used by the WSM to communicate with the frame relay network. Dynamic PVCs are usually data circuits, which are controlled by the frame relay network and not configured in advance. A logical frame relay DTE device such as the WSM does not create or control dynamic data VCs. It is informed of their status through periodic status updates from the frame relay network. Each virtual circuit is locally defined by a Data Link Connection Identifier (DLCI). The frame relay network (generally the carrier) assigns the DLCIs and informs the WSM about them.

The following steps describe the WSM self configuration of virtual circuits.

1. Cable is plugged into WSM.
2. WSM senses cable type and adjusts circuitry.

3. WSM send Status Enquiry message to the frame relay network. This message requests the DLCIs of all virtual circuits and their status (active/inactive).
4. Frame relay network returns a status message the includes the DLCIs for all the virtual circuits and their status.
5. WMS maps each DLCI (virtual circuit) to a virtual port within the switch. In this auto configuration mode all of these virtual ports are assigned to VLAN 1 (default VLAN).

Figure 237 shows a graphical representation of this process.




Figure 237. Virtual Circuit to Virtual Port Mapping

9.4 Congestion Control

Congestion control can be defined as a set of mechanisms incorporated to attain certain network performance objectives, particularly in the peak periods, while optimizing or improving the network resource requirements. It aims to minimize the number of occurrences of user-perceived congestion. Frame relaying networks should not allow users to monopolize network resource usage at the expense of other users. Congestion control includes both congestion avoidance and congestion recovery mechanisms.

9.4.1 FECN and BECN

The network is able to inform endstations about congestion by using two fields in the frame address field. For this purpose the forward explicit congestion notification (FECN) bit and the backward explicit congestion notification (BECN) bit have been reserved. The FECN bit will be set in frames flowing in the direction in which the network is experiencing congestion. The BECN field will be set in frames flowing in the opposite direction in which the network is experiencing congestion. The WSM supports BECN and FECN and can adjust traffic flow down to the Committed Information Rate (CIR) or below.

9.4.2 Comitted Information Rate and Burst Sizes

Congestion control based on the discarding of frames or the use of FECN/BE CN bits and relying on the "good behavior" of endstations has therefore been considered inadequate to networks providing a frame relaying service and additional provisions have been defined.

9.4.2.1 Comitted Information Rate and Burst Sizes

The maximum number of bits per seconds that an endstation can transmit into the network is bounded by the *access rate* of the user-network interface. The access rate is limited by the line speed of the user-network connection and established by subscription.

The maximum committed amount of data which a user may offer to the network is defined as committed burst size (B_c). B_c is a measurement for the volume of data for which the network will guarantee message delivery under normal conditions. It is measured during the committed rate measurement interval (T_c).

End stations are allowed to transmit data in excess of the committed burst rate. The excess burst size (B_e) has been defined as the allowed amount of data by which a user can exceed B_c during the committed measurement rate interval T_c . If spare capacity exists, the network will forward the data to its destination. The network however is free to mark the data as discard eligible (DE).

The committed information rate (CIR) has been defined as the allowed amount of data that the network is committed to transfer under normal conditions. The rate is averaged over a increment of time T_c . The CIR is also referred to as minimum acceptable throughput.

B_c and B_e are expressed in bits, T_c in seconds, the access rate and CIR in bits per second. B_c , B_e , T_c and CIR are defined per DLCI. The access rate is valid for the user-network interface. For B_c , B_e and CIR incoming and outgoing values can be distinguished. If the connection is symmetrical, the values in both directions are the same. For permanent virtual circuits, B_c (incoming and outgoing), B_e (incoming and outgoing) and CIR (incoming and outgoing) are defined at subscription time. They are negotiated for SVCs at call establishment time. T_c is calculated as depicted in Table 21.

Table 21. Measurement Interval Calculation			
CIR	B_c	B_e	Measurement Interval (T_c)
> 0	> 0	> 0	$T_c = B_c / CIR$
> 0	> 0	0	$T_c = B_c / CIR$
0	0	> 0	$T_c = (B_e / \text{Access Rate}) \bullet$
Notes:			
1. Table depicts the valid parameter configurations. Other configurations are for further study.			
2. When the two communicating endstations have different access rates, the network may define a smaller T_c value.			

Individual CIRs on a physical connection are always lower than the access rate; however, the sum of CIRs defined can be larger than the access rate. An example could be a network connection with an access rate of 256 kbps on which three virtual circuit have been defined two having a CIR of 128 kbps each, one having a CIR of 64 kbps.

Optimal values for the above parameters depend on network implementation, availability of spare network capacity, charging methods, type of user device and performance required. Only a number of considerations are mentioned and careful study is required.

Networks may mark frames above B_c with discard eligible (DE) but have plenty of spare capacity to transport the frame, or the reverse, have limited capacity and discard excessive frames immediately. Networks may mark frames above $B_c + B_e$ with discard eligible (DE), and possibly transport it, or just drop the frames as suggested by ITU-T I.370.

Network manager always try to balance costs and performance and have to examine the frame relay service provider charging schemes. Networks may implement fixed charging dependent on access rate, a scheme dependent on CIR, B_c and B_e or more sophisticated schemes for example charging on number of bits transported and charging progressively for data above B_c or $B_c + B_e$. Depending on the charging scheme employed subscribing to high values of CIR, B_c and B_e , may lead to high networking costs. It should be examined if the performance gain, if any, counterbalances the additional networking expenses.

Many devices have limited control over the volumes of data they send into the network. Assuming flow control mechanisms implemented on top of the layer two core functions are not inhibiting data transfer, data will be transmitted with a speed up to the network access rate. If the device has only one DLCI active, or has (temporarily) data to send for one DLCI only, the data rate on a single DLCI may be equal to the network access rate. If the sum of committed and excess burst size ($B_c + B_e$) is lower than the access rate times T_c , the network may decide to discard frames. In this situation it may be advisable to subscribe for all DLCIs to:

$$B_c + B_e = \text{Access rate} * T_c$$

Depending on functions implemented on top of the layer two core functions, lost frames may be quickly detected and recovered from. This may be a time-consuming activity severely impacting performance. In the latter case subscribing to high values of CIR, B_c and B_e is important.

As with compression, these values (Committed Information Rate (CIR), Committed Burst Rate (B_c), and Excess Burst Rate (B_e)) can be altered. The command `frmodify` is used, where the syntax is `frmodify <slot>/<port>/<DLCI>`. Figure 238 on page 263 shows the console screen where these parameters can be changed.


```

Modifying Frame Relay DLCI for Slot 3, Port 1, DLCI 17.
1) Administrative State .....= U
{(U)p, (D)own}
2) Committed Information Rate (CIR) in BPS ..... = 0
{0 through line speed in BPS}
3) Committed Burst Rate(Bc) ..... = 0
{0 through positive number in bits}
4) Excess Burst Rate(Be) ..... = 0
{0 through positive number in bits}
5) Compression Administrative Status ..... = Enabled
{(E)nabled, (D)isabled}
6) Compression PRetry Time ..... =3
{1..10}
7) Compression PRetry Count ..... =10
{3..255}
To change a value, enter the corresponding number, an '=', and the new
value. For example to set a new DLCI Active/Inactive Traps, use
: 5=d
When complete enter "save" to save all changes, or cancel or Ctrl-C to
cancel all changes. Enter ? to view the new configuration.

```

Figure 238. Congestion Control

9.5 WSM Routing and Interoperability Examples

The following examples show the setup and configuration of the RouteSwitches in a frame relay environment. The examples include back-to-back, with a frame relay group as well as IP routing with another RouteSwitch and an IBM 2210.

9.5.1 Back-to-Back Bridging/Routing Using a Frame Relay Group

This first example is of a self-configured pair of RouteSwitches. They are set up in a back-to-back configuration. Figure 239 shows the physical connection.




Figure 239. Back-to-Back Self Configuration

Two types of configurations are possible using this method: bridging or routing of IP. Both methods will be described

First described will be the bridging method. Very little user input has to be entered to configure this environment. As in most bridged environments the TCP/IP subnet remains the same. Both RouteSwitches will be configured using the same subnet with different host addresses. Figure 240 on page 264 shows

the configuration necessary to set the TCP/IP addresses. Subnet 10.2.2.x is used. Only switch one is shown.

```
/ % modvl 1
Current values associated with GROUP 1.1 are as follows:

1) GROUP Number          - 1:1
2) Description            - Default GROUP (#1)
IP parameters:
3) IP enabled             - Y
4) IP Network Address     - 192.168.10.1
5) IP Subnet Mask         - 255.255.255.0
6) IP Broadcast Address   - 192.168.10.255
7) Router Description     - GROUP #1.0 IP router vport
8) RIP Mode               - Silent
   {Active(a), Inactive(i), Deaf(d), Silent(s)}
9) Default Framing        - Ethernet II
   {Ethernet II(e), Ethernet 802.3(8), fddi(f),
    token ring(t), source route token ring(s)}
IPX parameters:
10) IPX enabled           - N

(save/quit/cancel)
: 4=10.2.2.1
New IP address generates new subnet and broadcast addresses.
Enter '?' to view the changes.
: 5=255.255.255.0
New mask caused change in broadcast address.
: save
```

Figure 240. Setting the TCP/IP Address

After both switches have been configured with their respective addresses and properly connected, either back-to-back with DCE/DTE cables or the E1/T1 connection the frs (frame relay service) and vas (view a service) commands can be used to check the status of the connections. Figure 241 on page 265 shows the output generated from these commands. Note that the switch recognized the DCE cable and has automatically assigned the service to virtual port 16.

```

/ % frs
Frame Relay Status for the Chassis:

```

Slot/Port	Admin/ Oper Status	Type	Intf BPS	Speed Clocking	Inactive
=====	=====	=====	=====	=====	=====
6/1	UP/DN	V35DCE	2048000	Split	0/0
6/2	UP/DN	*NONE*	EXT CLK	External	0/0
6/3	UP/DN	*NONE*	EXT CLK	External	0/0
6/4	UP/DN	*NONE*	EXT CLK	External	0/0
6/5	UP/DN	*NONE*	EXT CLK	External	0/0
6/6	UP/DN	*NONE*	EXT CLK	External	0/0
6/7	UP/DN	*NONE*	EXT CLK	External	0/0
6/8	UP/DN	*NONE*	EXT CLK	External	0/0

```

/ % vas

```

Frame-Relay Services						
Slot	Oper				Service	Service
Port	Sta.	VCs	Groups	Number	Vport	Description
=====	=====	=====	=====	=====	=====	=====
6/1	UP	32	1	1	16	Frame-Relay-BridgingBridging

Figure 241. Frame Relay Services

The setup for a routed environment is nearly the same. The same physical setup is used for the routed environment. The difference in this configuration is that a second group will be added to the RouteSwitches. In this configuration the second group will have a different TCP/IP address and all traffic will be routed across the frame relay or back-to-back connection. Effectively the wide area portion becomes a separate subnet. This allows multiple subnets to exist on each side of the wide area link. These subnets must, however, have different addresses. Again only one switch configuration will be shown.

Figure 242 on page 266 shows the new group being created with a different TCP/IP address. The `crgp` command is used to add and assign the address. The IP subnet 10.1.1.x will be used on the wide area link, where switch A will be assigned 10.1.1.1 and switch B 10.1.1.2. IPX was also configured to route across the link using subnet 100.

```

/ % crgp
GROUP Number ( 2 ) :
Description (no quotes) :
Enable FR Routing? (n): y
Enable IP (y) :
    IP Address : 100.1.1.1
    IP Subnet Mask (0xff000000) : 255.255.255.0
    IP Broadcast Address (100.1.1.255) :
    Description (30 chars max) :
    IP RIP mode {Deaf(d),
        Silent(s),
        Active(a),
        Inactive(i)} (a) :
Enable IPX? (y):
    IPX Network : 100
    Description (30 chars max) :
    IPX RIP and SAP mode {RIP and SAP active(a),
        RIP only active(r),
        RIP and SAP inactive(i)} (a) :

GROUP 2 has been added to the system.

```

Figure 242. Creating a Frame Relay Routing Group

Next the frame relay port needs to be altered to join the newly created group, Figure 243 on page 267 shows the command frmod used to alter its configuration.

Note: Option eleven is altered from its default to the new group two.

```

/ % frmod 6/1
Modify Frame Relay port for Slot: 6, Port: 1.

1) Speed in BPS ..... = 0
   {9600, 19200, 56000, 64000, 128000, 256000, 512000, 768000}
   {1024000, 1544000 2048000}
2) Clocking ..... = External
   {(I)nternal, (E)xternal, (S)plit}
3) DLCMI Type ..... = ANSI T1.617 Annex D
   {(L)MI Rev. 1.0, T1.617 Annex (D), Q.933 Annex (A), (N)one }
4) Polling Interval T391/nT1 in seconds ..... = 10
   {1 through 255 seconds}
5) Full Status Interval N391/nN1 ..... = 6
   {1 through 10}
6) Error Threshold N392/nN2 ..... = 3
   {1 through 10}
7) Monitored Events Counter N393/nN3 ..... = 4
   {1 through 10}
8) Administrative Status ..... = UP
   {(U)p, (D)own}
9) Default Bridging Group ..... = 1
   {1-65535}
10) Default Frame-Relay Bridging Mode ..... = Bridge All
    {Bridge (A)ll, (E)thernet only}
11) Default Routing Group ..... = 0
    {1-65535}
12) Default Compression Admin Status ..... = Enabled
    {(E)nable, (D)isable}
13) Default Compression PRetry Time ..... = 3
    {1-10}
14) Default Compression PRetry Count ..... = 10
    {3-255}
15) Description ..... =
    {Enter up to 30 characters}

(save/quit/cancel)
: 11=2
: save

```

Figure 243. Altering the Frame Relay Port

Lastly, by viewing the group and the service menu it can be seen that the switch automatically created the new service in group 2. Figure 244 on page 268 shows the newly created group and service.

/ % gp			
Group		NetworkProto/	
ID	Group Description	(IP Subnet Mask)	Encaps
(VLAN ID)	or (IPX Node Addr)		
=====			
1	Default GROUP (#1)	10.2.2.1	IP/
		(ff.ff.ff.00)	ETH2
2	New GROUP (#2)	10.1.1.1	
		(ff.ff.ff.00)	IP/
			1490
/ % vas			
Frame-Relay Services			
Slot	Oper	Service	Service
Port	Sta.	VCs	Groups
		Number	Vport
		Description	Type
=====			
6/1	UP	32	2
		2	24
		Frame-Relay-Routing	Routing

Figure 244. Group 2 Description

9.5.2 Frame Relay Routing with an IBM 2210

This example shows how the IBM 2210 can interoperate with the 8274's WSM module. In this example the frame relay network once again is used as a separated subnet. To begin the process the mpm.cfg and mpm.cnf files were deleted off of the RouteSwitches to set them to the default state. Figure 245 shows the physical network connections.




Figure 245. 8274 with an 2210

The first step is to configure the IP address on each switch. Switch A will be assigned 10.2.2.2, switch B 10.2.2.1 and the 2210 will use 10.2.2.3. Figure 246 on page 269 shows how to assign the IP address to the 8274. In this case no VLANs were configured and the IP address was assigned to the default VLAN. Only switch A's configuration is shown.

```

/ % modvl 1
Current values associated with GROUP 1.1 are as follows:

1) GROUP Number          - 1:1
2) Description            - Default GROUP (#1)
IP parameters:
3) IP enabled             - Y
4) IP Network Address     - 192.168.10.1
5) IP Subnet Mask         - 255.255.255.0
6) IP Broadcast Address   - 192.168.10.255
7) Router Description     - GROUP #1.0 IP router vport
8) RIP Mode               - Silent
   {Active(a), Inactive(i), Deaf(d), Silent(s)}
9) Default Framing        - Ethernet II
   {Ethernet II(e), Ethernet 802.3(8), fddi(f),
    token ring(t), source route token ring(s)}
IPX parameters:
10) IPX enabled           - N

(save/quit/cancel)
: 4=10.2.2.2
New IP address generates new subnet and broadcast addresses.
Enter '?' to view the changes.
: 5=255.255.255.0
New mask caused change in broadcast address.
: save

```

Figure 246. Changing the Default IP Address

Since this configuration will use routing through the frame relay network a new frame relay group must be created in each switch. Refer to Figure 242 on page 266 for the configuration of the group. Additionally the frame relay port must be configured into the new group. Figure 243 on page 267 shows the commands necessary to move the port into the new group.

Once the IP address is set and the ports have been modified, the switch learns that the DLCI is associated with its interfaces. This process is described in 9.3, "Virtual Circuits and DLCIs" on page 259. Once this process has been completed the the commands frs and vas were entered to view the frame relay configuration and status. Figure 247 on page 270 shows the status screen.

```

/ % frs
Frame Relay Status for the Chassis:

      Admin/      VC
Slot/Port Oper  Intf  Speed      Clocking  Active/
=====
6/1      UP/UP   V35DTE      EXT CLK   External 2/0
6/2      UP/DN   *NONE*      EXT CLK   External 0/0
6/3      UP/DN   *NONE*      EXT CLK   External 0/0
6/4      UP/DN   *NONE*      EXT CLK   External 0/0
6/5      UP/DN   *NONE*      EXT CLK   External 0/0
6/6      UP/DN   *NONE*      EXT CLK   External 0/0
6/7      UP/DN   *NONE*      EXT CLK   External 0/0
6/8      UP/DN   *NONE*      EXT CLK   External 0/0

/ % vas

                        Frame-Relay Services
Slot Oper      Service
Port Sta. VCs  Groups Number Vport Description      Service
=====
6/1  UP   102  2      1    16  Frame-Relay-Routing  Routing
6/1  UP   103  2      2    24  Frame-Relay-Routing  Routing

```

Figure 247. Frame Relay Status

Figure 248 shows switch C's configuration.

```

/ % frs
Frame Relay Status for the Chassis:

      Admin/      VC
Slot/Port Oper  Intf  Speed      Clocking  Active/
=====
3/1      UP/UP   V35DTE      EXT CLK   External 1/0
3/2      UP/DN   *NONE*      EXT CLK   External 0/0

/ % vas

                        Frame-Relay Services
Slot Oper      Service
Port Sta. VCs  Groups Number Vport Description      Service
=====
3/1  UP   102  2      1    16  Frame-Relay-Bridging  Routing

```

Figure 248. Frame Relay Status

For simplicity the qconfig command will be used to configure the 2210.

The following parameters were used in the configuration of the 2210 during qconfig:

- Interface 1 is WAN frame relay.
- Interface 1 cable type is V.35 modem.
- Configure Bridging = No

- Configure IPX = No
- Configure IP = Yes
- INT 0 IP= 10.2.2.3 M=255.255.255.0
- INT 1 IP= 100.1.1.4 M=255.255.255.0

Additional parameters must be changed that cannot be altered or set during the qconfig process. In this case the LMI type was set to ansi, the DLCI was added and the IP address set. Figure 249 shows the commands entered on the 2210 to set these parameters.

```
*t 6
Gateway user configuration
Config>net 1
Frame Relay user configuration

FR Config>set lmi-type ansi
FR Config>add permanent-virtual-circuit
Circuit number 16 ? 103
Committed Information Rate (CIR) in bps 64000 ?
Committed Burst Size (Bc) in bits 64000 ?
Excess Burst Size (Be) in bits 0 ?
Assign circuit name ?
FR Config>add protocol-address
Protocol name or number IP ?
IP Address 0.0.0.0 ? 10.2.2.2
Circuit number 16 ? 103

FR Config>
*
*restart
Are you sure you want to restart the gateway? (Yes or No): y
```

Figure 249. Configuring the 2210 for Frame Relay

The configuration parameters can be verified by listing the interface as shown in Figure 250 on page 272.

```

FR Config>li all

                                Frame Relay HDLC Configuration
Encoding      = NRZ IDLE      = Flag
Clocking      = External
Cable type    = V.35 DTE
Line access rate bps =      0 Interface MTU in bytes = 2048
Transmit delay =      0

                                Frame Relay Configuration
LMI enabled = Yes LMI DLCI = 0
LMI type = ANSI LMI Orphans OK = Yes
Protocol broadcast = Yes Congestion monitoring = Yes
Emulate multicast = Yes CIR monitoring = No
PVCs P1 allowed = 64 CIR monitor adjustment = 1
Timer T1 seconds = 10 Counter N1 increments = 6
LMI N2 error threshold = 3 LMI N3 error threshold window = 4
DECnet length field = No MIR % of CIR = 25
IR % Increment = 12 IR % Decrement = 25
Maximum PVCs allowable = 64
Total PVCs configured = 1

Circuit      Circuit      Circuit      CIR      Burst      Excess
Name          Number      Type      in bps    Size      Burst
-----
Unassigned    103      Permanent 64000     64000     0

                                Frame Relay Protocol Address Translations
Protocol Type  Protocol Address      Circuit Number
-----
IP             102.2.2              103

```

Figure 250. Verifying 2210 Setup

Chapter 10. Network Configuration Examples

This chapter examines and documents several network scenarios.

The configuration examples in this chapter utilize two IBM Nways RouteSwitches for demonstration purposes.

RouteSwitch A is an 8273 with an ATM submodule installed in slot 3 of the switch and a Fast Ethernet submodule installed in slot 4.

RouteSwitch B is a 5-slot chassis 8274 with an ATM module installed in slot 4 of the switch and a Fast Ethernet module installed in slot 5.

It is important to note the physical configuration setup of these switches as all of the following console screen examples refer to this setup. The relevant highlighted information in the console screens also makes reference to the setup of these switches.

10.1 Fast Ethernet Interconnected RouteSwitches

The objective of this example is to create multiple policy-based VLANs within one group in RouteSwitch A and interlink this switch via a Fast Ethernet link to RouteSwitch B. RouteSwitch B is also configured with multiple policy-based VLANs within one group. More than one policy-based VLAN rule will be used in the creation of VLANs in this example.

10.1.1 Network Topology

Shown in Figure 251 is the physical view of the network followed by the logical view in Figure 252 on page 274.

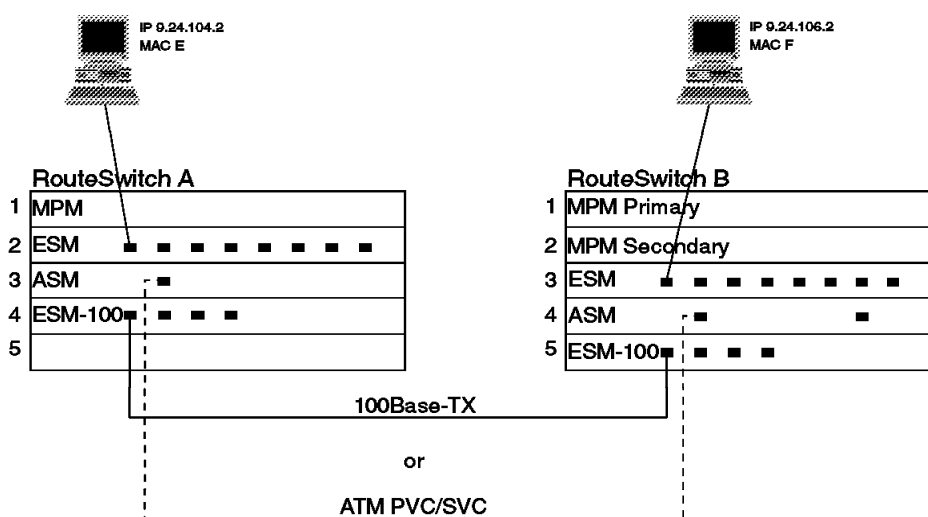


Figure 251. Physical View of Interconnected RouteSwitch Network

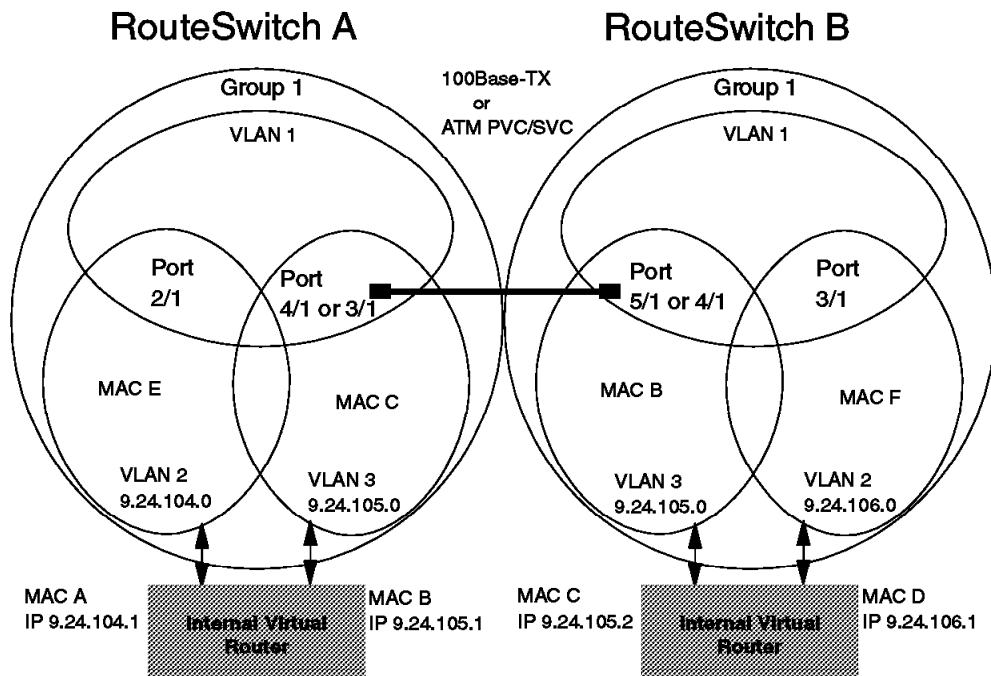


Figure 252. Logical View of Interconnected RouteSwitch Network

10.1.2 RouteSwitch Configuration

In all configuration examples, RouteSwitch A will be configured first and then RouteSwitch B.

In Figure 253 on page 275, the default VLAN 1 in RouteSwitch A was modified to disable IP. This was done for the following reasons:

- An internal virtual router arm is not going to be used in this configuration for the default VLAN.
- Management of the VLANs will not be done via the default VLAN but rather by the VLANs that will be created.
- To save IP address space since every virtual router port has to have its own IP subnet.
- The default VLAN should not be used for communication at all.
- Since there are no policies configurable for the default VLAN 1 all broadcast and unicast traffic to unknown destinations would be flooded out of all ports within the group.

```

SwitchA % modvl 1
Current values associated with GROUP 1.1 are as follows:

1) GROUP Number          - 1:1
2) Description           - Default GROUP (#1)
IP parameters:
3) IP enabled            - Y
4) IP Network Address    - 192.168.10.1
5) IP Subnet Mask        - 255.255.255.0
6) IP Broadcast Address  - 192.168.10.255
7) Router Description    - GROUP #1.0 IP router vport
8) RIP Mode              - Silent
   {Active(a), Inactive(i), Deaf(d), Silent(s)}
9) Routing disabled      - N
10) NHRP enabled         - N
11) Default Framing      - Ethernet II
   {Ethernet II(e), Ethernet 802.3(8), fddi(f),
    token ring(t), source route token ring(s)}
IPX parameters:
12) IPX enabled          - N

(save/quit/cancel)
: 3=n
: save

```

Figure 253. Modifying the Default VLAN in RouteSwitch A

IP was disabled by using the command 3=n. This disables IP within the default VLAN.

```

SwitchA % cratv1
Enter the VLAN Group id for this VLAN ( 1 ) :
Enter the VLAN Id for this VLAN ( 2 ) :
Enter the new VLAN's description: 9.24.104.0
Enter the Admin status for this vlan{(e)nable/(d)isable} (d):e
Select rule type:
  1. Port Rule
  2. MAC Address Rule
  3. Protocol Rule
  4. Network Address Rule
  5. User Defined Rule
  6. Binding Rule
  7. DHCP PORT Rule
  8. DHCP MAC Rule
Enter rule type (1): 4
Set Rule Admin Status to {(e)nable/(d)isable} (d):e
Select the Network Protocol:
  1. IP
  2. IPX
Enter protocol type: 1
Enter the IP Address: 9.24.104.0
Enter the IP Mask (255.0.0.0): 255.255.255.0
Configure more rules for this VLAN {y/n} (n): y
Select rule type:
  1. Port Rule
  2. MAC Address Rule
  3. Protocol Rule
  4. Network Address Rule
  5. User Defined Rule
  6. Binding Rule
  7. DHCP PORT Rule
  8. DHCP MAC Rule
Enter rule type (1): 1
Set Rule Admin Status to {(e)nable/(d)isable} (d): e
Enter the list of ports in Slot/Interface format: 2/1
Configure more rules for this VLAN {y/n} (n):
VLAN 1: 2 created successfully
Enable IP? (y):
  IP Address : 9.24.104.1
  IP Subnet Mask (0xff000000) : 255.255.255.0
  IP Broadcast Address (9.255.255.255 ) :
  Description (30 chars max) :
  IP RIP mode {Deaf(d),
    Silent(s),
    Active(a),
    Inactive(i)} (s) :
  Default framing type {Ethernet II(e),
    fddi(f),
    token ring(t),
    Ethernet 802.3(8),
    source route token ring(s)} (e) :
Created router port for VLAN 1: 2
Enable IPX? (y): n

```

Figure 254. Creating a VLAN on RouteSwitch A

VLAN 2 with the IP address of 9.24.104.0 was created according to the network address rule. The internal router arm for this VLAN was configured with the address 9.24.104.1. This will be the default gateway for any devices that will attach to this VLAN.

```

SwitchA % cratv1
Enter the VLAN Group id for this VLAN ( 1 ) :
Enter the VLAN Id for this VLAN ( 3 ) :
Enter the new VLAN's description: 9.24.105.0
Enter the Admin status for this vlan{(e)nable/(d)isable} (d):e
Select rule type:
  1. Port Rule
  2. MAC Address Rule
  3. Protocol Rule
  4. Network Address Rule
  5. User Defined Rule
  6. Binding Rule
  7. DHCP PORT Rule
  8. DHCP MAC Rule
Enter rule type (1): 4
Set Rule Admin Status to {(e)nable/(d)isable} (d): e
Select the Network Protocol:
  1. IP
  2. IPX
Enter protocol type: 1
Enter the IP Address: 9.24.105.0
Enter the IP Mask (255.255.255.0):
Configure more rules for this VLAN {y/n} (n): y
Select rule type:
  1. Port Rule
  2. MAC Address Rule
  3. Protocol Rule
  4. Network Address Rule
  5. User Defined Rule
  6. Binding Rule
  7. DHCP PORT Rule
  8. DHCP MAC Rule
Enter rule type (1): 1
Set Rule Admin Status to {(e)nable/(d)isable} (d): e
Enter the list of ports in Slot/Interface format: 4/1
Configure more rules for this vlan {y/n} (n):
VLAN 1: 3 created successfully
Enable IP? (y):
  IP Address : 9.24.105.1
  IP Subnet Mask (0xffffffff) :
  IP Broadcast Address (9.24.105.255 ) :
  Description (30 chars max) :
  IP RIP mode {Deaf(d),
    Silent(s),
    Active(a),
    Inactive(i)} (s) : a
  Default framing type {Ethernet II(e),
    fddi(f),
    token ring(t),
    Ethernet 802.3(8),
    source route token ring(s)} (e) :
Created router port for vlan 1: 3
Enable IPX? (y): n

```

Figure 255. Configuring the Fast Ethernet Inter-Switch Link on RouteSwitch A

VLAN 3 with the address 9.24.105.0 was created for the Fast Ethernet inter-switch link. The internal router arm with the address 9.24.105.1 (default gateway) was assigned to slot 4/2 for the Fast Ethernet inter-switch link of RouteSwitch A.

Configuration for RouteSwitch A is completed.

```

SwitchB % timeout 30
Auto logout time now 30 minutes
SwitchB % modvl 1
Current values associated with GROUP 1.1 are as follows:

1) GROUP Number          - 1:1
2) Description            - Default GROUP (#1)
IP parameters:
3) IP enabled             - Y
4) IP Network Address     - 192.168.10.1
5) IP Subnet Mask         - 255.255.255.0
6) IP Broadcast Address   - 192.168.10.255
7) Router Description     - GROUP #1.0 IP router vport
8) RIP Mode               - Silent
                          {Active(a), Inactive(i), Deaf(d), Silent(s)}
9) Routing disabled       - N
10) NHRP enabled          - N
11) Default Framing       - Ethernet II
                          {Ethernet II(e), Ethernet 802.3(8), fddi(f),
                           token ring(t), source route token ring(s)}
IPX parameters:
12) IPX enabled           - N

(save/quit/cancel)
: 3=n
: save

```

Figure 256. Modifying the Default VLAN in RouteSwitch B

The default VLAN was modified to disable IP.

The console timeout was changed to 30 minutes. This allows the operator more time to work away from the console while configuring VLANs without having the console session timing out due to keyboard inactivity.


```

SwitchB % cratv1
Enter the VLAN Group id for this VLAN ( 1 ) :
Enter the VLAN Id for this VLAN ( 2 ) :
Enter the new VLAN's description: 9.24.106.0
Enter the Admin status for this vlan{(e)nable/(d)isable} (d):e
Select rule type:
  1. Port Rule
  2. MAC Address Rule
  3. Protocol Rule
  4. Network Address Rule
  5. User Defined Rule
  6. Binding Rule
  7. DHCP PORT Rule
  8. DHCP MAC Rule
Enter rule type (1): 4
Set Rule Admin Status to {(e)nable/(d)isable} (d): e
Select the Network Protocol:
  1. IP
  2. IPX
Enter protocol type: 1
Enter the IP Address: 9.24.106.0
Enter the IP Mask (255.0.0.0): 255.255.255.0
Configure more rules for this vlan {y/n} (n): n
VLAN 1: 2 created successfully
Enable IP? (y):
  IP Address : 9.24.106.1
  IP Subnet Mask (0xff000000) : 255.255.255.0
  IP Broadcast Address (9.255.255.255 ) :
  Description (30 chars max) :
  IP RIP mode {Deaf(d),
    Silent(s),
    Active(a),
    Inactive(i)} (s) : a
  Default framing type {Ethernet II(e),
    fddi(f),
    token ring(t),
    Ethernet 802.3(8),
    source route token ring(s)} (e) :
Created router port for vlan 1: 2
Enable IPX? (y): n

```

Figure 257. Creating a VLAN on RouteSwitch B

VLAN 2 with the IP address of 9.24.106.0 was created according to the network address rule. This will be the default gateway for any devices attaching to this VLAN.

```

SwitchB % modatv1 1:2
VLAN 1: 2 is defined as:
  1. Description      = 9.24.106.0
  2. Admin Status    = Enabled
  3. Rule Definition
      Rule Num      Rule Type      Rule Status
      1              Net Addr Rule Enabled
Available options:
  1. Set VLAN Admin Status
  2. Set VLAN Description
  3. Add more rules
  4. Delete a rule
  5. Set rule Admin Status
  6. Quit
Option = 3
Select rule type:
  1. Port Rule
  2. MAC Address Rule
  3. Protocol Rule
  4. Network Address Rule
  5. User Defined Rule
  6. Binding Rule
  7. DHCP PORT Rule
  8. DHCP MAC Rule
Enter rule type (1): 1
Set Rule Admin Status to {(e)nable/(d)isable} (d): e
Enter the list of ports in Slot/Interface format: 3/1
  1. Description      = 9.24.106.0
  2. Admin Status    = Enabled
  3. Rule Definition
      Rule Num      Rule Type      Rule Status
      1              Net Addr Rule Enabled
      2              Port Rule      Enabled
Available options:
  1. Set VLAN Admin Status
  2. Set VLAN Description
  3. Add more rules
  4. Delete a rule
  5. Set rule Admin Status
  6. Quit
Option = 6

```

Figure 258. Modifying a VLAN on RouteSwitch B

Figure 258 shows how to add additional VLAN policies to an existing VLAN. VLAN 1:2 was modified because no internal router arm was configured for the 9.24.106.0 VLAN in the initial VLAN setup.

```

SwitchB % cratv1
Enter the VLAN Group id for this VLAN ( 1 ) :
Enter the VLAN Id for this VLAN ( 3 ) :
Enter the new VLAN's description: 9.24.105.0
Enter the Admin status for this vlan{(e)nable/(d)isable} (d):e
Select rule type:
  1. Port Rule
  2. MAC Address Rule
  3. Protocol Rule
  4. Network Address Rule
  5. User Defined Rule
  6. Binding Rule
  7. DHCP PORT Rule
  8. DHCP MAC Rule
Enter rule type (1): 4
Set Rule Admin Status to {(e)nable/(d)isable} (d): e
Select the Network Protocol:
  1. IP
  2. IPX
Enter protocol type: 1
Enter the IP Address: 9.24.105.0
Enter the IP Mask (255.255.255.0):
Configure more rules for this vlan {y/n} (n): y
Select rule type:
  1. Port Rule
  2. MAC Address Rule
  3. Protocol Rule
  4. Network Address Rule
  5. User Defined Rule
  6. Binding Rule
  7. DHCP PORT Rule
  8. DHCP MAC Rule
Enter rule type (1): 1
Set Rule Admin Status to {(e)nable/(d)isable} (d): e
Enter the list of ports in Slot/Interface format: 5/1
Configure more rules for this vlan {y/n} (n):
VLAN 1: 3 created successfully
Enable IP? (y):
  IP Address : 9.24.105.2
  IP Subnet Mask (0xffffffff) :
  IP Broadcast Address (9.24.105.255 ) :
  Description (30 chars max) :
  IP RIP mode {Deaf(d),
    Silent(s),
    Active(a),
    Inactive(i)} (s) : a
  Default framing type {Ethernet II(e),
    fddi(f),
    token ring(t),
    Ethernet 802.3(8),
    source route token ring(s)} (e) :
Created router port for vlan 1: 3
Enable IPX? (y): n

```

Figure 259. Creating the Inter-Switch link for RouteSwitch B

VLAN 3 with the address 9.24.105.0 was created for the Fast Ethernet inter-switch link. The address 9.24.105.2 was assigned to slot 5/1 for the Fast Ethernet inter-switch link of RouteSwitch B. The VLAN port rule was used so the link will always be up and not disconnected due to traffic inactivity across the inter-switch link.

Configuration for RouteSwitch B is completed.

The following console displays are used to confirm that the configuration of the RouteSwitches is correct. It also shows that the two RouteSwitches are communicating with each other.

These console displays are from RouteSwitch B.

```
SwitchB % ipr
```

IP ROUTING TABLE				
Network	Mask	Gateway	Metric	Group VLAN Id: Id
9.24.104.0	255.255.255.0	9.24.105.1	2	1:3
9.24.106.0	255.0.0.0	9.24.106.1	1	1:2
127.0.0.1	255.255.255.255	127.0.0.1	1	LOOPBACK
9.24.105.0	255.255.255.0	9.24.105.2	1	1:3

Figure 260. Displaying the Dynamic IP Routing Tables

The dynamic IP routing table of RouteSwitch B indicates that it is aware of an existing network on the remote side of its own Fast Ethernet link. This is indicated by the 9.24.104.0 IP address being listed in the Network column of the table.

```
SwitchB % vivl
```

Virtual Interface				VLAN Membership	
Slot/Intf/Service/Instance				Group	Member of VLAN#
1	/1	/Rtr	/1	1	2
1	/1	/Rtr	/2	1	3
3	/1	/Brg	/1	1	1 2
3	/2	/Brg	/1	1	1
3	/3	/Brg	/1	1	1
3	/4	/Brg	/1	1	1
3	/5	/Brg	/1	1	1
3	/6	/Brg	/1	1	1
3	/7	/Brg	/1	1	1
3	/8	/Brg	/1	1	1
3	/9	/Brg	/1	1	1
3	/10	/Brg	/1	1	1
3	/11	/Brg	/1	1	1
3	/12	/Brg	/1	1	1
4	/1	/Brg	/1	1	1
4	/2	/Brg	/1	1	1
5	/1	/Brg	/1	1	1 3

Figure 261. VLAN Membership within RouteSwitch B

This screen shows the VLAN membership of the slots in RouteSwitch B. Slot 5/1 can be seen to have become a member of both the inter-switch link VLAN and the router arm due to the port rule VLAN rule.

```
SwitchB % vi
Virtual Interface Summary Information- For All Interfaces
```

						Status				
Group	Intf	Slot/ Type/ Inst/Srvc	MAC Address	Prt	Encp	Admin	Oper	Spn Tr	Mode	
1	A11	Rtr/ 1	0020da:758ed0	IP	DFLT	Enabl	d	Active	N/A	
1	A11	Rtr/ 2	0020da:758ed3	IP	DFLT	Enabl	d	Active	N/A	
1	3/1	Brg/ 1/ na	0020da:73f1b0	Tns	DFLT	Enabl	d	Inactv	Disabl	
1	3/2	Brg/ 1/ na	0020da:73f1b1	Tns	DFLT	Enabl	d	Inactv	Disabl	
1	3/3	Brg/ 1/ na	0020da:73f1b2	Tns	DFLT	Enabl	d	Inactv	Disabl	
1	3/4	Brg/ 1/ na	0020da:73f1b3	Tns	DFLT	Enabl	d	Inactv	Disabl	
1	3/5	Brg/ 1/ na	0020da:73f1b4	Tns	DFLT	Enabl	d	Inactv	Disabl	
1	3/6	Brg/ 1/ na	0020da:73f1b5	Tns	DFLT	Enabl	d	Inactv	Disabl	
1	3/7	Brg/ 1/ na	0020da:73f1b6	Tns	DFLT	Enabl	d	Inactv	Disabl	
1	3/8	Brg/ 1/ na	0020da:73f1b7	Tns	DFLT	Enabl	d	Inactv	Disabl	
1	3/9	Brg/ 1/ na	0020da:73f1b8	Tns	DFLT	Enabl	d	Inactv	Disabl	
1	3/10	Brg/ 1/ na	0020da:73f1b9	Tns	DFLT	Enabl	d	Inactv	Disabl	
1	3/11	Brg/ 1/ na	0020da:73f1ba	Tns	DFLT	Enabl	d	Inactv	Disabl	
1	3/12	Brg/ 1/ na	0020da:73f1bb	Tns	DFLT	Enabl	d	Inactv	Disabl	
1	4/1	Brg/ 1/ 1	0020da:6fa8e0	Tns	DFLT	Enabl	d	Inactv	Disabl	
1	4/2	Brg/ 1/ 1	0020da:6fa8e1	Tns	DFLT	Enabl	d	Inactv	Disabl	
1	5/1	Brg/ 1/ na	0020da:76b210	Tns	DFLT	Enabl	d	Active	Fwdng	

Figure 262. Viewing Port Status

The Fast Ethernet inter-switch link that is connected to slot 5/1 can be seen in a forwarding state in this table. This is an indication that the Fast Ethernet link is operational on RouteSwitch B.

```
SwitchB % viatr1
VLAN  VLAN  Rule Rule      Rule      Rule
Group Id  Num  Type      Status    Definition
-----
1: 2      1  NET ADDR RULE Enabled  IP Addr = 9.24.106.0
                                IP Mask = 255.0.0.0
1: 2      2  PORT RULE   Enabled  3/1/Brg/1
1: 3      1  NET ADDR RULE Enabled  IP Addr = 9.24.105.0
                                IP Mask = 255.255.255.0
1: 3      2  PORT RULE   Enabled  5/1/Brg/1

/ %
Syncing configuration data with secondary 2 .. complete
```

Figure 263. Viewing VLAN Rules on RouteSwitch B

This screen indicates the different VLAN policy rules that are associated with the VLANs.

10.2 RouteSwitch Using Original Port Rule

In this example, a single RouteSwitch is using the original port rule. The original port rule specifies that once a port is assigned to a VLAN, all of the MAC addresses seen on this port will join the VLAN. This differs from the current port rule where the MAC addresses will not join the VLAN.

The first step is to edit the mpm.cmd and add the following line: `reg_port_rule=1` and reboot the RouteSwitch.

In this scenario we have two groups of workstations using a common server. Each workstations should be able to communicate with the server and at the same time workstations in one group should not be able to communicate with workstations in the other group.

In this example, the server is connected to port 2/1; workstations in group A are connected to port 2/2; and workstations in group B are connected to port 2/3




Figure 264. Logical View Using Original Port Rules

The configuration of the VLAN and the definition of the port rules for each VLAN do not change from the description in Chapter 3, "VLAN and Mobile Groups Concepts" on page 37.

10.3 RouteSwitch Interconnection Using MSS and 8260

In this example, the two RouteSwitches are configured to use ATM LAN Emulation (LANE). The RouteSwitches will connect via their ATM ports to an IBM Nways 8260 MultiProtocol Switching hub. The 8260 is connected to an IBM 8210 MultiProtocol Switching Services (MSS). The MSS is configured with an Emulated LAN (ELAN) and the RouteSwitches will connect to this ELAN.

The VLAN configurations for both the RouteSwitches stay the same as in the previous examples. The console screen captures in this example show how the RouteSwitches are configured to be LAN Emulation Clients (LECs) attached to the ELAN on the MSS.

10.3.1 Network Topology

Shown in Figure 265 is the physical view of the network followed by the logical view in Figure 266 on page 286.

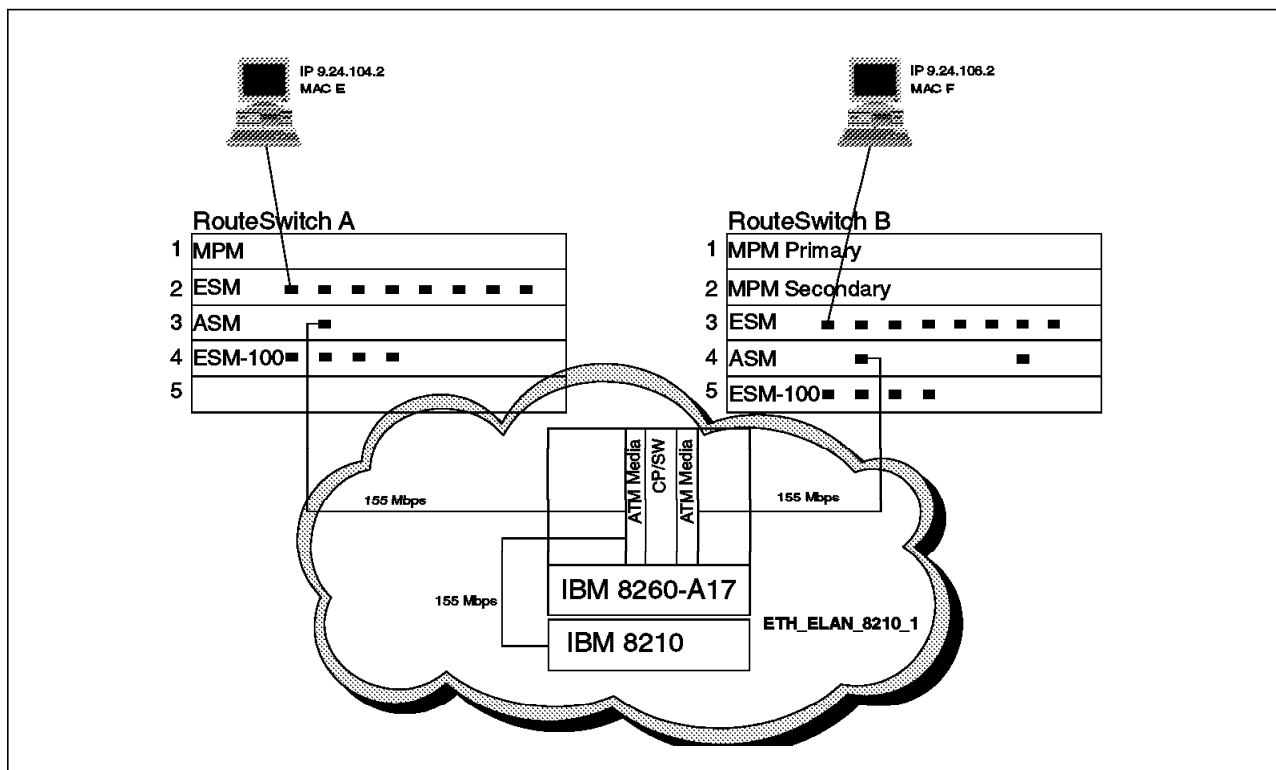


Figure 265. Physical View of Interconnection via ATM LANE

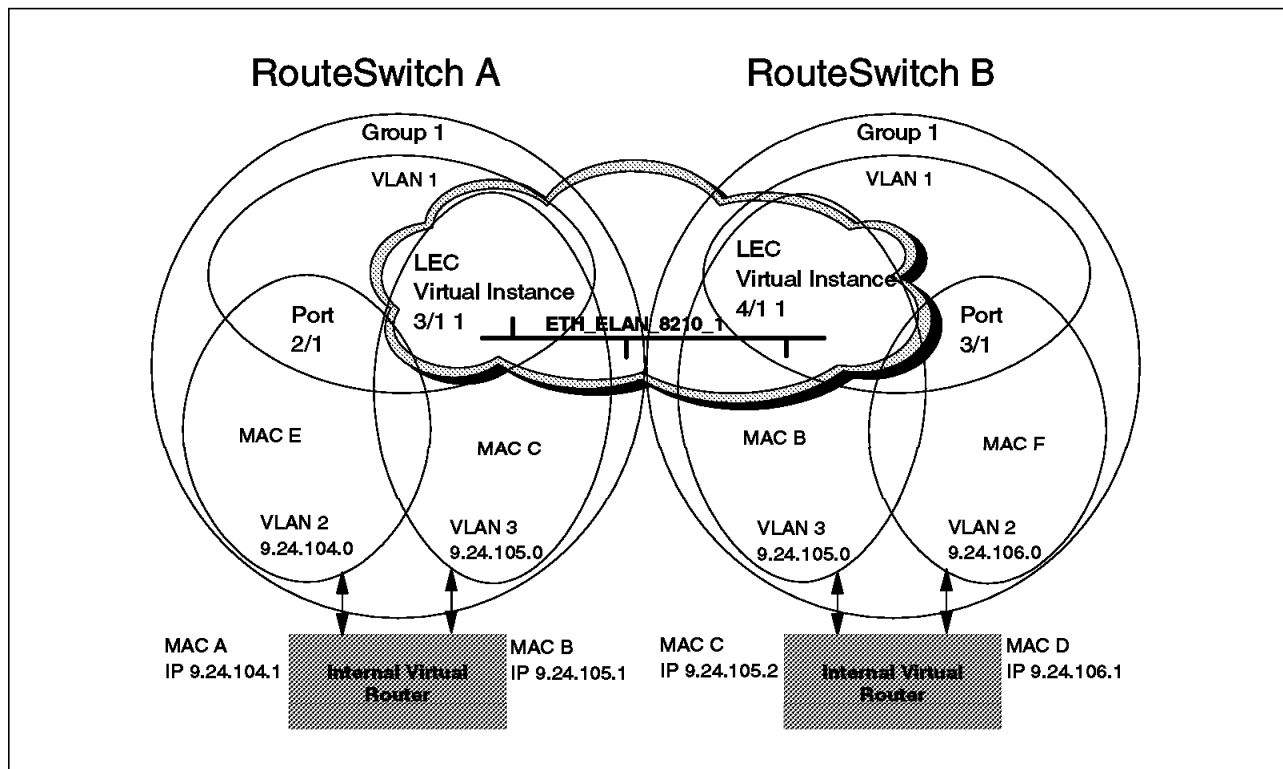


Figure 266. Logical View of Interconnection via ATM LANE

10.3.2 RouteSwitch Configuration

SwitchB % **map 4/1**

Slot 4 Port 1 Configuration

```
1) Description (30 chars max)           : ATM PORT
2) Conn Type { PVC(1), SVC(2) }         : PVC
3) Max VCCs (1-1023)                   : 1023
4) Max VCI bits (1..10)                 : 10
5) UNI Type { Pub(1), Priv(2) }         : Private
6) Tx Segment Size (4096-131072)        : 8192
7) Rx Segment Size (4096-131072)        : 8192
8) Tx Buffer Size (1800-8192)            : 4600
9) Rx Buffer Size (1800-8192)            : 4600
10) Pl Scramble {(False(1),True(2))}    : True
11) Timing Mode {(Loop(1),Local(2))}    : Local
12) Loopback Config { NoLoop(1), DiagLoop(2),
    LineLoop(3) }                       : NoLoop
13) Phy media { SONET(1),SDH(2)}        : SONET
Enter (option=value/save/cancel) : 2=2
```

Slot 4 Port 1 Configuration

```
1) Description (30 chars max)           : ATM PORT
2) Conn Type { PVC(1), SVC(2) }         : SVC

30) Sig version { 3.0(1) 3.1(2) }       : 3.0
31) Signaling VCI (0..1023)             : 5
32) ILMI Enable {(False(1),True(2))}    : True
33) ESI (12 hex-chars)                  : 0020da6fa8e0
34) ILMI VCI (0..1023)                  : 16

3) Max VCCs (1-1023)                   : 1023
4) Max VCI bits (1..10)                 : 10
5) UNI Type { Pub(1), Priv(2) }         : Private
6) Tx Segment Size (4096-131072)        : 8192
7) Rx Segment Size (4096-131072)        : 8192
8) Tx Buffer Size (1800-8192)            : 4600
9) Rx Buffer Size (1800-8192)            : 4600
10) Pl Scramble {(False(1),True(2))}    : True
11) Timing Mode {(Loop(1),Local(2))}    : Local
12) Loopback Config { NoLoop(1), DiagLoop(2),
    LineLoop(3) }                       : NoLoop
13) Phy media { SONET(1),SDH(2)}        : SONET
Enter (option=value/save/cancel) : save
```

```
Reset all services on slot 4 port 1 (n)? : y
Resetting port, please wait...
```

Figure 267. Modifying the ATM Port of RouteSwitch B for SVCs

The connection type is changed from PVC to SVC because LAN Emulation (LANE) is based on SVCs. The physical port is configured with SVC to match the service type, that is LANE.

```

SwitchB % cas 4/1

Slot 4 Port 1 Service 2 Configuration

1) Description (30 chars max)           : PTOP Bridging Service 2
2) Service type { LAN Emulation (1),
                  Trunking (4),
                  Classical IP(5),
                  PTOP Bridging(6),
                  VLAN cluster(7) } : PTOP Bridging
10) Encaps Type { Private(1),
                 RFC1483(2) } : Private
3) Connection Type { PVC(1),
                   SVC(2) } : PVC
4) PTOP Group : 1
5) PTOP connection : none
6) Admin Status { disable(1),
                 enable(2) } : Enable

Enter (option=value/save/cancel) : 2=1

Slot 4 Port 1 Service 2 Configuration

1) Description (30 chars max)           :LAN Emulation Service 2
2) Service type { LAN Emulation (1),
                  Trunking (4),
                  Classical IP(5),
                  PTOP Bridging(6),
                  VLAN cluster(7) } :LAN Emulation
21) LAN type { 802.3 (1),
               802.5 (2) } : 802.3
22) Change LANE Cfg { NO (1),
                    YES (2) } : NO
3) Connection Type { PVC(1),
                   SVC(2) } : SVC
30) SEL for the ATM address : 02
4) LAN Emulated Group : 1
5) LECS Address (40-char-hex) : 4700790000000000000000000000A03E0
6) Admin Status { disable(1),
                 enable(2) } : Enable

Enter (option=value/save/cancel) : 22=2

```

Figure 268. Changing the Service Type to LAN Emulation

The service type was changed from PTOP bridging to LAN Emulation. LANE will be used to attach to the MSS server.

```
Enter (option=value/save/cancel) : 22=2
```

Slot 4 Port 1 Service 2 LANE Configuration Parameters

```
1) Proxy { NO (1), YES (2) } : YES
2) Max Frame Size { 1516 (1), 4544 (2)
    9234 (3), 18190 (4) } : 4544
3) Use translation options{NO (1), YES (2) } : Yes
4) Use Fwd Delay time { NO (1), YES (2) } : NO
5) Use LE Cfg Server (LECS){ NO (1), YES (2)}: YES
6) Use Default LECS address { NO(1), YES (2)}: YES
7) Control Time-out (in seconds) : 120
8) Max Unknown Frame Count : 1
9) Max Unknown Frame Time (in seconds) : 1
10) VCC Time-out Period (in minutes) : 20
11) Max Retry Count : 1
12) Aging Time (in seconds) : 300
13) Expectd LE_ARP Resp Time (in seconds) : 1
14) Flush Time-out (in seconds) : 4
15) Path Switching Delay (in seconds) : 6
16) ELAN name (32 chars max) :
```

```
Enter (option=value/save/cancel) : 16=ETH_ELAN_8210_1
```

Slot 4 Port 1 Service 2 LANE Configuration Parameters

```
1) Proxy { NO (1), YES (2) } : YES
2) Max Frame Size { 1516 (1), 4544 (2)
    9234 (3), 18190 (4) } : 4544
3) Use translation options{NO (1), YES (2) } : Yes
4) Use Fwd Delay time { NO (1), YES (2) } : NO
5) Use LE Cfg Server (LECS){ NO (1), YES (2)}: YES
6) Use Default LECS address { NO(1), YES (2)}: YES
7) Control Time-out (in seconds) : 120
8) Max Unknown Frame Count : 1
9) Max Unknown Frame Time (in seconds) : 1
10) VCC Time-out Period (in minutes) : 20
11) Max Retry Count : 1
12) Aging Time (in seconds) : 300
13) Expectd LE_ARP Resp Time (in seconds) : 1
14) Flush Time-out (in seconds) : 4
15) Path Switching Delay (in seconds) : 6
16) ELAN name (32 chars max) : ETH_ELAN_8210_1
```

```
Enter (option=value/save/cancel) : save
```

```
Creating service, please wait...
```

```
Enabling service...
```

Figure 269. Definition of ELAN Name

The case-sensitive ELAN name on the MSS is specified so that the LAN Emulation Client service that was created on RouteSwitch B may join the ELAN of the MSS.

A maximum frame size of 4544 bytes is not possible in an Ethernet network. During the ELAN join phase, this value is reduced to the true Ethernet ELAN frame size configured at the LES. The value specified here has to be greater than the actual maximum frame size of the ELAN to join.

If option 5) is set to no, the ATM address of the LES has to be specified.

If option 6) is set to no, the ATM address of the LECS has to be specified in the previous screen. With the above setting, the ATM Forum well-known address for LECS is used to establish a connection to the LECS. Since this address may not be a valid LECS address, the ATM switch has to map it to a valid LECS address for the particular network. This is done by the CP/SW ATM switch.

```
SwitchB % das 4/1 1

                        ATM Services
Slot Port Serv      Service      Service
  Num Num      Description      Type
=====
4   1   1   PTOp Bridging Service 1   PTOp Priv
4   1   2   LAN Emulation Service 2   LANE

                        ATM Services
Slot Port Serv VC   Oper
  Num Num      Type Status SEL Groups Conn VCI's/Addresses
=====
4   1   1   PVC Enabled N/A 1      100
4   1   2   SVC LANE Op. 02 1      102 103 104 105

Remove ATM Slot 4 Port 1 Service 1 (n)? :y
Removing ATM Slot 4 Port 1 Service 1, please wait...

ATM Slot 4 Port 1 Service 1 removed
```

Figure 270. Deleting an ATM Service from the 4/1 RouteSwitch Port

PTOP bridging is deleted to ensure that the LAN Emulation service is used and not PTOp.

```
SwitchB % vas

                        ATM Services
Slot Port Serv      Service      Service
  Num Num      Description      Type
=====
4   1   2   LAN Emulation Service 2   LANE
4   2   1   PTOp Bridging Service 1   PTOp Priv

                        ATM Services
Slot Port Serv VC   Oper
  Num Num      Type Status SEL Groups Conn VCI's/Addresses
=====
4   1   2   SVC LANE Op. 02 1      102 103 104 105
                                107 106
4   2   1   PVC Disabled N/A 1      100
```

Figure 271. Viewing an ATM Service

The SVC-LAN Emulation service is configured and working while the original PVC PTOp bridging service is no longer operational.

SwitchB % ipr				
IP ROUTING TABLE				
Network	Mask	Gateway	Metric	Id: Id
9.24.104.0	255.0.0.0	9.24.105.1	2	1:3
9.24.106.0	255.0.0.0	9.24.106.1	1	1:2
127.0.0.1	255.255.255.255	127.0.0.1	1	LOOPBACK
9.24.105.0	255.255.255.0	9.24.105.2	1	1:3

Figure 272. Displaying the IP Routing Tables

Displaying the IP routing tables verifies the connection between the RouteSwitches.

10.4 IP Routing at the MSS

The following network example shows how to set up a RouteSwitch in an ATM MSS environment. In contrast to the prior example 10.3, “RouteSwitch Interconnection Using MSS and 8260” on page 284 where the internal virtual router of the RouteSwitch was used to do the IP routing, in this example the routing capabilities of the MSS will be used. Note that only a single LEC instance was used to route frames from one VLAN to another. The routing within a single ELAN was made possible by the MSS’s capability of supporting multiple IP addresses per LEC.

10.4.1 Network Topology

Shown in Figure 273 is the physical view of the network followed by the logical view in Figure 274 on page 292.

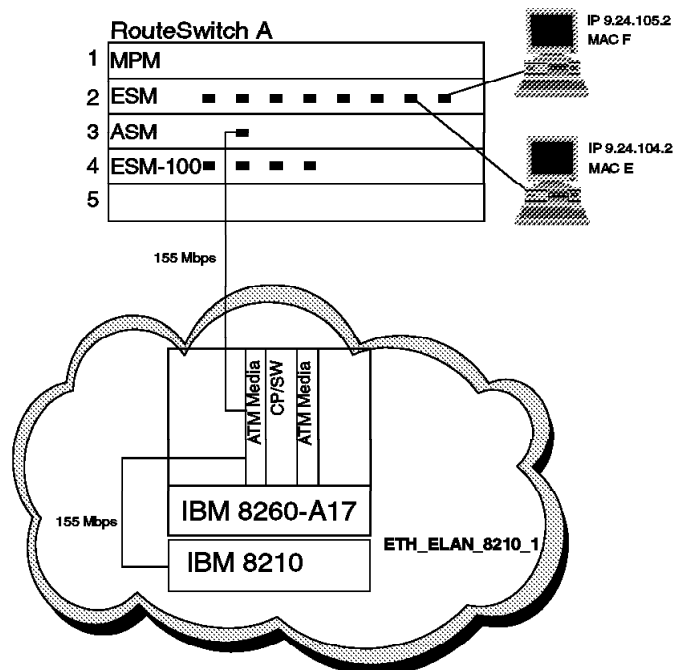


Figure 273. Physical View of Network Routed at MSS

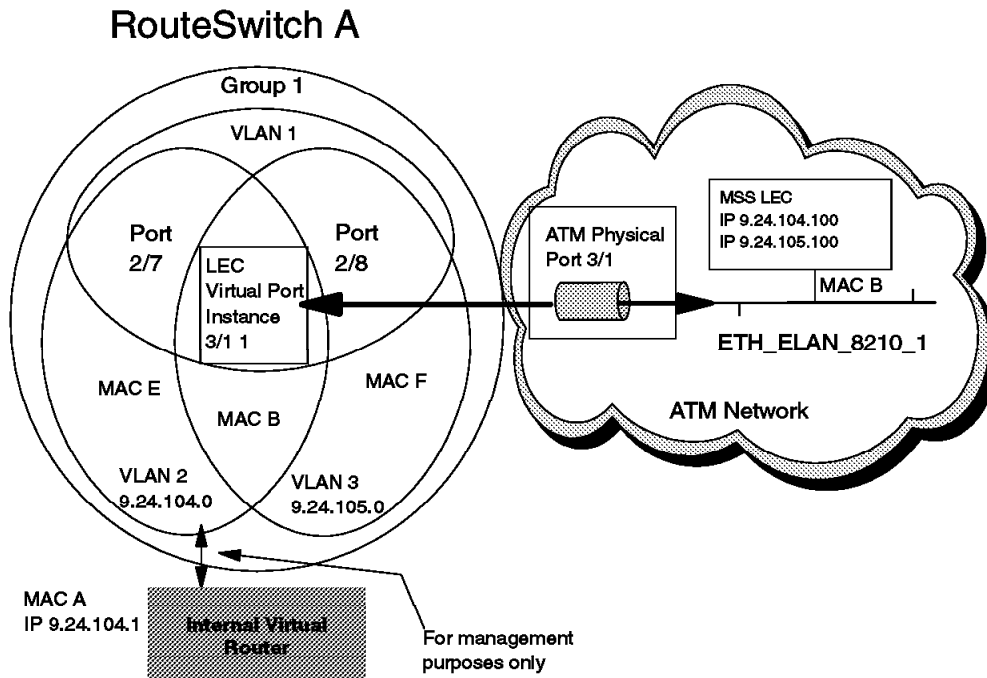


Figure 274. Logical View of Network Routed at MSS

10.4.2 RouteSwitch Configuration

```
SwitchA % modv1 1
Current values associated with GROUP 1.1 are as follows:

1) GROUP Number          - 1:1
2) Description            - Default GROUP (#1)
IP parameters:
3) IP enabled             - Y
4) IP Network Address     - 192.168.10.1
5) IP Subnet Mask         - 255.255.255.0
6) IP Broadcast Address   - 192.168.10.255
7) Router Description     - GROUP #1.0 IP router vport
8) RIP Mode               - Silent
                           {Active(a), Inactive(i), Deaf(d), Silent(s)}
9) Routing disabled       - N
10) NHRP enabled          - N
11) Default Framing       - Ethernet II
                           {Ethernet II(e), Ethernet 802.3(8), fddi(f),
                           token ring(t), source route token ring(s)}
IPX parameters:
12) IPX enabled           - N

(save/quit/cancel)
: 3=n
: save
```

Figure 275. Disabling the Virtual Router for the Default VLAN

In Figure 275, the default VLAN 1 in RouteSwitch A was modified to disable IP. This was done for the following reasons:

- The internal virtual router arm is not going to be used in this configuration for the default VLAN.
- Management of the RouteSwitch will not be done via the default VLAN but rather by the VLANs that will be created.
- Since there are no policies configurable for the default VLAN 1, all broadcast and unicast traffic to unknown destinations are flooded out all ports within the group.
- Save IP address space since every virtual router port has to have its own IP subnet.
- The default VLAN should not be used for communication at all.

```

SwitchA % cratvl
Enter the VLAN Group id for this VLAN ( 1 ) :
Enter the VLAN Id for this VLAN ( 2 ) :
Enter the new VLAN's description: ip_9.24.104.0
Enter the Admin status for this vlan [(e)nable/(d)isable] (d): e
Select rule type:
  1. Port Rule
  2. MAC Address Rule
  3. Protocol Rule
  4. Network Address Rule
  5. User Defined Rule
  6. Binding Rule
  7. DHCP PORT Rule
  8. DHCP MAC Rule
Enter rule type (1): 4
Set Rule Admin Status to &lrbk.(e)nable/(d)isable&rbrk. (d): e
Select the Network Protocol:
  1. IP
  2. IPX
Enter protocol type: 1
Enter the IP Address: 9.24.104.0
Enter the IP Mask (255.0.0.0): 255.255.255.0
Configure more rules for this vlan [y/n] (n):
VLAN 1: 2 created successfully
Enable IP? (y):
  IP Address : 9.24.104.1
  IP Subnet Mask (0xff000000) : 255.255.255.0
  IP Broadcast Address (9.24.104.255 ) :
  Description (30 chars max) :
  IP RIP mode {Deaf(d),
    Silent(s),
    Active(a),
    Inactive(i)} (s) :
  Default framing type {Ethernet II(e),
    fddi(f),
    token ring(t),
    Ethernet 802.3(8),
    source route token ring(s)} (e) :
Created router port for vlan 1: 2
Enable IPX? (y): n

```

Figure 276. Creating VLAN 2

In Figure 276, IP was enabled for management purposes only. No internal routing takes place because this is the only router arm configured.

```

SwitchA % cratv1
Enter the VLAN Group id for this VLAN ( 1 ) :
Enter the VLAN Id for this VLAN ( 3 ) :
Enter the new VLAN's description: ip_9.24.105.0
Enter the Admin status for this vlan (e)nable/(d)isable (d): e
Select rule type:
  1. Port Rule
  2. MAC Address Rule
  3. Protocol Rule
  4. Network Address Rule
  5. User Defined Rule
  6. Binding Rule
  7. DHCP PORT Rule
  8. DHCP MAC Rule
Enter rule type (1): 4
Set Rule Admin Status to [(e)nable/(d)isable] (d): e
Select the Network Protocol:
  1. IP
  2. IPX
Enter protocol type: 1
Enter the IP Address: 9.24.105.0
Enter the IP Mask (255.0.0.0): 255.255.255.0
Configure more rules for this vlan [y/n] (n):
VLAN 1: 3 created successfully
Enable IP? (y): n
Enable IPX? (y): n

```

Figure 277. Creating VLAN 3

In Figure 277, IP was disabled in order to suppress internal routing within the RouteSwitch.

```

SwitchA % map 3/1

Slot 3 Port 1 Configuration

1) Description (30 chars max)           : ATM PORT
2) Conn Type { PVC(1), SVC(2) }         : PVC
3) Max VCCs (1-1023)                   : 1023
4) Max VCI bits (1..10)                 : 10
5) UNI Type { Pub(1), Priv(2) }         : Private
6) Tx Segment Size (2048-131072)        : 16384
7) Rx Segment Size (2048-131072)        : 16384
8) Tx Buffer Size (1800-8192)            : 4600
9) Rx Buffer Size (1800-8192)            : 4600
10) Pl Scramble {(False(1),True(2))}    : True
11) Timing Mode {(Loop(1),Local(2))}    : Local
12) Loopback Config { NoLoop(1), DiagLoop(2),
    LineLoop(3) }                       : NoLoop
13) Phy media { SONET(1),SDH(2)}        : SONET

Enter (option=value/save/cancel) : 2=2

```

Figure 278. Changing Connection Type to SVC

In Figure 278, the connection type of the ATM port 3/1 was changed to SVC since LAN Emulation is based on SVCs. All the other values remained unchanged.


```

SwitchA % cas 3/1

Slot 3 Port 1 Service 2 Configuration

1) Description (30 chars max)      : PTOp Bridging Service 2
2) Service type { LAN Emulation (1),
                  Trunking (4),
                  Classical IP(5),
                  PTOp Bridging(6),
                  VLAN cluster(7) } : PTOp Bridging
10) Encaps Type { Private(1),
                 RFC1483(2) }      : Private
3) Connection Type { PVC(1),
                   SVC(2) }        : PVC
4) PTOp Group                      : 1
5) PTOp connection                 : none
6) Admin Status { disable(1),
                 enable(2) }       : Enable

Enter (option=value/save/cancel) : 2=1

```

Figure 279. Creating the LEC Instance

```

Slot 3 Port 1 Service 2 Configuration

1) Description (30 chars max)      : LAN Emulation Service 2
2) Service type { LAN Emulation (1),
                  Trunking (4),
                  Classical IP(5),
                  PTOp Bridging(6),
                  VLAN cluster(7) } : LAN Emulation
21) LAN type { 802.3 (1),
              802.5 (2) }          : 802.3
22) Change LANE Cfg { NO (1),
                    YES (2) }      : NO
3) Connection Type { PVC(1),
                   SVC(2) }        : SVC
30) SEL for the ATM address        : 02
4) LAN Emulated Group             : 1
5) LECS Address (40-char-hex)     : 4700790000000000000000000000A03E00000100
6) Admin Status { disable(1),
                 enable(2) }       : Enable

Enter (option=value/save/cancel) : 22=2

```

Figure 280. Modifying the LANE Configuration

In Figure 279 and Figure 280, the RouteSwitch automatically creates a service with the next available service instance. In this case, instance two is used since by default all service ports have a bridging service already configured.

Note: Each LEC service may only be assigned to one group.

Slot 3 Port 1 Service 2 LANE Configuration Parameters

1) Proxy { NO (1), YES (2) } : YES
2) Max Frame Size { 1516 (1), 4544 (2)
9234 (3), 18190 (4) } : 1516
3) Use translation options{NO (1), YES (2) : Yes (use Swch menu to set)
4) Use Fwd Delay time { NO (1), YES (2) } : NO
5) Use LE Cfg Server (LECS){ NO (1), YES (2)}: YES
6) Use Default LECS address { NO(1), YES (2)}: YES
7) Control Time-out (in seconds) : 10
8) Max Unknown Frame Count : 10
9) Max Unknown Frame Time (in seconds) : 1
10) VCC Time-out Period (in minutes) : 20
11) Max Retry Count : 2
12) Aging Time (in seconds) : 300
13) Expectd LE_ARP Resp Time (in seconds) : 1
14) Flush Time-out (in seconds) : 4
15) Path Switching Delay (in seconds) : 6
16) ELAN name (32 chars max) :

Enter (option=value/save/cancel) : 16=ETH_ELAN_8210_1

Figure 281. Definition of ELAN Name

In Figure 281, the ELAN name is specified. This name is case-sensitive.

A maximum frame size of 4544 bytes is not possible in an Ethernet network. During the ELAN join phase, this value is reduced to the true Ethernet ELAN frame size configured at the LES. The value specified here has to be greater than the actual maximum frame size of the ELAN to join.

If option 5) is set to no, the ATM address of the LES has to be specified.

If option 6) is set to no, the ATM address of the LECS has to be specified in the previous screen. With the above setting, the ATM Forum well-known address for LECS is used to establish a connection to the LECS. Since this address may not be a valid LECS address, the ATM switch has to map it to a valid LECS address for the particular network. This is done by the CP/SW ATM switch.

```

SwitchA % das 3/1 1

                                ATM Services

Slot Port Num      Service Description      Service Type
==== =====
3    1    1    PTOP Bridging Service 1    PTOP Priv
3    1    2    LAN Emulation Service 2    LANE

                                ATM Services

Slot Port Num  Serv VC  Oper  SEL Groups  Conn VCI's/Addresses
==== =====
3    1    1    PVC Disabled N/A 1    100
3    1    2    SVC Initial 02 1

Remove ATM Slot 3 Port 1 Service 1 (n)? : y

Removing ATM Slot 3 Port 1 Service 1, please wait...

ATM Slot 3 Port 1 Service 1 removed

```

Figure 282. Removing the Default Bridging Service

In Figure 282, the default bridging service that uses the VPI/VCI 0/100 and a PVC connection was removed because there are no other services needed for this example.

The operational status of service 2 shows Initial, which means that the LEC has not yet joined the ELAN.

Note: The service number is not necessarily equal to the instance number. In this case, the LAN Emulation service 2 becomes service instance 1 after removing of the bridging service.

```

SwitchA % modatv1 1:2
VLAN 1: 2 is defined as:
  1. Description      = ip_9.24.104.0
  2. Admin Status     = Enabled
  3. Rule Definition
      Rule Num      Rule Type      Rule Status
      1             Net Addr Rule Enabled
Available options:
  1. Set VLAN Admin Status
  2. Set VLAN Description
  3. Add more rules
  4. Delete a rule
  5. Set rule Admin Status
  6. Quit
Option = 3
Select rule type:
  1. Port Rule
  2. MAC Address Rule
  3. Protocol Rule
  4. Network Address Rule
  5. User Defined Rule
  6. Binding Rule
  7. DHCP PORT Rule
  8. DHCP MAC Rule
Enter rule type (1): 1
Set Rule Admin Status to [(e)nable/(d)isable] (d): e
Enter the list of ports in Slot/Interface format: 3/1
  1. Description      = ip_9.24.104.0
  2. Admin Status     = Enabled
  3. Rule Definition
      Rule Num      Rule Type      Rule Status
      1             Net Addr Rule Enabled
      2             Port Rule      Enabled

```

Figure 283. Adding the LEC Instance to VLAN 2

```

SwitchA % modatv1 1:3
VLAN 1: 3 is defined as:
  1. Description      = ip_9.24.105.0
  2. Admin Status     = Enabled
  3. Rule Definition
      Rule Num      Rule Type      Rule Status
      1             Net Addr Rule Enabled
Available options:
  1. Set VLAN Admin Status
  2. Set VLAN Description
  3. Add more rules
  4. Delete a rule
  5. Set rule Admin Status
  6. Quit
Option = 3
Select rule type:
  1. Port Rule
  2. MAC Address Rule
  3. Protocol Rule
  4. Network Address Rule
  5. User Defined Rule
  6. Binding Rule
  7. DHCP PORT Rule
  8. DHCP MAC Rule
Enter rule type (1): 1
Set Rule Admin Status to [(e)nable/(d)isable] (d): e
Enter the list of ports in Slot/Interface format: 3/1
  1. Description      = ip_9.24.105.0
  2. Admin Status     = Enabled
  3. Rule Definition
      Rule Num      Rule Type      Rule Status
      1             Net Addr Rule Enabled
      2             Port Rule      Enabled

```

Figure 284. Adding the LEC Instance to VLAN 3

In Figure 283 on page 298 and Figure 284, the LEC instance is added to both VLANs with an additional port rule. This prevents the LEC from being timed out and dropped from the VLAN membership if there is no traffic for a certain amount of time. With this additional rule, the LEC is always a member of both VLANs so that communication problems are avoided. For a further explanation, refer to 3.3.2, “Port Rules” on page 41.

```

SwitchA % viv1
          Virtual Interface VLAN Membership
          Slot/Intf/Service/Instance  Group  Member of VLAN#
          -----
1   /1   /Rtr   /1      1      2
2   /1   /Brg   /1      1      1
2   /2   /Brg   /1      1      1
2   /3   /Brg   /1      1      1
2   /4   /Brg   /1      1      1
2   /5   /Brg   /1      1      1
2   /6   /Brg   /1      1      1
2   /7   /Brg   /1      1      1
2   /8   /Brg   /1      1      1
2   /9   /Brg   /1      1      1
2   /10  /Brg   /1      1      1
2   /11  /Brg   /1      1      1
2   /12  /Brg   /1      1      1
3   /1   /Lne   /1      1      1 2 3
4   /2   /Brg   /1      1      1
SwitchA %

```

Figure 285. Viewing the VLAN Membership of Ports

As shown in Figure 285, there is only one virtual router arm defined and assigned to VLAN 2. Thus no internal routing takes place.

Because of the additional port rules, the virtual LEC instance one is also a member of VLANs two and three even though no frames were sent to the RouteSwitch at this time.

10.5 IP Routing at the MSS Using Multiple LECs

The following network example shows how to set up a RouteSwitch in an ATM MSS environment using separate LECs for each VLAN. Routing is done within the MSS.

This example is based on the physical network topology of Figure 273 on page 291.

10.5.1 Logical Network Topology

Shown in Figure 286 on page 301 is the logical view of the network.

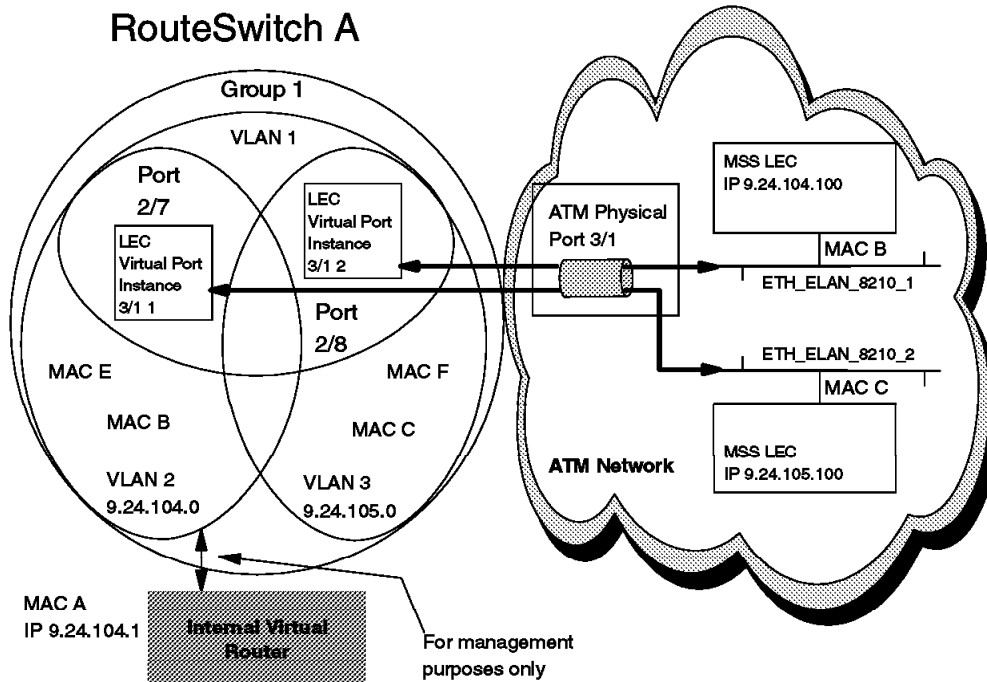


Figure 286. Logical View of Network with Multiple ELANs Routed at MSS

10.5.2 RouteSwitch Configuration

Follow the configuration steps starting in Figure 275 on page 292 until Figure 283 on page 298 for the basic VLAN and ATM LEC configuration.

```
SwitchA % cas 3/1

Slot 3 Port 1 Service 1 Configuration

1) Description (30 chars max)      : PTOP Bridging Service 1
2) Service type { LAN Emulation (1),
                  Trunking (4),
                  Classical IP(5),
                  PTOP Bridging(6),
                  VLAN cluster(7) } : PTOP Bridging
10) Encaps Type { Private(1),
                 RFC1483(2) }    : Private
3) Connection Type { PVC(1),
                    SVC(2) }      : PVC
4) PTOP Group                          : 1
5) PTOP connection                      : none
6) Admin Status { disable(1),
                  enable(2) }      : Enable

Enter (option=value/save/cancel) : 2=1
```

Figure 287. Creating the LEC Instance

```

Slot 3 Port 1 Service 1 Configuration

1) Description (30 chars max)           : LAN Emulation Service 1
2) Service type { LAN Emulation (1),
                  Trunking (4),
                  Classical IP(5),
                  PTOP Bridging(6),
                  VLAN cluster(7) } : LAN Emulation
21) LAN type { 802.3 (1),
               802.5 (2) }           : 802.3
22) Change LANE Cfg { NO (1),
                     YES (2) }      : NO
3) Connection Type { PVC(1),
                    SVC(2) }        : SVC
30) SEL for the ATM address           : 01
4) LAN Emulated Group                 : 1
5) LECS Address (40-char-hex)         : 4700790000000000000000000000A03E00000100
6) Admin Status { disable(1),
                 enable(2) }        : Enable

Enter (option=value/save/cancel) : 22=2

```

Figure 288. Modifying the LANE Configuration

In Figure 287 on page 301 and Figure 288, a second LEC instance was created. Note that each LEC instance gets its own ATM selector, which is the 20th byte of the ATM address.

Since all defined VLANs reside in group 1, this LEC is also assigned to group 1.

Note: Each LEC service may only be assigned to one group.

```

Slot 3 Port 1 Service 1 LANE Configuration Parameters

1) Proxy { NO (1), YES (2) }           : YES
2) Max Frame Size { 1516 (1), 4544 (2),
                   9234 (3), 18190 (4) } : 1516
3) Use translation options{NO (1), YES (2) } : Yes (use Swch menu to set)
4) Use Fwd Delay time { NO (1), YES (2) }   : NO
5) Use LE Cfg Server (LECS){ NO (1), YES (2)}: YES
6) Use Default LECS address { NO(1), YES (2)}: YES
7) Control Time-out (in seconds)           : 10
8) Max Unknown Frame Count                 : 10
9) Max Unknown Frame Time (in seconds)     : 1
10) VCC Time-out Period (in minutes)       : 20
11) Max Retry Count                       : 2
12) Aging Time (in seconds)                : 300
13) Expectd LE_ARP Resp Time (in seconds)  : 1
14) Flush Time-out (in seconds)            : 4
15) Path Switching Delay (in seconds)      : 6
16) ELAN name (32 chars max)              :

Enter (option=value/save/cancel) : 16=ETH_ELAN_8210_2

```

Figure 289. Definition of ELAN Name

This second LEC joins the other ELAN, which is ETH_ELAN_8210_2.

Note: Each LEC instance defined to one physical ATM port must join a different ELAN.


```

SwitchA % modatv1 1:3
VLAN 1: 3 is defined as:
  1. Description      = ip_9.24.105.0
  2. Admin Status     = Enabled
  3. Rule Definition
      Rule Num      Rule Type      Rule Status
      1             Net Addr Rule Enabled
Available options:
  1. Set VLAN Admin Status
  2. Set VLAN Description
  3. Add more rules
  4. Delete a rule
  5. Set rule Admin Status
  6. Quit
Option = 3
Select rule type:
  1. Port Rule
  2. MAC Address Rule
  3. Protocol Rule
  4. Network Address Rule
  5. User Defined Rule
  6. Binding Rule
  7. DHCP PORT Rule
  8. DHCP MAC Rule
Enter rule type (1): 1
Set Rule Admin Status to [(e)nable/(d)isable] (d): e
Enter the list of ports in Slot/Interface format: 3/1
3/1 has the following services configured:
Index  Service Type/Instance
=====
  1.    Lne/2
  2.    Lne/1
Enter the index of the service/instance to add: 1
  1. Description      = ip_9.24.105.0
  2. Admin Status     = Enabled
  3. Rule Definition
      Rule Num      Rule Type      Rule Status
      1             Net Addr Rule Enabled
      2             Port Rule      Enabled

```

Figure 290. Adding the Second LEC Instance to VLAN 3

Since multiple LEC instances are configured for group one, the RouteSwitch asks for the instance to choose for the VLAN assignment. Instance one is already assigned to VLAN two, thus instance two was chosen.

```

SwitchA % viv1
          Virtual Interface VLAN Membership
Slot/Intf/Service/Instance  Group  Member of VLAN#
-----
1  /1  /Rtr  /1      1      2
2  /1  /Brg  /1      1      1
2  /2  /Brg  /1      1      1
2  /3  /Brg  /1      1      1
2  /4  /Brg  /1      1      1
2  /5  /Brg  /1      1      1
2  /6  /Brg  /1      1      1
2  /7  /Brg  /1      1      1
2  /8  /Brg  /1      1      1
2  /9  /Brg  /1      1      1
2  /10 /Brg  /1      1      1
2  /11 /Brg  /1      1      1
2  /12 /Brg  /1      1      1
3  /1  /Lne  /1      1      1 2
3  /1  /Lne  /2      1      1 3
4  /2  /Brg  /1      1      1
SwitchA %

```

Figure 291. Viewing the VLAN Membership of Ports

Figure 291 shows the VLAN membership of both LEC instances. In this configuration, each VLAN has its own LEC and each LEC is connected to a different ELAN.

The internal virtual router is only connected to VLAN 2 thus no routing is performed within the RouteSwitch.

```

SwitchA % vas 3/1

          ATM Services

Slot  Port  Serv  Service  Service
====  ==  ===  =====  =====
3     1     1    LAN Emulation Service 1    LANE
3     1     2    LAN Emulation Service 2    LANE

          ATM Services

Slot  Port  Serv  VC  Oper  SEL Groups  Conn VCI's/Addresses
====  ==  ===  ==  =====  ==
3     1     1    SVC LANE Op. 01 1    225 226 227 228
                                     229
3     1     2    SVC LANE Op. 02 1    231 232 233 234
                                     235

SwitchA %

```

Figure 292. Verification of LANE Operation

The vas 3/1 command was used to verify the operational status of both LECs. In Figure 292, both LECs have joined the configured ELANs and established the VCCs to the LES/BUS of the MSS.

10.6 Configuration Example of the 8274 as an ATM Switch

The following example shows how the 8274 can be set up as an ATM switch. In the previous example 10.3, “RouteSwitch Interconnection Using MSS and 8260” on page 284, an 8260 is used as the ATM switch. In this example we use the 8274 as the ATM switch. We do not change the configuration of switch B or the MSS, which is now connected to CSM ports on the 8274.

Note: The MSS is used for LECS and LES/BUS functionality as the 8274 requires the code LSM.IMG in order to do LECS and LES/BUS internally. Please see Chapter 8, “LECS and LES/BUS Functionality” on page 213.

The physical view of the network is shown in Figure 265 on page 285.

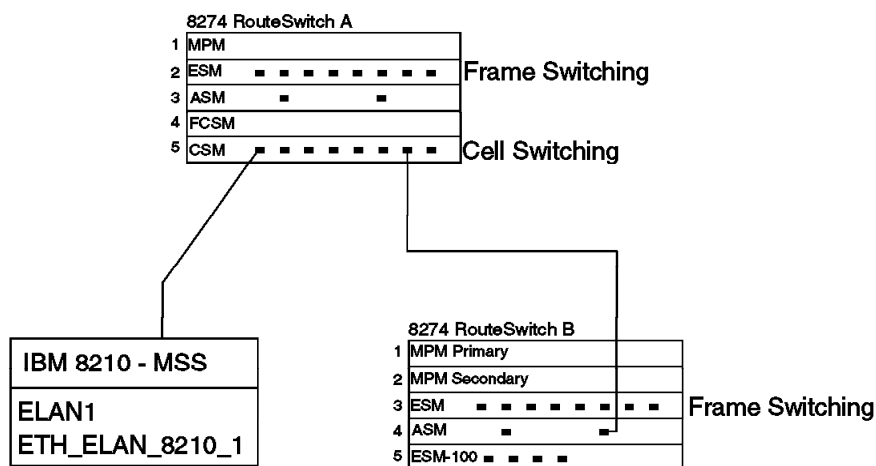


Figure 293. Physical View of Interconnection via ATM LANE

The configuration on switch B is unchanged from the previous example. Please see 10.3, “RouteSwitch Interconnection Using MSS and 8260” on page 284.

The IBM 8210 MSS configuration is also unchanged, and still has the ELAN ETH_ELAN_8210_1.

The configuration for switch A is as follows.

The internal port of the FCSM must be set to SVC. The FCSM module is in slot 4 of switch A, so the command is map 4/1. The configuration of this virtual port is similar to the configuration of the ASM port in switch B. The difference is that after saving the configuration shown in Figure 294 on page 306, the switch automatically continues with the configuration of the CSM portion of the virtual port 4/1, as seen in Figure 295 on page 306. No changes have to be made here and we just save the configuration. Please note that we are using UNI 3.0 in this example and thus leave option 7 in Figure 295 on page 306 unchanged.

```

Switch A / >map 4/1
Slot 4 Port 1 Configuration
1) Description (30 chars max)           : ATM PORT
2) Conn Type { PVC(1), SVC(2) }         : SVC
30) Sig version { 3.0(1) 3.1(2) }       : 3.0
31) Signaling VCI (0..1023)             : 5
32) ILMI Enable {(False(1),True(2))}    : True
33) ESI (12 hex-chars)                  : 0020da81d5e0
34) ILMI VCI (0..1023)                  : 16
3) Max VCCs (1-1023)                    : 1023
4) Max VCI bits (1..10)                  : 10
5) UNI Type                             : Private
6) Tx SAR Buffer Size (4096-131072)      : 16384
7) Rx SAR Buffer Size (4096-131072)      : 16384
8) Tx Frame Buffer Size (1800-16384)     : 4600
9) Rx Frame Buffer Size (1800-16384)     : 4600
10) Pl Scramble {(False(1),True(2))}    : True
11) Timing Mode {(Loop(1),Local(2))}    : Local
12) Loopback Config { NoLoop(1), DiagLoop(2), LineLoop(3) } : NoLoop
13) Phy media { SONET(1),SDH(2)}        : SONET
Enter (option=value/save/cancel) : save
Reset all services on slot 4 port 1 (n)? : y

```

Figure 294. Set the Internal Port of FCSM to SVC

```

Slot 4 Port 1 Configuration
1) Description (30 chars max)           : CSM PORT
2) Atm Address (40 hex-chars)           :
   00000000000000000000000000000000
3) Max VPI bits (1..12)                 : 2
4) Max VCI bits (1..12)                 : 10
5) I/F Type {Pub UNI(1), Pri UNI(2),
   PNNI(3), IISP(4)}                   : Private
6) Phy Protocol {SONET(1), SDH(2)}      : SONET
7) Signaling Ver {3.0(1), 3.1(2)}      : 3.0
8) ILMI Enable {False(1), True(2)}     : Enable

Enter (option=value/save/cancel) : save
Reset all connections on slot 4 port 1 (n)? : y
Resetting port, please wait...
send port change event
send port change event

Switch A / >

```

Figure 295. Configuration of the Internal FCSM Port

We now create a service on the virtual port in the FCSM. This configuration is exactly the same as configuring a service for the ASM port on switch B. The reason for this is that we are creating a LEC on the SAR in the FCSM virtual port that is like an ASM port, thus allowing for frame to cell conversion from switch A's frame switching modules. Please see Figure 174 on page 202 for a full description on the internal FCSM port.

```

Switch A / >cas 4/1
Slot 4 Port 1 Service 2 Configuration

1) Description (30 chars max) : LAN Emulation Service 2
2) Service type { LAN Emulation (1),
                  Trunking (4),
                  Classical IP(5),
                  PTOP Bridging(6),
                  VLAN cluster(7) } : LAN Emulation 1
21) LAN type { 802.3 (1),
               802.5 (2) } : 802.3 2
22) Change LANE Cfg { NO (1),
                     YES (2) } : NO
3) Connection Type { PVC(1),
                    SVC(2) } : SVC
30) SEL for the ATM address : 02
4) LAN Emulated Group : 1
5) LECS Address (40-char-hex) : 3
   4700790000000000000000000000A03E00000100
6) Admin Status { disable(1),
                 enable(2) } : Enable
Enter (option=value/save/cancel) :
Enter (option=value/save/cancel) : 22=2 4
Slot 4 Port 1 Service 2 LANE Configuration Parameters

1) Proxy { NO (1), YES (2) } : YES
2) Max Frame Size { 1516 (1), 4544 (2) 9234 (3), 18190 (4) } : 4544
3) Use translation options{NO (1), YES (2) : Yes (use Swch menu to set
4) Use Fwd Delay time { NO (1), YES (2) } : NO
5) Use LE Cfg Server (LECS){ NO (1), YES (2)}: YES
6) Use Default LECS address { NO(1), YES (2)}: YES
7) Control Time-out (in seconds) : 10
8) Max Unknown Frame Count : 10
9) Max Unknown Frame Time (in seconds) : 1
10) VCC Time-out Period (in minutes) : 20
11) Max Retry Count : 2
12) Aging Time (in seconds) : 300
13) Expected LE_ARP Resp Time (in seconds) : 1
14) Flush Time-out (in seconds) : 4
15) Path Switching Delay (in seconds) : 6
16) ELAN name (32 chars max) 5 : ETH_ELAN_8210_1
17) Ring Number (1 - 4095) : 0
18) Bridge Number (1 - 15) : 0

Enter (option=value/save/cancel) : save 6

```

Figure 296. Configuring a Service on FCSM Virtual Port

- 1** Create the service: LAN Emulation.
- 2** Select 802.3 because we are configuring a proxy LEC for the Ethernet ELAN on the MSS.
- 3** Enter the LECS address that is configured on the MSS.
- 4** Use the option 22=2 to enter the LANE configuration menu.
- 5** Set the ELAN name as used in the MSS.
- 6** Save the new configuration. This will also exit you from the configuration menu and enable the service.

10.7 Configuration of Two 8274s over MSS and Using DHCP Protocol Rule

The following example shows how two 8274s can be interconnected using MSS. The DHCP policy can be applied for the groups so that nodes that are configured for DHCP can boot with the IP address allocated by the DHCP server over ATM network.

In this example there are two 8274s: ibm8274a and ibm8274b. Both the 8274s are connected to the 8260 over two independent ATM ports. RouteSwitch ibm8274a has groups 2 and 3 in addition to default group. Group 2 has a network IP policy defined with the IP network 192.168.3.0 and net mask 255.255.255.0. Group 3 has two policies defined: network IP 192.168.4.0 with netmask 255.255.255.0 and DHCP.

RouteSwitch ibm8274b has only default group. The DHCP server will be connected to the ibm8274b RouteSwitch. There are two types of nodes in the network: one with the IP address defined and the others with the DHCP client enabled. These nodes are connected to the ibm8274a. All the nodes with the IP address in the 192.168.3.0 range will get connected to group 2 in the RouteSwitch. Nodes with the DHCP client enabled will get connected to group 3 and request for the IP address will be forwarded to the DHCP server connected to the other RouteSwitch ibm8274b through ATM switch. The DHCP server has the IP address range 192.168.4.0. Hence, the node remains in the group 3 of the ibm8274a as it also matches the second policy for that switch.

On both of the RouteSwitches LEC services are defined over ATM. There are two ELANs defined on the ATM switch: ELAN1 and ELAN2. Group 2 of ibm8274a is configured to get connected to the ELAN1, whereas group 3 of ibm8274a and default group of ibm8274b are configured to connect to ELAN2. Routing between the two ELANs is configured on the MSS.

10.7.1 Logical View of the Network

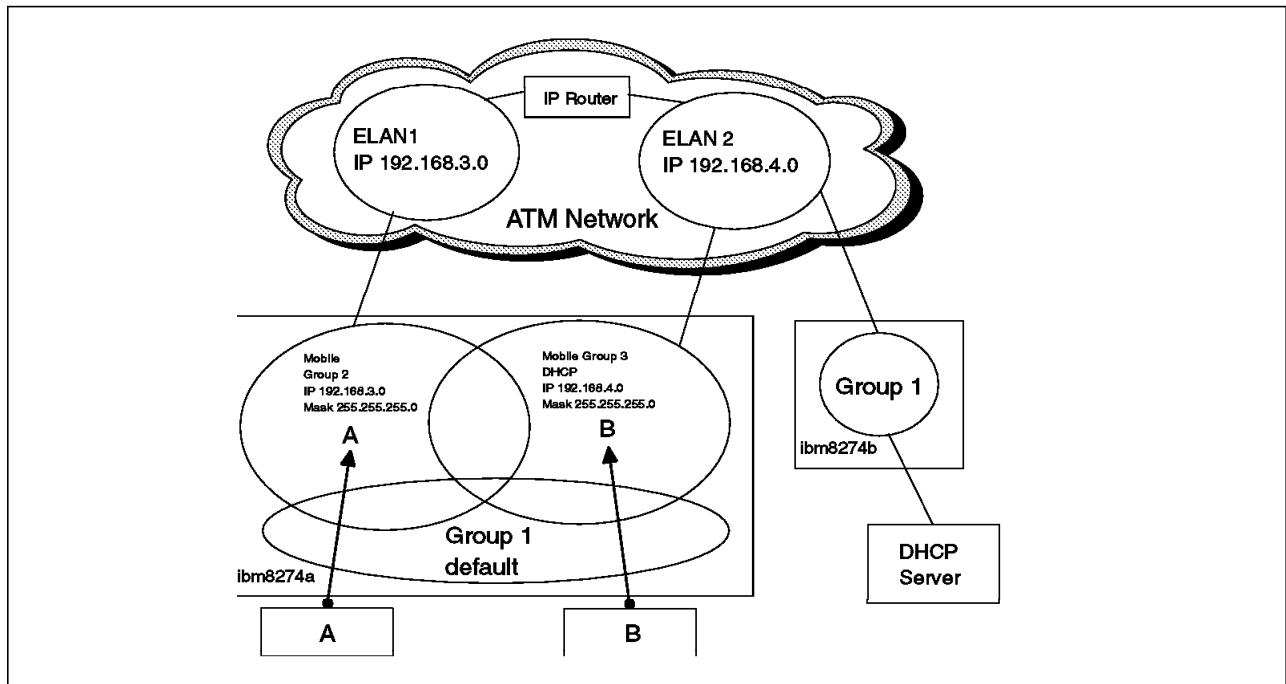


Figure 297. Logical View of the Network When Using DHCP

10.7.2 RouteSwitch Configuration

The following screens show the configuration steps necessary to implement the DHCP configuration.

1. Do the system configuration using the syscfg command.

```
/ % syscfg
System Contact           : Unset
System Name              : no_name
System Location          : Unset
System Description       : DESCRIPTION NOT SET.
Duplicate MAC Aging Timer : 0 (not configured)
Change any of the above {Y/N}? (N) : y
System Contact (Unset)   : John Parker
System Name (no_name)    : ibm8274a
System Location (Unset)  : LAB
System Description (DESCRIPTION NOT SET.) : RouteSwitch
Duplicate Mac Aging Timer (0) :
```

Figure 298. System Configuration of ibm8274a

2. Modify default VLAN configuration of the default group using the modvln command.

```

/ % modv1 1
Current values associated with GROUP 1.1 are as follows:

1) GROUP Number          - 1:1
2) Description            - Default GROUP (#1)
IP parameters:
3) IP enabled             - Y
4) IP Network Address     - 192.168.10.1
5) IP Subnet Mask         - 255.255.255.0
6) IP Broadcast Address   - 192.168.10.255
7) Router Description     - GROUP #1.0 IP router vport
8) RIP Mode               - Silent
   {Active(a), Inactive(i), Deaf(d), Silent(s)}
9) Routing disabled      - N
10) NHRP enabled         - N
11) Default Framing      - Ethernet II
   {Ethernet II(e), Ethernet 802.3(8), fddi(f),
    token ring(t), source route token ring(s)}
IPX parameters:
12) IPX enabled          - N

(save/quit/cancel)
: 4=9.24.105.99
New IP address generates new subnet and broadcast addresses.

```

Figure 299 (Part 1 of 2). Modifying the Default VLAN on ibm8274a


```

Enter '?' to view the changes.
: ?
1) GROUP Number          - 1:1
2) Description           - Default GROUP (#1)
IP parameters:
3) IP enabled            - Y
4) IP Network Address    - 9.24.105.99
5) IP Subnet Mask        - 255.0.0.0
6) IP Broadcast Address  - 9.255.255.255
7) Router Description    - GROUP #1.0 IP router vport
8) RIP Mode              - Silent
   {Active(a), Inactive(i), Deaf(d), Silent(s)}
9) Routing disabled      - N
10) NHRP enabled         - N
11) Default Framing      - Ethernet II
   {Ethernet II(e), Ethernet 802.3(8), fddi(f),
    token ring(t), source route token ring(s)}
IPX parameters:
12) IPX enabled          - N

: 5=255.255.255.0
New mask caused change in broadcast address.
: ?
1) GROUP Number          - 1:1
2) Description           - Default GROUP (#1)
IP parameters:
3) IP enabled            - Y
4) IP Network Address    - 9.24.105.99
5) IP Subnet Mask        - 255.255.255.0
6) IP Broadcast Address  - 9.24.105.255
7) Router Description    - GROUP #1.0 IP router vport
8) RIP Mode              - Silent
   {Active(a), Inactive(i), Deaf(d), Silent(s)}
9) Routing disabled      - N
10) NHRP enabled         - N
11) Default Framing      - Ethernet II
   {Ethernet II(e), Ethernet 802.3(8), fddi(f),
    token ring(t), source route token ring(s)}
IPX parameters:
12) IPX enabled          - N

: save

```

Figure 299 (Part 2 of 2). Modifying the Default VLAN on ibm8274a

3. Define static route parameters for the router port.

```

/ % aizr
Do you want to see the current route table? (y or n) (y) : y

                                IP ROUTING TABLE

Network      Mask      Gateway      Metric      Group VLAN
              Id: Id
-----
9.24.104.0    255.255.255.0    9.24.105.70      2          1:1
9.24.105.0    255.255.255.0    9.24.105.99      1          1:1
127.0.0.1     255.255.255.255  127.0.0.1        1          LOOPBACK
-----

Destination IP address of host or network : 0.0.0.0
IP address of next hop                     : 9.24.105.1
Route successfully added

```

Figure 300. Defining Static Route on ibm8274a

4. Disable IP from the default VLAN of the default group.

```

/ % modvl 1
Current values associated with GROUP 1.1 are as follows:

1) GROUP Number      - 1:1
2) Description        - Default GROUP (#1)
IP parameters:
3) IP enabled         - Y
4) IP Network Address - 9.24.105.99
5) IP Subnet Mask     - 255.255.255.0
6) IP Broadcast Address - 9.24.105.255
7) Router Description - GROUP #1.0 IP router vport
8) RIP Mode           - Silent
   {Active(a), Inactive(i), Deaf(d), Silent(s)}
9) Routing disabled   - N
10) NHRP enabled      - N
11) Default Framing   - Ethernet II
   {Ethernet II(e), Ethernet 802.3(8), fddi(f),
    token ring(t), source route token ring(s)}
IPX parameters:
12) IPX enabled       - N

(save/quit/cancel)
: 3=n
: ?
1) GROUP Number      - 1:1
2) Description        - Default GROUP (#1)
IP parameters:
3) IP enabled         - N
IPX parameters:
4) IPX enabled        - N

: save
/ %

```

Figure 301. Disabling IP from the Default VLAN

5. Enable group mobility on the switch.

```

/ % gmcfg
Group Mobility is Disabled. Enable Group Mobility ? yes/no (no): y
move_to_def is set to Disabled. Set to Enable ? yes/no (no):
def_group is set to Enable. Set it to Disable ? yes/no (no):
/ %

```

Figure 302. Enabling Group Mobility on ibm8274a

6. Create a new mobile group 2 for the IP network 192.168.3.0.

```

/ % crgp 2
GROUP Number ( 2 ) :
Description (no quotes) : IP Net 192.168.3.0
Enable WAN Routing? (n):
Enable ATM CIP? (n):
Enable IP (y) : n
Enable IPX? (y): n
Enable Group Mobility on this Group ? y/n(n): y
Enable User Authentication for this Group y/n(n):

Do you wish to configure the interface group for this Virtual LAN
at this time? (y) n

GROUP 2 has been added to the system.
You may add interfaces to this group using the addvp command at a later date.
For now, the GROUP is inactive until you add interfaces.
Configure Auto-Activated LEC service ? y/n(y): n

```

Figure 303. Creating Group 2 on ibm8274a

7. Add the AutoTracker network policy to the group.

```

Configure AutoTracker rules for this group y/n(y): y
Select rule type:
1. Port Rule
2. MAC Address Rule
3. Protocol Rule
4. Network Address Rule
5. User Defined Rule
6. Binding Rule
7. DHCP PORT Rule
8. DHCP MAC Rule
Enter rule type (1): 4
Set Rule Admin Status to (e)nable/(d)isable (d): e
Select the Network Protocol:
1. IP
2. IPX
Enter protocol type: 1
Enter the IP Address: 192.168.3.0
Enter the IP Mask (255.255.255.0): 255.255.255.0
Configure more rules for this vlan y/n (n):
VLAN 2: 1 created successfully
/ %

```

Figure 304. Adding AutoTracker Policy to Group 2

8. Create new group 3.

```

/ % crgp 3
  GROUP Number ( 3 ) :
  Description (no quotes) : DHCP and IP 192.168.4.0
  Enable WAN Routing? (n):
  Enable ATM CIP? (n):
  Enable IP (y) : n
  Enable IPX? (y): n
  Enable Group Mobility on this Group ? y/n(n): y
  Enable User Authentication for this Group y/n(n): n

  Do you wish to configure the interface group for this Virtual LAN
    at this time? (y) n

  GROUP 3 has been added to the system.
  You may add interfaces to this group using the addvp command at a later date.
  For now, the GROUP is inactive until you add interfaces.
  Configure Auto-Activated LEC service ? y/n(y): n

```

Figure 305. Creating Group 3 on ibm8274a

9. Add the AutoTracker DHCP policy to the mobile group 3.

```

Configure AutoTracker rules for this group y/n(y): y
Select rule type:
1. Port Rule
2. MAC Address Rule
3. Protocol Rule
4. Network Address Rule
5. User Defined Rule
6. Binding Rule
7. DHCP PORT Rule
8. DHCP MAC Rule
Enter rule type (1): 7
Set Rule Admin Status to (e)nable/(d)isable (d): e
Enter the list of ports in Slot/Interface format: 3/7-12
Configure more rules for this vlan y/n (n): y

```

Figure 306. Adding AutoTracker Policy to Group 3

10. Add another AutoTracker policy, IP network 192.168.4.0 to the mobile group 3.

```

Configure more rules for this vlan y/n (n): y
Select rule type:
  1. Port Rule
  2. MAC Address Rule
  3. Protocol Rule
  4. Network Address Rule
  5. User Defined Rule
  6. Binding Rule
  7. DHCP PORT Rule
  8. DHCP MAC Rule
Enter rule type (1): 4
Set Rule Admin Status to (e)nable/(d)isable (d): e
Select the Network Protocol:
  1. IP
  2. IPX
Enter protocol type: 1
Enter the IP Address: 192.168.4.0
Enter the IP Mask (255.255.255.0): 255.255.255.0
Configure more rules for this vlan y/n (n): n
VLAN 3: 1 created successfully
/ %

```

Figure 307. Adding Additional AutoTracker Policy to Group 3

11. View a service on the RouteSwitch ibm8274a using the vas command.

```

/ % vas

                                ATM Services

Slot Port Serv      Service      Service
  4   1   1   PTOP Bridging Service 1   PTOP Priv
  4   2   1   PTOP Bridging Service 1   PTOP Priv

                                ATM Services

Slot Port Serv VC   Oper
  4   1   1   PVC Enabled N/A 1
  4   2   1   PVC Disabled N/A 1
                                Conn VCI's/Addresses
  4   1   1   100
  4   2   1   100

FDDI Services do not exist!
/ %

```

Figure 308. Viewing Services on ibm8274a

12. Modify a port configuration of the switch using the map command.

```

Slot 4 Port 1 Configuration

1) Description (30 chars max) : ATM PORT
2) Conn Type { PVC(1), SVC(2) } : PVC
3) Max VCCs (1-1023) : 1023
4) Max VCI bits (1..10) : 10
5) UNI Type : Private
6) Tx SAR Buffer Size (2048-131072) : 16384
7) Rx SAR Buffer Size (2048-131072) : 16384
8) Tx Frame Buffer Size (1800-16384) : 4600
9) Rx Frame Buffer Size (1800-16384) : 4600
10) Pl Scramble {(False(1),True(2))} : True
11) Timing Mode {(Loop(1),Local(2))} : Local
12) Loopback Config { NoLoop(1), DiagLoop(2),
                    LineLoop(3) } : NoLoop
13) Phy media { SONET(1),SDH(2)} : SONET

Enter (option=value/save/cancel) :

```

13. Change the port connection type from PVC to SVC.

Figure 310. Modifying Connection Type

316 IBM Nways RouteSwitch Implementation Guide

```

Enter (option=value/save/cancel) : 32=1

Slot 4 Port 1 Configuration

1) Description (30 chars max) : ATM PORT
2) Conn Type { PVC(1), SVC(2) } : SVC

30) Sig version { 3.0(1) 3.1(2) } : 3.0
31) Signaling VCI (0..1023) : 5
32) ILMI Enable {(False(1),True(2))} : False
33) Net Prefix & ESI (38 hex-chars) : 000000000000000000000000000020da6fa8e1

3) Max VCCs (1-1023) : 1023
4) Max VCI bits (1..10) : 10
5) UNI Type : Private
6) Tx SAR Buffer Size (2048-131072) : 16384
7) Rx SAR Buffer Size (2048-131072) : 16384
8) Tx Frame Buffer Size (1800-16384) : 4600
9) Rx Frame Buffer Size (1800-16384) : 4600
10) Pl Scramble {(False(1),True(2))} : True
11) Timing Mode {(Loop(1),Local(2))} : Local
12) Loopback Config { NoLoop(1), DiagLoop(2),
LineLoop(3) } : NoLoop
13) Phy media { SONET(1),SDH(2)} : SONET

Enter (option=value/save/cancel) :

```

Figure 311. Disabling ILMI

15. Configure Net Prefix for the ATM port. It is a 38-hex character unique number for the port. At the end, save the configuration.

```

Enter (option=value/save/cancel) : 33=3911223344556677889900110140000008274a00

Slot 4 Port 1 Configuration

1) Description (30 chars max) : ATM PORT
2) Conn Type { PVC(1), SVC(2) } : SVC

30) Sig version { 3.0(1) 3.1(2) } : 3.0
31) Signaling VCI (0..1023) : 5
32) ILMI Enable {(False(1),True(2))} : False
33) Net Prefix & ESI (38 hex-chars) : 3911223344556677889900110140000008274a00

3) Max VCCs (1-1023) : 1023
4) Max VCI bits (1..10) : 10
5) UNI Type : Private
6) Tx SAR Buffer Size (2048-131072) : 16384
7) Rx SAR Buffer Size (2048-131072) : 16384
8) Tx Frame Buffer Size (1800-16384) : 4600
9) Rx Frame Buffer Size (1800-16384) : 4600
10) Pl Scramble {(False(1),True(2))} : True
11) Timing Mode {(Loop(1),Local(2))} : Local
12) Loopback Config { NoLoop(1), DiagLoop(2),
LineLoop(3) } : NoLoop
13) Phy media { SONET(1),SDH(2)} : SONET

Enter (option=value/save/cancel) : save

Reset all services on slot 4 port 1 (n)? : y
Resetting port, please wait...
/ %

```

Figure 312. Defining Net Prefix for ibm8274a

16. Create a new service for ELAN1 on port 4/1 using the cas command.

```

/ % cas 4/1

Slot 4 Port 1 Service 2 Configuration

1) Description (30 chars max) : PTOP Bridging Service 2
2) Service type { LANE client(1),
Trunking (4),
Classical IP(5),
PTOP Bridging(6),
VLAN cluster(7) } : PTOP Bridging
10) Encaps Type { Private(1),
RFC1483(2) } : Private
3) Connection Type { PVC(1),
SVC(2) } : PVC
4) PTOP Group : 1
5) PTOP connection : none
6) Admin Status { disable(1),
enable(2) } : Enable

Enter (option=value/save/cancel) :

```

Figure 313. Creating New Service 2 on ibm8274a

17. Change service type to LANE Client.


```

Enter (option=value/save/cancel) : 2=1

Slot 4 Port 1 Service 2 Configuration

1) Description (30 chars max)           : LAN Emulation Service 2
2) Service type { LANE client(1),
                  Trunking (4),
                  Classical IP(5),
                  PTOp Bridging(6),
                  VLAN cluster(7) } : LAN Emulation
21) LAN type { 802.3 (1),
               802.5 (2) }           : 802.3
22) Change LANE Cfg { NO (1),
                     YES (2) }      : NO
3) Connection Type { PVC(1),
                    SVC(2) }        : SVC
30) SEL for the ATM address           : 02
4) LAN Emulated Group                 : 1
5) LECS Address (40-char-hex)         : 4700790000000000000000000000A03E00000100
6) Admin Status { disable(1),
                  enable(2) }       : Enable

Enter (option=value/save/cancel) :

```

Figure 314. Modifying Service Type

18. Change LANE configuration.

```

Enter (option=value/save/cancel) : 22=2

Slot 4 Port 1 Service 2 LANE Configuration Parameters

1) Proxy { NO (1), YES (2) }           : YES
2) Max Frame Size { 1516 (1), 4544 (2),
                   9234 (3), 18190 (4) } : 1516
3) Use translation options{NO (1), YES (2) : Yes (use Swch menu to set)
4) Use Fwd Delay time { NO (1), YES (2) } : NO
5) Use LE Cfg Server (LECS){ NO (1), YES (2)}: YES
6) Use Default LECS address { NO(1), YES (2)}: YES
7) Control Time-out (in seconds)         : 10
8) Max Unknown Frame Count               : 10
9) Max Unknown Frame Time (in seconds)   : 1
10) VCC Time-out Period (in minutes)     : 20
11) Max Retry Count                      : 2
12) Aging Time (in seconds)              : 300
13) Expectd LE_ARP Resp Time (in seconds) : 1
14) Flush Time-out (in seconds)          : 4
15) Path Switching Delay (in seconds)    : 6
16) ELAN name (32 chars max)             :

Enter (option=value/save/cancel) :

```

Figure 315. Modifying LANE Configuration

19. Define the ELAN name to which this service will be connected.

```

Enter (option=value/save/cancel) : 16=ELAN1

Slot 4 Port 1 Service 2 LANE Configuration Parameters

1) Proxy { NO (1), YES (2) } : YES
2) Max Frame Size { 1516 (1), 4544 (2)
   9234 (3), 18190 (4) } : 1516
3) Use translation options{NO (1), YES (2) : Yes (use Swch menu to set)
4) Use Fwd Delay time { NO (1), YES (2) } : NO
5) Use LE Cfg Server (LECS){ NO (1), YES (2)}: YES
6) Use Default LECS address { NO(1), YES (2)}: YES
7) Control Time-out (in seconds) : 10
8) Max Unknown Frame Count : 10
9) Max Unknown Frame Time (in seconds) : 1
10) VCC Time-out Period (in minutes) : 20
11) Max Retry Count : 2
12) Aging Time (in seconds) : 300
13) Expectd LE_ARP Resp Time (in seconds) : 1
14) Flush Time-out (in seconds) : 4
15) Path Switching Delay (in seconds) : 6
16) ELAN name (32 chars max) : ELAN1

Enter (option=value/save/cancel) : save

Saving new LANE Configuration values

Slot 4 Port 1 Service 2 Configuration

1) Description (30 chars max) : LAN Emulation Service 2
2) Service type { LANE client(1),
   Trunking (4),
   Classical IP(5),
   PTOP Bridging(6),
   VLAN cluster(7) } : LAN Emulation
21) LAN type { 802.3 (1),
   802.5 (2) } : 802.3
22) Change LANE Cfg { NO (1),
   YES (2) } : NO
3) Connection Type { PVC(1),
   SVC(2) } : SVC
30) SEL for the ATM address : 02
4) LAN Emulated Group : 1
5) LECS Address (40-char-hex) : 4700790000000000000000000000A03E00000100
6) Admin Status { disable(1),
   enable(2) } : Enable

Enter (option=value/save/cancel) :

```

Figure 316. Modifying ELAN Name for the Service and Saving the Configuration

20. Define mobile groups that need to be added to this service. Here in this example group 2 will be added to the service which gets connected to the ELAN1. Define appropriate LECS address in option 5 and save the configuration.

```

Enter (option=value/save/cancel) : 4=2

Slot 4 Port 1 Service 2 Configuration

1) Description (30 chars max)      : LAN Emulation Service 2
2) Service type { LANE client(1),
                  Trunking (4),
                  Classical IP(5),
                  PTOP Bridging(6),
                  VLAN cluster(7) } : LAN Emulation
21) LAN type { 802.3 (1),
               802.5 (2) }         : 802.3
22) Change LANE Cfg { NO (1),
                     YES (2) }      : NO
3) Connection Type { PVC(1),
                    SVC(2) }        : SVC
30) SEL for the ATM address         : 02
4) LAN Emulated Group               : 2
5) LECS Address (40-char-hex)       : 4700790000000000000000000000A03E00000100
6) Admin Status { disable(1),
                  enable(2) }       : Enable

Enter (option=value/save/cancel) : save
Creating service, please wait...

Enabling service...
/ %

```

Figure 317. Adding Mobile Group to the LANE Group

21. Create service 3 on the port 4/1 of the ibm8274a switch.

```

/ % cas 4/1 3

Slot 4 Port 1 Service 3 Configuration

1) Description (30 chars max)      : PTOP Bridging Service 3
2) Service type { LANE client(1),
                  Trunking (4),
                  Classical IP(5),
                  PTOP Bridging(6),
                  VLAN cluster(7) } : PTOP Bridging
10) Encaps Type { Private(1),
                 RFC1483(2) }      : Private
3) Connection Type { PVC(1),
                    SVC(2) }        : PVC
4) PTOP Group                     : 1
5) PTOP connection                 : none
6) Admin Status { disable(1),
                  enable(2) }       : Enable

Enter (option=value/save/cancel) :

```

Figure 318. Creating an Additional Service on ibm8274a

22. Change service type to LANE Client.

```

Enter (option=value/save/cancel) : 2=1

Slot 4 Port 1 Service 3 Configuration

1) Description (30 chars max)           : LAN Emulation Service 3
2) Service type { LANE client(1),
                  Trunking (4),
                  Classical IP(5),
                  PTOP Bridging(6),
                  VLAN cluster(7) } : LAN Emulation
21) LAN type { 802.3 (1),
               802.5 (2) }           : 802.3
22) Change LANE Cfg { NO (1),
                     YES (2) }      : NO
3) Connection Type { PVC(1),
                    SVC(2) }        : SVC
30) SEL for the ATM address           : 03
4) LAN Emulated Group                 : 1
5) LECS Address (40-char-hex)         : 4700790000000000000000000000A03E00000100
6) Admin Status { disable(1),
                 enable(2) }        : Enable

Enter (option=value/save/cancel) :

```

Figure 319. Defining Service Type

23. Change LANE configuration for this service.

```

Enter (option=value/save/cancel) : 22=2

Slot 4 Port 1 Service 3 LANE Configuration Parameters

1) Proxy { NO (1), YES (2) }           : YES
2) Max Frame Size { 1516 (1), 4544 (2),
                   9234 (3), 18190 (4) } : 1516
3) Use translation options{NO (1), YES (2) : Yes (use Swch menu to set)
4) Use Fwd Delay time { NO (1), YES (2) } : NO
5) Use LE Cfg Server (LECS){ NO (1), YES (2)}: YES
6) Use Default LECS address { NO(1), YES (2)}: YES
7) Control Time-out (in seconds)         : 10
8) Max Unknown Frame Count               : 10
9) Max Unknown Frame Time (in seconds)   : 1
10) VCC Time-out Period (in minutes)     : 20
11) Max Retry Count                     : 2
12) Aging Time (in seconds)              : 300
13) Expectd LE_ARP Resp Time (in seconds) : 1
14) Flush Time-out (in seconds)          : 4
15) Path Switching Delay (in seconds)    : 6
16) ELAN name (32 chars max)            :

Enter (option=value/save/cancel) :

```

Figure 320. Changing Service Type

24. Define the name of the ELAN to which this service is going to be connected.
Here in this example it is the ELAN2.

```

Enter (option=value/save/cancel) : 16=ELAN2

Slot 4 Port 1 Service 3 LANE Configuration Parameters

1) Proxy { NO (1), YES (2) } : YES
2) Max Frame Size { 1516 (1), 4544 (2)
   9234 (3), 18190 (4) } : 1516
3) Use translation options{NO (1), YES (2) : Yes (use Swch menu to set)
4) Use Fwd Delay time { NO (1), YES (2) } : NO
5) Use LE Cfg Server (LECS){ NO (1), YES (2)}: YES
6) Use Default LECS address { NO(1), YES (2)}: YES
7) Control Time-out (in seconds) : 10
8) Max Unknown Frame Count : 10
9) Max Unknown Frame Time (in seconds) : 1
10) VCC Time-out Period (in minutes) : 20
11) Max Retry Count : 2
12) Aging Time (in seconds) : 300
13) Expectd LE_ARP Resp Time (in seconds) : 1
14) Flush Time-out (in seconds) : 4
15) Path Switching Delay (in seconds) : 6
16) ELAN name (32 chars max) : ELAN2

Enter (option=value/save/cancel) :

```

Figure 321. Defining ELAN Name

25. Save the configuration.

```

Enter (option=value/save/cancel) : save

Saving new LANE Configuration values

Slot 4 Port 1 Service 3 Configuration

1) Description (30 chars max) : LAN Emulation Service 3
2) Service type { LANE client(1),
   Trunking (4),
   Classical IP(5),
   PTOP Bridging(6),
   VLAN cluster(7) } : LAN Emulation
21) LAN type { 802.3 (1),
   802.5 (2) } : 802.3
22) Change LANE Cfg { NO (1),
   YES (2) } : NO
3) Connection Type { PVC(1),
   SVC(2) } : SVC
30) SEL for the ATM address : 03
4) LAN Emulated Group : 1
5) LECS Address (40-char-hex) : 4700790000000000000000000000A03E00000100
6) Admin Status { disable(1),
   enable(2) } : Enable

Enter (option=value/save/cancel) :

```

Figure 322. Saving the LANE Configuration

26. Add a mobile group to the LAN Emulated Group in option 4. Define the appropriate LECS address in option 5 and save the configuration.

```

Enter (option=value/save/cancel) : 4=3

Slot 4 Port 1 Service 3 Configuration

1) Description (30 chars max)      : LAN Emulation Service 3
2) Service type { LANE client(1),
                  Trunking (4),
                  Classical IP(5),
                  PTOP Bridging(6),
                  VLAN cluster(7) } : LAN Emulation
21) LAN type { 802.3 (1),
               802.5 (2) }         : 802.3
22) Change LANE Cfg { NO (1),
                     YES (2) }    : NO
3) Connection Type { PVC(1),
                    SVC(2) }      : SVC
30) SEL for the ATM address       : 03
4) LAN Emulated Group            : 3
5) LECS Address (40-char-hex)    : 4700790000000000000000000000A03E00000100
6) Admin Status { disable(1),
                 enable(2) }      : Enable

Enter (option=value/save/cancel) : save
Creating service, please wait...

Enabling service...
/ %

```

Figure 323. Adding Mobile Group to the LANE Group and Saving the Configuration

27. View the services using the vas command.

```

/ % vas

ATM Services

Slot Port Serv Service Service
Num Num Description Type
====
4 1 2 LAN Emulation Service 2 802.3 LEC
4 1 1 PTOP Bridging Service 1 PTOP Priv
4 1 3 LAN Emulation Service 3 802.3 LEC
4 2 1 PTOP Bridging Service 1 PTOP Priv

ATM Services

Slot Port Serv VC Oper
Num Num Type Status SEL Groups Conn VCI's/Addresses
====
4 1 2 SVC Initial 02 2
4 1 1 PVC Disabled N/A 1 100
4 1 3 SVC Initial 03 3
4 2 1 PVC Disabled N/A 1 100

FDDI Services do not exist!

/ %

```

Figure 324. Viewing Services on ibm8274a

28. Use the syscfg command to show or alter the configuration of the ibm8274b RouteSwitch.

```
/ % syscfg
System Contact           : Unset
System Name              : no_name
System Location          : Unset
System Description       : DESCRIPTION NOT SET.
Duplicate MAC Aging Timer : 0 (not configured)
Change any of the above {Y/N}? (N) : y
System Contact (Unset) : John Parker
System Name (no_name) : ibm8274b
System Location (Unset) : LAB
System Description (DESCRIPTION NOT SET.) : RouteSwitch
Duplicate Mac Aging Timer (0) :
```

Figure 325. System Configuration of the ibm8274b

29. Modify the default VLAN configuration of the default group using the modvl command.

```

/ % modv1 1
Current values associated with GROUP 1.1 are as follows:

1) GROUP Number          - 1:1
2) Description            - Default GROUP (#1)
IP parameters:
3) IP enabled             - Y
4) IP Network Address     - 192.168.10.1
5) IP Subnet Mask         - 255.255.255.0
6) IP Broadcast Address   - 192.168.10.255
7) Router Description     - GROUP #1.0 IP router vport
8) RIP Mode               - Silent
   {Active(a), Inactive(i), Deaf(d), Silent(s)}
9) Routing disabled       - N
10) NHRP enabled          - N
11) Default Framing       - Ethernet II
   {Ethernet II(e), Ethernet 802.3(8), fddi(f),
    token ring(t), source route token ring(s)}
IPX parameters:
12) IPX enabled           - N

(save/quit/cancel)
: 4=192.168.4.99
New IP address generates new subnet and broadcast addresses.
Enter '?' to view the changes.
: ?
1) GROUP Number          - 1:1
2) Description            - Default GROUP (#1)
IP parameters:
3) IP enabled             - Y
4) IP Network Address     - 192.168.4.99
5) IP Subnet Mask         - 255.255.255.0
6) IP Broadcast Address   - 9.168.4.255
7) Router Description     - GROUP #1.0 IP router vport
8) RIP Mode               - Silent
   {Active(a), Inactive(i), Deaf(d), Silent(s)}
9) Routing disabled       - N
10) NHRP enabled          - N
11) Default Framing       - Ethernet II
   {Ethernet II(e), Ethernet 802.3(8), fddi(f),
    token ring(t), source route token ring(s)}
IPX parameters:
12) IPX enabled           - N

: 5=255.255.255.0
New mask caused change in broadcast address.

```

Figure 326 (Part 1 of 2). Modifying Default VLAN of ibm8274b


```

: ?
1) GROUP Number      - 1:1
2) Description       - Default GROUP (#1)
IP parameters:
3) IP enabled        - Y
4) IP Network Address - 192.168.4.99
5) IP Subnet Mask    - 255.255.255.0
6) IP Broadcast Address - 192.168.4.255
7) Router Description - GROUP #1.0 IP router vport
8) RIP Mode          - Silent
   {Active(a), Inactive(i), Deaf(d), Silent(s)}
9) Routing disabled  - N
10) NHRP enabled     - N
11) Default Framing  - Ethernet II
   {Ethernet II(e), Ethernet 802.3(8), fddi(f),
    token ring(t), source route token ring(s)}
IPX parameters:
12) IPX enabled      - N

: save

```

Figure 326 (Part 2 of 2). Modifying Default VLAN of ibm8274b

30. Define static route parameters for the router port.

```

/ % a1sr
Do you want to see the current route table? (y or n) (y) : y

                        IP ROUTING TABLE

Network      Mask      Gateway      Metric      Group VLAN
-----
9.24.104.0    255.255.255.0    9.24.105.70      2          1:1
192.168.4.0   255.255.255.0    192.168.4.99     1          1:1
127.0.0.1     255.255.255.255  127.0.0.1        1          LOOPBACK
-----

Destination IP address of host or network : 0.0.0.0
IP address of next hop                    : 9.24.105.1
Route successfully added

```

Figure 327. Adding Static Route

31. Modify port parameters of the service port 4/1.

```

Slot 4 Port 1 Configuration

1) Description (30 chars max) : ATM PORT
2) Conn Type { PVC(1), SVC(2) } : PVC
3) Max VCCs (1-1023) : 1023
4) Max VCI bits (1..10) : 10
5) UNI Type : Private
6) Tx SAR Buffer Size (2048-131072) : 16384
7) Rx SAR Buffer Size (2048-131072) : 16384
8) Tx Frame Buffer Size (1800-16384) : 4600
9) Rx Frame Buffer Size (1800-16384) : 4600
10) Pl Scramble {(False(1),True(2))} : True
11) Timing Mode {(Loop(1),Local(2))} : Local
12) Loopback Config { NoLoop(1), DiagLoop(2),
                        LineLoop(3) } : NoLoop
13) Phy media { SONET(1),SDH(2)} : SONET

Enter (option=value/save/cancel) :

```

32. Change connection type to SVC.

```

Enter (option=value/save/cancel) : 2=2

Slot 4 Port 1 Configuration

1) Description (30 chars max) : ATM PORT
2) Conn Type { PVC(1), SVC(2) } : SVC

30) Sig version { 3.0(1) 3.1(2) } : 3.0
31) Signaling VCI (0..1023) : 5
32) ILMI Enable {(False(1),True(2))} : True
33) ESI (12 hex-chars) : 40000008274a
34) ILMI VCI (0..1023) : 16

3) Max VCCs (1-1023) : 1023
4) Max VCI bits (1..10) : 10
5) UNI Type : Private
6) Tx SAR Buffer Size (2048-131072) : 16384
7) Rx SAR Buffer Size (2048-131072) : 16384
8) Tx Frame Buffer Size (1800-16384) : 4600
9) Rx Frame Buffer Size (1800-16384) : 4600
10) Pl Scramble {(False(1),True(2))} : True
11) Timing Mode {(Loop(1),Local(2))} : Local
12) Loopback Config { NoLoop(1), DiagLoop(2),
LineLoop(3) } : NoLoop
13) Phy media { SONET(1),SDH(2)} : SONET

Enter (option=value/save/cancel) :

```

33. Change ILMI Enable to False.

```

Enter (option=value/save/cancel) : 32=1

Slot 4 Port 1 Configuration

1) Description (30 chars max) : ATM PORT
2) Conn Type { PVC(1), SVC(2) } : SVC

30) Sig version { 3.0(1) 3.1(2) } : 3.0
31) Signaling VCI (0..1023) : 5
32) ILMI Enable {(False(1),True(2))} : False
33) Net Prefix & ESI (38 hex-chars) : 3911223344556677889900110240000008274a

3) Max VCCs (1-1023) : 1023
4) Max VCI bits (1..10) : 10
5) UNI Type : Private
6) Tx SAR Buffer Size (2048-131072) : 16384
7) Rx SAR Buffer Size (2048-131072) : 16384
8) Tx Frame Buffer Size (1800-16384) : 4600
9) Rx Frame Buffer Size (1800-16384) : 4600
10) Pl Scramble {(False(1),True(2))} : True
11) Timing Mode {(Loop(1),Local(2))} : Local
12) Loopback Config { NoLoop(1), DiagLoop(2),
LineLoop(3) } : NoLoop
13) Phy media { SONET(1),SDH(2)} : SONET

Enter (option=value/save/cancel) :

```

Figure 330. Disabling ILMI

34. Change Net Prefix ATM address for the port and save the configuration.

```
Enter (option=value/save/cancel) : 33=3911223344556677889900110240000008274b
```

Slot 4 Port 1 Configuration

```
1) Description (30 chars max)           : ATM PORT
2) Conn Type { PVC(1), SVC(2) }         : SVC

30) Sig version { 3.0(1) 3.1(2) }       : 3.0
31) Signaling VCI (0..1023)             : 5
32) ILMI Enable {(False(1),True(2))}    : False
33) Net Prefix & ESI (38 hex-chars)    : 3911223344556677889900110240000008274b

3) Max VCCs (1-1023)                   : 1023
4) Max VCI bits (1..10)                 : 10
5) UNI Type                             : Private
6) Tx SAR Buffer Size (2048-131072)      : 16384
7) Rx SAR Buffer Size (2048-131072)      : 16384
8) Tx Frame Buffer Size (1800-16384)     : 4600
9) Rx Frame Buffer Size (1800-16384)     : 4600
10) Pl Scramble {(False(1),True(2))}    : True
11) Timing Mode {(Loop(1),Local(2))}    : Local
12) Loopback Config { NoLoop(1), DiagLoop(2),
    LineLoop(3) }                       : NoLoop
13) Phy media { SONET(1),SDH(2)}         : SONET
```

```
Enter (option=value/save/cancel) : save
```

```
Reset all services on slot 4 port 1 (n)? : y
Resetting port, please wait...
/ %
```

Figure 331. Defining Net Prefix for ibm8274b

35. Create a new service for the RouteSwitch ibm8274b.

```
/ % cas 4/1
```

Slot 4 Port 1 Service 2 Configuration

```
1) Description (30 chars max)           : PTOP Bridging Service 2
2) Service type { LANE client(1),
    Trunking (4),
    Classical IP(5),
    PTOP Bridging(6),
    VLAN cluster(7) } : PTOP Bridging
10) Encaps Type { Private(1),
    RFC1483(2) }      : Private
3) Connection Type { PVC(1),
    SVC(2) }          : PVC
4) PTOP Group         : 1
5) PTOP connection    : none
6) Admin Status { disable(1),
    enable(2) }       : Enable
```

```
Enter (option=value/save/cancel) :
```

Figure 332. Creating New Service on ibm8274b

36. Change service type to LANE Client.

```

Enter (option=value/save/cancel) : 2=1

Slot 4 Port 1 Service 2 Configuration

1) Description (30 chars max)           : LAN Emulation Service 2
2) Service type { LANE client(1),
                  Trunking (4),
                  Classical IP(5),
                  PTOp Bridging(6),
                  VLAN cluster(7) } : LAN Emulation
21) LAN type { 802.3 (1),
               802.5 (2) }           : 802.3
22) Change LANE Cfg { NO (1),
                     YES (2) }      : NO
3) Connection Type { PVC(1),
                    SVC(2) }        : SVC
30) SEL for the ATM address           : 02
4) LAN Emulated Group                 : 1
5) LECS Address (40-char-hex)         : 4700790000000000000000000000A03E00000100
6) Admin Status { disable(1),
                 enable(2) }        : Enable

Enter (option=value/save/cancel) :

```

Figure 333. Changing Service Type

37. Change LANE configuration for the RouteSwitch ibm8274b.

```

Enter (option=value/save/cancel) : 22=2

Slot 4 Port 1 Service 2 LANE Configuration Parameters

1) Proxy { NO (1), YES (2) }           : YES
2) Max Frame Size { 1516 (1), 4544 (2),
                   9234 (3), 18190 (4) } : 1516
3) Use translation options{NO (1), YES (2) : Yes (use Swch menu to set)
4) Use Fwd Delay time { NO (1), YES (2) } : NO
5) Use LE Cfg Server (LECS){ NO (1), YES (2)}: YES
6) Use Default LECS address { NO(1), YES (2)}: YES
7) Control Time-out (in seconds)         : 10
8) Max Unknown Frame Count                : 10
9) Max Unknown Frame Time (in seconds)    : 1
10) VCC Time-out Period (in minutes)      : 20
11) Max Retry Count                       : 2
12) Aging Time (in seconds)               : 300
13) Expectd LE_ARP Resp Time (in seconds) : 1
14) Flush Time-out (in seconds)           : 4
15) Path Switching Delay (in seconds)     : 6
16) ELAN name (32 chars max)              :

Enter (option=value/save/cancel) :

```

Figure 334. Modifying LANE Configuration

38. Configure the ELAN name for the service to which this service port will get connected.

```

Enter (option=value/save/cancel) : 16=ELAN2

Slot 4 Port 1 Service 2 LANE Configuration Parameters

1) Proxy { NO (1), YES (2) } : YES
2) Max Frame Size { 1516 (1), 4544 (2)
   9234 (3), 18190 (4) } : 1516
3) Use translation options{NO (1), YES (2) : Yes (use Swch menu to set)
4) Use Fwd Delay time { NO (1), YES (2) } : NO
5) Use LE Cfg Server (LECS){ NO (1), YES (2)}: YES
6) Use Default LECS address { NO(1), YES (2)}: YES
7) Control Time-out (in seconds) : 10
8) Max Unknown Frame Count : 10
9) Max Unknown Frame Time (in seconds) : 1
10) VCC Time-out Period (in minutes) : 20
11) Max Retry Count : 2
12) Aging Time (in seconds) : 300
13) Expectd LE_ARP Resp Time (in seconds) : 1
14) Flush Time-out (in seconds) : 4
15) Path Switching Delay (in seconds) : 6
16) ELAN name (32 chars max) : ELAN2

Enter (option=value/save/cancel) :

```

Figure 335. Configuring ELAN Name for the Service

39. Save the LANE configuration.

```

Enter (option=value/save/cancel) : save

Saving new LANE Configuration values

Slot 4 Port 1 Service 2 Configuration

1) Description (30 chars max) : LAN Emulation Service 2
2) Service type { LANE client(1),
   Trunking (4),
   Classical IP(5),
   PTOP Bridging(6),
   VLAN cluster(7) } : LAN Emulation
21) LAN type { 802.3 (1),
   802.5 (2) } : 802.3
22) Change LANE Cfg { NO (1),
   YES (2) } : NO
3) Connection Type { PVC(1),
   SVC(2) } : SVC
30) SEL for the ATM address : 02
4) LAN Emulated Group : 1
5) LECS Address (40-char-hex) : 4700790000000000000000000000A03E00000100
6) Admin Status { disable(1),
   enable(2) } : Enable

Enter (option=value/save/cancel) :

```

Figure 336. Save the LANE Configuration

40. Define the appropriate LECS address in option 5.

41. Save the service configuration.

```
Enter (option=value/save/cancel) : save  
Creating service, please wait...  
  
Enabling service...  
/ %
```

Figure 337. Saving the Service Configuration on ibm8274b

Appendix A. Correcting ARI/FCI Issues

The ARI/FCI issues came to light with the advent of the IBM 8272 Token-Ring switch.

A diskette was created that contained upgrades to various drivers to help alleviate the ARI/FCI issues.

The current diskette is available at the following Web site:
<http://www.networking.ibm.com>. Follow the link to the 8272 to find the file trdrvfix.exe.

Currently the URL to the file is:
<http://www.networking.ibm.com/nes/nesswite.htm#MISC>.

The diskette contains the following files:

- FILE_ID.DIZ
- READ.ME
- LSP138P.EXE
- PRODAID.TXT
- DXMC0MOD.SYS
- DXMC5MOD.SYS
- NETBEUI.OS2
- LANDD.OS2
- LANDD.NIF
- NET.EXE

A.1 Read Me File from the TRDRVFIX Diskette

=====

Token-Ring Protocol Driver Fixes, Version 1.20

March 25, 1996

=====

It has been revealed that some IBM token-ring drivers are susceptible to a latent condition that appears under certain configurations when an 8272 switch is in the client/server path. The affected IBM protocol drivers do not handle the ARI/FCI bit (Address Recognized and Frame Copied Indicator) correctly.

The symptoms of this condition are:

- 1) A client/server application fails to establish a session
- 2) A client/server session can be established but then fails when file transfer or other normal session activity occurs

- 3) TCP/IP sessions will log excessive errors but will not drop session links.

The condition does not occur in all client/server applications. In general, the symptoms will be apparent in different combinations of applications and network configurations.

Affected Configurations:

The condition is confined to these specific switch configurations:

- 1) Client or server stations must be on a ring that is directly attached to a switch port AND
- 2) There is an intermediate Source Route Bridge in the path

Applications That Are Affected:

The following applications are affected (driver dependent):

- 1) All DOS client/server applications that reside above the following DOS client device drivers:

JETBEUI Driver	DXMJOMOD.SYS
----------------	--------------

Native Token-Ring Support driver	DXMCOMOD.SYS DXMC1MOD.SYS DXMC5MOD.SYS
----------------------------------	--

DOS LAN Services (DLS) Driver	NET.EXE
-------------------------------	---------

- 2) All OS/2 client/server applications utilize the following:

NETBIOS Device Driver	NETBEUI.OS2
802.2 Device Driver	LANDD.OS2

- 3) All OS/2 and DOS TCP/IP client/server applications that utilize IBM TCP/IP drivers (false errors only)

Replacement drivers Available:

Replacement drivers are now available for the IBM drivers that are affected (see below).

Applications that are not affected (regardless of configuration):

The following applications have not shown to be affected to date:

- * Novell applications are not affected if they are using Novell drivers.
- * TCP/IP sessions can still be maintained, but false frame errors will be logged.
- * The condition has not been seen in RS/6000s.
- * 3745 TICs are not affected
- * The condition has not been observed with 3174 configurations

The following IBM protocol drivers are not affected:

- * The new IBM LAN CLIENT drivers are not affected (DOS, Windows, and Windows for Workgroup).

* The DOS IEEE 802.2 Protocol Driver (DXME0MOD.SYS) is not affected.

Configurations that are not affected:

The condition is confined to a limited set of 8272 and source route bridge configurations. The following configurations do not require the driver modifications:

- * Configurations that have a switch as a central backbone with source route bridges to other rings do not require driver updates (i.e., there is a bridge between every ring and the switch).
- * Dedicated, full-duplex server or client links are not affected

Configurations that may or may not be affected:

- * Configurations that have no source route bridges in the client/server path may not have the problem.

The impact of 8272 implementations in installations with other non-IBM drivers has not been determined. IBM will update the list of affected drivers if any others are discovered.

=====

To Obtain Replacement Drivers:

Replacement drivers are available for download that includes the updated drivers in a single package, which is this TRDRVFIX.ZIP.

Internal IBM users may obtain this driver package via TOOLCAT on LANPROD.

The driver package is also available via Internet access at:
<http://www.networking.ibm.com>

A BBS service is offered at (919)517-0001 (8-N-1).

=====

This archive includes the following updates for the IBM protocol drivers for DOS and OS/2 systems (except TCP/IP):

OS/2 802.2 (LANDD.OS2)

Some OS/2 802.2 applications, e.g. for access to a host, may inspect the Address Recognized/Frame Copied indicators to determine whether or not a frame was copied. This could lead to failure of the application, such as loss of a host session or failure to establish one. The 802.2 driver LANDD.OS2 can now be configured to ignore these indicators.

Copy LANDD.OS2 and LANDD.NIF to the IBMCOM\PROTOCOL or comparable LAPS or MPTS directory (make sure you remove, and not just rename, the old versions of these files in IBMCOM\PROTOCOL). To configure the driver to ignore the indicators run LAPS or MPTS, select Configure, then Edit the 802.2 driver to set NetFlags to 1000. Alternately, you can manually add the following line manually to the LANDD_nif section of PROTOCOL.INI:

NETFLAGS = 0x1000

The replacement driver can be used with any version of LAPS or MPTS at the level distributed with LAN Server 3.0 or higher. It will be included in future LAPS/MPTS CSD packages and in OS2 Warp Server.

Please note: The LANDD.NIF file is in English. Non-English versions of this file that have been updated with the NETFLAGS parameter are not yet available. This change will be included in a future CSD of LAPS/MPTS. If you cannot use the LANDD.NIF included in this package you will have to add the NETFLAGS line to PROTOCOL.INI manually, as it will not appear as a configuration option in LAPS/MPTS. Please also be aware that if you later run LAPS/MPTS in the future to make a configuration change, the NETFLAGS line will be removed from the LANDD section of PROTOCOL.INI -- since it is not in the NIF file LAPS/MPTS believes it is invalid and should be removed.

OS/2 LAN Requester (NETBEUI.OS2)

LAN Requester will enter a polling state when it believes the target station is not receiving frames, causing additional network traffic and possibly leading to loss of connection to the server. Simply copy NETBEUI.OS2 to the IBMCOM\PROTOCOL or comparable LAPS or MPTS directory to resolve this. The replacement driver can be used with any version of LAPS or MPTS at the level distributed with LAN Server 3.0 or higher. It will be included in future LAPS/MPTS CSD packages and in OS2 Warp Server.

TCP/IP for OS/2

The condition does not cause any traffic or connection failures for TCP/IP for OS/2 but results in an error counter being incorrectly incremented (reported by NETSTAT -- number of packets dropped). These have been reported in APARs PN77766 for version 2.0 and IC11795 for version 3.0 shipped with OS/2 Warp Connect. A future CSD for these packages will include a fix not to increment this counter for this condition.

DOS LAN Services (NET.EXE)

A protocol-layer driver loaded by NET.EXE can experience problems if the indicators are not set, which can result in loss of connection to the server or failure to connect to the server at all. The replacement driver can be used with any level of DLS 4.x.

LAN Support Program (Version 1.38P, full release)

The DXMCOMOD.SYS, DXMC1MOD.SYS, and DXMJOMOD.SYS drivers of LAN Support Program Version 1.38 have been patched to handle the condition gracefully; other drivers in the package such as DXMEOMOD.SYS and DXMTOMOD.SYS are unaffected. If you are using the full version of LAN Support Program, use the LOADDISKF.EXE program to create the LAN Support Program 1.38P diskette (3.5-inch 1.44M formatted) from the diskette image:

LOADDISKF LSP138P.DSK A:

LAN Support Program Custom (DXMCOMOD.SYS, DXMCSMOD.SYS)

This archive also includes replacement driver DXMCOMOD.SYS for LAN Support Custom shipped with some IBM Token-Ring adapters such as Auto 16/4 and ISA-16, and DXMCSMOD.SYS for the LAN Support Custom shipped with the Auto 16/4 Token-Ring Credit Card adapter.

=====

1/05/96

Initial release, version 1.00.

Includes the following updated drivers/products:

DXMCOMOD.SYS	11,648	1-05-96
DXMCSMOD.SYS	23,629	1-05-96
LSP138P.DSK	615,610	1-05-96
NETBEUI.OS2	114,676	1-05-96

1/25/96

Version 1.10. Expanded READ.ME to include more background and details on the problem. Added the following:

LANDD.NIF	10,603	1-25-96
LANDD.OS2	115,236	1-23-96
NET.EXE	352,999	12-06-95

3/25/96

Version 1.20. Fixed a problem in LANDD.NIF that could cause the NETFLAGS settings to be lost on subsequent LAPS/MPTS configuration.

LANDD.NIF	10,602	3-25-96
-----------	--------	---------

=====

Appendix B. Ether-Types and SAP Listings

The following listings contain all possible Ether-Types and SAP values that may be used to configure protocol-based VLANs in the RouteSwitch.

The SNAP values are the same as the Ether-type values preceeded by six zeroes. The SNAP value for Novell NetWare would then be 0000008137.

B.1 Ether-Type Listing

Value	Description
0000-05DC	IEEE 802.3 Length Fields
0101-01FF	Experimental (for development) -- Conflicts with 802.3 Length Fields
0200	Xerox PUP -- Conflicts with 802.3 Length Field
0201	PUP Address Translation -- Conflicts with 802.3 Length Fields
0600	Xerox XNS IDP
0800	DOD IP
0801	X.75 Internet
0802	NBS Internet
0803	ECMA Internet
0804	CHAOSnet
0805	X.25 Level 3
0806	ARP (for IP and CHAOS)
0807	Xerox XNS Compatibility
081C	Symbolics Private
0888-088A	Xyplex
0900	Ungermann-Bass network debugger
0A00	Xerox 802.3 PUP
0A01	Xerox 802.3 PUP Address Translation
0A02	Xerox PUP CAL Protocol (unused)
0BAD	Banyan Systems, Inc.
1000	Berkeley Trailer negotiation
1001-100F	Berkeley Trailer encapsulation for IP
1066	VALIS Systems
1600	VALID Systems
3C01-3C0D	3Com Corporation
3C10-3C14	3Com Corporation
4242	PCS Basic Block Protocol
5208	BBN Simnet Private
6000	DEC Unassigned
6001	DEC MOP Dump/Load Assistance
6002	DEC MOP Remote Console
6003	DEC DECnet Phase IV
6004	DEC LAT
6005	DEC DECnet Diagnostic Protocol: DECnet Customer Use
6007	DEC DECnet LAVC
6008	DEC Amber
6009	DEC MUMPS
6010-6014	3Com Corporation
7000	Ungermann-Bass download
7001	Ungermann-Bass NIU

7002	Ungermann-Bass diagnostic/loopback
7007	OS/9 Microware
7020-7028	LRT (England)
7030	Proteon
7034	Cabletron
8003	Cronus VLN
8004	Cronus Direct
8005	HP Probe protocol
8006	Nestar
8008	AT&T
8010	Excelan
8013	SGI diagnostic type (obsolete)
8014	SGI network games (obsolete)
8015	SGI reserved type (obsolete)
8016	SGI "bounce server" (obsolete)
8019	Apollo
802E	Tymshare
802F	Tigan, Inc.
8035	Reverse ARP (RARP)
8036	Aeonic Systems
8038	DEC LANBridge
8039	DEC DSM
803A	DEC Aragon
803B	DEC VAXELN
803C	DEC NSMV
803D	DEC Ethernet CSMA/CD Encryption Protocol
803E	DEC DNA
803F	DEC LAN Traffic Monitor
8040	DEC NetBIOS
8041	DEC MS/DOS
8042	DEC Unassigned
8044	Planning Research Corporation
8046	AT&T
8047	AT&T
8049	ExperData (France)
805B	VMTP (Versatile Message Transaction Protocol, RFC-1045, Stanford)
805C	Stanford V Kernel production, Version 6.0
805D	Evans & Sutherland
8060	Little Machines
8062	Counterpoint Computers
8065	University of Massachusetts, Amherst
8066	University of Massachusetts, Amherst
8067	Veeco Integrated Automation
8068	General Dynamics
8069	AT&T
806A	Autophon (Switzerland)
806C	ComDesign
806D	Compugraphic Corporation
806E-8077	Landmark Graphics Corporation
807A	Matra (France)
807B	Dansk Data Elektronik A/S (Denmark)
807C	Merit Intermodal
807D	VitaLink Communications
807E	VitaLink Communications
807F	VitaLink Communications
8080	VitaLink Communications bridge
8081	Counterpoint Computers

8082	Counterpoint Computers
8083	Counterpoint Computers
8088	Xyplex
8089	Xyplex
808A	Xyplex
809B	AppleTalk and Kinetics Appletalk over Ethernet
809C	Datability
809D	Datability
809E	Datability
809F	Spider Systems, Ltd. (England)
80A3	Nixdorf Computer (West Germany)
80A4-80B3	Siemens Gammasonics, Inc.
80C0	Digital Communication Associates
80C1	Digital Communication Associates
80C2	Digital Communication Associates
80C3	Digital Communication Associates
80C6	Pacer Software
80C7	Applitek Corporation
80C8-80CC	Integraph Corporation
80CD	Harris Corporation
80CE	Harris Corporation
80CF-80D2	Taylor Inst.
80D3	Rosemount Corporation
80D4	Rosemount Corporation
80D5	IBM SNA Services over Ethernet
80DD	Varian Associates
80DE	Integrated Solutions TRFS
80DF	Integrated Solutions
80E0-80E3	Allen-Bradley
80E4-80F0	Datability
80F2	Retix
80F3	Kinetics, AppleTalk ARP (AARP)
80F4	Kinetics
80F5	Kinetics
80F7	Apollo Computer
80FF-8103	Wellfleet Communications
8107	Symbolics Private
8108	Symbolics Private
8109	Symbolics Private
8130	Waterloo Microsystems
8131	VG Laboratory Systems
8137	Novell (old) NetWare IPX
8138	Novell
8139-813D	KTI
9000	Loopback (Conifguration Test Protocol)
9001	Bridge Communications XNS Systems Management
9002	Bridge Communications TCP/IP Systems Management
9003	Bridge Communications
FF00	BBN VITAL LANBridge cache wakeup

B.2 SAP Listing

Address	Assignment
00	Null LSAP
02	Individual LLC Sublayer Management Function
03	Group LLC Sublayer Management Function
04	IBM SNA Path Control (individual)
05	IBM SNA Path Control (group)
06	ARPANET Internet Protocol (IP)
08	SNA
0C	SNA
0E	PROWAY (IEC955) Network Management & Initialization
10	Novell and SDLC Link Servers
18	Texas Instruments
20	CLNP ISO OSI
34	CLNP ISO OSI
42	IEEE 802.1 Bridge Spanning Tree Protocol
4E	EIA RS-511 Manufacturing Message Service
7E	ISO 8208 (X.25 over IEEE 802.2 Type 2 LLC)
80	Xerox Network Systems (XNS)
86	Nestar
8E	PROWAY (IEC 955) Active Station List Maintenance
98	ARPANET Address Resolution Protocol (ARP)
AA	SNAP
BC	Banyan VINES
AA	SubNetwork Access Protocol (SNAP)
E0	Novell NetWare
EC	CLNP ISO OSI
F0	IBM NetBIOS
F4	IBM LAN Management (individual)
F5	IBM LAN Management (group)
F8	Remote Program Load
FA	Ungermann-Bass
FC	IBM Remote Program Load
FE	ISO Network Layer Protocol
FF	Global LSAP

Appendix C. Sample GATED Configuration File

This gated file is a sample taken from an RS/6000. It shows the various settings that are available in a UNIX environment. The RouteSwitch architecture and gated function is based on UNIX. This file is not meant to be a full example but explains some of the functions used in a gated configuration file.

```
# @(#)61      1.6  src/tcpip/etc/gated.conf, tcpip, tcpip411, GOLD410
#            12/6/93 14:23:11
#
# COMPONENT_NAME: TCPIP gated.conf
#
# FUNCTIONS:
#
# ORIGINS: 27
#
# (C) COPYRIGHT International Business Machines Corp. 1985, 1989
# All Rights Reserved
# Licensed Materials - Property of IBM
#
# US Government Users Restricted Rights - Use, duplication or
# disclosure restricted by GSA ADP Schedule Contract with IBM Corp.
#
#
# gated configuration file
#
#####
#
# How to configure this file for your system:
#
# Statement classes.
#   There are eight classes of statements. The first two may
#   be specified in the configuration file in any order.
#
#       directives
#           These statements are immediately acted upon by the
#           parser. They are used to specify included files and
#           the directory in which the reside. Unlike other
#           statements which terminate with a semi-colon (;),
#           directive statements terminate with a newline
#           character.
#
#       trace
#           These statements control tracing options.
#
#   The six remaining classes must be specified in order:
#
#       options
#           These statements allow specification of some global
#           options.
#
#       interface
#           These statements specify interface options.
#
#       definition
#           These statements options, the autonomous system and
#           martian networks.
#
#       protocol
#           These statements enable or disable protocols and set
#           protocol options.
#
#       route
```

```

#           These statements define static routes.
#           control
#           These statements define routes that are imported from
#           routing peers and routes that are exported to these
#           peers.
#
#           Statements within a class may be listed in any order.
#
#####
#
# 1) Directive Statements
#     Set optional directive statements.
#
#     %directory "<path_name>"
#
#     Sets the current directory to <path_name>. This is the
#     path that gated uses to look for included files that do
#     not begin with "/".
#
#     %include "filename"
#
#     Causes the specified file to be parsed completely before resuming
#     this file. Nesting up to 10 levels is supported.
#
#####
#
# 2) Trace Statements
#     Set optional trace statements.
#
#     tracefile ["filename" [replace]]
#     [size <size> [k|m] files <files>] ;
#
#     Specifies the file to contain tracing output. If a filename is
#     specified, trace information is appended to this file unless "replace"
#     specified.
#     If specified, <size> and <files> cause the trace file to be limited
#     to <size>, with <files> files kept (including the active file). The
#     backup file names are created by appending a period and a number to
#     the trace file name, starting with ".0". The minimum size that can
#     be specified is 10k, the minimum number of files that can be specified
#     is 2. The default is not to rotate log files.
#
#     traceoptions <traceoption> [<traceoption> [ ... ]]
#     [except <traceoption>
#     [<traceoption> [ ... ]]];
#
#     Where "traceoption" is:
#
#           all, general, internal, external, nostamp, kernel, mark, task,
#           timer, parse, route, kernel, bgp, egp, rip, hello, icmp, snmp,
#           protocol, or update.
#
#####
#
# 3) Options Statements
#     options <option_list> ;
#           Sets gated options:
#
#           noinstall Do not change kernel's routing table.

```

```

# Useful for verifying configuration
# files.
#
# gendefault
# BGP and EGP neighbors should cause the
# internal generation of a default route
# when up. This route will not be
# installed in the kernel's routing table,
# but may be announced by other protocols.
# Announcement is controlled by
# referencing the special protocol
# "default".
#
# nosend Do not send any packets. This allows
# running gated on a live network to test
# protocol interactions without actually
# participating in the routing protocols.
# The packet traces in the gated log can
# be examined to verify that gated is
# functioning properly. This is most
# useful for RIP and HELLO and possibly
# the SMUX SNMP interface. This does not
# yet apply to BGP packets.
#
# noresolv Do not try to resolv symbolic names into
# IP addresses by using the host/network
# tables or Domain Name System. This is
# intended for systems where a lack of
# routing information could cause a DNS
# lookup to hang.
#
# syslog Controls the amount of data gated logs
# via syslog on systems where setlogmask()
# is supported. The log_levels and other
# terminology are as defined in the
# setlogmask() man page. The default is
# equivalent to "syslog upto info".
#
#####
#
# 4) Interface Statements
# interfaces {
#     options [strictifs] [scaninterval <time>] ;
#     interface <interface_list> <interface_options> ;
#     define <address> [broadcast <broadaddr>|pointopoint
#         <lcladdr>] [netmask <netmask>]
#         [multicast] ;
# } ;
#
# options Sets some global options related to interfaces.
#
# Options are:
#
# strictifs Indicates that it is a fatal error to
# reference an interface in the
# configuration file that is not listed in
# a define statement or not present when
# gated is started. Without this option a
# warning message will be issued and gated

```

```

#                               will continue.
#
#
#       scaninterval <time>
#           Specifies how often gated scans the
#           kernel interface list for changes. The
#           default is every 15 seconds on most
#           systems, 60 seconds on systems that pass
#           interface status changes through the
#           routing socket (i.e. BSD 4.4). Note
#           that gated will also scan the interface
#           list on receipt of a SIGUSR2.
#
#       define    Defines interfaces that may not be present when
#                 gated is started. Gated considers it an error to
#                 reference a non-existent interface in the config
#                 file. This clause allows specification of that
#                 interface so it can be referenced in the config
#                 file.
#
#                 Definition keywords are:
#
#       broadcast <broad_addr>
#           Defines the interface as broadcast
#           capable (i.e. Ethernet and Token Ring)
#           and specifies the broadcast address.
#
#       pointopoint <local_addr>
#           Defines the interface as a point to
#           point interface (i.e. SLIP and PPP) and
#           specifies the address on the local side.
#           For this type of interface the
#           <interface_addr> specifies the address
#           of the remote host.
#
#       An interface not defined as broadcast or
#       pointopoint is assumed to be non-broadcast
#       multiaccess (NBMA), such as an X.25 network.
#
#       netmask <subnetmask>
#           Specifies the non-standard subnet mask
#           to be used on this interface. Note that
#           this currently ignored on pointopoint
#           interfaces.
#
#       multicast Specifies the interface is multicast
#           capable.
#
#       interface Sets interface options on the specified
#           interfaces. An interface list is "all" or a list
#           of interface names (see warning about interface
#           names), domain names, or numeric addresses.
#
#
#       Options are:
#
#       preference <pref>
#           Sets the preference for routes to this
#           interface when it is up, defaults to 0.
#

```

```

#
# down preference <pref>
#       Sets the preference for routes to this
#       interface when gated believes it to be down
#       due to lack of received routing information,
#       defaults to 120.
#
# passive
#       Prevents gated from changing the preference
#       of the route to this interface if it is
#       believed to be down due to lack of received
#       routing information.
#
# simplex
#       Defines an interface as unable to hear it's
#       own broadcast packets. Currently defining an
#       interface as simplex is functionally
#       equivalent to defining it as passive.
#
# reject
#       Specifies that the address loopback
#       interfaces which match these criteria will be
#       used as the local address when installing
#       reject routes in the kernel. Should only
#       used with systems based on BSD 4.3 Tahoe or
#       earlier which have installed a
#       reject/blackhole pseudo interface.
#
# blackhole
#       Specifies that the address loopback
#       interfaces which match these criteria will be
#       used as the local address when installing
#       blackhole routes in the kernel. Should only
#       used with systems based on BSD 4.3 Tahoe or
#       earlier which have installed a
#       reject/blackhole pseudo interface.
#
#####
#
# 5) Definition Statements
#       autonomoussystem <autonomous system> ;
#       Sets the autonomous system of this router to be
#       <autonomous system>. This option is required if
#       BGP or EGP are in use.
#
#       routerid <host> ;
#       Sets the router identifier for use by the BGP and
#       OSPF protocols. The default is the address of the
#       first interface encountered by gated. The address
#       of a non-POINTOPOINT interface is preferred over
#       the local address of a POINTOPOINT interface and
#       an address on a loopback interface that is not the
#       loopback address (127.0.0.1) is most preferred.
#
#       martians {
#           <martian_list>
#       } ;
#
#       Defines a list of martian addresses about which

```

```

#           all routing information is ignored. The
#           <martian_list> is a semi-colon separated list of
#           symbolic or numeric hosts with optional masks.
#           See dest_mask. Also, the allow parameter may be
#           specified to explicitly allow a subset of a range
#           that was disallowed.
#
#####
#
# 7) Protocol Statements
#       Enables or disables use of a protocol and controls protocol
#       options. These may be specified in any order.
#
#       For all protocols, "preference" controls the choice of
#       routes learned via this protocol or from this autonomous
#       system in relation to routes learned from other
#       protocols/autonomous systems. The default metric used when
#       propagating routes learned from other protocols is specified
#       with "defaultmetric" which itself defaults to the highest
#
#       valid metric for this protocol, for many protocols this
#       signifies a lack of reachability.
#
#       For distance vector IGPs with no explicit connections or
#       authentication (RIP and HELLO) and redirects (ICMP), the
#       "trustedgateways" clause supplies a list of gateways
#       providing valid routing information; routing packets from
#       other gateways are ignored. This defaults to all gateways
#       on the attached networks.
#
#       Routing packets may be sent not only to the remote end of
#       point-to-point links and the broadcast address of
#       broadcast-capable interfaces, but also to specific gateways
#       if they are listed in a "sourcegateways" clause and "yes" or
#       "on" is specified. If "nobroadcast" is specified, routing
#       updates will be sent only to gateways listed in the
#       "sourcegateways" clause, and not at all to the broadcast
#       address. Disabling the transmission and reception of
#       routing packets for a particular protocol may be specified
#       with the "interface" clause. An "interface" clause which
#       disables sending or receiving protocol packets may be
#       overridden for specific peers using the "trustedgateways"
#       and "sourcegateways" clauses.
#
#       For exterior protocols (BGP, EGP), the autonomous system
#       advertised to the peer is specified by the global
#       "autonomoussystem" clause unless overridden by the "asout"
#       parameter. The incoming autonomous system number is not
#       verified unless "peeras" is specified. Specifying
#       "metricout" fixes the outgoing metric for all routes
#       propagated to this peer. If the peer does not share a
#       network, "interface" can be used to specify which interface
#       address to use when communicating with this peer and
#       "gateway" can be used to specify the next hop to use for all
#       routes learned from this peer. An internal default is
#       generated when routing information is learned from a peer
#       unless the "nogendefault" parameter is specified.
#
#       Any protocol can have a "traceoptions" clause, which enables

```



```

#      tracing for a particular protocol, group or peer.  The
#      allowable protocol-specific options are: all, general,
#      internal, external, route, update, task, timer, protocol, or
#      kernel.
#
#      rip yes|no|on|off [ {
#          broadcast ;
#          nobroadcast ;
#          nocheckzero;
#          preference <preference> ;
#          defaultmetric <metric> ;
#          interface <interface_list> [noripin]
#          [noripout]
#          [metricin <metric>] [metricout <metric>]
#
#          [version 1] | [version 2
#              [multicast|broadcast]];
#          ...
#          trustedgateways <gateway_list> ;
#          sourcegateways <gateway_list> ;
#          traceoptions <traceoptions> ;
#      } ] ;
#
#      If the rip clause is not specified the default is "rip
#      on".  "Nobroadcast" specifies that RIP packets will
#      only be sent to gateways listed in the "sourcegateways"
#      clause, if there are any.  If "yes" or "on" is
#      specified, RIP will assume "nobroadcast" if there is
#      only one interface and "broadcast" if there is more
#      than one.  "Broadcast" specifies that RIP packets will
#      always be generated.  "Nocheckzero" specifies that RIP
#      should not make sure that the reserved fields in RIP
#      packets are zero.
#
#      Note that using "broadcast" with only one
#      interface is useful only when propagating static
#      routes or routes learned from another protocol.
#      This will cause data packets to travel across the
#      same network twice, which may be tolerable in
#      certain configurations.
#
#      The default metricout is zero, the default metricin is
#      the kernel interface metric plus 1 (the default RIP hop
#      count).
#
#      If the version is specified as or defaults to 1, RIP
#      version 2 packets will never be sent except in response
#      to a v2 POLL packet.  If the version is specified as 2,
#      RIP version 2 packets will be sent to the RIP multicast
#      address if possible, or to the broadcast addresss,
#      unless the method is explicitly specified.
#
#      The default metric is 16; the default preference is
#      100.
#
#      hello yes|no|on|off [ {
#          broadcast ;
#          nobroadcast ;
#          preference <preference> ;

```

```

#         defaultmetric <metric> ;
#         interface <interface_list> [nohelloin]
#             [nohelloout]
#             [metricin <metric>]
#             [metricout <metric>];
#         ...
#         trustedgateways <gateway_list> ;
#         sourcegateways <gateway_list> ;
#         traceoptions <traceoptions> ;
#     } ] ;
#
#     If "yes" or "on" is specified, HELLO will assume
#     "nobroadcast" if there is only one interface and
#     "broadcast" if there is more than one. If the HELLO
#     clause is not specified the default is "hello off".
#     "Broadcast" specifies that HELLO packets will be
#     generated. "Nobroadcast" specifies that HELLO packets
#     will only be sent to gateways listed in the
#     "sourcegateways" clause, if there are any.
#
#     Note that using "broadcast" with only one
#     interface is useful only when propagating static
#     routes or routes learned from another protocol.
#     This will cause data packets to travel across the
#     same network twice, which may be tolerable in
#     certain configurations.
#
#     The default metricout is zero, the default metricin is
#     a translation of the kernel interface metric into a
#     hello metric plus 100 (the default HELLO hop count).
#
#     The default metric is 30000; the default preference is
#     90.
#
ospf yes|no|on|off [ {
#     [ defaults {
#         preference <preference> ;
#         cost <cost> ;
#         tag [<tag> | as [<as_tag>]] ;
#         type <1|2> ;
#     } ] ;
#     [exportlimit <routes> ;]
#     [exportinterval <time> ;]
#     [traceoptions <traceoptions> ;]
#     [monitorauthkey <authkey> ;]
#     [area <area> {
#         authtype <0|1|none|simple> ;
#         stub [cost <cost>];
#         networks {
#             network [mask <mask>] ;
#         } ;
#         stubhosts {
#             <host> cost <cost> ;
#         } ;
#         interface <interface> [cost <cost>] {
#             [enable|disable] ;
#             retransmitinterval <time> ;
#             transitdelay <time> ;
#             priority <priority> ;

```

```

#             hellointerval <time> ;
#             routerdeadinterval <time> ;
#             authkey <auth_key> ;
#         } ;
#     interface <interface> nonbroadcast
#         [cost|<cost>]{
#             pollinterval <time> ;
#             routers {
#                 <gateway> [eligible] ;
#                 ...
#             } ;
#             [enable|disable] ;
#             retransmitinterval <time> ;
#             transitdelay <time> ;
#             priority <priority> ;
#             hellointerval <time> ;
#             routerdeadinterval <time> ;
#             authkey <auth_key> ;
#         } ;
#     } ; ]
# [ backbone {
#     authtype <0|1|none|simple> ;
#     networks {
#         network [mask <mask>] ;
#     } ;
#     subhosts {
#         <host> cost <cost> ;
#     } ;
#     interface <interface> [cost <cost>] {
#         [enable|disable] ;
#         retransmitinterval <time> ;
#         transitdelay <time> ;
#         priority <priority> ;
#         hellointerval <time> ;
#         routerdeadinterval <time> ;
#         authkey <auth_key> ;
#     } ;
#     . . .
#     interface <interface> nonbroadcast
#         [cost <cost>]{
#             pollinterval <time> ;
#             routers {
#                 <gateway> [eligible] ;
#                 ...
#             } ;
#             [enable|disable] ;
#             retransmitinterval <time> ;
#             transitdelay <time> ;
#             priority <priority> ;
#             hellointerval <time> ;
#             routerdeadinterval <time> ;
#             authkey <auth_key> ;
#         } ;
#     . . .
#     virtuallink neighborid <host> transitarea <area> {
#         [enable|disable] ;
#         retransmitinterval <time> ;
#         transitdelay <time> ;
#         priority <priority> ;

```

```

#             hellointerval <time> ;
#             routerdeadinterval <time> ;
#             authkey <auth_key> ;
#         } ;
#         . . .
#     } ; ]
# } ] ;
#
# interface
#     An interface is specified with an address, a name,
#     a wildcard name (name without any number), or
#     "all". Multiple interface clauses may be
#     specified with different parameters, the
#     parameters used are accumulated from the interface
#     clauses. If a parameter is specified more than
#     once the instance with the most specific interface
#     reference is used. The order of precedence is
#     address, name, wildcard name, "all".
#
# cost A number between 0 and 65535 specifying an OSPF
#     internal cost.
#
# tag The OSPF tag (an unsigned 31-bit number) to be
#     placed on all routes exported by gated into OSPF.
#
# as_tag
#     The OSPF-BGP tag (an unsigned 12-bit number) to be
#     placed on all routes export by gated into OSPF.
#     When "tag as [<as_tag>]" is used, tag fields are
#     automatically generated and the as_tag field is
#     assigned if specified.
#
# metric
#     A number between 0 and 16777215 specifying an OSPF
#     external (ASE) cost.
#
# area A dotted quad or a number between 1 and
#     4294967295. Area 0 is always referred to as the
#     "backbone".
#
# auth_key
#     One to eight decimal digits separated by periods,
#     a one to eight byte hexadecimal string preceded by
#     "0x", or a one to eight character string in double
#     quotes.
#
# priority
#     A number between 0 and 255 specifying the priority
#     of becoming the designated router on this
#     interface.
#
# OSPF inter and intra area are always imported into the
#     gated routing table with a preference of 10. It would
#     be a violation of the protocol to do otherwise so it is
#     not possible to override this. OSPF Autonomous System
#     External (ASE) routes are imported with a preference of
#     150. This default may be changed with the preference
#     keyword in the defaults section. ASE routes are
#     imported at a rate of 100 ASEs every 1 second, these

```

```

# parameters can be tuned with the "exportlimit" and
# "exportinterval" parameters.
#
# Gated routes are exported to OSPF as ASEs with a
# default cost of 0 and a type of 1. By default, the tag
# is calculated from the AS path of the route being
# exported (tag as). These may all be changed in the
# defaults section.
#
# OSPF areas may be specified in any order, but the
# "backbone" area must be specified last.
#
# Reconfiguration (SIGHUP) is currently disabled when
# OSPF is enabled. This will hopefully be fixed in a
# future release.
#
# egp yes|no|on|off [ {
#     [preference <preference> ;]
#     [defaultmetric <metric> ;]
#     [packetsize <maxpacketsize> ;]
#     [traceoptions <traceoptions> ;]
#     [group          [peeras <autonomous system>]
#                     [localas <autonomous system>]
#                     [maxup <number>]
#                     [preference <preference>]
#     {
#         neighbor <host>
#             [metricout <metric>]
#             [nogendefault]
#             [importdefault]
#             [exportdefault]
#             [gateway <gateway>]
#             [lcladdr <local_address>]
#             [sourcenet <network>]
#             [minhello <min_hello>]
#             [minpoll <min_poll>]
#             [traceoptions <traceoptions>] {
#             ;
#         ...
#     } ;
#     ...]
# } ] ;
#
# "Packetsize" specifies the size, in bytes, of the
# largest EGP packet that will be accepted or sent. A
# "group" lists a group of EGP peers in one autonomous
# system. "Maxup" specifies the maximum number of peers
# that will be maintained in the Up state.
# "Importdefault" and "exportdefault" tell gated to
# import or export the default route (0.0.0.0) in updates
# exchanged with an EGP neighbor. If not specified, the
# the default network is ignored when exchanging EGP
# updates. "Sourcenet" specifies the network to query in
# EGP Poll packets, this is normally the shared network.
# The minimum EGP hello and poll intervals acceptable may
# be specified with the "minhello" and "minpoll"
# arguments, respectively. These are both specified as a
# time in seconds, minutes:seconds or
# hours:minutes:seconds. Any number of "group" clauses

```

```

# may be specified containing any number of "neighbor"
# clauses. Any parameters from the "neighbor" clause may
# be specified on the "group" clause to provide defaults
# for the group.
#
# The "local_address" is used to set the address the
# local address to be used when there is a choice of
# interfaces. If not specified it defaults to whichever
# interface is shared with the neighbor. If a network is
# not shared with the neighbor, "gateway" may be used to
# specify the next-hop gateway to use when installing
# routes learned from this neighbor. In this case the
# default interface is the one shared with the specified
# gateway.
#
# The default metric is 255; the default preference is
# 200.
#
# bgp yes|no|on|off [ {
#     [preference <preference> ;]
#     [defaultmetric <metric> ;]
#     [traceoptions <traceoptions> ;]
#     [group type external|internal|igp|test peeras <peeras>
#         [metricout <metric>]
#         [localas <localas>]
#         [nogendefault]
#         [gateway <gateway>]
#         [preference <preference>]
#         [lcladdr <local_address>]
#         [holdtime <time>]
#         [traceoptions <traceoptions>]
#         [version <version>]
#         [passive]
#         [importdefault]
#         [exportdefault]
#         [sendbuffer <bufsize>]
#         [recvbuffer <bufsize>]
#         [spoolbuffer <bufsize>]
#         [keepall]
#         {
#             [allow { dest_mask ... } ;]
#             [peer <host>
#                 [metricout <metric>]
#                 [localas <localas>]
#                 [nogendefault]
#                 [gateway <gateway>]
#                 [preference <preference>]
#                 [lcladdr <local_address>]
#                 [holdtime <time>]
#                 [traceoptions <traceoptions>]
#                 [version <version>]
#                 [passive]
#                 [importdefault]
#                 [exportdefault]
#                 [sendbuffer <bufsize>]
#                 [recvbuffer <bufsize>]
#                 [spoolbuffer <bufsize>]
#                 [keepall]
#             ]
#         }
#     ;]

```

```

#
#           } ;
#           ....]
#       } ] ;
#
#       BGP peers are assigned to groups based on the type and
#       peeras, it is not possible to have two groups with the
#       same type and peeras. Peer specifies the address of
#       each BGP peer. Group options provide the defaults for
#       all peers within that group.
#
#       "Peeras" is the autonomous system expected from a peer.
#       "Metricout" is the default metric to use when sending
#       to this peer. "Localas" specifies the autonomous
#       system advertised to this peer, the default is that
#       which has been set globally. "Nogendefault" specifies
#       that this peer should not cause the automatic default
#       to be generated.
#
#       The "local_address" specifies the address to be used on
#       the local end of the TCP connection with the peer. For
#       "external" peers the local address must be on an
#       interface which is shared with the peer (or for a non-
#       local peer's configured next-hop gateway when the
#       "gateway" option is used to specify this) and a session
#       with the peer will be opened only when an interface
#       with the appropriate local address through which the
#       peer (gateway) address is directly reachable is
#       operating. For other types of peers a peer session
#       will be maintained when any interface with the
#       specified local address is operating. In either case
#       incoming connections will only be recognized as
#       matching a configured peer if they are addressed to the
#       configured local address.
#
#       "Holdtime" specifies the BGP holdtime to use with this
#       peer. Traceoptions specify tracing options for this
#       peer (and are not yet implemented). Version specifies
#       the version of the BGP protocol to use with this peer.
#       If not specified, the highest supported version is used
#       first and version negotiation is attempted. "Passive"
#       specifies that active opens to this peer should not be
#       attempted. "Importdefault" and "exportdefault" control
#       whether the default network (0.0.0.0) can be exchanged
#       with this peer. "Keepall" is used to retain routes
#       learned from a peer that contain one of our autonomous
#       system numbers in their path.
#
#       "Sendbuffer" and "Recvbuffer" control the amount of
#       buffering asked of the kernel, the default is to
#       configure the maximum supported, up to 65535 bytes.
#       "Spoolbuffer" is used to indicate that BGP should
#       buffer data for peers when the kernel queues are full,
#       the default is to break the connection. These options
#       are normally not needed on properly functioning
#       systems..
#
#       If a metric is not specified, the default is not to
#       send a metric. The default preference is 170, the

```

```

#           default holdtime is 180 and the default version is 3.
#
#   redirect yes|no|on|off [ {
#       preference <preference> ;
#       interface <interface_list> [noredirects] ;
#       trustedgateways <gateway_list> ;
#       traceoptions <traceoptions> ;
#   } ] ;
#
#       Controls whether gated makes routing table changes
#       based on ICMP redirects when not functioning as a
#       router. When functioning as a router (i.e. any
#       interior routing protocols (RIP, HELLO, OSPF) are
#       participating in routing on any interface, ICMP
#       redirects are disabled. When ICMP redirects are
#       disabled, gated must actively remove the effects of
#       redirects from the kernel as the kernel always
#       processes ICMP redirects.
#
#       The default preference is 30.
#
#   snmp yes|no|on|off [ {
#       preference <preference> ;
#       traceoptions <traceoptions> ;
#       port <port> ;
#   } ] ;
#
#       Controls whether gated tries to contact the SMUX SNMP
#       daemon to register supported variables. The default is
#       "on". The default preference is 50. The default port
#       is 199 (SMUX).
#
#####
# 8) Route/Static Statements
#       Static routes are specified with "static" clauses.
#       static {
#           <dest_mask> gateway <gateway> [<gateway2>
#               [<gateway3> [...]]]
#               [interface <interface_list>]
#               [preference <preference>]
#               [retain] [reject]
#               [blackhole] [noinstall] ;
#       ...
#       <dest_mask> interface <interface> [preference
#           <preference>] [retain]
#           [reject] [blackhole]
#           [noinstall] ;
#       ...
#       } ;
#
#       Any number of "static" statements may be specified,
#       each containing any number of static route definitions.
#       The first form defines a static route through one or
#       more gateways. If multiple gateways are specified,
#       they are limited by the number of multipath
#       destinations supported (on Unix this is almost always
#       one). Only gateways on interfaces that are configured
#       and up are used.
#

```



```

# The second defines a static interface route which is
# used for primitive support of multiple networks on one
# interface.
#
# The interface list on the first form restricts static
# routes to a specific set of interfaces.
#
# "Retain" causes the route to be retained in the kernel
# after gated is shut down. "Reject" causes all packets
# to this route to be rejected. "Blackhole" causes all
# packets to this route to be silently discarded.
# "Reject" and "blackhole" are not supported by all
# systems. "Noinstall" is used to prevent this route
# from being installed in the kernel
#
# The preference for static routes defaults to 60.
#
#####
#
# 9) Control Statements
# Importation of routes from routing protocol peers and
# exportation of routes to routing protocol peers are
# controlled by "import" and "export" clauses.
#
# import proto bgp|egp as <autonomous system> restrict ;
#
# import proto bgp|egp as <autonomous system>
# [preference <preference>] {
# <import_list>
# } ;
#
# import proto bgp aspath <aspath_spec> restrict ;
#
# import proto bgp aspath <aspath_spec>
# [preference <preference>] {
# <import_list>
# } ;
#
# import proto rip|hello|redirect restrict ;
#
# import proto rip|hello|redirect
# [preference <preference>] {
# <import_list>
# } ;
#
# import proto rip|hello|redirect interface
# <interface_list> restrict ;
#
# import proto rip|hello|redirect interface <interface_list>
# [preference <preference>] {
# <import_list>
# } ;
#
# import proto rip|hello|redirect gateway <gateway_list>
# restrict;
#
# import proto rip|hello|redirect gateway <gateway_list>
# [preference <preference>] {
# <import_list>
#

```

```

#           } ;
#
#           import proto ospfase [tag <ospf_tag>] restrict ;
#
#           import proto ospfase [tag <ospf_tag>]
#               [preference <preference>] [{
#                   <import_list>
#               }] ;
#
#           If an OSPF type is specified, only routes of that
#           type will be considered for import, otherwise
#           either type will be considered. If an ospf_tag
#           specification is given, only routes matching that
#           tag specification will be considered, otherwise
#           any tag will be considered. An OSPF tag
#           specification may be a decimal, hexadecimal or
#           dotted quad number.
#
#           If more than one import statement relevant to a
#           protocol is specified, they are processed most
#           specific to least specific (i.e. for RIP and
#           HELLO, gateway, interface and protocol), then in
#           the order specified in the config file.
#
#           import_list
#           An import_list consists of zero or more
#           destinations (with optional mask). One of two
#           parameters may be specified, "restrict" to prevent
#           a set of destinations from being imported or a
#           specific preference for this set of destinations.
#
#           <dest_mask> [[restrict] |
#               [preference <preference>]] ;
#
#           Note that the contents of an import_list are
#           sorted internally so that entries with the most
#           specific masks are examined first. The order in
#           which dest_mask entries are specified does not
#           matter.
#
#           If no import list is specified, all routes will be
#           accepted. If an import list is specified, the
#           import list is scanned for a match. If no match
#           is found, the route is discarded. Rephrased, a
#           "all restrict" entry is assumed in an import list.
#
#           export proto bgp|egp as <autonomous system> restrict ;
#
#           export proto bgp|egp as <autonomous system>
#               [metric <metric>] {
#                   <export_list>
#               } ;
#
#           export proto rip|hello restrict ;
#
#           export proto rip|hello [metric <metric>] {
#               <export_list>
#           } ;

```

```

#
# export proto rip|hello interface <interface_list> restrict ;
#
# export proto rip|hello interface <interface_list>
#     [metric <metric>] {
#     <export_list>
#     } ;
#
# export proto rip|hello gateway <gateway_list> restrict ;
#
# export proto rip|hello gateway <gateway_list>
#     [metric <metric>] {
#     <export_list>
#     } ;
#
# export proto ospfase [type 1|2]
#     [tag <ospf_tag>] restrict ;
#
# export proto ospfase [type 1|2]
#     [tag <ospf_tag>]
#     [cost <ospf_cost>] {
#     <export_list>
#     } ;
#
# export_list
#     The export list specifies exportation based on the
#     origin of a route to a destination:
#
#     proto bgp|egp as <autonomous system>
#         [restrict] | [metric <metric>]
#         [ {
#         <announce_list>
#         } ] ;
#
#     proto rip|hello|direct|static|default
#         [restrict] | [metric <metric>]
#         [ {
#         <announce_list>
#         } ] ;
#
#     proto rip|hello|direct|static|default interface
#         <interface_list>
#         [restrict] | [metric <metric>]
#         [ {
#         <announce_list>
#         } ] ;
#
#     proto rip|hello gateway <gateway_list>
#         [restrict] | [metric<metric>]
#         [ {
#         <announce_list>
#         } ] ;
#
#     proto ospf [restrict] | [metric
#         <metric>] [ {
#         <announce_list> ;
#         } ] ;
#
#     proto ospfase [restrict | metric <metric>]]

```

```

#           [ {
#             <announce_list> ;
#           } ] ;
#
# proto <proto> aspath <aspath_spec>
#           [restrict] | [metric <metric>]
#           [ {
#             <announce_list>
#           } ] ;
#
# proto <proto> tag <tag>
#           [restrict] | [metric <metric>]
#           [ {
#             <announce_list>
#           } ] ;
#
# If a tag is specified, only routes with that tag
# will be considered, otherwise any tag will be
# considered. An OSPF tag on an export statement
# may be a decimal, hexadecimal, or "AS" to generate
# a tag based on the AS path of route being
# announced. An OSPF tag on an export list is just
# an 31 bit number that is matched against the tag
# present (if any) on that route.
#
# If more than one export statement relevant to a
# protocol is specified, they are processed most
# specific to least specific (i.e. for RIP and
# HELLO, gateway, interface and protocol), then in
# the order specified in the config file.
#
# By default interface routes are exported to all
# protocols. RIP and HELLO also export their own
# routes. An export specification with just a
# restrict will prevent these defaults from being
# exported. Note that it is not possible to change
# the metric RIP and HELLO use for their own routes;
# any attempt to override it will be silently
# ignored.
#
# Any protocol may be specified for import lists
# referring to aspaths and tags. AS paths are most
# meaningful with BGP and OSPF ASE routes, but are
# generated for all routes. Tags are currently only
# meaningful for OSPF ASE routes.
#
# announce_list
# An announce_list consists of zero or more
# destinations (with optional mask). One of two
# parameters may be specified, "restrict" to prevent
# a set of destinations from being exported or a
# specific metric for this set of destinations.
#
# <dest_mask> [[restrict] |
#               [metric <metric>]] ;
#
# Note that the contents of an announce_list are sorted
# internally so that entries with the most specific masks

```

```

#         are examined first. The order in which dest_mask
#         entries are specified does not matter.
#
#         If no announce_list is specified, all destinations
#         are announced. If an announce list is specified,
#         an "all restrict" is assumed. Therefore, an empty
#         announce list is the equivalent of "all restrict".
#
#         Note that to announce routes which specify a next
#         hop of the loopback interface (i.e. static and
#         internally generated default routes) via RIP or
#         HELLO it is necessary to specify the metric at
#         some level in the propagate clause. Just setting
#         a default metric for RIP or HELLO is not
#         sufficient.
#
#         aspath_spec
#         An AS path specification is used to match one or
#         more AS paths.
#
#         aspath <regexp> origin [igp]egp[incomplete]any]
#
#         where the regexp is a regular expression over the
#         set of AS numbers as defined in RFC-1164 section 4.2.
#
#####
#
# Sample setups.
#
#####
#
# Simple RIP quiet
# This configuration runs RIP in quiet mode, it only listens to
# packets, no matter how many interfaces are configured.
# It traces all trace output to a file called /var/tmp/rip.quiet.trace.
# The tracefile options allow the creation of 4 files to rotate trace
# information. Each of these trace files will grow to about 50K bytes
# before rotating to the next trace file. The trace files will be called,
# /var/tmp/rip.quite.trace, /var/tmp/rip.quite.trace.0,
# /var/tmp/rip.quite.trace.1, and /var/tmp/rip.quite.trace.2.
#
# traceoptions all ;
#
# tracefile "/var/tmp/rip.quiet.trace" replace size 50k files 4 ;
#
# rip yes {
#     nobroadcast ;
# } ;
#
#####
#
# Simple RIP supplier
# Run as a RIP supplier, do not supply RIP packets to
# tr0, and only listen to RIP packets from 192.100.110.1.
#
# rip yes {
#     broadcast ;
#     interface tr0 noripout ;
#     trustedgateways 192.100.110.1 ;

```

```

# } ;
#
#####
#
# Simple EGP
#   This host is in autonomous system 283. Verify neighbor
#   192.35.82.100 is in autonomous system 145. Do not
#   generate a default route from EGP information learned.
#
# autonomoussystem 283;
#
# egp yes {
#   group peeras 145 {
#     neighbor 192.100.110.100 nogendefault ;
#   } ;
# } ;
#
#####
#
# Simple BGP
#   This host is in autonomous system 283. Verify peer
#   192.100.110.1 is in autonomous system 145.
#
# autonomoussystem 283;
#
# bgp yes {
#   group type
#     External peeras 145 {
#       peer 192.100.110.1 ;
#     } ;
# } ;
#
#####

```

Appendix D. Special Notices

This publication is intended to help customers, system engineers, service specialists and marketing specialists understand VLANs and the IBM Nways RouteSwitch architecture for planning, installation, and support purposes. The information in this publication is not intended as the specification of any programming interfaces that are provided by Nways RouteSwitches or the associated software. See the PUBLICATIONS section of the IBM Programming Announcement for IBM Nways RouteSwitch for more information about what publications are considered to be product documentation.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The information about non-IBM ("vendor") products in this manual has been supplied by the vendor and IBM assumes no responsibility for its accuracy or completeness. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any pointers in this publication to external Web sites are provided for convenience only and do not in any manner serve as an endorsement of these Web sites.

Any performance data contained in this document was determined in a controlled environment, and therefore, the results that may be obtained in other

operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.

Reference to PTF numbers that have not been released through the normal distribution process does not imply general availability. The purpose of including these reference numbers is to alert IBM customers to specific information relative to the implementation of the PTF when it becomes available to each customer according to the normal IBM PTF distribution process.

The following terms are trademarks of the International Business Machines Corporation in the United States and/or other countries:

IBM	Nways
NetView	AIX
AT	Current
OS/2	IBM

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc.

Java and HotJava are trademarks of Sun Microsystems, Incorporated.

PC Direct is a trademark of Ziff Communications Company and is used by IBM Corporation under license.

Pentium, MMX, ProShare, LANDesk and ActionMedia are trademarks or registered trademarks of Intel Corporation in the United States and other countries.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Microsoft, Windows, and the Windows 95 logo are trademarks or registered trademarks of Microsoft Corporation.

Other company, product and service names may be trademarks or service marks of others.

Appendix E. Related Publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

E.1 International Technical Support Organization Publications

For information on ordering these ITSO publications see "How to Get ITSO Redbooks" on page 369.

- *Understanding and Using the IBM MSS Server*, SG24-4915
- *IBM 8260 As a Campus ATM Switch*, SG24-5003
- *Internetworking over ATM: An Introduction*, SG24-4699

E.2 Redbooks on CD-ROMs

Redbooks are also available on CD-ROMs. **Order a subscription** and receive updates 2-4 times a year at significant savings.

CD-ROM Title	Subscription Number	Collection Kit Number
System/390 Redbooks Collection	SBOF-7201	SK2T-2177
Networking and Systems Management Redbooks Collection	SBOF-7370	SK2T-6022
Transaction Processing and Data Management Redbook	SBOF-7240	SK2T-8038
Lotus Redbooks Collection	SBOF-6899	SK2T-8039
Tivoli Redbooks Collection	SBOF-6898	SK2T-8044
AS/400 Redbooks Collection	SBOF-7270	SK2T-2849
RS/6000 Redbooks Collection (HTML, BkMgr)	SBOF-7230	SK2T-8040
RS/6000 Redbooks Collection (PostScript)	SBOF-7205	SK2T-8041
RS/6000 Redbooks Collection (PDF Format)	SBOF-8700	SK2T-8043
Application Development Redbooks Collection	SBOF-7290	SK2T-8037

E.3 Other Publications

These publications are also relevant as further information sources:

- *Nways RouteSwitch User's Guide*, GA27-4166
- *Nways RouteSwitch Network Manager Software User's Guide*, GC30-3871

How to Get ITSO Redbooks

This section explains how both customers and IBM employees can find out about ITSO redbooks, CD-ROMs, workshops, and residencies. A form for ordering books and CD-ROMs is also provided.

This information was current at the time of publication, but is continually subject to change. The latest information may be found at <http://www.redbooks.ibm.com/>.

How IBM Employees Can Get ITSO Redbooks

Employees may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **Redbooks Web Site on the World Wide Web**

<http://w3.itso.ibm.com/>

- **PUBORDER** — to order hardcopies in the United States

- **Tools Disks**

To get LIST3820s of redbooks, type one of the following commands:

```
TOOLCAT REDPRINT
TOOLS SENDTO EHONE4 TOOLS2 REDPRINT GET SG24xxxx PACKAGE
TOOLS SENDTO CANVM2 TOOLS REDPRINT GET SG24xxxx PACKAGE (Canadian users only)
```

To get BookManager BOOKs of redbooks, type the following command:

```
TOOLCAT REDBOOKS
```

To get lists of redbooks, type the following command:

```
TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET ITSOCAT TXT
```

To register for information on workshops, residencies, and redbooks, type the following command:

```
TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ITSOREGI 1998
```

- **REDBOOKS Category on INEWS**

- **Online** — send orders to: USIB6FPL at IBMMAIL or DKIBMBSH at IBMMAIL

Redpieces

For information so current it is still in the process of being written, look at "Redpieces" on the Redbooks Web Site (<http://www.redbooks.ibm.com/redpieces.html>). Redpieces are redbooks in progress; not all redbooks become redpieces, and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

How Customers Can Get ITSO Redbooks

Customers may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **Online Orders** — send orders to:

In United States:
In Canada:
Outside North America:

IBMMAIL
usib6fpl at ibmmail
caibmbkz at ibmmail
dkibmbsh at ibmmail

Internet
usib6fpl@ibmmail.com
lmannix@vnet.ibm.com
bookshop@dk.ibm.com

- **Telephone Orders**

United States (toll free)
Canada (toll free)

1-800-879-2755
1-800-IBM-4YOU

Outside North America
(+45) 4810-1320 - Danish
(+45) 4810-1420 - Dutch
(+45) 4810-1540 - English
(+45) 4810-1670 - Finnish
(+45) 4810-1220 - French

(long distance charges apply)
(+45) 4810-1020 - German
(+45) 4810-1620 - Italian
(+45) 4810-1270 - Norwegian
(+45) 4810-1120 - Spanish
(+45) 4810-1170 - Swedish

- **Mail Orders** — send orders to:

IBM Publications
Publications Customer Support
P.O. Box 29570
Raleigh, NC 27626-0570
USA

IBM Publications
144-4th Avenue, S.W.
Calgary, Alberta T2P 3N5
Canada

IBM Direct Services
Sortemosevej 21
DK-3450 Allerød
Denmark

- **Fax** — send orders to:

United States (toll free)
Canada
Outside North America

1-800-445-9269
1-403-267-4455
(+45) 48 14 2207 (long distance charge)

- **1-800-IBM-4FAX (United States) or USA International Access Code -408-256-5422 (Outside USA)** — ask for:

Index # 4421 Abstracts of new redbooks
Index # 4422 IBM redbooks
Index # 4420 Redbooks for last six months

- **On the World Wide Web**

Redbooks Web Site
IBM Direct Publications Catalog

<http://www.redbooks.ibm.com/>
<http://www.elink.ibm.link.ibm.com/pbl/pbl>

Redpieces

For information so current it is still in the process of being written, look at "Redpieces" on the Redbooks Web Site (<http://www.redbooks.ibm.com/redpieces.html>). Redpieces are redbooks in progress; not all redbooks become redpieces, and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

IBM Redbook Order Form

Please send me the following:

Title	Order Number	Quantity

First name	Last name
------------	-----------

Company

Address

City	Postal code	Country
------	-------------	---------

Telephone number	Telefax number	VAT number
------------------	----------------	------------

• Invoice to customer number _____

• Credit card number _____

Credit card expiration date	Card issued to	Signature
-----------------------------	----------------	-----------

We accept American Express, Diners, Eurocard, Master Card, and Visa. Payment by credit card not available in all countries. Signature mandatory for credit card payment.

List of Abbreviations

APA	all points addressable	MSB	most significant bit
ARE	all route explorer	MTU	maximum transmission unit
ATM	asynchronous transfer mode	NetBIOS	network basic input output system
BCM	broadcast manager	NIC	network interface card
BPDU	bridge protocol data unit	PDU	protocol data unit
BUS	broadcast and unknown server	PROFS	professional office system
DA	destination address	PTOP	point to point
DSAP	destination service access point	PVC	permanent virtual circuit
ELAN	emulated LAN	REM	ring error monitor
FDDI	fiber distributed data interface	RIF	routing information field
IBM	International Business Machines Corporation	RII	routing information indicator
ILMI	interim local management interface	SA	source address
IPX	internetwork packet exchange	SAP	service access point
ITSO	International Technical Support Organization	SNA	systems network architecture
LAN	local area network	SNAP	sub-network access protocol
LANE	LAN emulation	SNMP	simple network management protocol
LEC	lan emulation client	SPX	sequenced packet exchange
LECS	lan emulation configuration server	SRB	source route bridge
LED	light emitting diodes	SRF	specifically routed frame
LES	lan emulation server	SR-TB	source route to transparent bridge
LLC	logical link control	SSAP	source service access point STE
LSB	least significant bit		spanning tree explorer
LUNI	LAN emulation user to network interface	SVC	switched virtual circuit
MAC	medium access control	TB	transparent bridge
MPM	management processor module	TCN	topology change notification
		UNI	user to network interface
		VLAN	virtual LAN

Index

Numerics

802.10 185, 186
802.1d 110
802.3 based formats 159
8210 284
8260 284
8274 LAN RouteSwitch 2
8274 Model Descriptions 3

A

abbreviations 373
acronyms 373
Active 129
Address Registration 196
address resolution 194
Address Resolution Protocol 156
address resolution requests 196
addvp command 162
all route explorer 95
any-to-any switching 13, 150, 157, 159
APPLETALK 75, 154
ARE 95
ARP 156
ARP bit swapping 158
ATM 273
 addresses 191
 cells 189
 LAN Emulation 189, 284
 selector 302
 trunking 123
autocaps 161
AutoEncaps 161
AutoTracker 71

B

backplane 13
backup file 21
Banyan Vines 154, 160
Basic Console Setup 14
basic information 16
Basic Setup 12
bibliography 367
binding based VLAN 78
binding rules 40
bit swap 156
blocking 108
BPDU 59, 98, 100, 186
Bridge Aging Timer 105
bridge management 23
bridge priority 105, 112
bridging 89

Broadcast and Unknown Server 189, 194, 197
broadcast domain 12
broadcasts 37, 52, 194
Burst Sizes 261
 committed 261
 excess 261
 Measurement Interval 261
BUS 189, 192, 194, 197

C

CAM 53
canonical 156
cas command 171
cell switching 7
Cell Switching Modules (CSM) 2
Checksum field 154
checksums 160
CIP 69
CIR 261
collision domain 38
community name 14
Configuration Fallback 25
configuration files 20
configuration settings 20
configuring VLANs 67
connection management 189
Console 14
console port 28
Content Addressable Memory 4, 53
control direct VCC 195
control distribute VCC 195
cratvl 71, 128
CRC-flawed packets 4
crgp command 69, 162

D

data 151
data direct VCC 195
data link layer 189
Deaf status 128
DECNET 75
default group 13, 49
default VLAN 13, 49, 274
DHCP MAC address-based rule 80
DHCP MAC rules 40
DHCP port based rule 79
DHCP port rules 40
diagnostic check 12
disabled 108
discard frames 262
DSAP 75
dynamic IP routing table 282

E

- ELAN 189, 194, 197, 284
- ELAN_NAME 190
- Emulated LAN 189, 284
- Encapsulation 151, 161
- Encapsulation Transformation Rules 154
- encapsulation transformations 155
- End System Identifier 191
- Ether-Type 75, 154, 155, 158
- Ether-Type encapsulated frame 158
- Ether-Type/SNAP transformations 154
- Ethernet
 - 802.3 129
 - LANs 90
 - switching module 4
 - Version II 129, 160

F

- factory settings 14
- Fast Ethernet 273
- FDDI 129, 158, 159
 - trunking 123
 - trunking commands 187
 - virtual bridge port 185
- filtering database 90
- Firmware-Version 26
- flood 52
- flooded frames 160
- forwarding 108
- forwarding state 283
- fragmentation/reassembly 156
- frame flooding 54
- frame path 4
- frame relay 69
- frame size requirements 156
- frame type 13
- frame-to-cell conversion 189
- Frame-to-Cell Switching Module (FCSM) 2, 7
- FTP 30, 129
- ftp commands 31
- fwc command 185

G

- global network-wide addresses 156
- graphical applications 14
- group
 - configuration 67
 - identifiers 47
 - multiplexing on FDDI 185, 186
 - number 50
- GROUP can 12

H

- hardware forwarding engine 54

- hardware routing engine (HRE) 6
- Hello bridge protocol data unit 100
- High-Speed Switching Module 4
- hot-swappable modules 3

I

- IBM spanning tree protocol 105
- ICMP 129, 159
- ICMP-based MTU discovery 159
- IEEE 802.2 160
- IEEE 802.2 IPX 163
- IEEE 802.2 IPX Pass Through 163
- IEEE 802.3 154, 158
- IEEE 802.5 196
- ILMI 192
- ILMI MIB 192
- Inactive 129
- Initial 297
- instance 50
- inter-switch link 13
- Interim Local Management Interface 192
- IP 75
 - fragmentation 159
 - subnet 17
- IPX 75
- IPX frame types 155
- IPX packet size negotiation 160

L

- LAN Emulation 159, 284, 294
 - Clients 189, 194, 197, 284
 - Configuration Server 189, 197
 - Server 189, 197
 - summary 197
 - user-to-network interface 192
- LANE 284
- Last Topology Change 105
- LE Service components 189
- learning 108
- LEC 189, 194
- LECS 189, 192, 193, 284, 290, 296
- LECS ATM addresses 192
- LES 189, 192, 197
- listening 108
- LLC 155, 158, 160
- LLC encapsulation 154
- loading boot parameters 34
- loading from disk 35
- LSB 152
- LUNI 192

M

- MAC address rules 40
- MAC header 151
- MAC-based VLAN 72

- Management Bus 4, 5
- Management Processor Module (MPM) 2, 4, 11, 22
- management station 17
- maximum frame size 158, 289, 296
- maximum frame size negotiation 156
- Maximum Transmission Unit 159
- MBUS 5
- minimum frame size 158
- misaligned packets 4
- modatvl 71
- modem port 28
- module slot 50
- modvp command 162
- MPM 2, 14, 22, 52
- MPM Changeover 25
- MSB 152
- MSS server 190
- MTU 159
- multicast 194
- multicast forward VCC 196
- multicast frames 160
- multicast send VCC 196
- multicast VLANs 48
- multiplex VLAN 169
- MultiProtocol Switching hub 284
- MultiProtocol Switching Services 284

N

- native 161
- network
 - address rules 40, 276, 279
 - address-based VLAN 75
 - header 151
 - level address 156
 - number 157
 - policy-based VLAN 57
 - prefix 191
 - switch module 4
- next hop resolution 142
- NHRP 142
- node address 157
- non-canonical 156
- non-IPX Ether-Type encapsulated frames 158
- non-native 161
- Novell 154
 - 4.1 163
 - network header 154
 - Proprietary 154, 155, 158
 - Raw 154
- Nways RouteSwitch Network Manager 11
- Nways RouteSwitch Software Program 11
- Nways RouteTracker Manager 11

O

- odd-byte padding 159
- OK1 LED 12

- OK2 27
- OK2 LED 12
- Optimized device switching 11, 59

P

- padding 158
- password 14
- path cost 108
- PDU fragmentation/reassembly 159
- physical port numbers 50
- policies 47, 49
- policy-based ELANs 193
- policy-based VLANs 273
- port assignment 70
- port rules 40, 72, 299
- port timeout 53
- port-based translation 160
- port-based VLANs 47
- PRI LED 25
- priority 108
- protocol data unit (PDU) 159
- protocol identifier 155
- protocol rules 40
- protocol-based VLAN 73
- proxy 189
- PS1 LED 12
- PTOP 161

R

- Raw 802.3 154
- redundancy 25
- redundant MPMs 24
- REFERENCE CHECKING 27
- reset 20
- RFC 1490 69
- RIF 94, 151
- RIP 129
- RIP update 157
- Rlogin 130
- rmgp 70
- root bridge 112
- route designator 94
- router 38
- RouteSwitch Architecture 4
- RouteSwitch group 54
- RouteSwitch Network Manager 18
- RouteSwitch Software 11
- RouteTracker Manager 18
- RouteTracker VLANs 48
- routing 89
 - control field 94
 - functions 23
 - information field 94
 - protocols 13
- rts command 187
- runs 4

S

- SAP 129
- SEC LED 25
- Selector byte 191
- serial port specification 14
- service 50
- Silent 128
- SLIP 130
- SNAP 75, 154, 155
- SNAP encapsulated protocols 158
- SNAP frames 160
- SNMP 129
- SNMP management 23
- socket number 157
- Software Fallback 25
- software version 25
- source route bridges 89
- source route token-ring 129
- source route transparent bridges 90
- spanning tree 89, 170, 186
 - algorithm 90, 112
 - explorer 95
 - parameters 103
 - setting parameters 103
- SPARC RISC 3, 5
- specifically routed frame 95
- SRF 95
- SSAP 75
- static route 17
- STE 95
- store-and-forward 4
- stpc 108
- stps 108
- subnetwork MAC address 156
- SVC 196
- SVCs 294
- swap state 27
- swapping 26
- switching 89
- system contact 14
- system location 14
- system name 14

T

- TCN BPDU 103
- TCP 129
- TCP/IP address 14
- TELNET 130
- telnet client 17
- terminal emulation 14
- timeout 278
- token-ring 129, 158
- Token-Ring IPX Switching 158
- topology change notification 103
- transfer methods 30
- translational bridges 89

- transparent bridge 89, 90
- transparent bridges 89
- trunk 161
- trunk ports 170
- trunking 89, 169, 185
 - configuration 170
 - hello message 186
 - over ATM 187
 - ports 161

U

- UDP 129
- unicast destination address 160
- unicast floods 59
- unicasts 52
- universally administered 191
- user interface 29
- user-defined based VLAN 77
- user-defined rules 40
- user-to-network interface 192
- using ZMODEM 32

V

- vas command 173
- VBUS 5
- VCC 192
- VCCs 194
- VCI 187
- vi command 162, 185
- via 70
- virtual circuit identifier 187
- Virtual LANs (VLANs) 12, 37
- virtual port numbers 50
- virtual rings 120
- virtual router 125
- virtual router port 50
- virtual transparent bridge port 185
- VLAN 189
 - bus 5
 - group number 170
 - groups 169
 - membership 58
 - policies 40
 - policy rules 283
 - timers 52
 - trunking 169
- vportEncapsulation 162
- vportSwitchDefaultTable 162
- vportSwitchTable 162
- vportTable 162

Z

- ZMODEM protocol 32

ITSO Redbook Evaluation

IBM Nways RouteSwitch Implementation Guide
SG24-4881-01

Your feedback is very important to help us maintain the quality of ITSO redbooks. **Please complete this questionnaire and return it using one of the following methods:**

- Use the online evaluation form found at <http://www.redbooks.ibm.com>
- Fax this form to: USA International Access Code 914 432 8264
- Send your comments in an Internet note to redbook@us.ibm.com

Which of the following best describes you?

☐ **Customer** ☐ **Business Partner** ☐ **Solution Developer** ☐ **IBM employee**
☐ **None of the above**

Please rate your overall satisfaction with this book using the scale:
(1 = very good, 2 = good, 3 = average, 4 = poor, 5 = very poor)

Overall Satisfaction _____

Please answer the following questions:

Was this redbook published in time for your needs? Yes_____ No_____

If no, please explain:

What other redbooks would you like to see published?

Comments/Suggestions: (THANK YOU FOR YOUR FEEDBACK!)

