


Understanding and Using MSS Release 1.1 and 2.0

Erol Lengerli, David Mindel, Volkert Kreuk, Randel T. Sink



International Technical Support Organization

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

This book was printed at 240 dpi (dots per inch). The final production redbook with the RED cover will be printed at 1200 dpi and will provide superior graphics resolution. Please see "How to Get ITSO Redbooks" at the back of this book for ordering instructions.



International Technical Support Organization

SG24-2115-01

**Understanding and Using
MSS Release 1.1 and 2.0**

April 1998

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix A, "Special Notices" on page 225.

Second Edition (April 1998)

This edition applies to Release 1.1 of the IBM MSS Server.

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. MSS Functional Overview	1
1.1 New in MSS Release 1.1	2
1.2 New in MSS 2.0	5
 Chapter 2. MSS 1.1 Hardware Overview	7
 Chapter 3. MSS 2.0 Hardware Overview	9
 Chapter 4. MSS R1.1 - New Interfaces	11
4.1 FDDI Interface	11
4.1.1 Benefits of the FDDI Interface	11
4.2 ATM Virtual Interfaces (AVIs)	11
4.2.1 Description of ATM Virtual Interface	13
4.2.2 Benefits of Using ATM Virtual Interfaces	13
4.3 IBM-Compliant LAN Emulation Client	15
4.3.1 ILEC Routing and Bridging Overview	16
4.3.2 Benefits of ILECs	17
4.3.3 Limitations of ILECs	18
4.3.4 Quick Overview of IBM-Compliant LAN Emulation	18
 Chapter 5. RFC 1483 Bridging over ATM	21
5.1 RFC 1483 Bridging over ATM with MSS 1.0 and 1.1	21
5.1.1 Configuration Issues	25
5.2 RFC 1483 Bridging over ATM with MSS 2.0	26
 Chapter 6. MSS R1.1 - Bridging and Routing	27
6.1 Real, Virtual and Logical Interfaces	27
6.2 Bridging Overview	28
6.2.1 Protocol Support	29
6.3 Routing Overview	29
6.3.1 IP Routing	30
6.3.2 IPX Routing	31
6.3.3 AppleTalk Routing	31
6.4 Concurrent Bridging and Routing	32
6.4.1 Routing and Bridging on ARIs	33
6.4.2 Routing and Bridging on ELANs	33
 Chapter 7. MSS 2.0 Routing Enhancements	35
7.1 IP Routing Enhancements	35
7.2 IPX Routing Enhancements	35
7.3 Banyan Vines Routing	35
7.4 APPN Routing	36
7.5 SR-TB Duplicate MAC Address Support	36
 Chapter 8. APPN	39
8.1 MSS As a Network Node	39
8.2 End Node (EN)	40

8.3 Low-Entry Network End Nodes (LEN)	40
8.4 Dependent LU Requester (DLUR)	41
8.5 APPN Connection Networks (CN)	41
8.6 High-Performance Routing (HPR)	41
Chapter 9. Configurable Quality of Service for LAN Emulation	43
9.1 Overview	43
9.2 Benefits	44
9.2.1 Migration Issues	44
9.3 Functional Description	45
9.3.1 Using Configurable QoS Parameters	45
9.3.2 LAN Emulation Configuration Server (LECS) Behavior	45
9.3.3 LAN Emulation Server (LES) Behavior	45
9.3.4 LAN Emulation Client (LEC) Behavior	45
9.3.5 Configurable QoS Parameters of an LE Client	46
9.4 LANE QoS Negotiation Algorithm	52
9.5 Benefits	52
Chapter 10. SuperVLANs	55
10.1 Shortcut Bridging (SCB)	56
10.2 Broadcast Management	62
10.2.1 Broadcast Management (BCM)	63
10.2.2 Bridge Broadcast Management (BBCM)	66
10.2.3 Dynamic Protocol Filtering (DPF) for PVLANS	66
10.2.4 IP Cut-Through Support for Zero-Hop Routing	71
10.2.5 Configuring DPF and/or IP Cut-Through	73
10.2.6 Configuring DPF for IPX	74
10.2.7 Configuring DPF for NetBIOS	76
10.3 Multiple SuperVLAN Instances	76
10.3.1 Relationship between SuperVLANs and LNNI	78
10.4 MSS and IBM 8274/3 Broadcast Control	78
10.4.1 Overview	78
10.4.2 Things You Should Know	79
10.4.3 An Example	83
Chapter 11. Next Hop Resolution Protocol Support for IP	87
11.1 ELAN Support	90
11.2 MSS 2.0 Zero-Hop Route Server (RouteSwitching)	95
11.2.1 Overview of RouteSwitching Clients	96
11.2.2 RouteSwitching Environments for Clients	97
11.2.3 RouteSwitching Environments With MSS	98
11.2.4 RouteSwitching Flows	100
Chapter 12. Default IP Gateway Redundancy and LAN Emulation	103
12.1 Configuration Issues	106
Chapter 13. Classical IP Redundancy and Multicast Features	107
13.1 MSS 1.1	108
13.1.1 ATMARF Server Redundancy	108
13.1.2 ATMARF Server/Default IP Gateway Redundancy	109
13.1.3 IP Addresses	110
13.1.4 ATM Addresses	111
13.1.5 Configuration Issues	112
13.2 MSS 2.0	112
13.2.1 Distributed ATM ARP Server	112

13.2.2 IP Multicast over ATM	116
Chapter 14. BUS Performance Enhancements	123
14.1 Benefits of the BUS Performance Enhancements	124
14.2 Limitations of the BUS Performance Enhancements	124
14.3 Details on BUS Modes	125
14.3.1 System Mode	125
14.3.2 Adapter Mode	125
14.3.3 VCC Splicing Mode	126
Chapter 15. Miscellaneous Enhancements MSS 1.1	129
15.1 Sharing of RFC 1483 PVCs	129
15.2 Multiple Internal LE Clients on the Same ELAN	129
15.2.1 ATM Interface VCC Binding	130
15.3 Dynamic Allocation of LE_ARP Cache Entries	130
15.3.1 Recommendation	130
15.4 SRAM Size Increase	131
15.5 New ESI Selector LECS Assignment Policy	131
15.5.1 Maximum Size of BCM's IPX Forwarding List	131
15.5.2 Persistent SET Support for LE Service MIBs	132
15.6 SR-TB Translational Bridging Support For IP	132
15.7 Association of Voice Files with Load Image	132
Chapter 16. Miscellaneous Enhancements MSS 2.0	133
16.1 Dynamic Reconfiguration	133
16.2 Dynamic Linking and Loading	134
16.3 Time Activated Reboot	134
Chapter 17. Configuration Examples MSS Version 1.1	137
17.1 Preparation	138
17.1.1 Connecting MSS Server to the ATM Network	139
17.1.2 Allow Dynamic Access to the LECS	140
17.1.3 Establishing IP Connectivity	141
17.1.4 Configuring SNMP	143
17.2 Selecting and Transferring a Sample Configuration	144
17.2.1 Terminology	144
17.2.2 Opening a Configuration	145
17.2.3 Optional Customization	145
17.2.4 Downloading a Configuration	146
17.2.5 Loading a Configuration Using XModem	147
17.3 Example - Basic Connectivity	149
17.3.1 Configuration Details	149
17.3.2 Logical IP Layout	151
17.3.3 Verifying the Basic Configuration	152
17.4 Example - Shortcut Bridging	154
17.4.1 Configuration Details	155
17.4.2 Required Customization	157
17.4.3 Logical IP Layout	157
17.4.4 Verifying Shortcut Bridging	158
17.5 Example - IP Cut-Through	164
17.5.1 Configuration Details	164
17.5.2 Logical IP Layout	166
17.5.3 Verifying IP Cut-Through	167
17.6 Example - Backward Connectivity	171
17.6.1 Configuration Details	171

17.6.2 IBM 8260 PVC Definitions	173
17.6.3 IBM 8274 Definitions	173
17.6.4 Logical IP Layout	174
17.6.5 Verifying Backward Connectivity	174
17.7 Example - Redundancy	178
17.7.1 LES/BUS Service Distribution	181
17.7.2 Configuration Details	181
17.7.3 Customization	186
17.7.4 Logical IP Layout	188
17.7.5 Verifying Redundancy	189
Chapter 18. Configuration Examples For MSS Version 2.0	201
18.1 Example - Zero-Hop IP Routing (RouteSwitching)	202
18.1.1 Configuration Details	203
18.1.2 Verifying Zero-Hop Routing	205
18.2 Example - IP Multicasting in a LIS Using MARS	208
18.2.1 Setting Up Classical IP via the Configurator	210
18.2.2 Configuring MARS	211
18.2.3 Verifying MARS And MCS	212
18.3 Example - Distributed ATMARP Servers in a LIS	213
18.3.1 Distributed and Redundant ATMARP Server Definition	213
18.3.2 Verifying Distributed and Redundant ATMARP Server	216
18.4 APPN Configuration	218
18.4.1 Configuring MSS Server for APPN	218
18.4.2 Configuration Steps for APPN	218
18.4.3 APPN Monitoring	219
18.5 APPN Configuration Example	219
18.6 Example - RFC 1483 SVC Bridging	220
18.6.1 Configuration Details	220
18.6.2 Verifying RFC 1483 SVC Bridging	222
18.6.3 Verifying the Source Routing Information Field (RIF) in the RFC 1483 SVC Bridging Example	224
Appendix A. Special Notices	225
Appendix B. Related Publications	227
B.1 International Technical Support Organization Publications	227
B.2 Redbooks on CD-ROMs	227
B.3 Other Publications	227
How to Get ITSO Redbooks	229
How IBM Employees Can Get ITSO Redbooks	229
How Customers Can Get ITSO Redbooks	230
IBM Redbook Order Form	231
List of Abbreviations	233
Index	237
ITSO Redbook Evaluation	239

Preface

This redbook will help you to understand and use the functions offered by MSS Server Releases 1.1 and 2.0. Proper understanding of the technology and following the configuration examples given in this publication will be helpful in designing and building an ATM network using IBM's Multiprotocol Switching Services.

An overview of the functions provided by MSS Server Release 1.1 and a short description of each of the new functions is included.

With MSS R2.0, new functions in addition to the new hardware are delivered. The major functions of MSS R2.0 are RouteSwitching, APPN routing, IP multicast over ATM, RFC 1577+ support, duplicate MAC address in SR-TB networks. These and other new functions are detailed in this redbook.

Finally, due to its versatility, getting started with the MSS Server can sometimes create difficulties. Therefore, a number of configuration examples are provided. They can also be downloaded from the Internet at <ftp://www.redbooks.ibm.com/redbooks/SG242115>. A description is included on how to activate and verify each of the configurations.

Have a look at the MSS R1.0 redbook as well

We recommend that you read this publication in conjunction with *Understanding and Using the MSS Server*, SG24-4915.

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.

Erol Lengerli is an Advisory ITSO Specialist for Networking at the Systems Management and Networking ITSO Center, Raleigh. He has 13 years of experience in WAN design/installations. His area of expertise include ATM switches, WAN switches, and NetView. Before joining the ITSO, Erol worked as a Systems Engineer and Networking Specialist in IBM Turkey.

David Mindel is an I/T specialist in IBM USA. He joined IBM in 1982; since 1986 he has supported technical marketing of networked solutions. His current responsibilities include designing/installing ATM-based networks, architecting TCP/IP networks, leading design workshops, and developing solutions for Web server load balancing, network application testing and networked multimedia. David has also been involved in creating IBM services and hands-on sessions such as LAN Doctor and ATM training.

Volkert Kreuk is a Senior Networking Specialist working at the Advanced Technical Support (ATS) organization in Raleigh. His current responsibility is to advise on the design and implementation of ATM campus networks. He writes extensively and teaches IBM classes world-wide on SNA/APPN, TCP/IP, frame relay, and ATM. Before joining the ATS he worked at the ITSO and IBM Global Networking in the Netherlands. He received BS and MS degrees in Physics from the Vrije Universiteit in Amsterdam.

Randel T. Sink is an I/T Program Mgr/Consultant for Networking Hardware Division's Installation Support Department in Raleigh. He has 24 years of experience in networking design/installations. His area of expertise include ATM, LAN, and WAN systems. Randel is a registered Professional Engineer and consults with IBM customers worldwide.

The authors of the first edition of this redbook were:

Cedel Alexander
Chief Architect
MSS Development

Heiner Heng
IBM Germany

Volkert Kreuk
Advanced Technical Support
IBM USA

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

Dee Vig
MSS Development

Network Services and Support
IBM USA

Roger Lowrie
IBM Canada

Bernie Grunwald
IBM USA

Mohammad Shabani
IBM UK

Paul Braun
Kathryn Casamento
Mike Haley
John Parker
Linda Robinson
Gail Wojton
ITSO Raleigh

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 239 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. MSS Functional Overview

The IBM Multiprotocol Switched Services Server is now being shipped in MSS 2.0. The following sections describe the functional overview of the previous release (MSS 1.1) and this current release (MSS 2.0). With the new functions introduced in MSS Server Release 1.1 (see 1.1, "New in MSS Release 1.1" on page 2), and MSS Server 2.0 (see 1.2, "New in MSS 2.0" on page 5), the MSS Server now supports the following functions:

1. ATM Forum-compliant LAN Emulation (LANE)
2. IBM LAN Emulation Client (ILEC)
3. Classical IP (CIP)
4. Bridging
5. Routing
6. SuperVLAN
7. Next Hop Resolution Protocol (NHRP)

The ATM Forum LAN emulation support includes:

- LAN Emulation Client (LEC)

LE clients are required to enable the MSS routing and bridging for emulated LANs and to provide access to the MSS management facilities. Support is available for ATM Forum and IBM-compliant LE clients.

- LAN Emulation Services
 - LAN Emulation Server (LES)
 - Broadcast and Unknown Server (BUS)
 - LAN Emulation Configuration Server (LECS)

LE server functions are required to establish ATM Forum emulated LANs. No server functions for IBM-compliant LANs are offered and an external NetWare server is required to provide IBM-compliant server functions.

- IBM value-add extensions for:

- Congestion control
- Redundancy

Redundancy features have been implemented to enable dynamic backup for LECS, LES/BUS, ARP server, and default IP gateway functions.

- Security

The LECS/LES security interface enables authentication of LE clients joining an ELAN.

- Broadcast management

Broadcast management facilities are available for the BUS function, and the bridging function. For details see Chapter 10, "SuperVLANs" on page 55.

- Enhanced network management including Classical IP MIB support

The IBM LAN Emulation Client enables bridging and routing between IBM ELANs, ATM Forum ELANs, and ATM connections using RFC 1483.

RFC 1483 encapsulation can be used as an alternative or in addition to LAN emulation. RFC 1483 is an encapsulation mechanism that enables multiple protocols sharing a single VCC. MSS 2.0 brings SVC support for RFC 1483 which

greatly enhances its utility while relaxing configuration tasks. A distinction is made between the bridged and routed frame format.

The routed frame format is used by MSS for the transport of either IP or IPX. IP transport using RFC 1483 is referred to as Classical IP. The Classical IP support includes both the ATMARP client and the ATMARP server function. The bridged frame format is used to encapsulate and transport legacy LAN frames over ATM.

Bridging is supported over emulated Ethernet and token-ring interfaces, either IBM or ATM Forum-compliant, and RFC 1483 connections. Support is included for transparent bridging, source-route bridging, and translational bridging.

Routing support is provided for IP, IPX, AppleTalk and Banyan Vines. Several routing enhancements have been added in MSS 2.0 making it more robust than ever. IP routing is supported between Classical IP, ATM Forum LANE, IBM-compliant LANE subnets and FDDI subnets. IPX routing is supported over LAN Emulation interfaces, FDDI interface and RFC 1483 connections to other routers. AppleTalk routing is supported over LAN Emulation and FDDI interfaces.

APPN routing has been added in MSS 2.0 as well as the ability to distribute the ARP server function and provide redundancy using RFC 1577+. Other enhancements include VCC multiplexing on the ATM interfaces, SR-TB duplicate MAC address support, dynamic reconfiguration, linking and loading functions and Ethernet support for out-of-band management.

Other MSS-specific functions such as a zero-hop routing server and a multicast address resolution server (RFC 2022) have also been added. For details on these and other bridging and routing functions see Chapter 6, "MSS R1.1 - Bridging and Routing" on page 27.

1.1 New in MSS Release 1.1

MSS Server Release 1.1 adds the following new features to the functions already introduced in Release 1.0:

1. New interfaces

The new types of interfaces introduced are:

- FDDI interface (see Figure 4 on page 12)

MSS R1.1 introduces FDDI support for the IBM 8210 MSS Server (note, *not* for the IBM 8260 MSS Server module). For a detailed description of the FDDI functions on the IBM 8210 see 4.1, "FDDI Interface" on page 11.

- ATM Virtual Interface (AVI) (see Figure 5 on page 14)

MSS provides an ATM real interface for each ATM attachment. MSS R1.1 introduces the concept of ATM virtual interface (AVI) (see 4.2, "ATM Virtual Interfaces (AVIs)" on page 11) to increase the number of Classical IP and RFC 1483/IPX connections, and improve the support for IP multicast over Classical IP.

2. IBM LAN Emulation Client (ILEC) (see Figure 6 on page 16)

MSS R1.0 provides extensive routing and bridging support for ATM Forum-compliant emulated LANs. MSS R1.1 adds the same routing and bridging support for attachments to IBM-compliant emulated LANs. For a

detailed description see 4.3, “IBM-Compliant LAN Emulation Client” on page 15.

3. Quality of Service (QoS) for LAN emulation clients (see Figure 18 on page 43)

MSS R1.0 LAN emulation functions assumed the use of best-effort data direct VCCs between LAN emulation clients. MSS R1.1 enables the use of Quality of Service (QoS) for the LE-client_to_LE-client data direct VCC. For a detailed description see Chapter 9, “Configurable Quality of Service for LAN Emulation” on page 43.

4. SuperVLAN support

MSS R1.0 provides broadcast management functions (see Figure 26 on page 64) and backup features that are unique in the industry. The SuperVLAN functions (for details see Chapter 10, “SuperVLANs” on page 55) introduced in MSS R1.1 further expand the possibility to manage and reduce ELAN broadcast traffic, to enable service distribution, and to implement dynamic backup facilities.

The SuperVLAN package consists of the following functions:

- Shortcut Bridging (SCB) (see Figure 23 on page 60)

SCB enables the establishment of inter-ELAN data direct VCC, that is VCCs between LE clients connected to different ELANs. For details see 10.1, “Shortcut Bridging (SCB)” on page 56.

- Bridge Broadcast Management (BBCM) (see Figure 28 on page 67)

BBCM is an enhancement to the MSS bridging function. By spoofing the (IP and NetBIOS) broadcast traffic between ELANs, BBCM learns about the location of IP, IPX, and NetBIOS station. Using this acquired knowledge BBCM converts broadcast messages into unicast traffic. The unicast traffic is directed only the target station which avoids unnecessary traffic and station perturbation. For details see 10.2.2, “Bridge Broadcast Management (BBCM)” on page 66.

- Dynamic Protocol Filtering (DPF) (see Figure 29 on page 68)

DPF introduces another enhancement to the MSS bridging function. Based on the pre-configured and/or learned protocol VLAN (PVLAN) definitions, the MSS R1.1 bridging function will limit protocol broadcasts within the PVLAN. For details see 10.2.3, “Dynamic Protocol Filtering (DPF) for PVLANS” on page 66.

The SuperVLAN package is a major contributor to the implementation of IBM’s switched virtual networking architecture. By bypassing the MSS bridging function for unicast traffic SCB increases network performance and avoids potential capacity problems. BBCM helps to decrease the impact of broadcast traffic on the the ELANs bridged by MSS *and* the legacy LAN bridged/switched to these ELANs.

Switched Virtual Networking

IBM switched virtual networking (SVN) architecture is based on the concept of using LAN and ATM switching rather than using router solutions. A switching solution is more scalable and less expensive. Better performance can be expected due to the lower latency on ATM and LAN switches, and the use of small fixed-size cells, rather than variable length packets, in the ATM backbone.

Essential for a switched solution is proper broadcast management as offered by MSS. Avoiding unnecessary broadcast enables network administrators to build larger flat networks. Flat networks are more simple to build, easier to manage, and less error prone.

Note: A flat network is a single broadcast domain with all stations using the same addressing structure (for example, a single IP subnet).

Each function within the SuperVLAN package is independent and can either be used individually or in conjunction with other SuperVLAN functions. The SCB function currently applies to transparently bridged ELANs only. BBCM and PDF can be used both in transparent and source-route bridged environments. BBCM applies to IP and NetBIOS, while PDF provides support for IP, IPX, and NetBIOS.

5. Next Hop Resolution Protocol (see Figure 42 on page 88)

In a traditional ELAN and/or Classical IP environment inter-subnet traffic requires an IP router interconnecting the subnets. The Next Hop Resolution Protocol (NHRP), however, provides capability for stations to determine the Non-Broadcast Multi-Access (NBMA) address of a destination station and to establish shortcut VCCs over CIP or ELAN connections.

The NHRP support introduced in MSS R1.1 enables data direct VCCs to be established between ATM-attached IP endstations and or routers that bypass intermediate router(s). This function increases performance and avoids capacity problems.

For a detailed description see Chapter 11, "Next Hop Resolution Protocol Support for IP" on page 87.

6. RFC 1483 support for bridging over ATM (see Figure 9 on page 21)

MSS R1.1 introduces the possibility to exchange bridged traffic using RFC 1483 encapsulation over pre-configured PVCs, perform bridging functions between RFC 1483 bridged connections and vice versa, and between RFC 1483 bridged connections and emulated LANs (both ATM Forum and IBM-compliant). For details see Chapter 5, "RFC 1483 Bridging over ATM" on page 21.

7. Redundant default IP gateway for ELAN (see Figure 54 on page 103)

MSS R1.1 provides enhanced support for a redundant default IP gateway. Similar to the MSS R1.0 implementation, the backup MSS Server will only function as the operational default IP gateway when the primary gateway is not available; however, support has been added to enable the backup gateway to establish an active IP subnet attachment while being in normal (primary gateway operational) mode. This avoids possible routing problems when running in normal mode, especially when using dynamic routing protocols.

See Chapter 12, “Default IP Gateway Redundancy and LAN Emulation” on page 103.

8. Classical IP redundancy (see Figure 57 on page 110)

MSS R1.1 enhances the ARP server redundancy feature introduced in MSS R1.0. MSS now allows a pre-configured primary/backup relation. The backup node will only become the operational ARP server if the primary node is inactive.

The enhanced ARP server redundancy feature can also be used to provide default IP gateway redundancy for your Classical IP subnets. Similar to the support for ELANs (see 7 on page 4), the backup IP gateway only becomes active when the primary fails; however, during normal mode, the backup node has an active interface to the subnet thus avoiding routing problems when using dynamic routing protocols.

See Chapter 13, “Classical IP Redundancy and Multicast Features” on page 107.

9. BUS performance enhancements

MSS R1.1 enables the user to specify in which mode the Broadcast and Unknown Server (BUS) will run. Options are:

- System mode (see Figure 66 on page 125)
- Adapter mode (see Figure 66 on page 125)
- VCC splicing mode (see Figure 66 on page 125)

In system mode the BUS operates similar to the MSS R1.0 BUS function. In adapter and VCC splicing mode BUS intelligence is traded for BUS forwarding power. Maximum BUS performance can be reached with VCC splicing mode. The BUS mode can be assigned individually for each of your defined BUSes.

See Chapter 14, “BUS Performance Enhancements” on page 123.

10. Miscellaneous enhancements

See Chapter 15, “Miscellaneous Enhancements MSS 1.1” on page 129.

1.2 New in MSS 2.0

MSS Server 2.0 adds the following new features to the functions already introduced in Release 1.0 and Release 1.1:

1. Advanced Peer-to-Peer Networking

Advanced Peer-to-Peer Networking (APPN) extends the SNA architecture by enabling indirectly connected Type 2.1 (T2.1) nodes to communicate without requiring the services of a SNA host computer. In a way, this dramatically reduces the amount of predefinition required.

2. Distributed ARP server

The Distributed ARP server eliminates the single point of failure for Classical in conjunction with enhanced Classical IP ARP clients (RFC 1577+). Users will have logical IP subnets (LIS) in the event of an ARP server failure.

3. RFC 1577+ Client

LIS clients must be RFC 1577+ capable. This allows them to recognize when connection to the ARP server is not operational and be able to switch to an alternate server.

4. RFC 1483 bridging using SVCs

RFC 1483 bridging support uses the bridged frame format for encapsulation. This, along with the advantages of using switched virtual circuits, significantly reduces the configuration efforts required to set up these functions on ATM switches.

5. Improved IP routing support

This release of MSS now supports RIP Version 2 (RFC 1723). RIP2 includes variable subnet mask support, route advertisement to well-known IP multicast addresses (instead of broadcasting), use of authentication keys, and backward compatibility with RIP V1. Improvements in IP routing also include multiple static routes (≤ 4) per destination, IP route filtering, OSPF demand circuits (RFC 1793), and TraceRoute enhancements.

6. Improved IPX routing support

Improved IPX support includes IPX static routes and services, default routes, IPX ping enhancements, and the introduction of IPX RecordRoute.

7. Banyan Vines routing

The support for Banyan Vines includes the Vines IP protocol, RIP, ARP, and Internet Control Protocol (ICP).

8. Other routing enhancements

In addition to the above routing enhancements, this release of MSS also includes functions for routing between ATM FC ELANs and IBM ELANs.

9. LAN Emulation enhancements

LANE enhancements in this release of MSS include a completely dynamic LE_ARP cache for LECs, a random delay for LECs trying to join a LES/BUS during a congested period.

10. ATM interface enhancements

MSS now has the ability to automatically multiplex SVCs and PVCs on native ATM interfaces for the RFC 1483 ATM protocols.

11. SR-TB duplicate MAC address support

This feature provides redundancy and load balancing for Ethernet-attached devices by allowing the cache of Ethernet MAC addresses with a specific RIF for a given token-ring address.

12. Dynamic reconfiguration

Dynamic reconfiguration allows users to change the configuration of the MSS Server without requiring a reboot to activate the changes.

13. Dynamic linking and loading

This feature allows selective loading of functions from the operational code image into processor memory.

14. MSS as a zero-hop routing server for IP

The term zero-hop routing refers to mechanisms that allow legacy LAN-attached devices to participate in NHRP.

15. MSS as an ATM Multicast Address Resolution Server (RFC 2022)

The Multicast Address Resolution Server (MARS) is an extended analog of the ATM ARP server that provides the necessary connection and addressing services required by layer 3 IP multicast clients.

Chapter 2. MSS 1.1 Hardware Overview

MSS consists of hardware and software. MSS 1.1 is a software-only release that builds on the hardware of Release 1.0.

The hardware comes in two forms:

1. The 8210 stand-alone box (see Figure 1)

The 8210 stand-alone box provides two high-speed adapter slots which can be used for ATM and/or FDDI attachments. The ATM and FDDI adapters use the same adapter slots. The FDDI adapter must be installed in slot #2. Using the FDDI adapter limits the number of ATM attachments to one.

For the 155 Mbps ATM attachments you have the option of using single-mode or multimode fiber. Both use SC connectors. The 8210 conforms to ATM Forum standards UNI 3.0 and 3.1 and can be connected to any ATM switch that supports these standards. By using two ATM adapters maximum redundancy can be achieved.

For FDDI a dual or single ring (DAS/SAS) attachment, using either a fiber or copper adapter, is provided. Using an FDDI adapter provides an easy migration path from an FDDI backbone to ATM.




Figure 1. IBM 8210 MSS Server Model 001

2. The MSS Server Module (see Figure 2 on page 8)

The MSS Server module occupies two slots in an IBM 8260 hub. It connects to the ATM network using ATM Forum standards UNI 3.0 and 3.1 via a single backplane attachment.

The module in the IBM 8260 takes advantage of the following dependability and fault tolerance options provided by the IBM 8260: hot swap-ability, intelligent power management, distributed management, and redundancy for critical components including power supplies and switches.




Figure 2. MSS Server Module Model 001

Functionally Equivalent

The IBM 8210 MSS Server and the MSS Server module are functionally equivalent. They both connect to the ATM network as an end device. The FDDI functions, however, are only available on the IBM 8210.

Both products are based on 100 MHz PowerPC 603e platforms that include 512 KB of L2 Cache, 32 MB of processor DRAM, 12 MB flash memory, and 8 MB of buffer memory per ATM interface. In addition they have:

- A standard serial service port that supports asynchronous communication for configuration and maintenance. The port conforms to the EIA 232 standard, with a male, 9-pin, D-shell connector and is capable of operating at up to 38.4 Kbps. The port supports direct attachment or modem attachment and auto answering.
- A slot for a PCMCIA hard disk for code and configuration storage, first failure data capture, logging, trace and dump data. The hard disk is mandatory. In case of a hard disk malfunction the MSS Server will be able to startup or continue to operate with limited functionality, as long as no disk accesses are required.
- A PCMCIA modem slot (which may contain a Data/Fax modem or, in the US and Canada, a Voice/Data/Fax modem) for remote installation, network management, and service access.

Chapter 3. MSS 2.0 Hardware Overview

MSS Release 2.0 consists of hardware changes and software changes. MSS 2.0 is a multiproduct release that includes a new MSS Server Module for the 8260/8265 hubs and changes to the 8210 stand-alone box.

The hardware comes in two forms:

1. The 8210-002 stand-alone box

The 8210-002 stand-alone box provides the same two high-speed adapter slots as the model 001.

2. The MSS Server Module (see Figure 3)

The MSS Server module occupies one slot in an IBM 8260/8265 hub. It connects to the ATM network using ATM Forum standards UNI 3.0 and 3.1 via a single backplane attachment.

As with MSS 1.1 the MSS Server module in the IBM 8260/8265 takes advantage of the dependability and fault tolerance options provided by the IBM 8260/8265.




Figure 3. Single-Slot MSS Server Module

Functionally Equivalent

The IBM MSS Server module and stand-alone 8210-002 have a 10 Mbps RJ45 Ethernet service port for use with out-of-band management of the MSS Server. This allows for fast upgrades of operational code and configurations. They both connect to the ATM network as an end device. The FDDI functions, however, are still only available on the IBM 8210.

Both products are based on 166 MHz PowerPC 603eV platforms that include 512 KB of L2 cache, 64 MB of processor EDO DRAM, 12 MB flash memory, and 8 MB of buffer memory per ATM interface. They also have two PCMCIA slots which can be used for a hard disk, a voice/data/fax or flash memory. In addition they have:

- Management/Configuration Interfaces
 1. A standard serial service port that supports asynchronous communication for configuration and maintenance. The port conforms to the EIA 232 standard, with a male, 9-pin, D-shell connector and is capable of operating at up to 38.4 kbps. The port supports direct attachment or modem attachment and auto answering.
 2. 10Base-T RJ45 Ethernet port on the single-slot MSS module, and 8210-002 that can be installed in the 8260/8265.
- A built-in 1.6 GB IDE hard drive for operational code and configuration storage, first failure data capture, logging, trace and dump data. In case of a hard disk malfunction the MSS Server will still be able to startup or continue to operate.

Note: The plug-in PCMCIA disk required in Model 001 can still be used for migration, startup, or continued operation with limited functionality, as long as no disk accesses are required.
- A PCMCIA modem slot (which may contain a data/fax modem or, in the U.S. and Canada, a voice/data/fax modem) for remote installation, network management, and service access.

Chapter 4. MSS R1.1 - New Interfaces

This chapter details the three new types of interfaces introduced in MSS Release 1.1:

1. FDDI interface

FDDI support has been introduced for the IBM 8210. For details on this new type of (physical) interface, see 4.1, "FDDI Interface."

2. ATM virtual interface (AVI)

ATM virtual interfaces are a new type of (logical) interfaces that extend the routing capabilities of the MSS Server. For details, see 4.2, "ATM Virtual Interfaces (AVIs)."

3. IBM-compliant LAN emulation interface

By providing support for IBM-compliant LAN emulation clients a new type of (logical) interface has been introduced. For details, see 4.3, "IBM-Compliant LAN Emulation Client" on page 15.

4.1 FDDI Interface

Support for an FDDI router interface has been added in MSS Release 1.1. The following routing protocols, see Figure 4 on page 12, will be supported by the FDDI interface: IP, IPX, and AppleTalk. Bridging is not supported over the FDDI interface in Release 1.1 of the MSS Server. Note that bridged frames can be exchanged using bridge tunneling over IP.

The FDDI interface is implemented with a standard PCI adapter. The FDDI adapter must be inserted in slot #2 on the IBM 8210. Using an FDDI adapter limits the number of ATM attachments to one.

The FDDI interface is only supported on the stand-alone (8210) version of the MSS Server; it is not supported on the MSS Server blade for the 8260.

4.1.1 Benefits of the FDDI Interface

This new interface enables the MSS Server to provide routed connectivity between high-speed FDDI and ATM backbone networks. This connectivity provides a path for a migration from existing FDDI backbones to an ATM backbone.

4.2 ATM Virtual Interfaces (AVIs)

The MSS bridging and routing functions associate a single ATM real interface (ARI) with each ATM adapter that is physically installed. This physical interface can be used for IP routing, IPX routing, and bridging function all using RFC 1483 encapsulation.

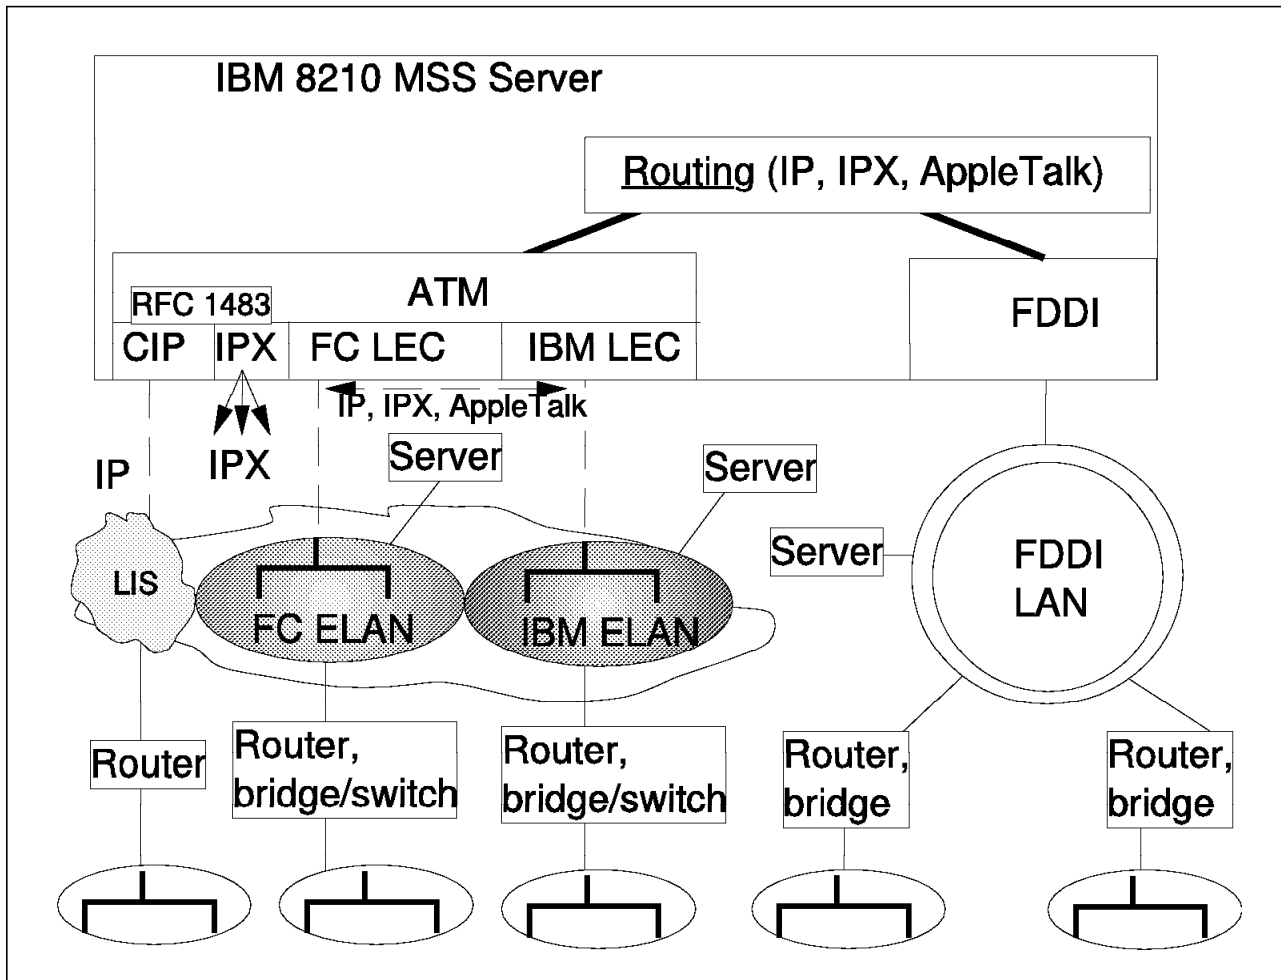


Figure 4. IBM 8210 - FDDI Routing

Notes:

1. IP routing using RFC 1483 encapsulation is generally referred to as Classical IP.
2. AppleTalk routing requires an FDDI or a LAN emulation interface.

The maximum number of ATM real interfaces (ARIs) on the IBM 8210 MSS Server is two. The number is limited to one on the IBM 8260 MSS Server module. Using this limited number of ARIs for all RFC 1483 bridging and routing functions creates some potential problems:

1. Number of Classical IP subnet attachments

The number of Classical IP addresses that can be defined per real interface is 32. Therefore, per ARI you can have a maximum of 32 Classical IP subnet attachments.

2. IP multicast

The IP routing code will not forward an IP multicast on the same interface as the multicast message has been received. This means when connected to multiple subnets on the same interface, a multicast message received on one subnet will never be forwarded to the other subnets on the same interface.

3. Number of RFC 1483/IPX subnet attachments

Only a single IPX network number can be defined per real interface.

4. Interference of bridging and routing

Adding an IP and/or IPX address to a real interface effectively disables the bridging function for IP and/or IPX respectively, for bridged connections on the same interface. As an example, having defined a Classical IP client and an RFC 1483 bridging connection on the same real interface means that the bridging function is disabled for IP.

To enable customers to bypass above limitations MSS Server Release 1.1 introduces the concept of ATM virtual interfaces (AVI).

4.2.1 Description of ATM Virtual Interface

An ATM virtual interface (AVI) is a feature used to create the look of multiple ATM interfaces when, in fact, there is only one physical ATM interface. Each AVI is associated with a real interface; however, you can configure multiple AVIs for each physical ATM interface.

ARIs and AVIs, see Figure 5 on page 14, provide identical routing functions. The routing protocols differentiate between interfaces by the interface number that is assigned at the time the interface is defined.

Note: MSS Release 1.1 does not support bridging over ATM virtual interfaces.

MSS assigns the interface numbers for the real and virtual (and also LAN emulation) interfaces sequentially and the interface number assigned depends on the order in which you configure your interface. A real interface must be defined prior to configuring a virtual interface associated with the real interface. Currently, the total number of interfaces (real, virtual, and LAN emulation) you can define on a single MSS server is 253. Be aware that the actual number of AVIs that can be configured is limited by physical resources, such as the amount of memory/buffers that are available.

Protocols can be configured on each ATM interface (real or virtual) independent of other interfaces. Hence, for example, one can configure IP on interface 0 (which is a real ATM interface) with IP address 9.1.1.1 and another instance of IP with address 9.2.1.1 on interface 1 (which is an AVI). Whether an interface is a real ATM interface or a virtual interface configured on a real interface makes no difference to IP. In addition, in the case above, whether virtual interface 1 is configured on top of real ATM interface 0 or some other physical ATM interface is also transparent to the protocols.

4.2.2 Benefits of Using ATM Virtual Interfaces

The key advantage of using the ATM virtual interface feature is the improved scalability in terms of the number of protocol instances that can be supported on a physical ATM interface.

The increase in number of protocol instances leads to the following advantages:

1. Increase of Classical IP subnet attachments

The limitation of having 32 Classical IP subnets attachment per virtual or real interface remains; however, because you can define multiple virtual interfaces the total number of Classical IP subnets increases considerably.

2. Enhanced IP multicast support

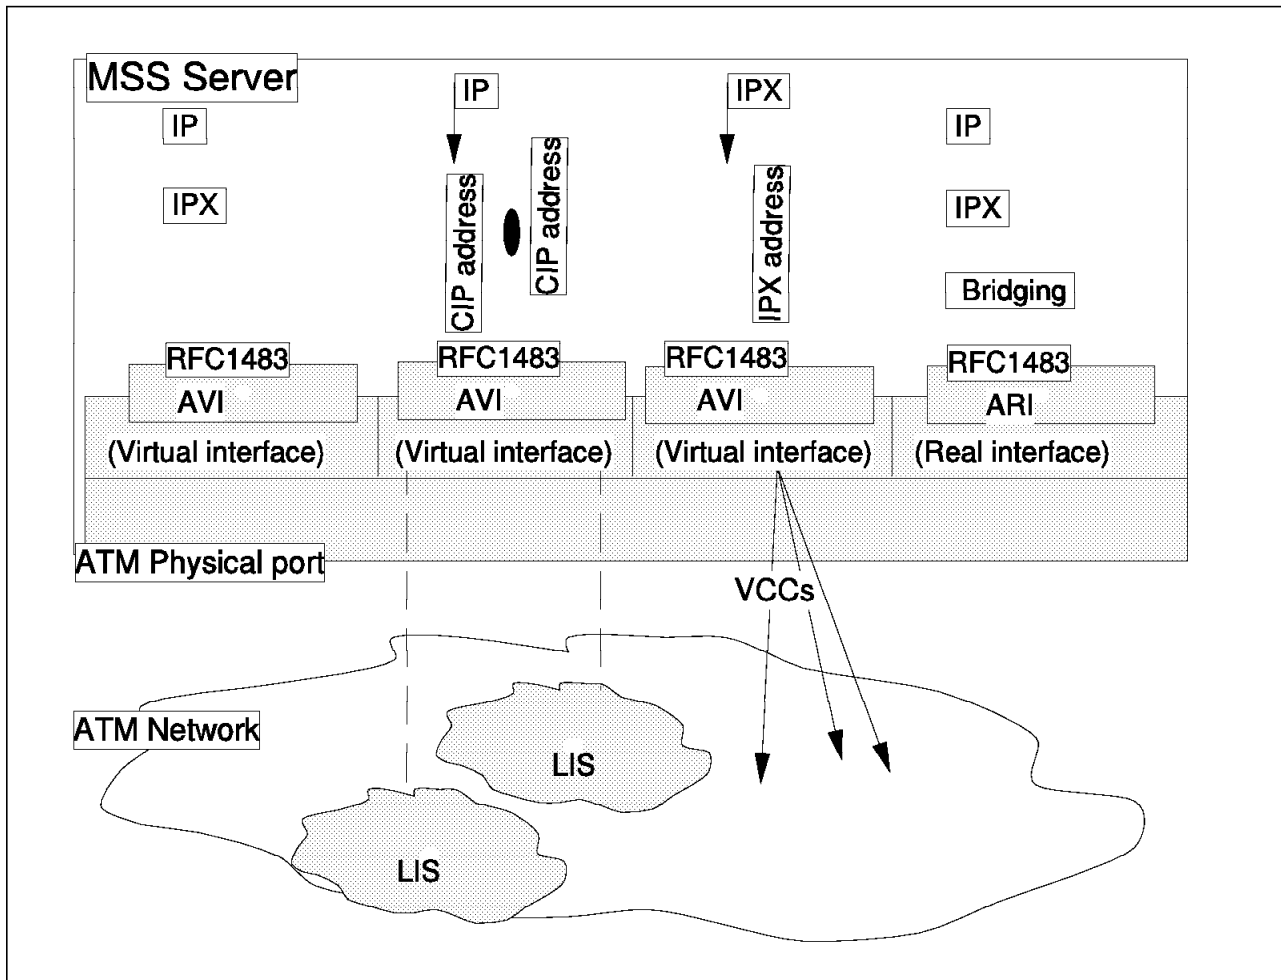


Figure 5. ATM Real and Virtual Interfaces

The ATM Virtual Interface feature is crucial for supporting multicast routing protocols (such as MOSPF and DVMRP) when using Classical IP. By using a different, real or virtual, interface for each of your Classical IP subnets multicast messages on one subnet will be properly forwarded to all other subnets.

Note: An alternative to the network-layer multicast is to use level 2 IP multicast. ATM is a non-broadcast media and requires address resolution support for layer 2 forwarding of IP multicast packets. The IP multicast packets must be mapped to ATM's point to multipoint connection service. This requires that ATM based IP multicast hosts and routers use a multicast address resolution server (MARS), see *RFC 2022, Support for Multicast over UNI 3.0/3.1 based ATM Networks*. Support for RFC 2022 can be expected in a future release of MSS.

3. Increase of RFC 1483 IPX addresses

Rather than using a single real interface for your RFC 1483 IPX connections, you can now use a mixture of real and virtual interfaces. Because you can define multiple virtual interfaces per real interface the number of IPX networks (one per, real or virtual, interface) increases considerably. per, real or virtual, interface.

4. Interference of bridging and routing can be avoided

By using a different interface for your bridging and routing functions the interference between routing and bridging functions described in item 4 on page 13, can be avoided.

An important point to note is that AVIs do not have any physical resources of their own. Hence, each virtual interface will not have as many virtual connections (VCs) as a single physical interface. Rather, the available resources (in this example VCs) will get partitioned between the different virtual interfaces configured on a single real interface. In the current implementation, there is no explicit allocation of physical resources to each virtual interface and the resource allocation is on demand. Each physical ATM interface has a pool of resources which are available for use by all AVIs and the single ARI itself.

The use of AVIs can also lead to a reduction in packet fragmentation and re-assembly. When multiple subnets are configured on a single physical ATM interface, the interface will have to reduce the maximum transmission unit or MTU (the maximum packet size that can be sent or received over that interface) to the minimum of the MTUs of all subnets sharing the same interface. However, if multiple AVIs are created on that ARI and each IP subnet is configured on a different AVI, every subnet can continue to use its existing MTU size without getting affected by other subnets configured on the same physical ATM interface. This avoids possible reductions in throughputs and delays due to packet fragmentation and re-assembly when MTU size is reduced. A second performance improvement can be achieved by distributing the number of protocol addresses configured on a physical interface over different virtual interfaces configured on the same physical interface.

4.3 IBM-Compliant LAN Emulation Client

— ATM Forum-compliant versus IBM-compliant LAN Emulation —

IBM considers the use of ATM Forum-compliant LAN emulation more strategic and recommends its use. The main reason to introduce the IBM-compliant LAN emulation client functions is to simplify the migration of networks (see Figure 7 on page 17) using IBM-compliant LAN emulation to ATM Forum-compliant LAN emulation. IBM provides ATM Forum-compliant LAN emulation drivers for most of its adapters.

As a pioneer of ATM technology, IBM defined one of the early LAN Emulation architectures. Although IBM submitted this architecture to the ATM Forum, the final ATM Forum LAN Emulation Specification diverged in a number of ways. Consequently, IBM, like several other vendors, now has products that implement two different LAN Emulation architectures: some earlier products implement the IBM LAN Emulation Architecture, while other, more recently developed products, implement the ATM Forum LAN Emulation specification.

Although ATM Forum LAN Emulation is the strategic direction for multi-vendor interoperability, IBM continues to support its customers that have installed products implementing the IBM LAN Emulation Architecture. These customers will generally fall into one of the following two categories: (1) customers who wish to continue using IBM LAN Emulation, which offers some technical advantages relative to ATM Forum LAN Emulation, or (2) customers who wish to migrate to ATM Forum LAN Emulation. By adding a client for the IBM LAN

Emulation Architecture, the MSS Server can provide routing and bridging services that support both purposes.

4.3.1 ILEC Routing and Bridging Overview

From routing and bridging perspectives, IBM LECs (ILECs) are functionally equivalent to ATM Forum LECs; both provide emulated Ethernet or token-ring interfaces with operational characteristics of real interfaces. Thus, all the (IP, IPX, and AppleTalk) routing and bridging functions previously described for ATM Forum LECs also apply to IBM LECs.

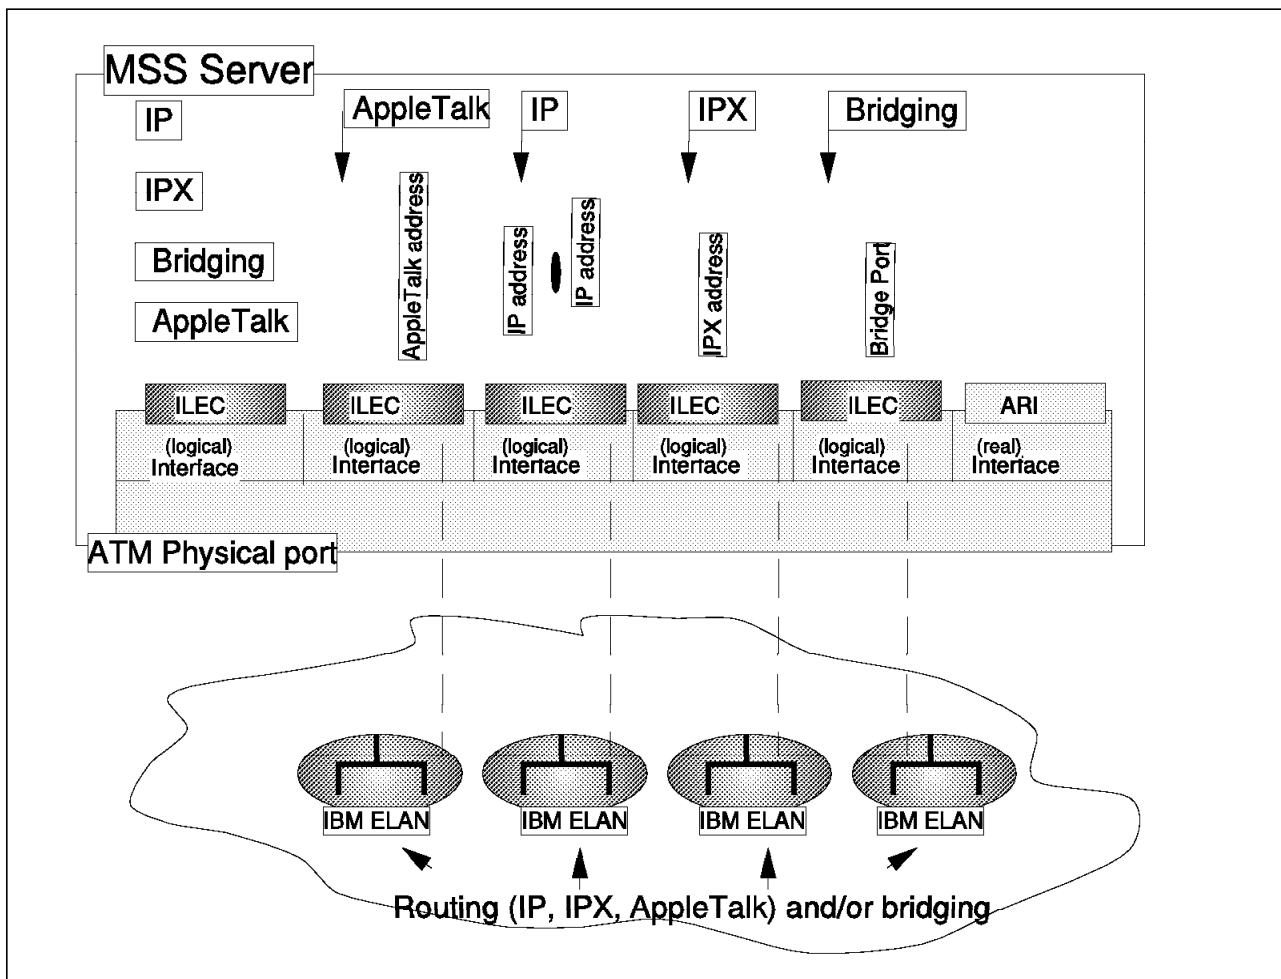


Figure 6. IBM-Compliant LAN Emulation

Figure 6 illustrates some of these capabilities, which include:

1. Bridging between IBM ELANs
2. Bridging between IBM ELANs and bridged connections using RFC 1483 encapsulation on ATM real or virtual interfaces
3. Bridging between IBM ELANs and ATM Forum ELANs
4. IP, IPX, and AppleTalk routing between IBM ELANs
5. IP, IPX, and AppleTalk routing between IBM ELANs and ATM Forum ELANs
6. IP, IPX, and AppleTalk routing between IBM ELANs and FDDI (IBM 8210 only)

7. IP and IPX routing between IBM ELANs and IP and IPX using RFC 1483 encapsulation on ATM real or virtual interfaces

Each IBM LAN emulation client is associated with a real interface and all traffic to and from the LE client will use this interface only. The real interface has to be defined prior to configuring the LE client.

All bridging and routing functions can be combined on a single LE client or you can use multiple LE clients. Be aware of the consequences when defining routing and bridging functions on the same interface (see also 6.4, "Concurrent Bridging and Routing" on page 32).

The bridging and routing protocols interact with the LE client through a logical interface. Each (real, virtual, and logical) interface has a unique interface number. MSS assigns the interface numbers in sequence starting with zero for the first real interface defined. When defining an LE client the next numeric value available is used.

4.3.2 Benefits of ILECs

As described in the preceding section, adding ILEC support to the MSS Server provides routing, bridging, and migration support for customers who are currently running IBM LAN Emulation (see Figure 7). By having an IBM ELAN and an ATM Forum ELAN operational at the same time, and providing bridging functions between the ELANs end stations can gradually migrate to the ATM Forum ELANs.

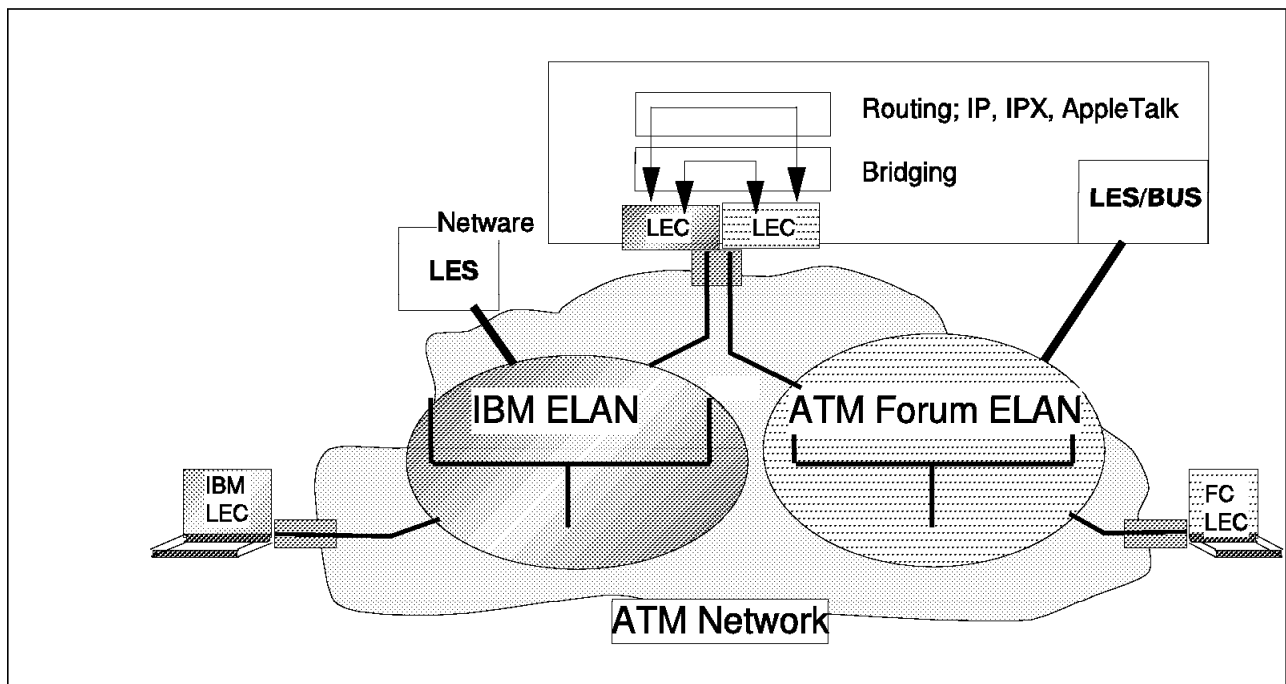


Figure 7. Migrating IBM-compliant to ATM Forum LAN Emulation

Section 4.3.4, "Quick Overview of IBM-Compliant LAN Emulation" on page 18 provides an overview of IBM-compliant LAN Emulation and highlights the main differences between ATM Forum and IBM-compliant LAN emulation. A more detailed description is given in *SG24-5002 - ATM Campus Network Design*.

4.3.3 Limitations of ILECs

Table 1 summarizes ILEC capacity constraints for three different hardware configurations. Upper bounds are identified for the number of ILEC and ILEC bridge instances. The source of the constraint, either 8260 point-to-multiPoint (PtMP) VCCs or MSS Server design, is listed with each upper bound. Each capacity characteristic was analyzed independently, and all the maximums cannot be supported simultaneously. The analysis assumes that all of the 8260's PtMP VCC resources are devoted to this function.

Table 1. ILEC Capacity Constraints			
Capacity Characteristic	Hardware Configuration		
	(1)	(2)	(3)
Internal ILECs	127 8260 PtMP VCC	252 8260 PtMP VCC	252 MSS Server Design
Internal Bridge ILECs	63 8260 PtMP VCC	126 8260 PtMP VCC	252 MSS Server Design
Note: Hardware Configurations			
1. MSS Server blade, or 8210 with single ATM adapter, connected to 8260 ATM port			
2. 8210 with two ATM adapters connected to two different 8260s			
3. 8210 connected to ATM switch network with no VCC resource constraints			

4.3.4 Quick Overview of IBM-Compliant LAN Emulation

IBM-compliant LAN emulation supports the following Ethernet and token-ring frame types:

- Ethernet (Version 2 (DIX) and IEEE 802.3)
- IEEE 802.5

Note: The data transported on the emulated LAN is transparent to the LE emulation layer.

The client-server relation within IBM-compliant LAN emulation is depicted in the upper part of Figure 8 on page 19. The lower part of the figure indicates the VCCs required to maintain client-server connectivity and establish client-client connectivity.

The LES can be implemented in a stand-alone system or can be incorporated into a network switch. The LES provides initialization, registration, address resolution, and data forwarding services to the LECs. It does this using a modified version of the UNI interface called LAN Emulation UNI or LUNI.

MSS

MSS provides IBM LAN emulation client functions only. For the LES function you need a NetWare server.

During the *initialization* phase LE clients establish a *default VCC* with the LE server. To do so, they must know the ATM address of the LE server. In the current IBM products, this address needs to be manually configured. The VPI and VCI used will depend on the VPs and VCCs available at call setup. After the connection to the LE server has been established, the LE clients enter the registration phase.

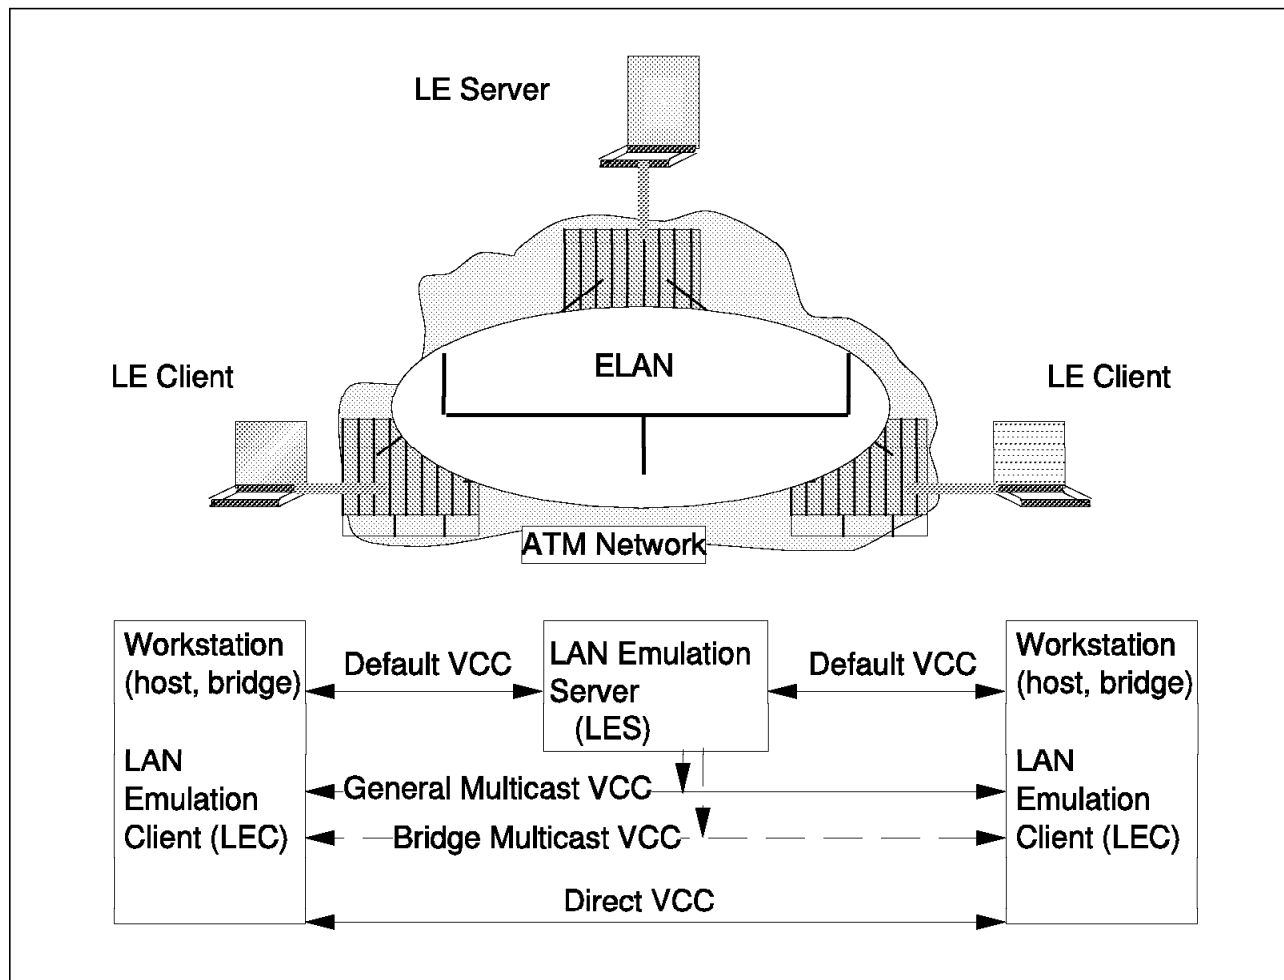


Figure 8. IBM-Compliant LAN Emulation Components

LE clients have to register themselves for every emulated LAN they want to participate in. Only after an LE client has been registered, can data transfer be initiated.

4.3.4.1 Comparing IBM and ATM Forum-Compliant LAN Emulation

The main difference between IBM and ATM Forum-compliant LAN emulation, are the LAN Emulation Connection Server (LECS) and a Broadcast and Unknown Server (BUS). The BUS which is identified within the ATM Forum specification as a separate entity, is, of course, also present within IBM-compliant LAN emulation but fully integrated within the LE server. Because of the BUS function, an ATM Forum ELAN requires more VCCs.

Within IBM-compliant LAN emulation there is no LECS. The ATM Forum LECS concept may simplify the management of ELANs significantly. Learning server addresses dynamically (LECS via ILMI, LES from LECS, BUS from LES), relieves LE clients from hard-coding server addresses, simplifies LE server backup, and lays the foundation for LE server distribution.

Note: ATM Forum-compliant LAN emulation implementations, may decide not to implement LECS at all.

The number of required VCCs is much lower in the IBM ELAN implementation than in the ATM Forum ELAN. The component that requires the highest number of VCCs is the shared LES/BUS for the ATM Forum ELAN.

Chapter 5. RFC 1483 Bridging over ATM

The MSS Server provides a wide range of bridging capabilities including: transparent bridging, source-route bridging, SRT bridging, and SR-TB translation.

5.1 RFC 1483 Bridging over ATM with MSS 1.0 and 1.1

In Release 1.0, the MSS Server supported bridging over emulated Ethernet and token-ring interfaces. In Release 1.1, support has been added for RFC 1483 bridge ports as illustrated in Figure 9. The emulated LANs can be ATM Forum or IBM-compliant. Bridging is supported between any of the emulated LANs and the RFC 1483 bridged connections.

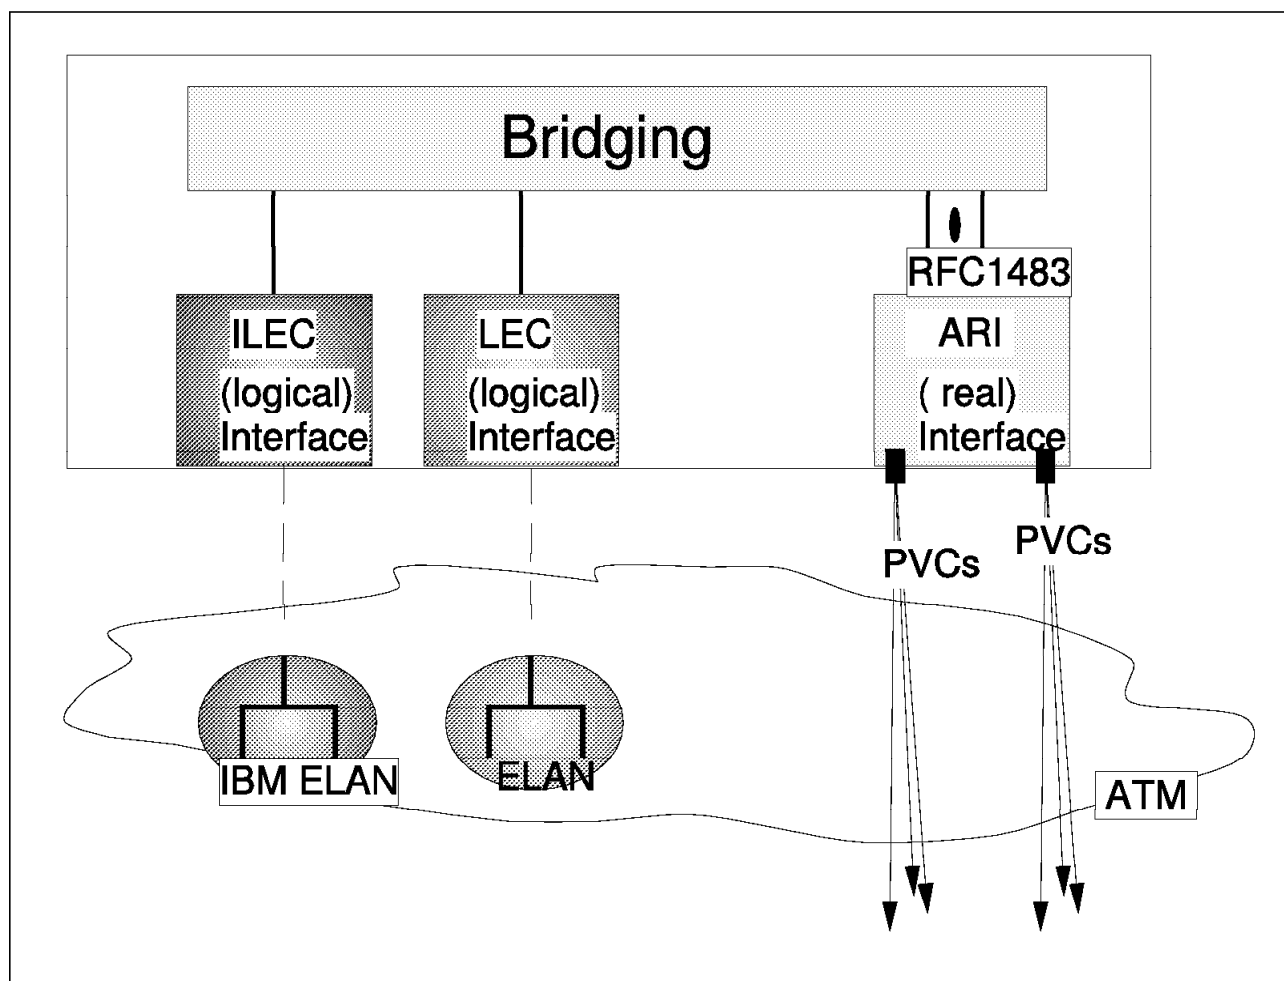


Figure 9. Bridging Using LAN Emulation and RFC 1483 PVCs

RFC 1483, multiprotocol encapsulation over ATM adaptation layer 5 (AAL5), specifies encapsulations for both bridged and routed format frames. The bridged format carries MAC frames, while the routed format does not include a MAC header. Encapsulated MAC frames received on an RFC 1483 port are passed to the bridge as if they had been received on a LAN (or emulated LAN) interface. Similarly, the bridge uses the same criteria to forward frames out both RFC 1483 ports and LAN ports (that is, in accordance with the bridging protocol/configuration in effect).

The main difference in configuring a bridge port on an ATM interface versus a LAN interface is that a PVC to be associated with the port must be specified in the ATM case. The PVC is specified via a VPI/VCI combination that must be unique across all RFC 1483 bridge ports defined on the ATM interface. Another difference is that multiple bridge ports may be defined on a single ATM attachment. MSS R1.1 supports RFC 1483 bridged connection on ATM real interfaces (ARIs). ATM virtual interfaces (AVIs) can be used for IP and IPX routing only. For configuration details see 5.1.1, "Configuration Issues" on page 25.

Once a bridging port has been configured on an ATM interface, it functions as a normal bridging port on a LAN, and the association of the port with an ATM interface PVC is transparent to the bridging function. Thus, all the normal bridging configuration options are applicable.

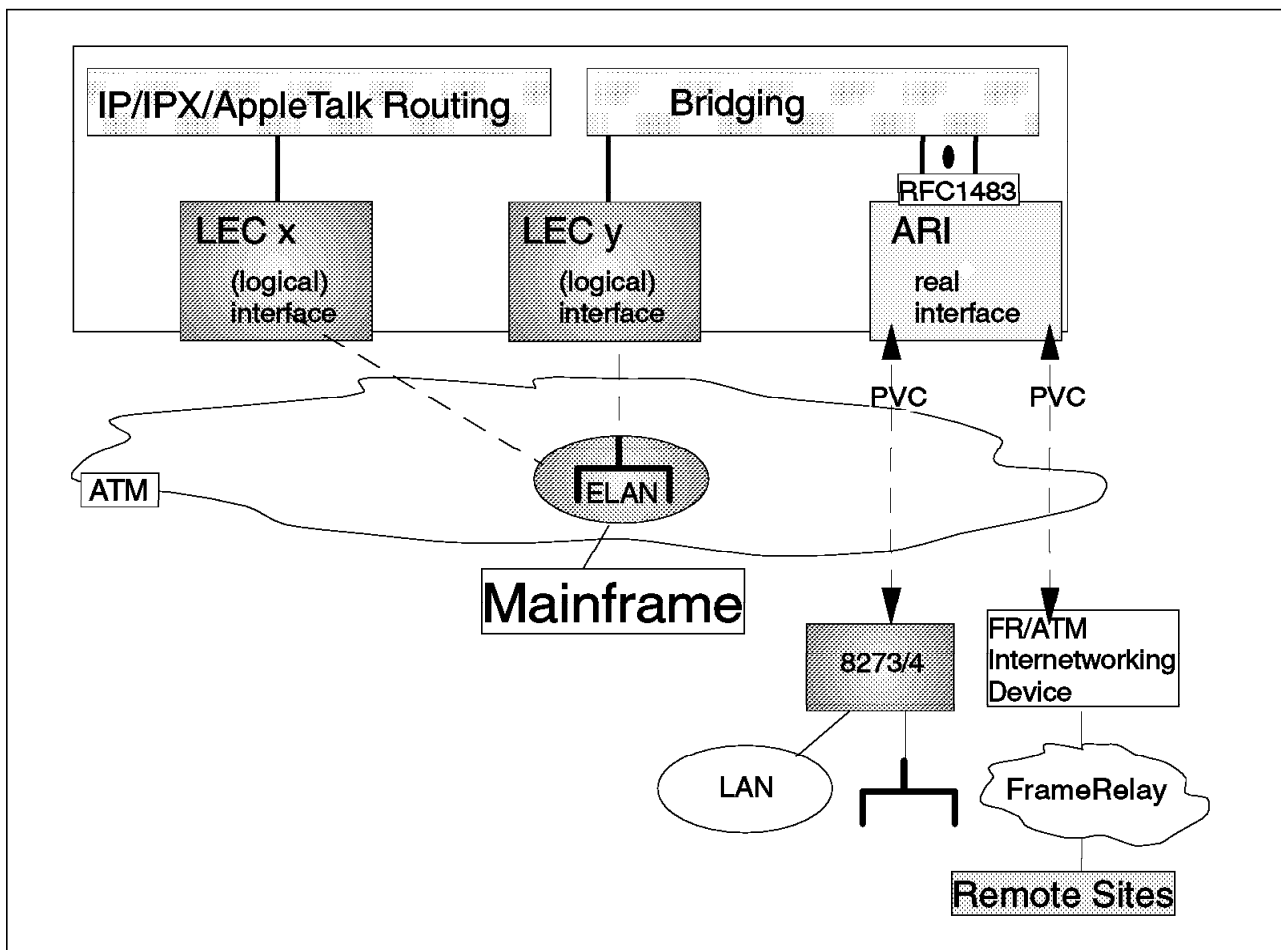


Figure 10. MSS Server Support for Frame Relay/ATM Service Interworking

RFC 1483 bridging support is important because it enhances the internetworking capabilities of the MSS Server, and thereby provides opportunities for the MSS Server to play a role in additional networking environments. Two example are depicted in Figure 10.

An RFC 1483 PVC can be defined to connect to an IBM 8274 RouteSwitch. As the MSS Bridging Broadcast Management and Dynamic Protocol Filtering (see Chapter 10, "SuperVLANs" on page 55) can be used on any bridge port the VLAN capabilities of the IBM 8274 and MSS can be shared.

Frame relay and ATM service interworking is another example. Figure 10 depicts a network where remote sites are connected to an ATM campus network via frame relay. For this example, assume that stations on remote LANs are communicating with mainframe servers using SNA. The conversion between frame relay and ATM is performed by the frame relay to ATM interworking switch at the edge of the campus network (using one-to-one mappings between frame relay DLCIs and ATM VPI/VCIs). The frame relay to ATM interworking switch converts between RFC 1490 encapsulation on the frame relay side and RFC 1483 encapsulation on the ATM side, while the MSS Server bridges RFC 1483 frames to/from the ELAN(s) that the mainframe servers are members of.

Note: There is an OC-3 ATM/FR Interworking module for the IBM 2225. Use feature number #8064 for the multi-mode fiber (MMF) module and feature number #8065 for single mode fiber (SMF) module. There are also redundant versions of the two modules.

The key points of this example are that frame relay is very prevalent in wide-area networks and that a PVC carrying RFC 1483 encapsulated frames may be the only ATM interface supported by the Frame Relay/ATM Interworking Switch. It is also important to point out that the Bridging Broadcast Manager (BBCM) function described in 10.2.2, "Bridge Broadcast Management (BBCM)" on page 66 is applicable to RFC 1483 ports, since BBCM makes the network configuration of Figure 10 on page 22 viable for broadcast-oriented protocols, like IP, by directing broadcasts to local servers (consolidated in the data center) instead of needlessly flooding low-speed WAN links.

In addition to the ELAN bridging function, the MSS Server can also provide routing services for stations supported by RFC 1483 PVCs. In the case of bridged format frames (where there is no downstream router), the MSS Server routes to/from an ELAN that is bridged to the RFC 1483 PVC as illustrated in Figure 11 on page 24.

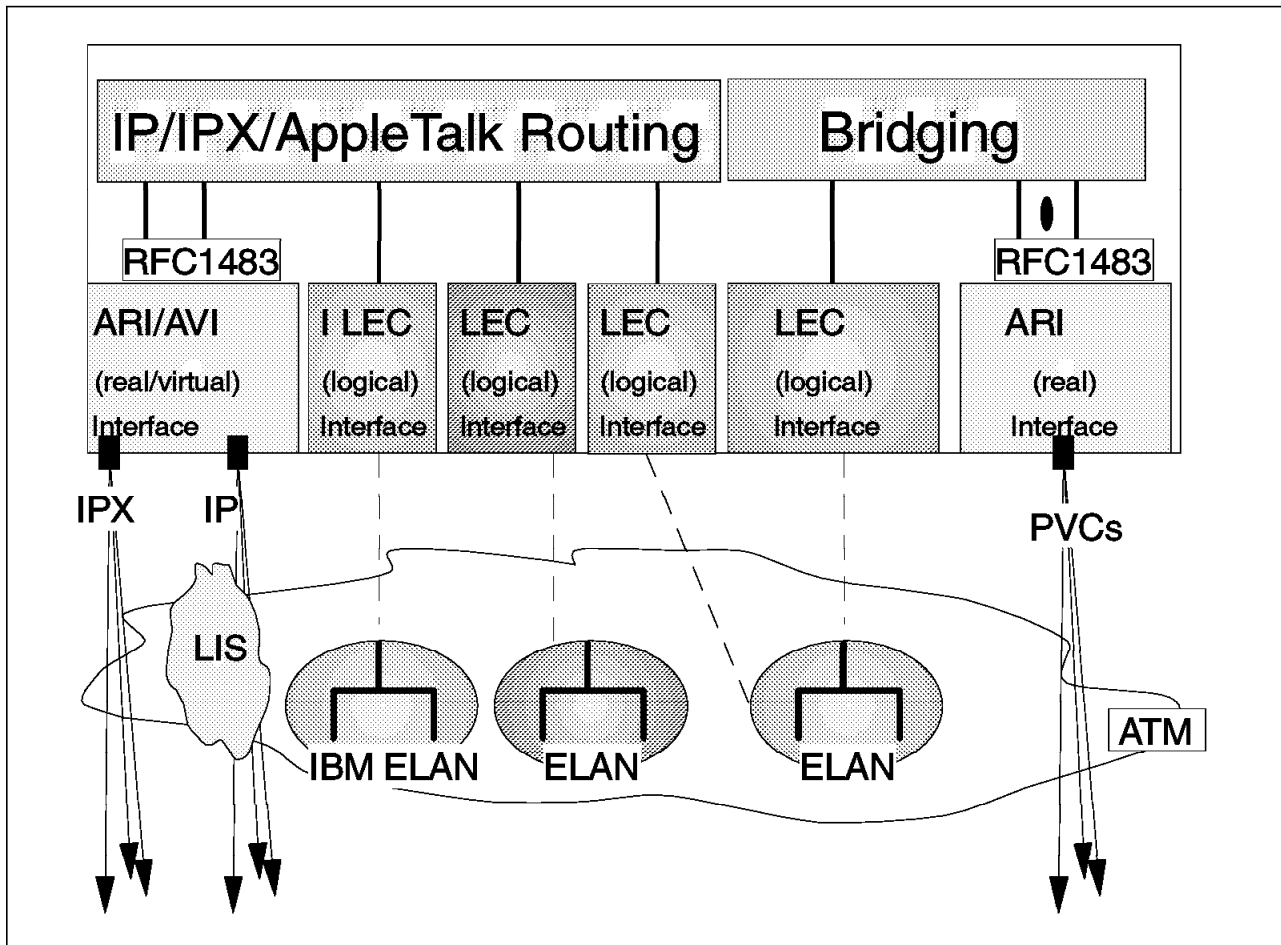


Figure 11. MSS Server Routing Support for RFC 1483 Bridged Format Frames

In the case of RFC 1483 routed format frames, the MSS Server provides native routing services for IP and IPX. Furthermore, the bridged and routed traffic may share a single RFC 1483 PVC, which is useful in multiprotocol environments, like the one shown in Figure 12 on page 25, where some protocols are bridged and others are routed (see 15.1, "Sharing of RFC 1483 PVCs" on page 129 for more information on RFC 1483 PVC sharing).

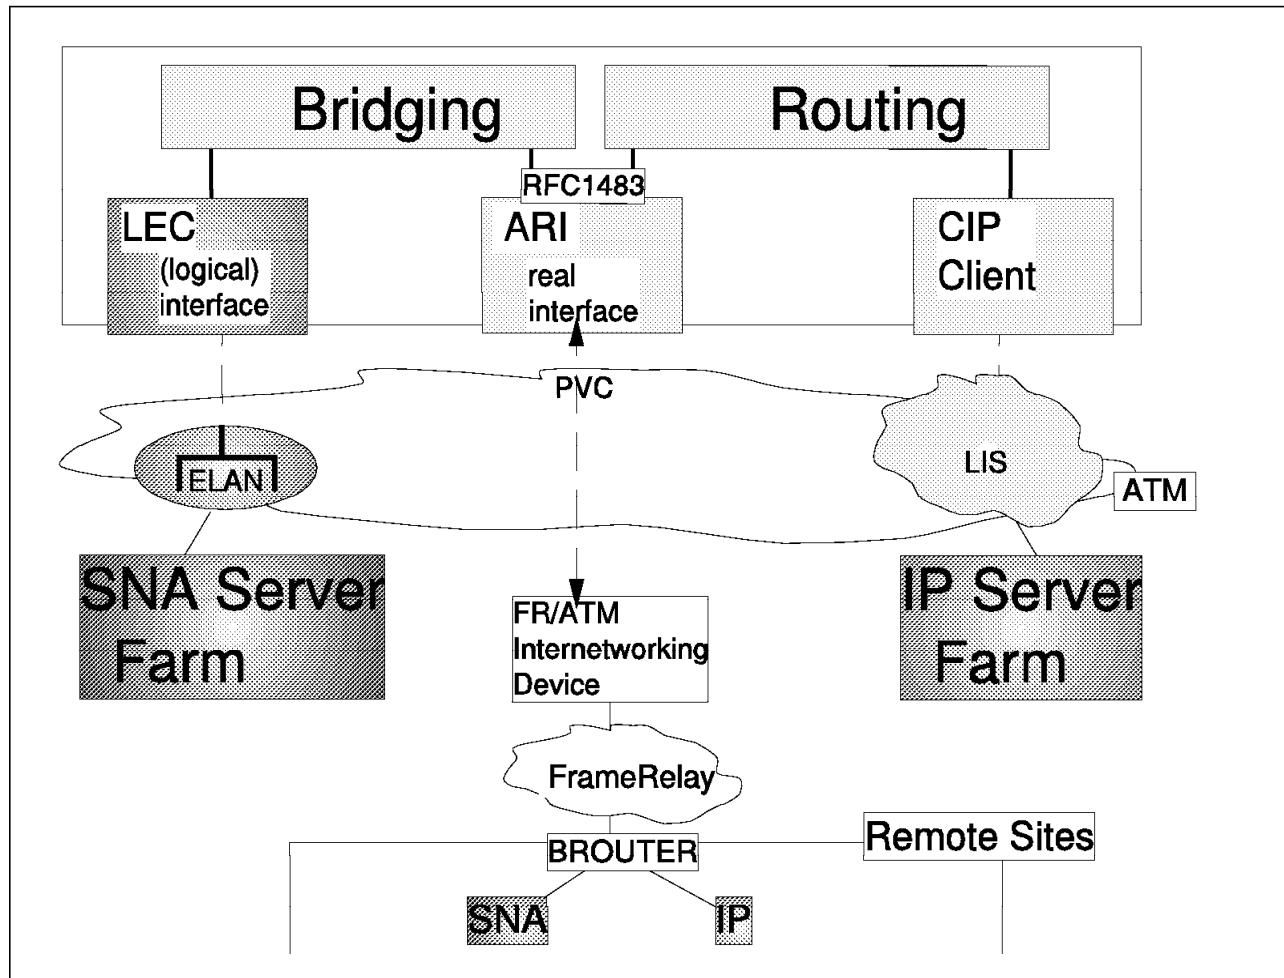


Figure 12. Multiplexing Support for RFC 1483 Bridged and Routed Format Frames

As a final point, note that in Release 1.0 of the MSS Server, the bridging versus routing decision was made on a global basis, where a given protocol was routed (and not bridged) on all interfaces if an address for that protocol was defined on any interface. In Release 1.1 of the MSS Server, this has been changed such that the bridging versus routing decision is made on a local (that is, port) basis. With this change, a protocol is bridged unless an address for that protocol has been defined on the bridge port interface. Thus, if an RFC 1483 encapsulated protocol is to be both bridged and routed, the protocol address and bridge port must be associated with different ATM interfaces (one or both of which may be virtual ATM interfaces).

5.1.1 Configuration Issues

When defining RFC 1483 bridging ports be aware of the following:

- ### 1. Interaction between bridging and routing function

RFC 1483 bridging ports have to be defined on ATM real interfaces. Defining IP and/or IPX addresses on the same interface effectively disables the bridging function for IP and/or IPX on the shared interface. Therefore, it may be considered to use a virtual ATM interface for the RFC 1483 routing function (see 4.2, “ATM Virtual Interfaces (AVIs)” on page 11).

- ## 2. Spanning tree

The current MSS Server runs a single spanning tree protocol for all its bridge ports and, by default, MSS uses the MAC address of the lowest numbered port as the bridge ID. As RFC 1483 bridge ports do not have MAC addresses; the bridge ID must be explicitly configured in the case that the RFC 1483 bridged port is the lowest numbered bridge port. Use the **set bridge** line command for this purpose.

3. Number of bridging ports

By using multiple RFC 1483 or ELAN connections one can define up to 254 bridge ports. A maximum of 32 bridging ports may be configured on a single ATM real interface.

4. PVCs only

RFC 1483 bridging ports in MSS V1.1 can be defined using PVCs only.

5.2 RFC 1483 Bridging over ATM with MSS 2.0

In MSS Version 2.0, RFC 1483 bridging operates as in MSS 1.1 except that it supports SVCs; this may allow for less definitions on some ATM switches when there are multiple switches along the path between two MSSs. MSS uses signalling to find the destination ATM address of the partner RFC 1483 bridge. Using RFC 1483 bridge ports with SVCs requires defining the partner bridge destination ATM addresses instead of the VPI/VCI. It is also necessary to configure the ESI and selector used by the local RFC 1483 bridge port. The local selector cannot be generated at runtime, since the address is predefined in the partner RFC 1483 bridge.

An MSS RFC 1483 bridge port will remain disabled until the SVC associated is set up and active. After the SVC is set up, the bridge port will progress through appropriate initialization (blocking, learning and forwarding). If the SVC is torn down for any reason, the port moves to a state of postconfiguring. All RFC 1483 SVC ports in this state are retried at fixed intervals, and the port will reactivate when the SVC is set up again.

Each of the two MSSs bridging over an RFC 1483 link, sets up a call with the other, and transmits data to the destination port. There are two channels then between two MSSs on an RFC 1483 link. Each RFC 1483 bridge port only transmits over 1 channel and only receives on the other.

If a channel associated with a bridge port is deleted (T 5, p arp, delete), the RFC 1483 bridge port will stay disabled until the MSS is reset.

Chapter 6. MSS R1.1 - Bridging and Routing

This chapter briefly overviews the bridging and routing functions of MSS Server Release 1.1. Before that we introduce the concept of interfaces.

6.1 Real, Virtual and Logical Interfaces

Depending on the type of MSS Server (module or stand-alone) the MSS Server can be equipped with up to two physical ATM ports, or one ATM port and an FDDI attachment.

The bridging and routing functions of the MSS Server do not interact directly with the physical ports but instead the concept of *interface* is used.

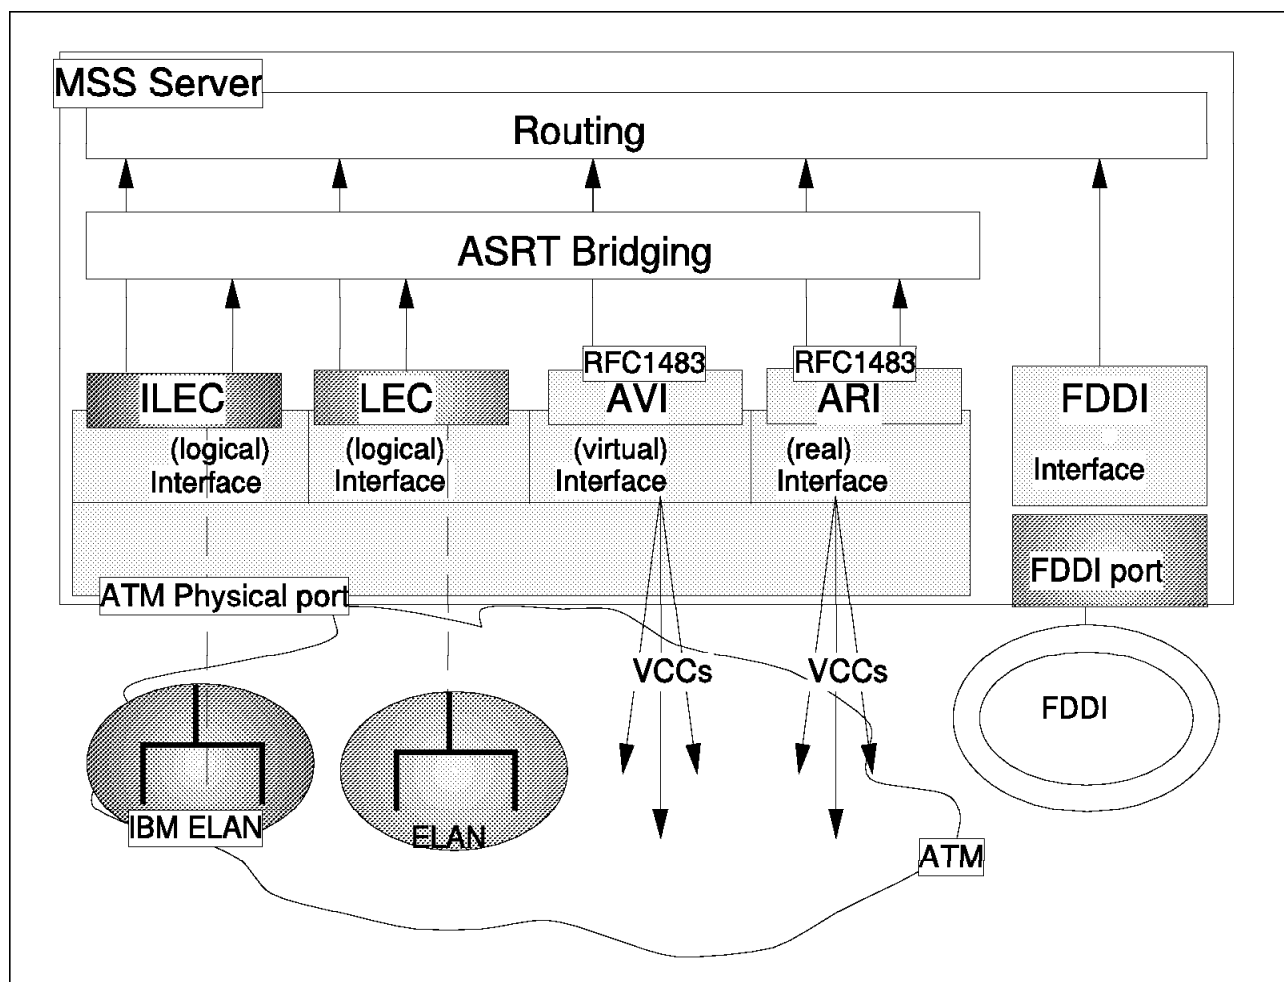


Figure 13. Types of Interfaces

Figure 13 overviews the type of interfaces possible:

1. FDDI real interface

When adding an FDDI attachment an FDDI real interface results. An FDDI interface is supported on the IBM 8210 only (for details see 4.1, "FDDI Interface" on page 11).

2. ATM real interface (ARI)

When adding an ATM attachment an ATM real interface (ARI) results. The IBM 8210 supports up to two ATM ports. The MSS Server module supports a single ATM port.

3. ATM virtual interface (AVI)

The concept of ATM virtual interface (ARI) has been introduced to increase the number of routing ports and improve IP multicast support. For details see 4.2, "ATM Virtual Interfaces (AVIs)" on page 11.

Each virtual interface is associated with a real ATM attachment and all routing traffic using the virtual interface will flow over it. For the routing functions there is no distinction between an ATM real and virtual interface.

Note: MSS R1.1 does not support bridging over AVIs.

4. Logical interface

For each LAN emulation client defined a logical interface results. LAN emulation clients are either IBM-compliant (ILECs) or ATM Forum-compliant. For the routing and bridging functions there is no distinction between either type of logical interface. Each logical interface is associated with a real ATM attachment and all routing and bridging traffic using the logical interface will flow over it.

The bridging and routing functions distinguish between the various interfaces by the interface number that is assigned to it. The MSS Server assigns the interface number sequentially when defining an interface, starting with zero¹. The maximum number of (FDDI, ATM real, ATM virtual and logical) interfaces that can be defined is 253.

6.2 Bridging Overview

The MSS Server is a single group any-to-any port bridge. It supports source-route bridging (SRB), transparent (TB), source-route transparent (SRT), and source-route translational (SR-TB) bridging. One or multiple logical bridge ports can be defined on any of the interfaces. Based on the attributes defined on the bridge ports the MSS Server dynamically decides, therefore, adaptive source-route transparent bridging (ASRT), what type of bridge behavior is required.

The interfaces on which bridging ports can be defined are:

1. ATM real interfaces (ARIs)

Bridge ports defined on ARIs use RFC 1483 encapsulation (see Chapter 5, "RFC 1483 Bridging over ATM" on page 21). Only PVCs can be defined. Per ARI you can define up to 32 bridge ports. Each bridge port can be defined as either an Ethernet or a token-ring bridge port. The type of bridge attributes (TB, SRB, SRT) that can be assigned per port depend on the type of port.

2. Logical interfaces

On each logical interface (that is, LE client) you can define one bridge port. The type of bridge attributes that can be assigned to this port depend on the type of emulated LAN attachment (token-ring or Ethernet).

¹ As a consequence the interface number of the ATM real interface of the first ATM attachment added is always zero.

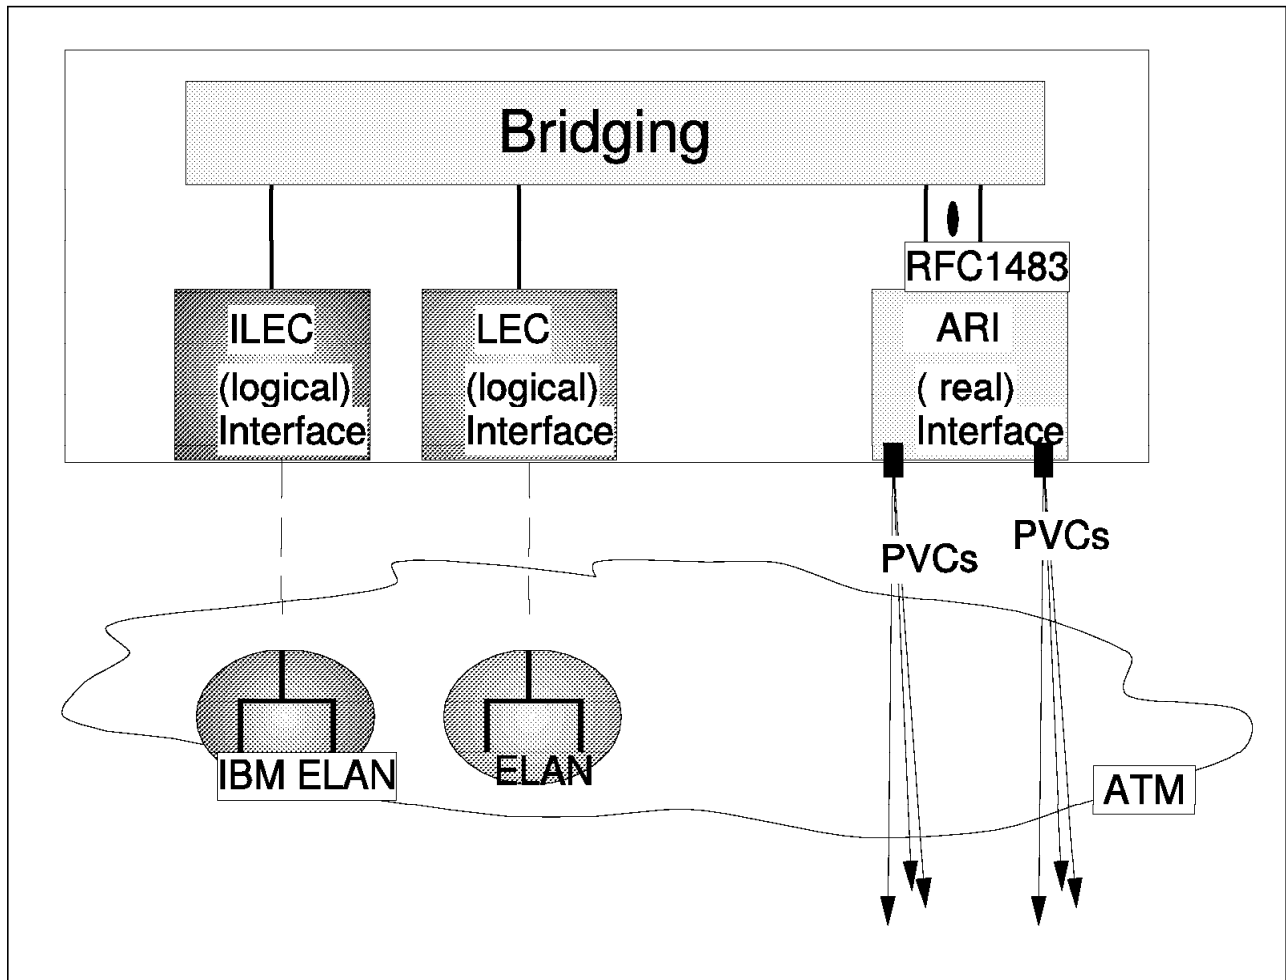


Figure 14. MSS R1.1 - Bridging Overview

Note: There is no bridging support for the FDDI interface.

The number of bridging ports per MSS Server is 254.

6.2.1 Protocol Support

The TB, SRB and SRT bridging functions are layer three protocol-independent. Bridging for networking protocols such as IP, IPX, NetBIOS, SNA, Banyan Vines, AppleTalk, etc. are all supported. Translational bridging (SR-TB), however, is network protocol dependent. The MSS Server provides support for NetBIOS, SNA, and IP. SR-TB translation support for IP is new in Release 1.1 and is discussed in 15.6, "SR-TB Translational Bridging Support For IP" on page 132.

6.3 Routing Overview

The MSS Server provides routing functions for IP, IPX, and AppleTalk. The next section gives a short overview of the functions for each of the protocols. An overview is depicted in Figure 15 on page 30. Further information can be found in the *Multiprotocol Switched Services (MSS) Server Command Line Interface Volume 2*.

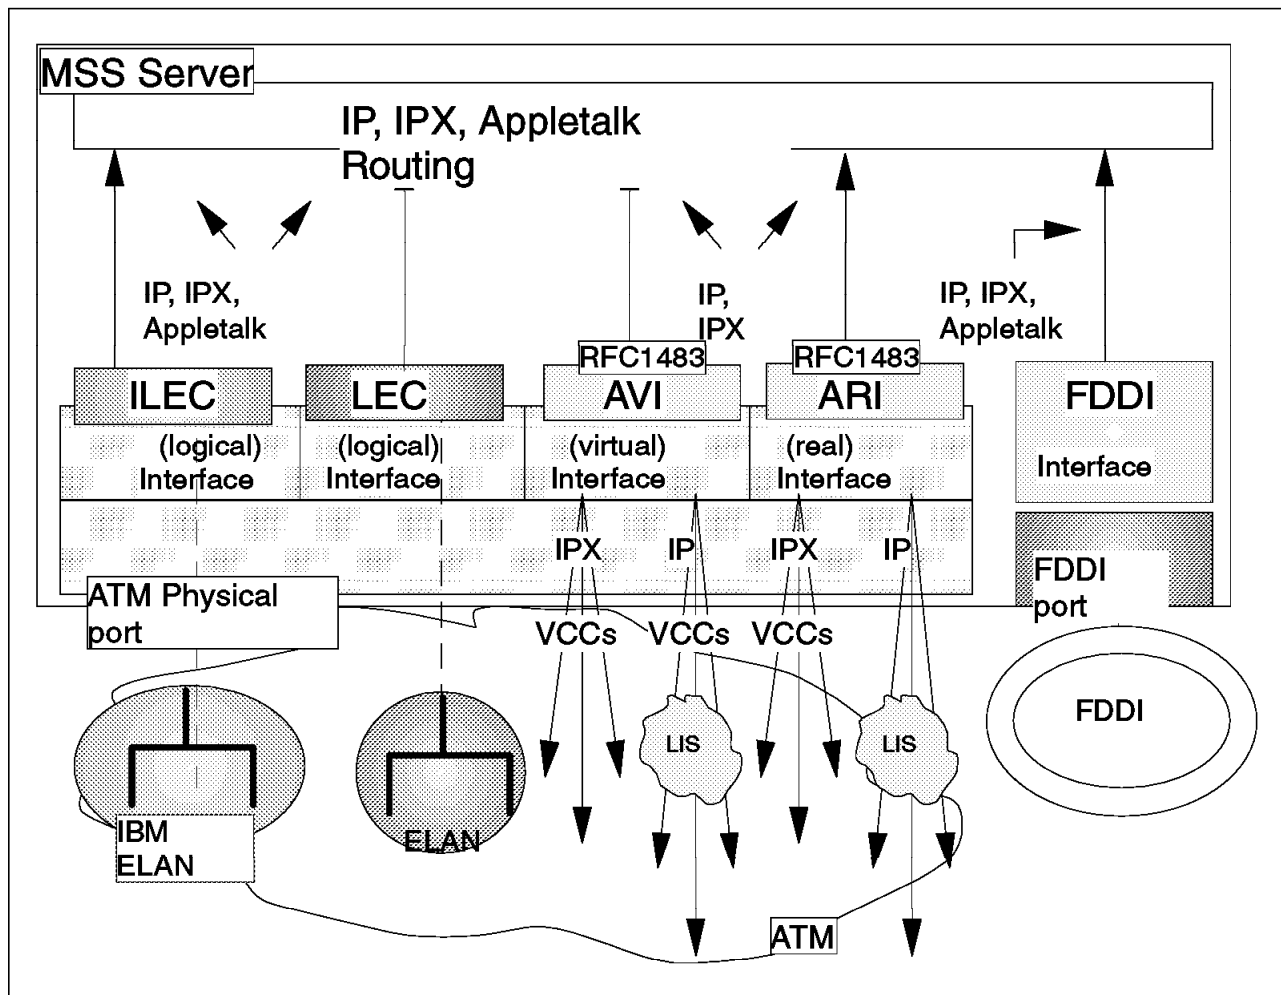


Figure 15. MSS R1.1 - Routing Overview

6.3.1 IP Routing

The interfaces that can be used for IP routing are:

1. FDDI interface

The IP routing over FDDI is standard-compliant.

2. ATM real and virtual interfaces (ARI/AVI)

IP routing on ARIs and AVIs uses RFC 1483 encapsulation and is referred to as Classical IP. Per AVI/ARI you can define up to 32 ARP servers and/or clients connecting to up to 32 logical IP subnets (LISes) per interface. Both PVC and SVC support is provided. The SVCs can be set up dynamically, using an ARP server, or can be pre-configured. PVC connections need to be pre-configured.

3. Logical interfaces

The IP routing makes no distinction between being connected between an IBM or an ATM-compliant ELAN and is independent of the type of emulated LAN attachment (token-ring or Ethernet).

Per interface you can define up to 32 IP addresses each connecting to a different IP subnet. The MSS Server provides IP routing between any of the IP subnets to which it connects.

IP support on the MSS Server is standard-compliant. It is important to realize that the support for Classical IP is similar to the support for conventional IP routers on X.25 or frame-relay networks while the ELAN functions are essentially the same as used by conventional IP routers on legacy LANs.

The MSS Server can use RIP, OSPF Version 2, and/or BGP Version 4 to dynamically exchange routing information. RIP, however, is not supported on Classical IP subnets. Support is provided for the multicast extensions to OSPF (MOSPF).

The MSS Server provides extensive support for IP filtering.

6.3.2 IPX Routing

The interfaces that can be used for IPX routing are:

1. FDDI interface

The IPX routing over FDDI is standard-compliant.

2. ATM real and virtual interfaces (ARI/AVI)

IPX routing on ARIs and AVIs uses RFC 1483 encapsulation. Both PVC and SVC support is provided. Contrary to Classical IP, IPX using RFC 1483 does not provide an ARP server function, and all VCCs need to be pre-configured. Both SVCs and PVCs are supported.

3. Logical interfaces

The IPX routing makes no distinction between being connected between an IBM or an ATM-compliant ELAN and is independent of the type of emulated LAN attachment (token-ring or Ethernet).

Per interface you define a single IPX network number. The MSS Server provides IPX routing between any of the IPX networks it connects to.

IPX support on the MSS Server is compatible with all previous Novell NetWare version environments. It is also compatible with the bridging function in a NetWare file server and supports the Novell NetBIOS emulator. It is important to realize that IPX routing using RFC 1483 is similar to the support for conventional IPX routers on X.25 or frame-relay networks while the ELAN routing functions are essentially the same as used by conventional IPX routers on legacy LANs.

IPX uses the routing information protocol (RIP) to exchange routing information on a NetWare internetwork and the service advertising protocol (SAP) to exchange information about available services on a NetWare network.

The MSS Server provides extensive support for IPX filtering.

6.3.3 AppleTalk Routing

In Release 1.0, the MSS Server provided support for the IP and IPX routing protocols. (That is, the MSS Server could function as a router for these protocols.) In MSS Server Release 1.1, support has been added for the AppleTalk routing protocol. Thus, the MSS Server is now capable of functioning as an AppleTalk router. (In Release 1.0, the MSS Server could bridge, but not route, AppleTalk.)

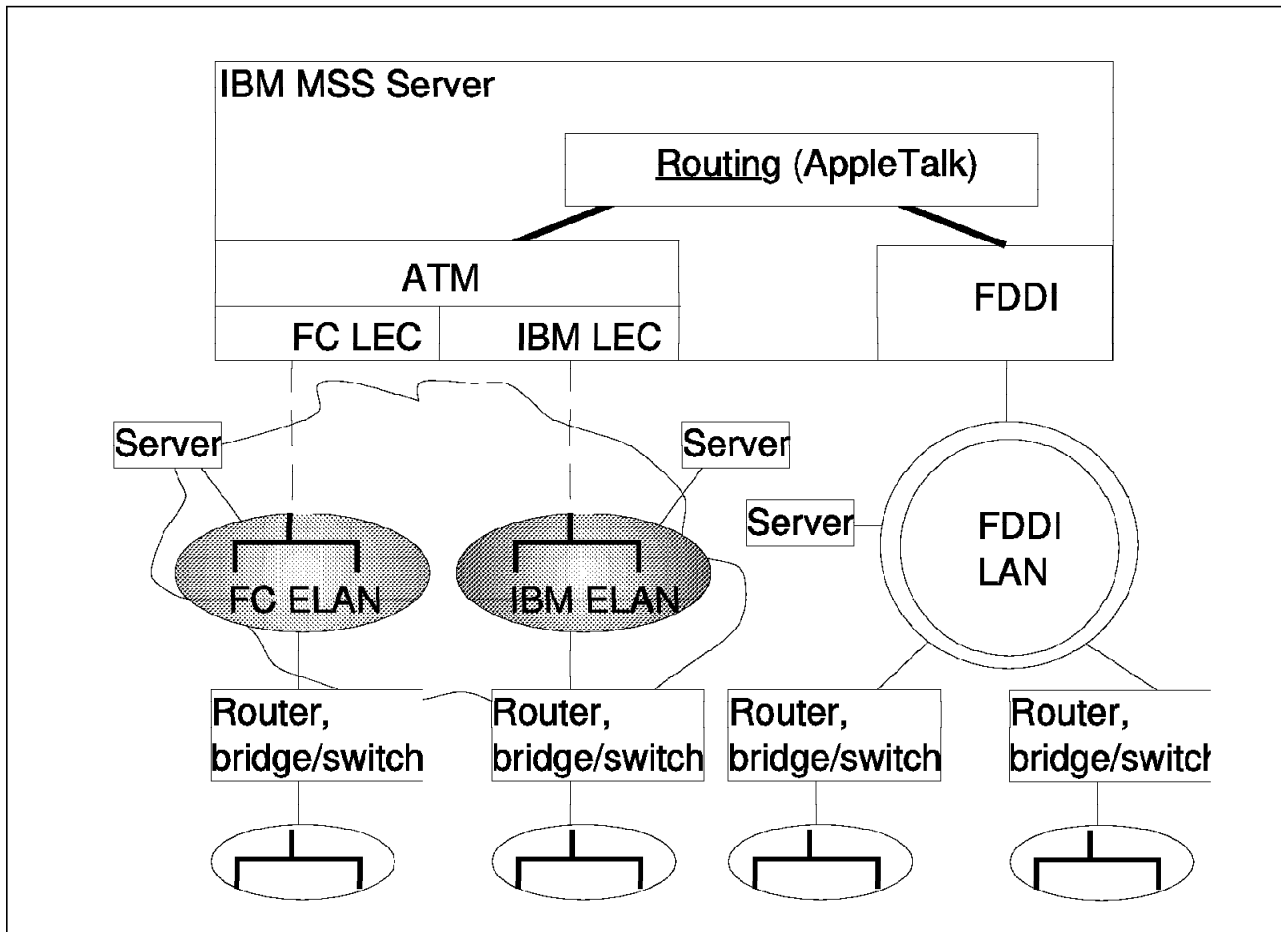


Figure 16. AppleTalk Routing

The MSS Server supports the same AppleTalk routing functions as the IBM 2210 Nways Multiprotocol Router and the IBM 2216 Nways Multiaccess Connector. In Release 1.1 of the MSS Server, AppleTalk routing functions are supported over emulated LAN (token-ring or IEEE 802.3/Ethernet) and FDDI interfaces as shown in Figure 3 below. The emulated LANs can be ATM Forum or IBM-compliant.

6.3.3.1 Configuration Issue

It is important to note that the use of AppleTalk network and zone filters is limited to interfaces 0 through 31. (That is, each network interface is assigned a unique interface number, and AppleTalk network and zone filters are only supported on network interfaces that have been assigned an interface number that is less than 32.) For details on interface numbers, see 6.1, "Real, Virtual and Logical Interfaces" on page 27.

6.4 Concurrent Bridging and Routing

In Release 1.0 of the MSS Server, the bridging versus routing decision was made on a global basis, where a given protocol was routed (and not bridged) on all interfaces if an address for that protocol was defined on any interface.

In Release 1.1 of the MSS Server, this is being changed such that the bridging versus routing decision is made on an interface basis. With this change, a protocol is bridged unless an address for that protocol has been defined on the bridge port interface.

6.4.1 Routing and Bridging on ARIs

New in MSS Server Release 1.1 is the ability to perform bridging on ATM real interfaces (see Chapter 5, “RFC 1483 Bridging over ATM” on page 21). If you have classical IP or IPX using RFC 1483 encapsulation, and are defining an RFC 1483 bridging PVC to a different network, ensure that your bridging port uses a different interface than your (Classical) IP and IPX connections using RFC 1483 encapsulation. Otherwise, no IPX or IP traffic will be bridged, even though this may be needed. Use the AVI for routed IP or IPX over RFC 1483 in this scenario.

6.4.2 Routing and Bridging on ELANs

The recommendation of using different interface for bridging and routing applies to emulated LANs as well. New in MSS Server Release 1.1 is that you can define multiple LE clients (that is, interfaces) to the same ELAN, simplifying the definition process when the MSS Server is routing and bridging on the same ELAN. For details see 15.2, “Multiple Internal LE Clients on the Same ELAN” on page 129.

Chapter 7. MSS 2.0 Routing Enhancements

Release 2.0 builds upon MSS routing capabilities by the addition and enhancement of available routing protocols. With the latest suite of available routing and route switching functions, MSS can be fully integrated into existing traditional routed networks, newer layer 3 switched networks, or be used as the building block for distributed network interconnection in a LAN, WAN and ATM environment.

7.1 IP Routing Enhancements

RIP Version 2 support has been added, compliant with RFC 1723. This extension of the RIP protocol allows routers to share information such as a subnet mask for each destination route and an authentication key for additional security between routers. The subnet mask per destination route allows for the use of variable subnet masking with RIP, allowing more flexibility in IP addressing. RIP Version 2 also reduces interruptions to other network devices by advertising routes to an IP multicast address instead of broadcasting them. It is also compatible with RIP protocol Version 1.

To aid in troubleshooting, the IP TraceRoute function has been enhanced. It now also allows source IP address, data size, number of probes to send, delay between probes, and the packets' maximum Time-To-Live (TTL) to be specified.

7.2 IPX Routing Enhancements

To aid in troubleshooting, the IPX PING function has been enhanced. IPX PING is similar to IP PING, which tests reachability between two network devices. IPX PING now also allows source network, source node, data size, and delay between probes.

IPX TraceRoute is a new function which is similar to IP TraceRoute. It allows the tracing of the routed path taken by an IPX PING packet. The source address, data size, number of probes to send, delay between probes and the maximum number of router hops may be specified.

IPX RecordRoute is a new function that allows tracing of the both forward and reverse paths to/from the destination device (since these may be different). Specified parameters are source address (and destination), delay between packets, and the number of requests to send.

7.3 Banyan Vines Routing

Banyan VINES routing support is provided for Ethernet and token-ring Emulated LANs. Virtual Networking System (VINES) consists of four main protocols at the network layer, all of which are supported by MSS. VINES Internet Protocol (VINES IP) is used to route packets; VINES Routing Update Protocol (VINES RTP) is used to exchange routing information between VINES routers; VINES Address Resolution Protocol (VINES ARP) is used to assign Internet addresses to clients; VINES Internet Control Protocol (VINES ICP) provides diagnostic and support functions.

7.4 APPN Routing

APPN routing support is provided for Ethernet and token-ring Emulated LANs, and natively via RFC 1483 encapsulation. The implementation supports both Intermediate Session Routing (ISR) and High-Performance Routing (HPR).

Chapter 8, "APPN" on page 39 details the function and configuration of APPN.

7.5 SR-TB Duplicate MAC Address Support

Source Route-Transparent Bridge (SR-TB) translational bridging is used when communication is required between token-ring and Ethernet devices and routing is not possible or desirable (due to protocols such as NetBEUI, or issues such as subnet masks).

Previously, SR-TB bridging did not work in networks with duplicate token-ring MAC addresses and source routing (SR) bridges. However, duplicate MAC addresses are often used in SNA to provide redundancy and load balancing across 374X, 3172, and 2216 mainframe gateways. MSS 2.0 SR-TB now supports up to seven duplicate MAC address pairs on token-ring. There must be no transparent bridge path between the two addresses, and the SR paths to each duplicate address must be distinct.

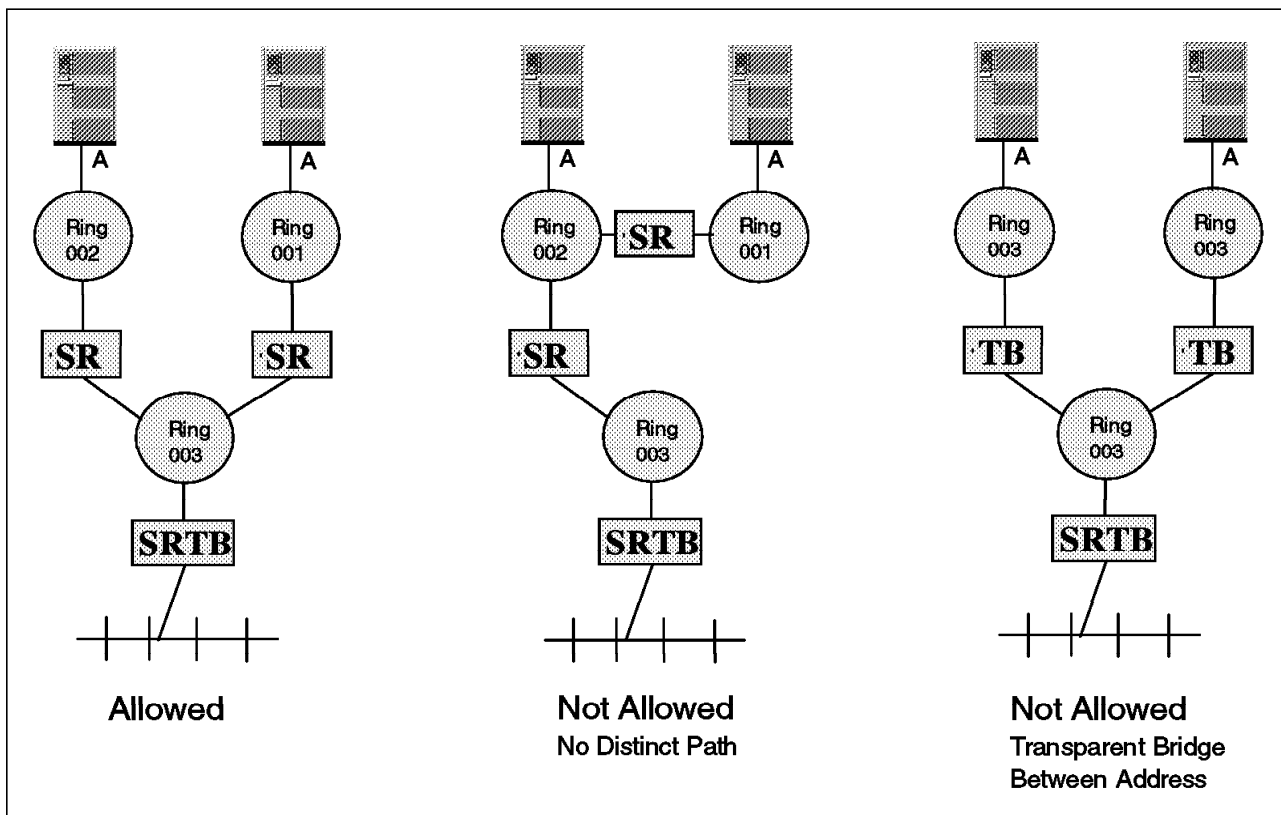


Figure 17. Examples of Duplicate Token-Ring Address Configurations in MSS 2.0

Notice in the left example, there is both source route bridging between the duplicate addresses and the paths are distinct (R#3-R#1, R#3-R#2). In the middle example, there is an SR bridge between the duplicate addresses, but the paths are not distinct (R#3-R#2, R#3-R#2-R#1). Thus if a session is set up to the right gateway on Ring 3, the left gateway must examine and process each

packet, even though it is not meant for that device. In the right example, there is a transparent bridge path between the duplicate addresses.

Duplicate MAC address support must be configured with the MAC addresses. This can be done with the MSS V2.0 configurator, and can be found in NW;Network Device/Protocols/Bridging/General in the Duplicate MACs tab.

To enable duplicate address support, left-mouse button click on the **Enable Duplicate MAC Addresses** check box. By selecting the **Load Balancing** check box below it, MSS will send an explorer frame from an Ethernet device to the duplicate MAC address destination with the RIF that was least recently used.

Chapter 8. APPN

Advanced Peer-to-Peer Networking (APPN) is used to extend the SNA architecture by enabling Type 2.1 (T2.1) nodes to communicate directly without requiring the services of an SNA host computer.

T2.1 nodes can activate connections with other T2.1 nodes and establish LU-LU sessions with these nodes. The relationship between a pair of T2.1 nodes is referred to as a *peer relationship* because either side can initiate communications.

8.1 MSS As a Network Node

The MSS is defined as a network node and supports connections with the following node types:

- APPN Network Nodes (NN)
- APPN End Nodes (EN)
- Low Entry Networking (LEN) End Nodes

The MSS is a network node (NN) and cannot function as an end node for APPN. The MSS provides directory services, routing services and management services to APPN end nodes (EN) and low-entry networking (LEN) end nodes. Some of the major functions performed by the MSS network node (NN) are:

1. Intermediate session routing (ISR)
2. High-performance routing (HPR)

Two of the major functions of HPR include:

- Rapid Transport Protocol (RTP)
- Automatic Network Routing (ANR)

3. PU 2.0 nodes supported by DLUR
4. APPN connection networks (CN)

CN enables direct communications between nodes connected to the same shared access transport facility (SATF).

5. Branch Extender:

- A new SNA/APPN architecture for very large networks
- Subdivides the large APPN network
- Overcomes the problems of:
 - Too many NNs
 - Too much unintelligent searching
 - Too much topology traffic
- Best for large intranets with many branches
- Behaves as:
 - An EN on UPlinks
 - An NN on DOWNlinks
 - APPN HPR over ATM

MSS as a network node supports DLCs as token-ring and Ethernet ELANs.

The following Logical Data Link Control (LDLC) support for native ATM connections is offered:

- RFC 1483 encapsulation is used.
- AAL5 with null SSCS.
- PVC and SVC support.
- HPR only (no ISR).

The MSS NN exchanges information about the topology of the network. This information is exchanged each time NNs establish a connection or when there is a change in the topology of the network (when a network node is deactivated, brought on line, or when a link fails or is congested). When a topology change is detected the NN sends a broadcast with this information to the other active network nodes that it has CP-CP sessions.

The MSS can act as an intermediate node, receiving session data from one adjacent node and passing that data on to the next adjacent node.

8.2 End Node (EN)

The MSS acts as a network node server for the EN. It both registers the EN resources and forwards requests for unknown resources. This is a locate search request, and uses the NNs directory and routing services. This will allow the NN to determine the path from the EN to the destination.

The EN and the network node server establish CP-CP sessions. An EN may be connected to a number of network nodes, but only one of these nodes acts as the ENs server at any one time.

The EN forwards requests for unknown resources to the network node server. The network node server then uses its search facilities to locate the requested resource and calculate the route from the EN to the resource.

8.3 Low-Entry Network End Nodes (LEN)

A LEN node is a T2.1 node without APPN extensions. A LEN node can establish peer connections with other LEN nodes, ENs and NNs as long as all of the required destination LUs are registered with the LEN node. A LEN node can also serve as a gateway between an APPN network and an SNA subarea network.

The LEN node is incapable of CP-CP sessions with the NN. The LEN node must predefine the LUs of other non-adjacent nodes with the CP name of the NN. Now that the LEN node has predefined LUs, when the LEN node sends a bind session activation to the NN, which then locates the destination and sends the Bind along the calculated route to the destination.

When configuring the MSS NN you can specify the names of LUs that are associated with an attached LEN node. These LU names reside in the NNs local directory. If the NN receives a request to search for one of these LEN nodes resources, it will be able to find the LU in its local directory and return a positive response to the node originating the search. To reduce the number of LU names you need to specify for an attached LEN node, the MSS supports the use of

generic LU names, which allow a wildcard character to represent a portion of an LU name.

8.4 Dependent LU Requester (DLUR)

A PU 2.0 node is a T2.0 node containing dependent LUs. PU 2.0 nodes are supported by the dependent LU requestor (DLUR) function, which is implemented by an EN or NN. PU 2.0 nodes require the services of a System Services Control Point (SSCP), which is made available through the DLUR-enabled APPN node.

Note that APPN nodes can contain dependent LUs supported by an internal DLUR function. APPN nodes that do not contain the DLUR function, but do contain dependent LUs, require DLUR support in an adjacent network node to allow those LUs to communicate across an APPN network. In such a case, the APPN node will also have a PU 2.0-like appearance to the DLUR support in the adjacent APPN NN.

8.5 APPN Connection Networks (CN)

With connection network support an APPN NN (MSS) can establish a direct connection with any APPN node without having all the nodes defined to the shared access transport facility (SATF). CN is supported on token-ring and Ethernet ELANs. A fully qualified name must be specified for the connection network name parameter. The network name of CN must be the same as the network identifier for the MSS APPN NN. To become a member of the CN, an APPN node's port must be attached to the CN by defining a connection network interface. When the port is defined, a connection network TG is created by the APPN component to identify the direct connection from the port to the SATF (that is, connection network).

8.6 High-Performance Routing (HPR)

HPR is an enhancement to APPN architecture that provides better performance over high-speed, low-error rate links using existing hardware. HPR replaces the normal APPN intermediate session routing (ISR) with a Network Control Layer (NCL) containing a new type of source routing function called Automatic Network Routing (ANR). The complete HPR route is contained in the ANR packet allowing intermediate routing nodes to route the packets with less processing overhead and storage.

HPR also eliminated the error recovery and flow control (session-level pacing) procedures for each link between nodes and moves the error recovery and flow/congestion control procedures to the endpoints of an HPR connection. A transport layer using a new error recovery procedure called Rapid Transport Protocol (RTP) is used by the endpoints of the HPR connection. HPR intermediate nodes have no session or RTP connection awareness.

HPR uses APPN network control functions including class of service (COS)-based least-weight route calculation and transmission priority.

- The network automatically adapts to the presence of HPR-capable nodes and HPR-enabled links.
- An APPN network can have any mix of ISR and HPR links, although the greatest benefit of HPR is realized when the network has three or more

HPR-enabled nodes with two or more HPR-capable links back-to-back. This allows the middle HPR node to be an HPR intermediate node and use only ANR routing, allowing session data to be routed through the middle node using only NCL.

- A given session route can be made up of a combination of ISR and HPR links.
- HPR uses the same TG and node characteristics for least-weight route calculation as APPN ISR. No special consideration is given to HPR-capable nodes or links other than their potentially improved characteristics (such as higher effective capacity if a higher speed link).

8.6.1.1 Traffic Types

APPN ISR uses the IEEE 802.2 LLC Type 2 protocol for token-ring and Ethernet APPN HPR, which is supported on token-ring and Ethernet, does not use LLC Type 2 protocol, but does use some functions of an APPN link station for XID and inactivity timeout. A single APPN link station is therefore used for ISR or HPR. Different mechanisms are used to distinguish between ISR and HPR traffic depending upon the DLC type.

Each protocol that uses a port must have a unique SAP address. A unique SAP address identifies the APPN link station for HPR traffic (local HPR SAP address parameter). If ISR traffic is destined for a station, then a different SAP address (local APPN SAP address parameter) must be used. The ISR traffic uses LLC Type 2 LAN frames. The HPR traffic is handled in similar fashion to LLC Type 1 LAN frames and must have a different SAP address.

The default SAP address for HPR traffic is X'C8'. If X'C8' has already been used by another protocol on a port, the default must be overridden.

There is only one APPN link station even though APPN ISR and HPR traffic use different SAP addresses.

Chapter 9. Configurable Quality of Service for LAN Emulation

The current LAN Emulation standard (LANE 1.0) does not specify how Quality of Service (QoS) can be used for the connections set up by an LE client. Consequently, most LANE implementations currently use best-effort connections. In order to leverage the benefits of ATM QoS capabilities for LAN Emulation, MSS Release 1.1 provides a configurable QoS service for data direct connections between LE clients.

9.1 Overview

An LE client will use configured or learned QoS parameters for its data direct connections. The user has the flexibility to configure QoS parameters on an ELAN and/or LE client basis. LE clients may use a negotiation algorithm to pick the parameters best-suited for the pair of communicating stations.

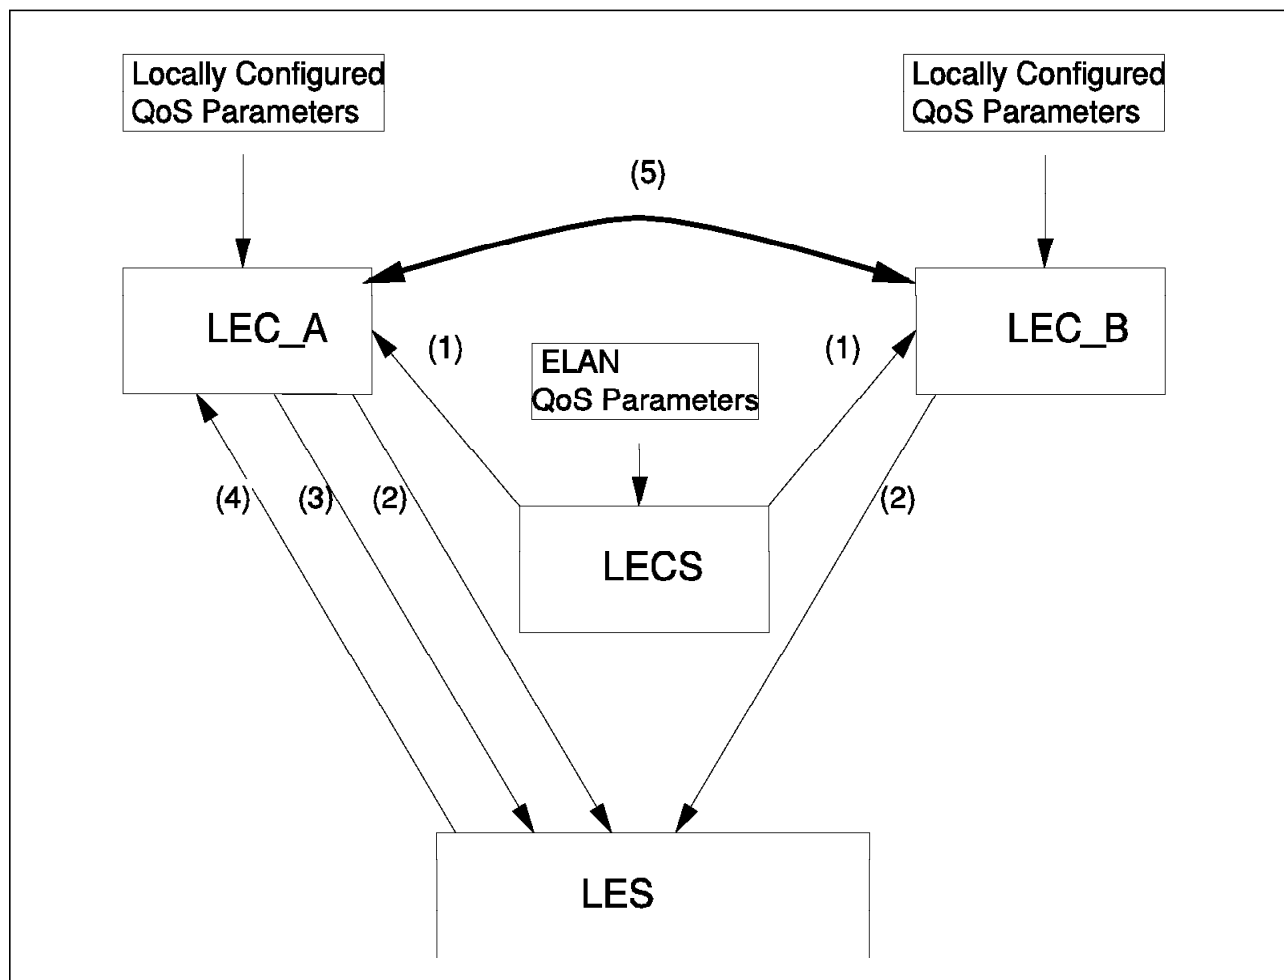


Figure 18. Quality of Service for LAN Emulation Clients

With ELAN-based QoS, the LECS provides the same set of QoS parameters to all LE clients on an ELAN during the configuration phase. This enables traffic for specific ELANs to receive preferential treatment.

LE client-based QoS differs in that QoS parameters may be configured individually at particular LE clients. This enables traffic for specific LE clients to receive preferential treatment. LE client-based QoS may be used in conjunction with ELAN-based QoS or independently. When used together, LE clients with individually configured QoS parameters should be configured to ignore QoS parameters received from the LECS. (By default, parameters received from the LECS take preference over locally defined values.)

A QoS negotiation mechanism is provided so that an LE client preparing to set up a data direct VCC may be made aware of the partner LE clients QoS parameters. Each QoS-capable LE client registers (see step 2 in Figure 18 on page 43) its QoS parameters at the LES. When requesting the ATM address of a partner LE client (step 3) the LES informs the requesting LE client also about the QoS parameters of the partner LE client (step 4). Using this information the LE client establishes a data direct VCC (step 5) using the QoS parameters that are best suited for communication between these two LE clients.

9.2 Benefits

By configuring QoS for an ELAN, all data direct VCCs established between any pair of LE clients on that ELAN are QoS-enabled. This allows QoS to be configured for selective ELANs. As an example, consider the case of a multiprotocol network containing multiple ELANs, where the traffic on one of the ELANs is SNA. In this case, it may be desirable to configure QoS parameters that guarantee some minimum bandwidth for communications between LE clients on the SNA ELAN as a precaution against LLC timeouts during bursty periods or simply to ensure timely delivery of critical data.

LE client-based QoS provides a finer level of granularity in configuring QoS. It addresses some situations where a particular LE client is required to be configured differently than the other LECs of the same ELAN. For example, QoS parameters might be configured for LE clients that provide access to servers with high-bandwidth requirements. A customer can provide better service for all transactions with these servers by configuring QoS for its LE clients to ensure acceptable performance levels.

9.2.1 Migration Issues

To maximize the use of configurable QoS, LECS, LES and LE clients all need to be QoS-enabled. MSS Server Release 1.1 has implemented all configurable QoS parameter features.

The configurable QoS feature is completely compliant with the LANE Version 1.0 specification and interoperable with LECS, LES, and LE clients that do not support the feature. If a reserved bandwidth call is rejected due to the traffic parameters or QoS class, the call is retried as a best-effort connection. Thus, all the LE clients on an ELAN do not have to support the QoS functions in order to realize benefits; however, the benefits increase with the breadth of support.

An LE client may participate in configurable QoS at several levels. At the most basic level, the LE client will accept reserved bandwidth connections; this allows the LE client to receive QoS-enabled calls. At the next level, the LE client must be capable of placing reserved bandwidth calls according to a configured set of parameters. To participate in ELAN-based QoS, the LE client must be able to accept QoS parameters returned by the LECS, and to fully participate in

configurable QoS the LE client must also implement the QoS negotiation algorithm.

9.3 Functional Description

Several entities participate in providing configurable QoS service. For minimal support, LECs must be able to use configured QoS parameters for establishing data direct connections and accept QoS VCCs for data direct connections. To extend it a step further, LECs must be capable of receiving QoS parameters from the LE Configuration Server (LECS). Complete support requires the participation of the LE Server (LES) to enable QoS negotiation.

9.3.1 Using Configurable QoS Parameters

The set of QoS parameters that are configured for an entity include the traffic type, sustained cell rate (SCR), peak cell rate (PCR), and the QoS class. Also, a customer can control the level of participation of an LE client by configuring the maximum reserved bandwidth (MRB) for a data direct connection; this parameter protects the LE client from accepting/establishing data direct connections whose traffic characteristics it cannot support. All of the QoS parameters are described in detail in 9.3.5, “Configurable QoS Parameters of an LE Client” on page 46.

The QoS parameters configured at the LECS for an ELAN apply to all LE clients belonging to the ELAN. QoS parameters configured for an LE client pertain to that client only. LE clients can be configured to enable/disable QoS negotiation, which is used only if the LES is participating in configurable QoS. LE clients can also be configured to enable/disable acceptance of QoS parameters from the LECS. If this parameter is enabled, then the QoS parameters from the LECS will override the locally configured parameters.

9.3.2 LAN Emulation Configuration Server (LECS) Behavior

The LECS provides the user with an interface to configure QoS parameters for the supported ELAN. Whenever an LE client sends an `LE_CONFIGURE_REQUEST`, the LECS provides the QoS information to the LE client as TLVs in the `LE_CONFIGURE_RESPONSE` control message.

9.3.3 LAN Emulation Server (LES) Behavior

The LES plays a role in configurable QoS service by supporting QoS negotiation. Each LE client that is participating in the negotiation protocol registers its QoS parameters with the LES by sending them in an `LE_JOIN_REQUEST` control message. During the `LE_ARP` process, the LES provides the requesting LE client with the TLVs registered by the target LEC as TLVs in the `LE_ARP_RESPONSE` control message. If the source or the destination LE client has not registered any parameters, then TLVs are not forwarded in the `LE_ARP_RESPONSE` control message.

9.3.4 LAN Emulation Client (LEC) Behavior

The LE client is the main entity that enables configurable QoS service. The LEC can participate in the service at different levels. The minimum requirement from an LE client is to be a passive participant where it can accept QoS data direct connections. An LE client can actively participate in one of the following two modes:

1. Configurable QoS

In this mode, the LE client participates by being able to set up and accept QoS VCCs for data direct connections. The QoS parameters are available to the LE client as locally configured values or as TLVs from the LECS. The LE client uses the QoS parameters by default whenever it has to set up a data direct connection. The LE client uses a local policy for retrying a connection request that failed due to the QoS parameters; the policy must include the use of best-effort traffic parameters at some point. If QoS parameters are not available, then the LE client uses the best-effort traffic parameters for its data direct connections.

2. Configurable QoS with Negotiation

An LE client can be configured to participate in QoS negotiation. In this mode, the LEC registers its QoS parameters with the LES by including QoS parameters in the LE_JOIN_REQUEST control message. Whenever an LE client registers its parameters with the LES, the LES may provide the traffic parameters of the target LE clients in LE_ARP_RESPONSE control messages. The LE client uses a local policy (based on its parameters and the destinations parameters) to resolve the negotiated parameters used for the connection to the address-resolved destination.

Figure 19 on page 47 overviews three examples of how traffic parameters are negotiated for data direct VCCs. Three LE clients (A, B, C) with pre-configured QoS parameters connect to LE client D, for which the QoS parameters have also been defined. The resulting (actual) QoS parameters are also displayed. As can be seen, different types of QoS parameters result for each of the connections.

The QoS parameters, their meaning and default values, and how actual values are negotiated are discussed in the following section.

9.3.5 Configurable QoS Parameters of an LE Client

The keywords to control an LE client's QoS parameters during the acceptance of *incoming* data direct VCC are:

1. Validate Peak Cell Rate (PCR) (see 9.3.5.1, "Validate Peak Cell Rate (PCR)" on page 48)

This keyword controls the acceptance of best-effort VCCs. Possible values are TRUE and FALSE. When TRUE, VCCs will be rejected if the signalled forward PCR exceeds the line rate of the LE client's ATM port.

2. Maximum Reserved Bandwidth (see 9.3.5.2, "Maximum Acceptable Reserved Bandwidth" on page 49)

This parameter defines the maximum acceptable sustained cell rate (SCR) for a data direct VCC; if SCR is not specified on the incoming call then this parameter defines the maximum acceptable PCR for a data direct VCC with reserved bandwidth.

Note: This parameter also plays a role in determining the traffic characteristics of outgoing calls.

The following characteristics are associated with all the data direct VCCs established by the LE client:

- Bandwidth is not reserved for best-effort traffic.
- Traffic parameters apply to both forward and backward directions.

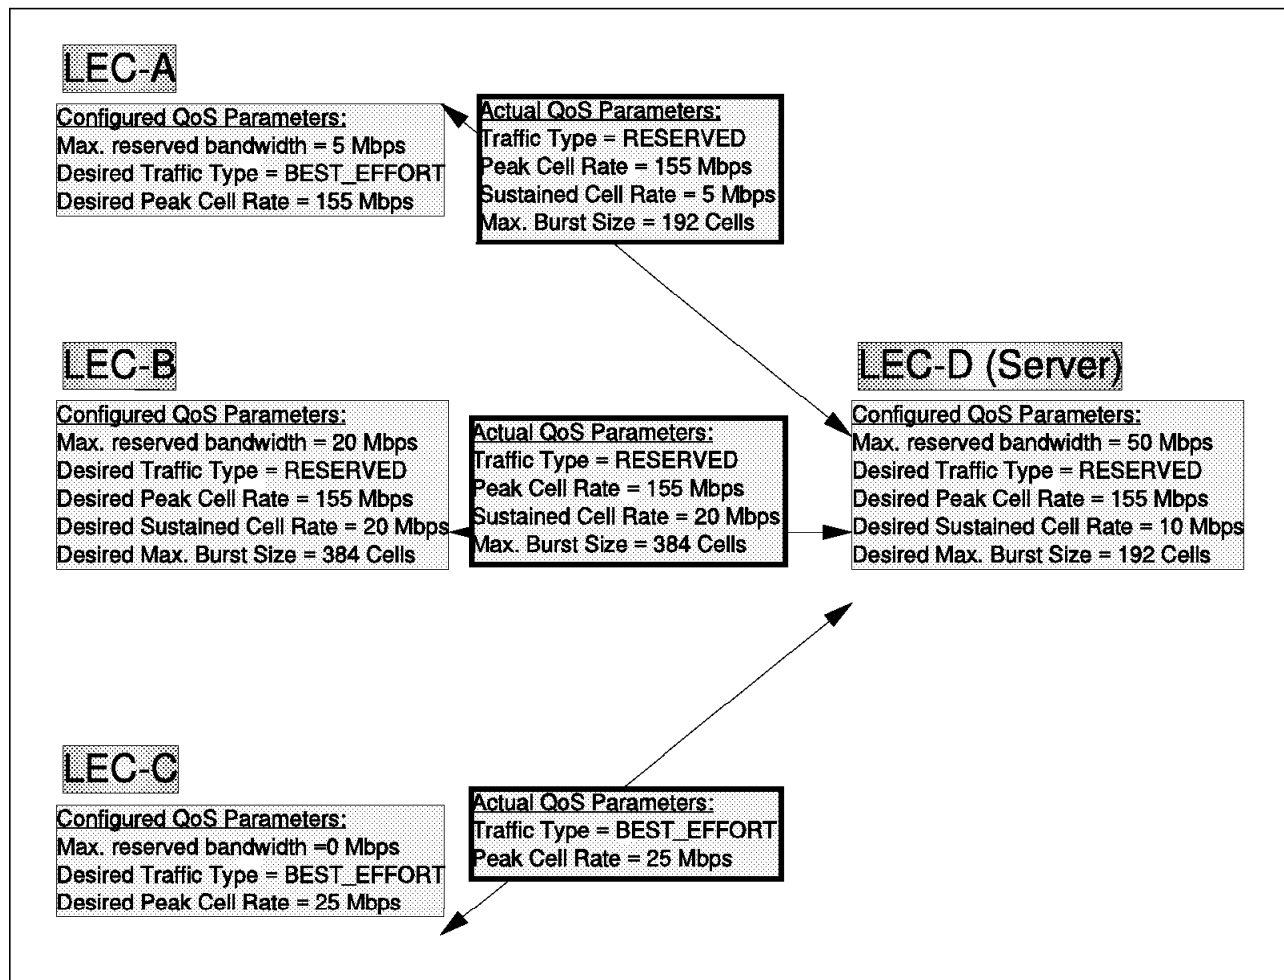


Figure 19. QoS Negotiation

- When a reserved bandwidth connection is rejected due to traffic parameters or QoS class, the call is retried as a best-effort connection with the configured peak cell rate.
- When a best-effort connection is rejected due to the PCR, the call may be automatically retried with a lower PCR; retries may be performed with:
 - if (rejected) PCR > 100 Mbps, retry with PCR = 100 Mbps
 - if (rejected) PCR > 25 Mbps, retry with PCR = 25 Mbps

During the establishment of the VCCs either pre-configured QoS parameters, values received from LECS, and/or values received from the LES are used. The receipt of values received from LECS and/or LES is controlled by the following keywords:

1. Accept QoS parameters from LECS (9.3.5.9, "Accept QoS Parameters from LECS" on page 51)

This parameter controls the interaction between an LE client and its QoS-enabled LECS. When enabled, the LE client overwrites its pre-configured QoS parameters with values received from the LECS.

2. Enable QoS parameter negotiation (9.3.5.8, "Enable QoS Parameter Negotiation" on page 51)

This parameter controls the interaction between an LE client and its QoS-enabled LES. When enabled, the LE client learns the QoS parameters of its partner LE client from the LES, and establishes a data direct VCC with matching traffic parameters.

If not enabled, the pre-defined traffic parameters on the LE client are used.

The QoS parameters used to control the traffic parameters of data direct VCCs are: maximum reserved bandwidth (RB), traffic type, peak cell rate (PCR), sustained cell rate (SCR), and QoS class. The interaction between the QoS parameters entered on LE clients and/or LECS definitions is:

1. Maximum reserved bandwidth (RB) (see 9.3.5.2, “Maximum Acceptable Reserved Bandwidth” on page 49)

This parameter sets an upper bound on the amount of reserved bandwidth that may be requested. The traffic type and peak cell rate are dependent upon this parameter.

2. Traffic type (see 9.3.5.3, “Desired Traffic Type” on page 49)

This parameter controls whether best-effort or reserved bandwidth connections are used. A reserved bandwidth connection will only result when requested by both LE clients.

3. Peak cell rate (see 9.3.5.4, “Desired Peak Cell Rate (PCR)” on page 49)

This parameter controls the desired PCR for data direct VCCs. The actual value depends on whether QoS negotiation is active and the PCR used by the partner LE client.

4. Sustained cell rate (SCR) (see 9.3.5.5, “Desired Sustained Cell Rate (SCR)” on page 50)

This parameter controls the desired SCR for data direct VCCs. The actual value depends on whether QoS negotiation is active and the SCR used by the partner LE client. The SCR is only signalled when PCR is signalled.

SCR only applies to reserved bandwidth connections.

5. Desired maximum burst size (MBS) (see 9.3.5.6, “Desired Maximum Burst Size (MBS)” on page 50)

This parameter controls the desired MBS for data direct VCCs. The actual value depends on whether QoS negotiation is active and the MBS used by the partner LE client. The MBS is only signalled when SCR is signalled.

MBS only applies to reserved bandwidth connections.

6. QoS class (see 9.3.5.7, “Desired QoS Class for Reserved Bandwidth Calls” on page 50)

This parameter controls the desired QoS class for data direct VCCs. The actual value depends on whether QoS negotiation is active and the QoS class used by the partner LE client.

9.3.5.1 Validate Peak Cell Rate (PCR)

This is an LE client parameter that controls the acceptance of incoming best-effort VCCs. Possible values are TRUE and FALSE.

When FALSE, best-effort VCCs will be accepted without regard to the signalled forward PCR. Accepting best-effort VCCs with forward PCRs that exceed the line

rate can result in poor performance due to excessive retransmissions; however, rejecting these VCCs can result in interoperability problems.

When TRUE, best-effort VCCs will be rejected if the signalled forward PCR exceeds the line rate of the LE client's ATM device. Calls will not be rejected due to the backward PCR. The signalled backward PCR will be honored if it does not exceed the line rate; otherwise, transmissions to the caller will be at line rate.

The TRUE setting is useful when callers retry with a lower PCR following call rejection due to unavailable cell rate.

9.3.5.2 Maximum Acceptable Reserved Bandwidth

This parameter controls the sustained cell rate (SCR) on both data direct VCC calls received by the LE client and data direct VCC calls placed by the LEC. Values between 0 and the line speed can be coded.

For incoming calls, this parameter defines the maximum acceptable SCR for a data direct VCC; if SCR is not specified on the incoming call, then this parameter defines the maximum acceptable PCR for a data direct VCC with reserved bandwidth. This parameter is applied to both forward and backward directions. Calls received with traffic parameters specifying higher rates will be released. If SCR is specified on the incoming call, the call will not be rejected due to the PCR or maximum burst size. The constraint imposed by this parameter is not applicable to best-effort connections.

For outgoing calls, this parameter sets an upper bound on the amount of reserved bandwidth that may be requested for a data direct VCC; thus, traffic type and peak cell rate are dependent upon this parameter.

9.3.5.3 Desired Traffic Type

This parameter controls whether best-effort or reserved bandwidth connections are desired.

If QoS parameters are not negotiated, then this parameter specifies the type of traffic characteristics to be associated with the data direct VCC calls placed by the LEC.

If QoS parameters are negotiated, this parameter specifies the desired type of traffic characteristics for data direct VCCs. If either the source or target LE client desires a reserved bandwidth connection and both LE clients support reserved bandwidth connections, then an attempt will be made to establish a reserved bandwidth data direct VCC between the two LE clients. Otherwise, the data direct VCC will be a best-effort connection.

9.3.5.4 Desired Peak Cell Rate (PCR)

This parameter (*integer* kbps) controls the desired PCR for data direct VCCs. Values between 0 and the line speed can be coded.

If QoS parameters are not negotiated, then this parameter specifies the PCR traffic parameter for data direct VCC calls placed by the LEC.

If QoS parameters are negotiated, this parameter specifies the desired PCR traffic parameter for data direct VCCs. The smallest of the desired PCRs of the two LE clients is used for negotiated best-effort VCCs. When a reserved

bandwidth VCC is negotiated and only one of the LE clients requests a reserved bandwidth connection, then the desired PCR of that LE client is used for the data direct VCC subject to the upper bound imposed by the line rate of the local ATM device; if both LE clients request a reserved bandwidth connection, then the largest of the desired PCRs of the LE clients is used for the data direct VCC subject to the upper bound imposed by the line rate of the local ATM device.

9.3.5.5 Desired Sustained Cell Rate (SCR)

This parameter (*integer* kbps) controls the desired SCR for data direct VCCs. Values between 0 and the minimum of MRB and PCR can be coded.

If QoS parameters are not negotiated, then this parameter specifies the SCR traffic parameter for data direct VCC calls placed by the LEC.

If QoS parameters are negotiated, this parameter specifies the desired SCR traffic parameter for data direct VCCs. When a reserved bandwidth VCC is negotiated and only one of the LE clients requests a reserved bandwidth connection, then the desired SCR of that LE client is used for the data direct VCC (subject to the upper bound imposed by the maximum RB parameter of the other LEC); if both LE clients request a reserved bandwidth connection, then the maximum of the desired SCRs of the LE clients is used for the data direct VCC (subject to the upper bound imposed by the maximum RB parameters of both LE clients).

In any case (negotiation or not), if the SCR that is to be signalled equals the PCR that is to be signalled, then the call is signalled with PCR only.

9.3.5.6 Desired Maximum Burst Size (MBS)

This parameter (*integer* of cells) controls the desired MBS for data direct VCCs. Values between 0 and the maximum data frame size plus eight, divided by 53 $((MFS+8)/53)$ can be coded.

If QoS parameters are not negotiated, then this parameter specifies the MBS traffic parameter for data direct VCC calls placed by the LEC.

If QoS parameters are negotiated, this parameter specifies the desired maximum burst size traffic parameter for data direct VCCs. When a reserved bandwidth VCC is negotiated and only one of the LE clients requests a reserved bandwidth connection, then the desired MBS of that LE client is used for the data direct VCC; if both LE clients request a reserved bandwidth connection, then the maximum of the desired MBSes of the LE clients is used for the data direct VCC.

In any case (negotiation or not), the MBS is signalled only when SCR is signalled. Although this parameter is expressed in units of cells, it is recommended that it be presented to the user as an integer multiple of the maximum data frame size (with a lower bound of 1).

9.3.5.7 Desired QoS Class for Reserved Bandwidth Calls

This parameter (*integer*) controls the desired QoS Class. Accepted values are: unspecified, Class 1, Class 2, Class 3 or Class 4.

If QoS parameters are not negotiated, then this parameter specifies the QoS class to be used for reserved bandwidth data direct VCC calls placed by the LEC.

If QoS parameters are negotiated, this parameter specifies the QoS class that is desired for data direct VCCs.

Unspecified QoS class is always used on best-effort calls. Specified QoS classes define objective values for ATM performance parameters such as cell loss ratio and cell transfer delay. The UNI specification states that:

1. QoS Class 1

Should yield performance comparable to current digital private line performance.

2. QoS Class 2

Intended for packetized video and audio in teleconferencing and multimedia applications.

3. QoS Class 3

Intended for interoperation of connection-oriented protocols, such as frame relay.

4. QoS Class 4

Intended for interoperation of connectionless protocols, such as IP or SMDS.

LE clients must be able to accept calls with any of the above QoS classes. When QoS parameters are negotiated, the configured QoS classes of the two LE clients are compared, and the QoS class with the more stringent requirements is used.

9.3.5.8 Enable QoS Parameter Negotiation

This parameter controls the interaction between an LE client and its QoS-enabled LES. Possible values are TRUE and FALSE.

QoS Negotiation

Enable only when connecting to an IBM MSS LES.

When enabled the LE client will include traffic parameters (maximum RB, traffic type, PCR, SCR, QoS class) in LE_JOIN_REQUEST and LE_ARP_RESPONSE frames sent to the LES. These traffic parameters may also be included in an LE_ARP_RESPONSE returned to the LE client by the LES. If there are no traffic parameters included in an LE_ARP_RESPONSE received by the LEC, then the local configuration parameters must be used to set up the data direct VCC. If traffic parameters are included in an LE_ARP_RESPONSE, the LE client must compare the values with the corresponding local values to determine the negotiated or best set of parameters acceptable to both parties before signalling for the data direct VCC.

The algorithm used by the LE client to determine the signalled parameters is given in 9.4, "LANE QoS Negotiation Algorithm" on page 52.

9.3.5.9 Accept QoS Parameters from LECS

This parameter controls the interaction between an LE client and its QoS-enabled LECS. Possible values are TRUE and FALSE.

When enabled the LE client should use the QoS parameters obtained from the LECS in the LE_CONFIGURE_RESPONSE frames. Any QoS parameters received from the LECS override the locally configured QoS parameters.

If disabled, the LE client will ignore any QoS parameters received from the LECS.

9.4 LANE QoS Negotiation Algorithm

Details of the QoS negotiation algorithm, in the form of a pseudo-code description, are provided below.

```
struct QoS_Tlv {
    /* QoS TLV exchanged by LECs and LES */
    unsigned long maxReservedBW; /* Maximum Reserved Bandwidth Acceptable */
    unsigned long trafficType; /* Desired Traffic Type */
    unsigned long pcr; /* Desired Peak Cell Rate */
    unsigned long scr; /* Desired Sustained Cell Rate */
    unsigned long qosClass; /* Desired QoS Class */
    unsigned long maxBurstSize; /* Desired Maximum Burst Size */
} Source, Target;
/* Source is included in LE_JOIN_REQUEST */
/* Target is returned in LE_ARP_RESPONSE */

if ((QoS parameter negotiation is disabled) || (no TLV is included in LE_ARP_RESPONSE))
{ trafficType = Source.trafficType; pcr = Source.pcr; scr = Source.scr; }
else if (Target.trafficType == BEST_EFFORT)
{
    if ((Source.trafficType == BEST_EFFORT) || (Target.maxReservedBW == 0))
    { trafficType = BEST_EFFORT; pcr = min(Source.pcr, Target.pcr); }
    else
    {
        trafficType = RESERVED_BANDWIDTH; pcr = Source.pcr;
        scr = min(Source.scr, Target.maxReservedBW); maxBurstSize = Source.maxBurstSize;
    }
}
else if (Source.trafficType == BEST_EFFORT)
{
    if (Source.maxReservedBW == 0)
    { trafficType = BEST_EFFORT; pcr = min(Source.pcr, Target.pcr); }
    else
    {
        trafficType = RESERVED_BANDWIDTH;
        pcr = min(Target.pcr, local ATM line rate);
        scr = min(Target.scr, Source.maxReservedBW);
        maxBurstSize = Target.maxBurstSize;
    }
}
else
{
    trafficType = RESERVED_BANDWIDTH;
    pcr = min(max(Source.pcr, Target.pcr), local ATM line rate);
    scr = min(max(Source.scr, Target.scr), Source.maxReservedBW, Target.maxReservedBW);
    maxBurstSize = max(Source.maxBurstSize, Target.maxBurstSize);
}
if (trafficType == BEST_EFFORT)
    qosClass = UNSPECIFIED_QOS_CLASS;
else
{
    qosClass = min(Source.qosClass, Target.qosClass);
    if (qosClass == UNSPECIFIED_QOS_CLASS)
        qosClass = max(Source.qosClass, Target.qosClass);
}
if ((trafficType == BEST_EFFORT) || (scr == pcr))
    do not signal SCR or maxBurstSize
```

9.5 Benefits

Standard UNI 3.0/3.1 signalling does not provide the capability to negotiate traffic parameters. If a call is rejected due to the traffic parameters or QoS class, the station may retry with different parameters, but must do so without knowing what parameter values are acceptable at the called station, which is both suboptimal and inefficient.

The situation is better with UNI 4.0, which supports peak cell rate negotiation or an alternate set of traffic descriptor parameters. However, the configurable QoS feature provides greater flexibility than UNI 4.0. With UNI 4.0 sustained cell rate negotiation is not supported, and it is not possible to negotiate cell rates upward

based on the traffic characteristics desired by the called station, as might be advantageous when a client connects to a server. Furthermore, the configurable QoS feature is compatible with the large installed base of UNI 3.0/3.1 installations.

While the preceding discussion has demonstrated the value of the configurable QoS feature, it must be noted that the benefits will initially be limited to backbone connections between MSS Servers, since the internal MSS Server LECs will be the only LAN Emulation clients that implement the feature when Release 1.1 is made available.

Chapter 10. SuperVLANs

Although virtual LANs (VLANs) and the benefits they render have received considerable attention throughout the networking community, the basic concepts are briefly reviewed to provide a framework for articulating the value of the MSS Server's SuperVLAN feature.

A VLAN is a logical grouping of hosts that form a broadcast domain. The logical grouping is independent of the physical network topology with VLAN membership governed by a set of rules or policies. VLANs have several advantageous characteristics:

1. VLANs imply broadcast control (see 10.2, "Broadcast Management" on page 62).

The scope of a broadcast frame originating within a VLAN is limited to LAN (or ELAN) segments containing hosts that are members of the VLAN. Most protocols use broadcasts for address resolution; therefore, if the scope of the broadcast is limited, the scope of host-to-host communication is also limited.

2. IntraVLAN communications is efficient.

IntraVLAN communications is typically very efficient with VLAN members communicating directly over a switched hardware path.

3. It is easy to move, add, and change VLAN members.

Host moves, adds, and changes are simplified since VLAN membership is independent of physical topology. No reconfiguration is required to retain VLAN memberships when stations move to new physical locations. Similarly, no wiring modifications are needed to move stations from one VLAN to another; instead, these operations can be performed at a management console, often via a drag-and-drop graphical user interface.

VLANs are generally associated with LAN switch products, which are, fundamentally, bridges with efficient hardware data paths. Analogously, MSS Server VLANs are based on ELAN bridging. The MSS Server implements a particular type of VLAN called protocol virtual LAN (or PVLAN). Membership in a PVLAN is determined by network protocol (for example, NetBIOS) or the combination of protocol and network address (for example, a particular IP subnet or IPX network). The MSS Server creates broadcast domains by dynamically learning the set of PVLANs active on each ELAN segment. Broadcasts are then limited to ELAN segments containing stations that are members of the PVLAN, and the shortcut bridging component of the SuperVLAN feature enables efficient intraVLAN communications by allowing direct ATM connections to be established between clients residing on different ELAN segments.

The net result is that VLANs are important because they enable a switched infrastructure to be partitioned into broadcast domains in a manner that simplifies network administration with regard to mobility. Carving the switched infrastructure into controlled broadcast domains is a significant step toward network scalability, but interVLAN communications must also be addressed to complete the solution. The SuperVLAN is an eclectic feature that was designed to address interVLAN communication requirements in addition to providing VLAN function. The SuperVLAN package is comprised of four complementary functions:

1. Shortcut Bridging (SCB) to distribute the LAN Emulation services

2. Bridging BroadCast Manager (BBCM)
3. Dynamic Protocol Filtering (DPF) support for PVLANS
4. IP cut-through support for zero-hop routing

These functions provide the versatility needed to support both large flat networks, as well as more traditional subnetted designs.

Flatter network designs offer the dual benefits of simplicity and more efficient utilization of high-speed switching infrastructures. When the number of subnets is reduced, configuration is simplified, which naturally eases network management. Furthermore, larger subnets tend to increase intraVLAN traffic within the enterprise, resulting in more direct switched connections, with a corresponding reduction in routing bandwidth requirements for interVLAN traffic.

While flatter networks are advantageous, security concerns or existing addressing conventions may dictate a more highly subnetted design. The MSS Server provides the flexibility to support either network design model, depending upon customer needs. Two mechanisms enable direct switched connections for intersubnet traffic. One of the mechanisms is the IP cut-through function of the SuperVLAN, while the other is the Next Hop Resolution Protocol (NHRP) support that is described in Chapter 11, "Next Hop Resolution Protocol Support for IP" on page 87.

NHRP is appropriate when endstations cannot be modified, as it can provide one-hop routing by establishing shortcut VCCs on behalf of the endstations. NHRP is also appropriate when NHRP Client software may be installed on endstations. In this case, NHRP enables zero-hop routing, where all routers are removed from the steady-state data path. Unfortunately, NHRP Client software is not yet widely available, and the initial support is limited to ATM-attached hosts in Classical IP environments.

IP cut-through complements the NHRP support by enabling zero-hop routing for stations accessible via emulated LANs (that is, ATM-attached hosts with LANE interfaces or legacy LAN stations represented by a LEC in an ATM bridge or LAN switch). In either form (NHRP or IP cut-through), zero-hop routing requires endstation changes. With IP cut-through, subnet address masks are modified at selected hosts so that those hosts will send IP ARPs for destinations that were previously reached through a router. The selective nature of the changes is significant, since the modifications may be limited to a few servers, where the benefits are most pronounced (assuming that the majority of the traffic flows from servers to clients), or gradually applied throughout the network as an incremental migration strategy to a flatter subnet model.

The preceding overview has attempted to highlight the value of the MSS Server's SuperVLAN feature. With this background, a more detailed description of the principal SuperVLAN components is presented in the following subsections.

10.1 Shortcut Bridging (SCB)

The ATM Forum is currently defining a standard for distributing LAN Emulation Services. This standard is called the LANE Network-to-Network Interface (LNNI) Specification. The ATM Forum began the LNNI work over two years ago in response to requirements from LANE users who need the reliability and scalability improvements associated with distributed LE services. For pragmatic reasons, the schedule for the LNNI Specification was recently decoupled from

the schedule for Version 2.0 of the LANE User-to-Network Interface (LUNI) Specification, in order to expedite LUNI V2.0. As a result, it now seems likely that the LNNI Specification will not be completed until late '97 or early '98.

Fortunately, the MSS Server's SuperVLAN feature provides distributed LE service functionality now in a standards-based fashion. The SuperVLAN does not rely on a proprietary protocol between service entities; instead, the distribution is accomplished by a new bridging variant called shortcut bridging. Shortcut bridging is a value-add enhancement to standard ELAN bridges, providing a convenient foundation for distributing the LE services. A simple ELAN bridging configuration is illustrated in Figure 20. Note that independent LES/BUS instances service each ELAN segment, and that stations on different ELANs communicate by establishing data direct VCCs to the bridge, which forwards interELAN traffic.

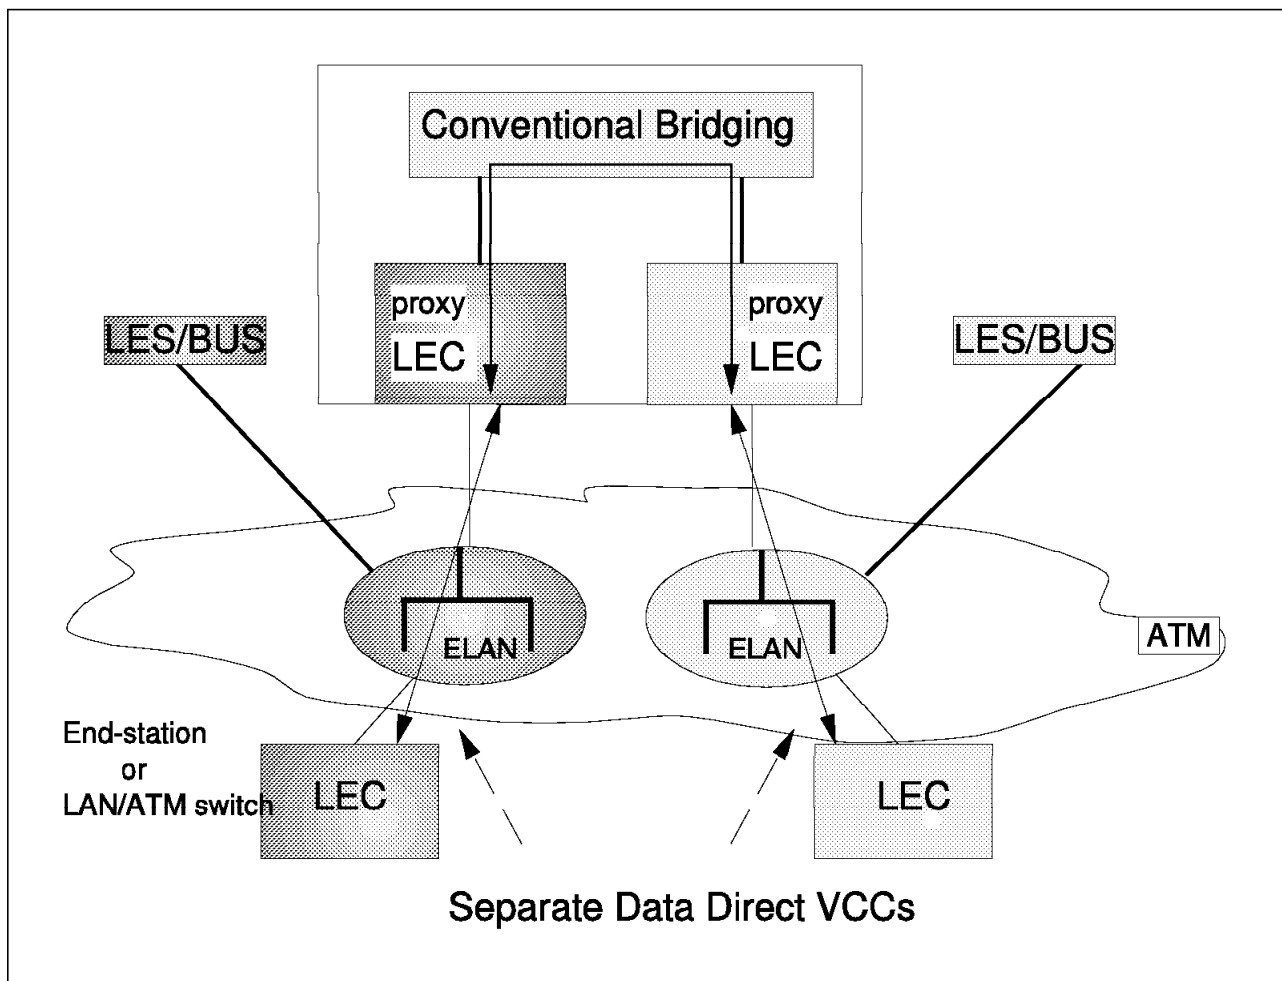


Figure 20. Conventional ELAN Bridging

Two aspects of conventional ELAN bridges are particularly significant from a LANE reliability perspective. First, if one of the LES/BUS instances fails, clients of the failed LES/BUS may be reassigned to any of the bridged ELAN segments without loss of connectivity. The MSS Server LECS facilitates the reassignment when a primary and backup LEC are configured for an ELAN. The LECS heuristically alternates client assignment between the primary and backup LES whenever a client returns to the LECS within a five minute period of its last configuration request. Thus, each of the LES/BUS instances may be

simultaneously serving clients for which it is the primary and acting as a backup for other ELAN segments.

The second significant reliability characteristic is that standard spanning tree protocols protect against failure of the bridge itself. Figure 21 depicts a configuration with redundant ELAN bridges. If the active bridge in MSS Server 1 fails, MSS Server 2 takes over the ELAN bridging function. During normal operation the spanning tree algorithm blocks parallel paths to avoid loops.

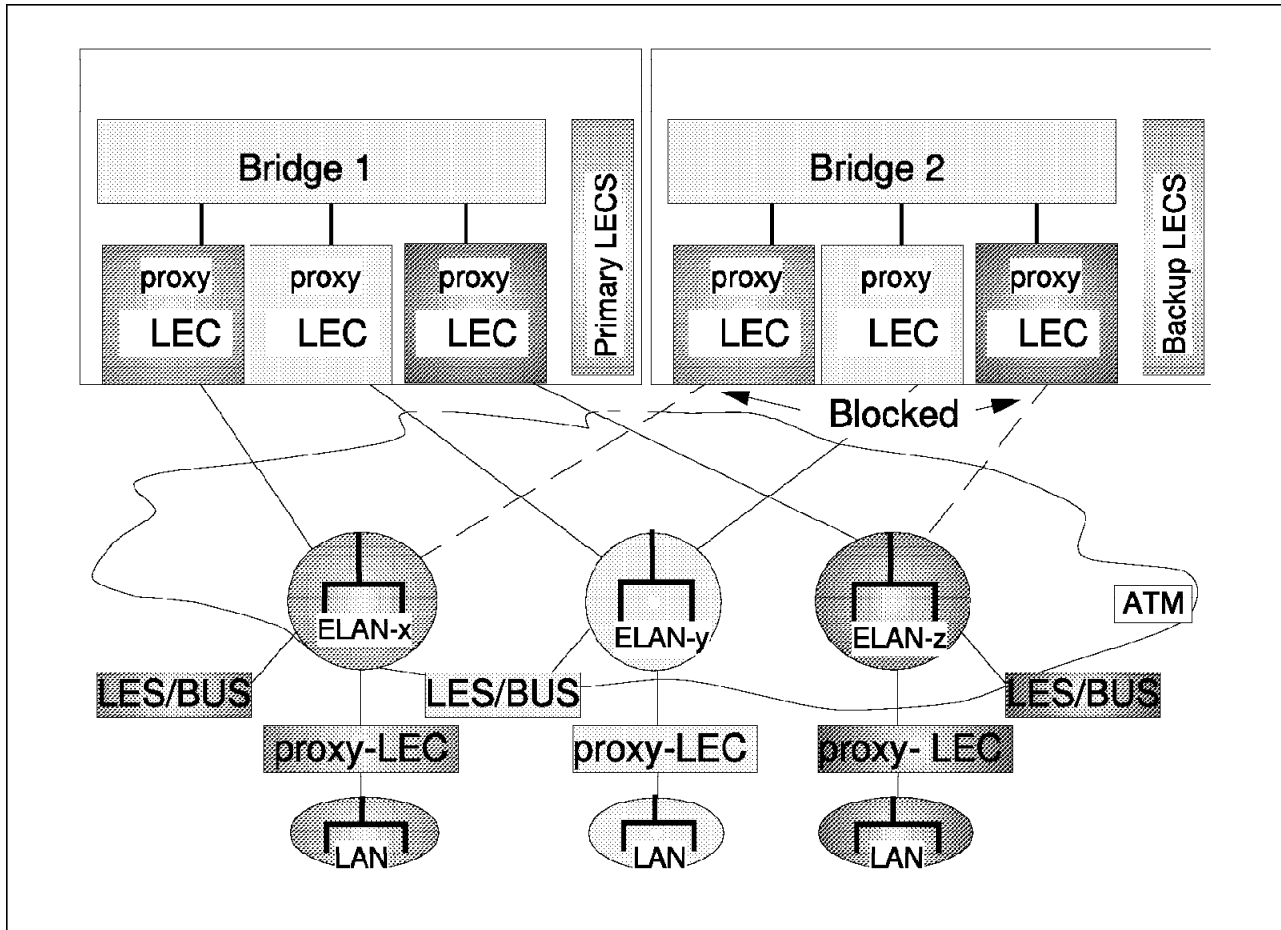


Figure 21. ELAN Bridging Redundancy

Distributing the LE services via ELAN bridging also addresses resource-related scalability issues that can hinder deployment of large-scale LANE networks. A single LES/BUS instance may be limited in terms of processing power, forwarding throughput, memory availability, or VCC capacity. However, with independent LES/BUSs serving each ELAN segment, the cumulative resources can be scaled up to meet the needs of the installation. Additionally, distributing the LE services can decrease client join times during high-activity periods, such as power-on sequences, by partitioning the signalling load across the ATM network.

Although ELAN bridging offers a solid foundation for distributing the LE services, it does not provide the entire answer. Shortcut bridging augments conventional ELAN bridging by enabling data direct VCCs to be established between clients on any of the bridged ELANs, which eliminates the bridge as a potential forwarding bottleneck for unicast traffic. By permitting clients to communicate directly over the switched ATM network, shortcut bridging makes it possible for

the collection of bridged ELANs to function as a single SuperELAN. In fact, the essence of shortcut bridging is its role as a robust control mechanism for distributing the LE services in a standards-based manner.

Figure 22 depicts a simple shortcut bridging configuration. Suppose Station x wishes to communicate with Station y, and Station x issues an LE_ARP_REQUEST for the MAC address of Station y. MSS Server LEC x receives the LE_ARP_REQUEST, and either answers the request (if it can) or forwards the request to the LES for ELAN y through proxy-LEC y. Assume LES y answers the LE_ARP_REQUEST. The LE_ARP_RESPONSE is received by MSS Server LEC y, and forwarded to LES x by proxy-LEC x. LES x forwards the LE_ARP_RESPONSE on to Station x, and Station x then sets up a data direct VCC to Station y.

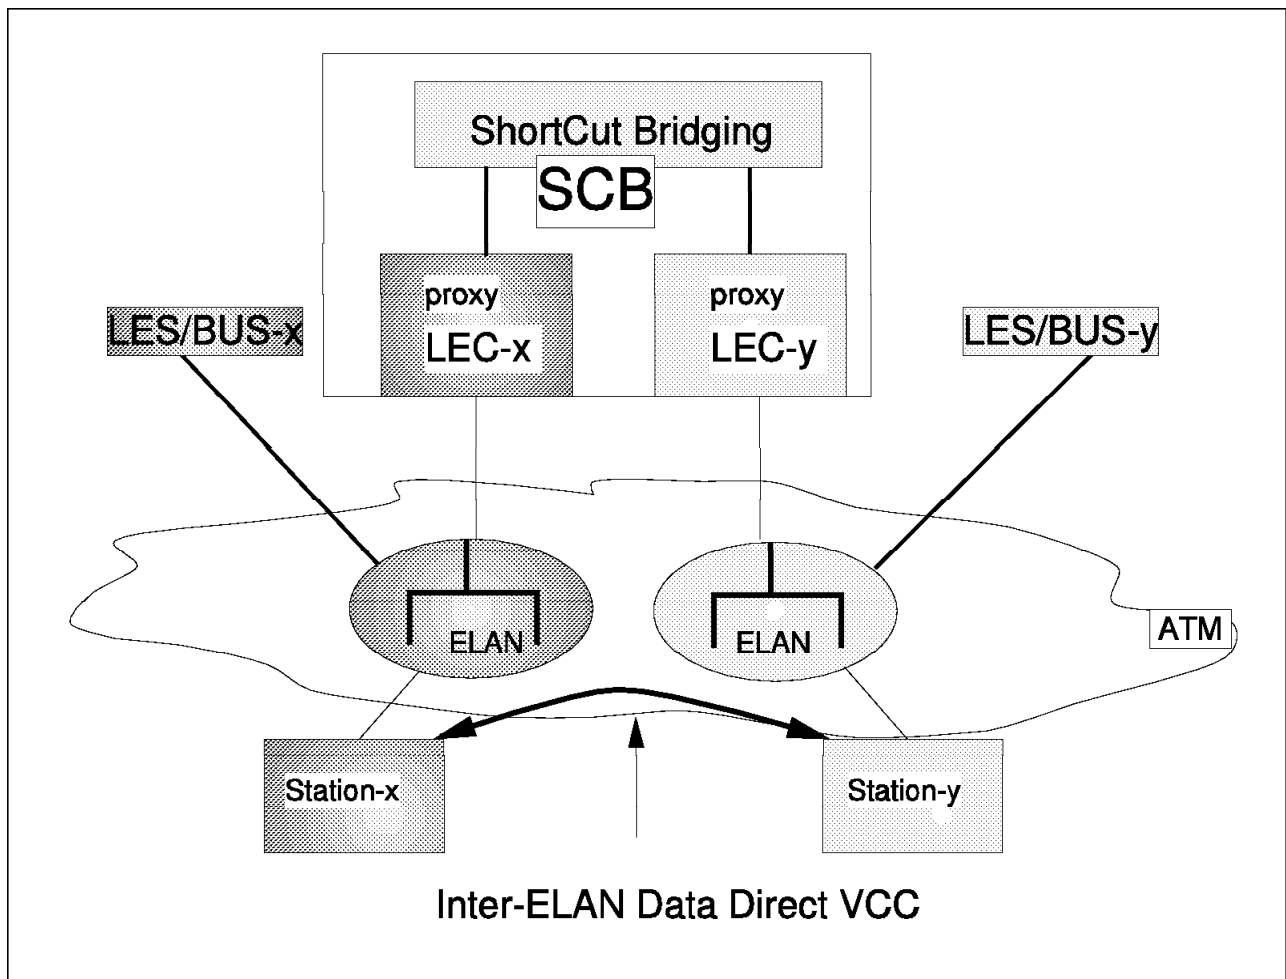


Figure 22. Simple Shortcut Bridging

The operation of the shortcut bridge differs from a conventional ELAN bridge in that it returns the ATM address associated with the destination station, while a conventional ELAN bridge would have returned the ATM address of proxy LEC x in response to Station x's LE_ARP_REQUEST. The shortcut bridge uses the bridge's MAC address database in conjunction with the LECs' LE ARP caches to answer LE_ARP_REQUESTs. The shortcut bridge also includes provisions to ensure that other LANE control frames (for example, LE_FLUSH_REQUESTs, LE_FLUSH_RESPONSEs, and LE_NARP_REQUESTs) are delivered to the

appropriate destinations. When forwarding data frames, shortcut bridge operation is similar to that of a transparent bridge.

Figure 23 contains a slightly more complex view of shortcut bridging. In this example, the shortcut bridge must be running the spanning tree protocol. To prevent a loop, the spanning tree protocol will deactivate one of the shortcut bridge ports on one of the MSS Servers. Assume that MSS Server 1 has been elected root of the spanning tree, that LEC f on MSS Server 3 has been deactivated by the spanning tree protocol, and that Station x wishes to communicate with Station y. As in the previous example, Station x sends a LE_ARP_REQUEST for Station y, and may send some initial unicast data frames destined for Station y to BUS x. If this were a transparently bridged network, all data frames from Station x to Station y would pass through MSS Server 1 and MSS Server 2. However, since it is a shortcut bridged network, the unicast data that traverses this path is limited to a few initial frames. After receiving the LE_ARP_RESPONSE, Station x will set up a data direct VCC to Station y, and subsequent unicast communications bypass the bridges.

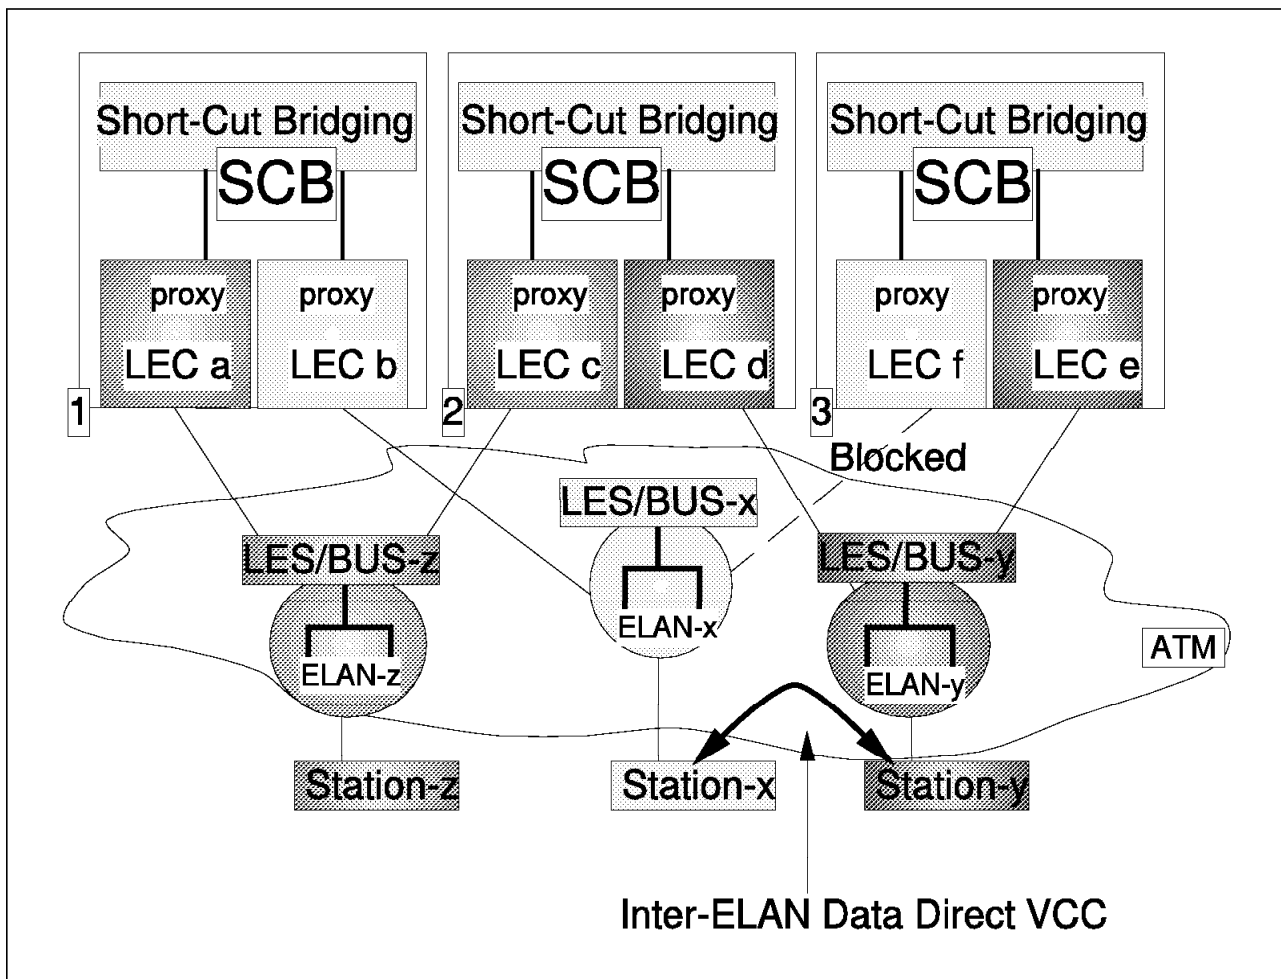


Figure 23. Shortcut Bridging

Since a SuperELAN is a collection of independent emulated LANs, LECIDs are not, in general, guaranteed to be unique across the constituent ELANs. It is indeed likely that distinct LESs will assign LECIDs from the same range (for example, starting at 1). This creates the potential for interoperability problems when two LECs with the same LECID attempt to communicate. The potential for

a problem arises because the LANE Version 1.0 Specification states that a LEC may discard frames received on any data direct VCC if the LECID in the frame matches the LEC's own LECID. Although this clause was undoubtedly intended to support filtering of frames echoed by the BUS, a compliant LEC might apply the same filtering rules to data direct VCCs. However, since LECID filtering slows data transfer and is completely unnecessary on data direct VCCs (as evidenced by the fact that LECID filtering on data direct VCCs is prohibited in the LANE Version 2.0 specification), it is unlikely that this option has been widely implemented. Nonetheless, the MSS Server LES may be configured with a range of LECIDs. As clients join the ELAN, they are assigned a LECID within the configured range. Thus, when MSS LESs serve each ELAN segment, the potential interoperability problem may be eliminated by configuring a unique LECID range at each LES.

In Release 1.1 of the MSS Server, a protocol is routed, and not bridged, if an address for the protocol has been defined on the bridge port interface.

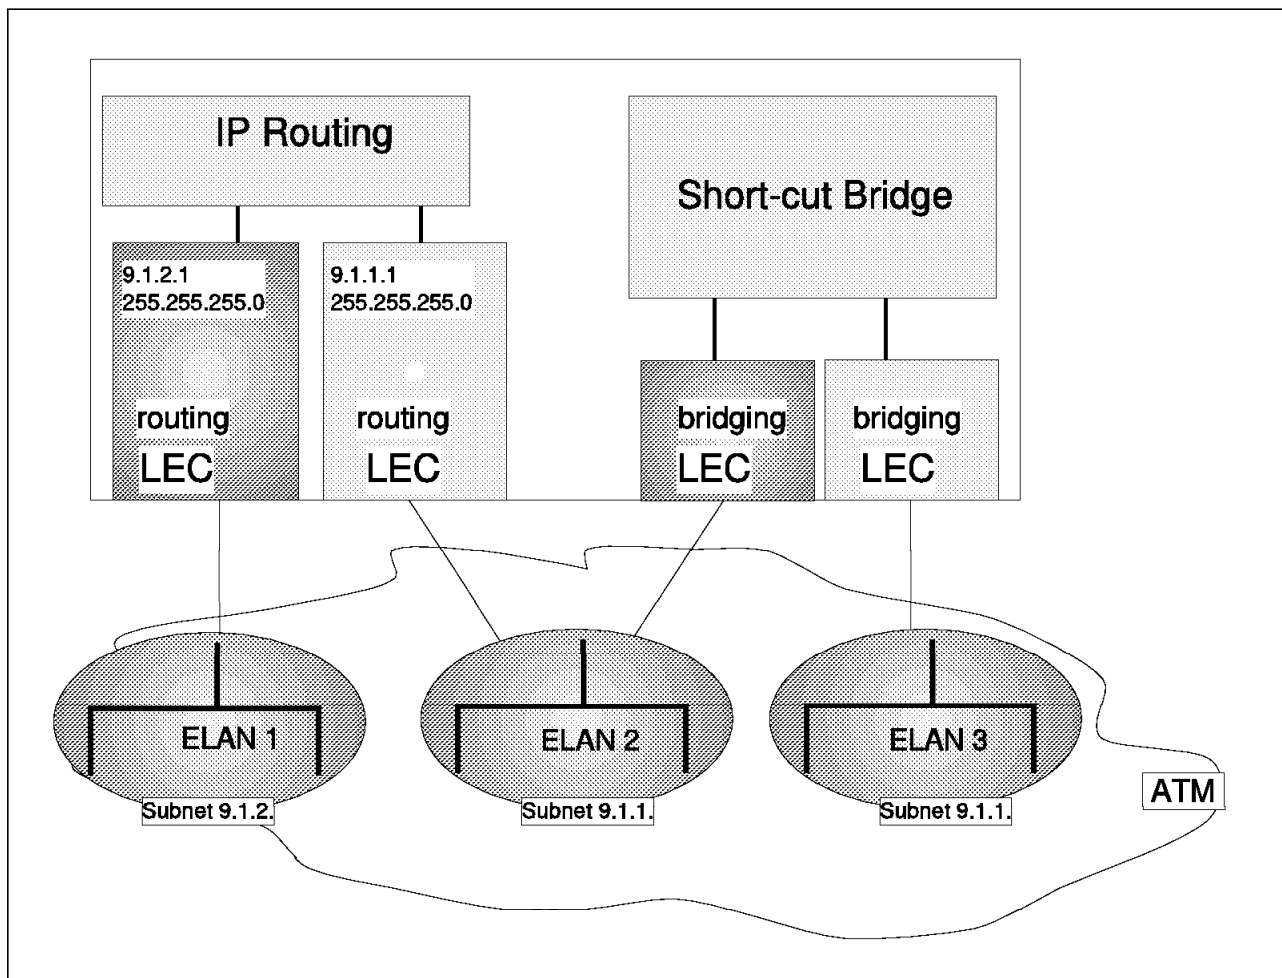


Figure 24. Concurrent Routing and Bridging

Important

Either all or no shortcut bridging LECs should be given an address for any given protocol; having some shortcut bridging LECs route a protocol and while others bridge can create serious connectivity problems.

Figure 24 provides a simple example where an additional LEC has been added to perform routing for an IP subnet that spans multiple ELAN segments. As a further example, assume that ELAN 1 also carries traffic for IPX Network 1, that ELAN 2 carries traffic for IPX Network 2, and that the MSS Server is to route, and not bridge, IPX between ELAN 1 and ELAN 2, while still bridging IP. This can be accomplished by defining IPX protocol addresses on the bridging LECs, which will transform these LECs into routing LECs for IPX.

Release 1.1 of the MSS Server provides shortcut bridging support for transparently bridged Ethernet and token-ring ELAN networks; Release 1.1 does not include shortcut bridging support for source-routed networks.

10.2 Broadcast Management

Two factors have traditionally limited the scalability of local area networks. Physical characteristics have constrained the size of any single segment, and broadcasts have limited the number of segments that can be effectively bridged together. ATM technology has addressed the physical constraints, and the MSS Server has been equipped with multiple functions to combat the effects of broadcast frames.

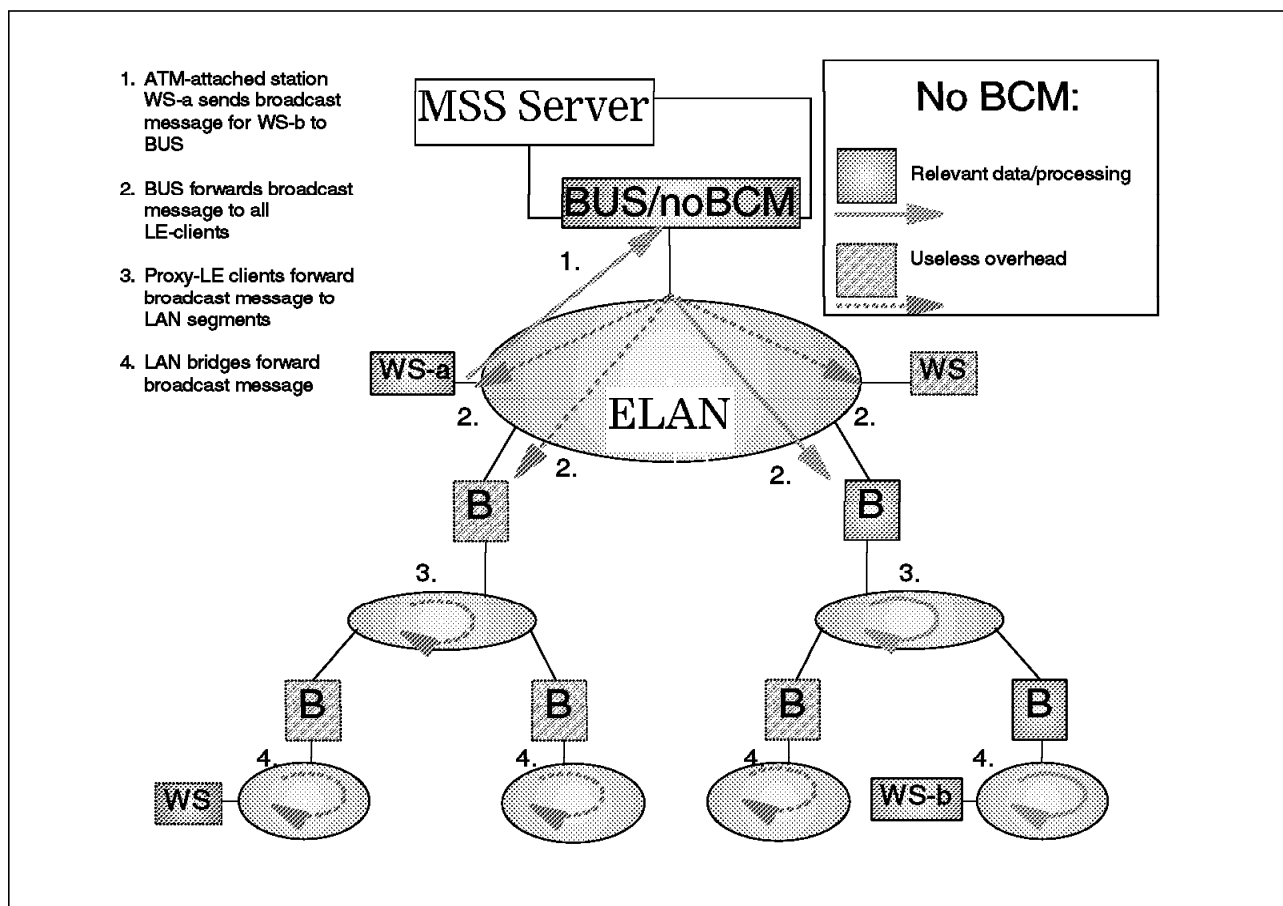


Figure 25. No Broadcast Management

10.2.1 Broadcast Management (BCM)

An IBM extension to the LANE BUS is the *broadcast manager (BCM)*. BCM will operate independently within each ELAN by intercepting broadcast and multicast frames sent to the BUS. BCM examines broadcast frames sent to the BUS. If possible, BCM transforms the broadcast frames into unicast frames and transmits the frames directly to the destination clients over multicast send VCCs. When BCM manages broadcast frames, both network bandwidth consumption and endstation perturbations are reduced.

10.2.1.1 Details on BCM

Today's LAN-based applications depend on broadcast techniques to locate partner applications before connections can be established. Unless a form of broadcast management is used, broadcast messages will be flooded to all of the (legacy and emulated) LAN segments that comprise the broadcast domain. In addition, the broadcast messages have to be processed on all of the stations attached to the broadcast domain. As can be seen from Figure 25 on page 62, this leads to much unnecessary traffic and station perturbation.

The principal goals of the BCM are two-fold: improve overall performance and efficiency by reducing network traffic and endstation processing associated with filtering nuisance frames, and enable practical deployment of larger ELANs. Broadcast management is especially useful when ATM connections traverse a WAN, where bandwidth is more precious, and when bridging legacy LAN segments to your ELANs. As the name implies the primary function of the BCM is to reduce the number of broadcast frames received by disinterested LECs and endstations. As can be seen from Figure 26 on page 64, BCM avoids most of the unnecessary traffic and station perturbation.

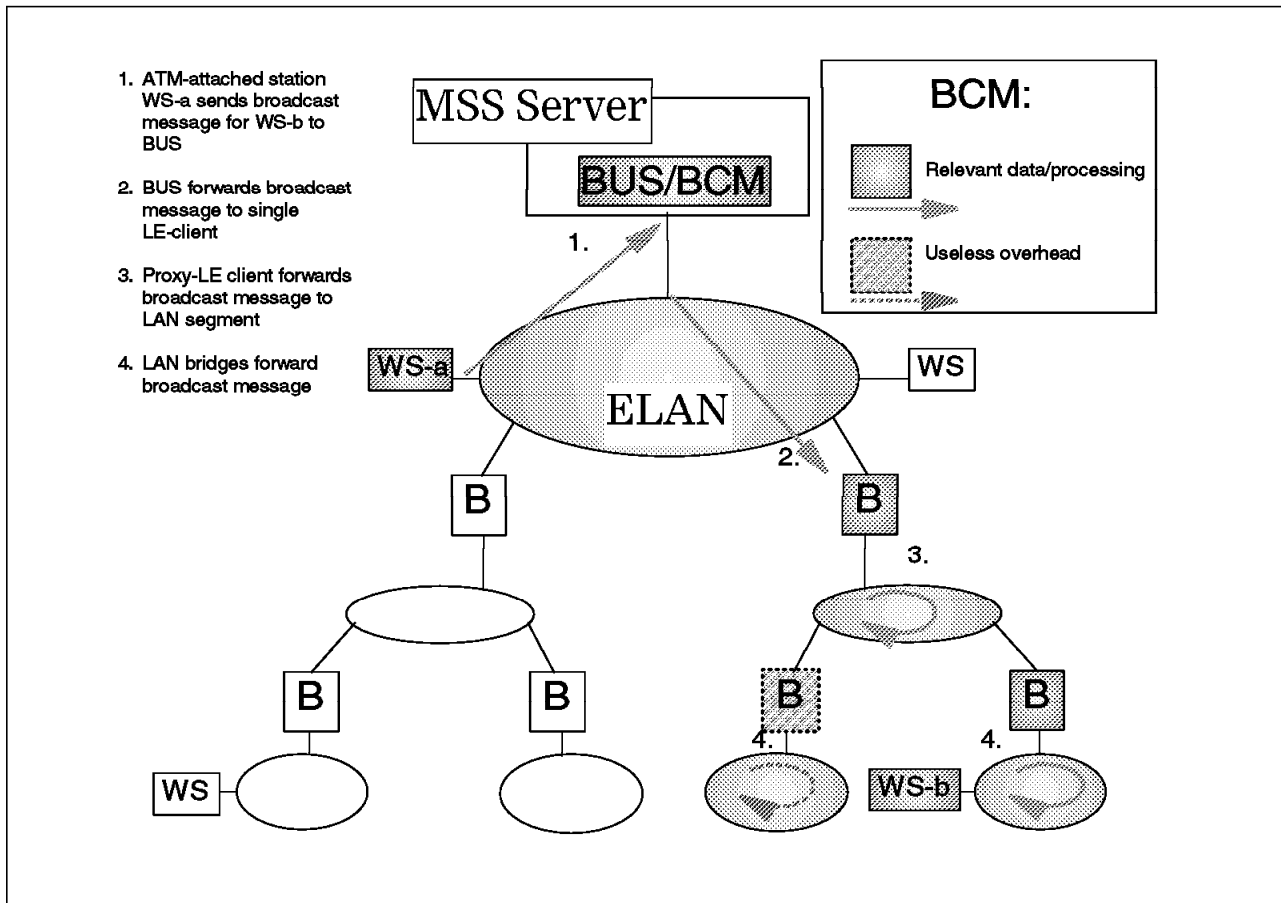


Figure 26. Using Broadcast Management

Broadcast manager operates independently within each ELAN by intercepting broadcast and multicast frames sent to the BUS. It performs a minimum amount of decoding of layer 2 and layer 3 information in order to perform its function. Support is provided for IP, IPX, and NetBIOS on Ethernet/802.3 or 802.5 ELANs.

1. IP Support

The broadcast manager scans all IP ARP requests and replies for the purpose of learning the location of IP addresses in the subnet containing this ELAN. The goal is for broadcast manager to take the workstation's broadcast ARP request frame and forward it as a unicast directly to the target station. Without broadcast manager, every device in the ELAN and bridged legacy LANs would unnecessarily have to receive and process all ARP broadcasts transmitted by every IP workstation.

2. IPX Support

Periodically, IPX routers and servers will broadcast their known network and service information. These periodically broadcasted packets, called informational RIP and SAP response packets, need only be received by other IPX routers and servers. Broadcast manager accomplishes this by intelligently forwarding the broadcast frames to these IPX devices. Without broadcast manager, every device in the ELAN and bridged legacy LANs would unnecessarily have to receive and process all RIP and SAP broadcasts transmitted by every IPX router and server.

3. NetBIOS

Many NetBIOS protocol frames are multicast to the NetBIOS functional address and must be processed by every NetBIOS device. Also, many NetBIOS devices tend to repeat the transmission of certain frame types (as much as ten times). This is a historical artifact to ensure that all devices received the frame in the event of bridge congestion or a heavily loaded endstation.

Broadcast manager attempts to associate unique NetBIOS names with MAC addresses and LECs by learning names from NetBIOS frames sent to the BUS. After a name is learned, all NetBIOS broadcasts for that name are forwarded to the appropriate LAN emulation client as a unicast frame. Broadcast manager also filters certain NetBIOS frames that are repeated when broadcast.

10.2.1.2 Source Route Management (SRM)

Source route management (SRM) is an additional function to BCM for 802.5 ELANs. Broadcast manager learns the bridged token-ring topology behind each LEC by learning the Routing Information Field (RIF) of frames received by the BUS. Broadcast manager uses this knowledge to specifically source route all route explorer (ARE) or spanning tree explorer (STE) frames through a bridged token-ring network. This reduces the amount of unwanted explorer traffic transmitted on the token-ring network from the ATM emulated LAN. As can be seen from Figure 27, SRM avoids all unnecessary traffic and station perturbation.

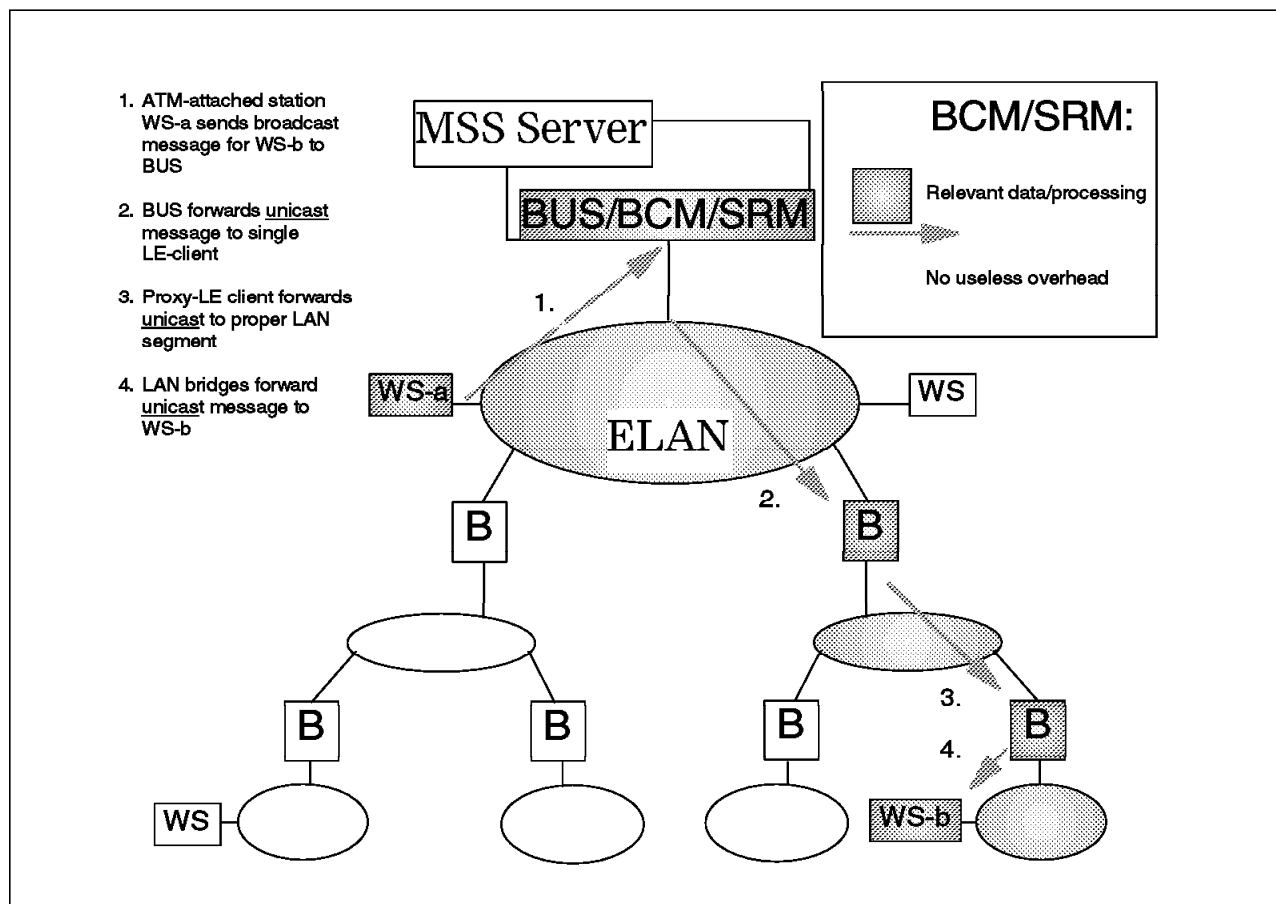


Figure 27. Using Source Route Management (SRM)

10.2.2 Bridge Broadcast Management (BBCM)

In Release 1.1 of the MSS Server, broadcast management function has been added to the MSS bridging function. Bridging BroadCast Manager (BBCM) can transform many broadcast frames received by the bridge into unicast frames, thus lessening their effects on network performance. Like BCM, BBCM dynamically learns bindings between protocol-layer addresses and MAC addresses. Subsequent broadcasts destined for learned protocol addresses can then be transformed into unicast frames directed toward a single ELAN. Thus, BBCM enables hierarchical broadcast management, with BCM managing broadcasts within ELANs and BBCM managing broadcasts between ELANs.

BBCM applies when a NetBIOS network or an IP subnet spans multiple bridged segments. A simple example of a bridged IP subnet is shown in Figure 28 on page 67. Assume Server z is heavily accessed and that BBCM and BCM 2 have both already learned the IP-to-MAC address mapping for Server z. Now suppose Client x ARPs for Server z. First, BBCM records the mapping between IP address 9.1.1.1 and MAC address MAC x for future use. Then, since BBCM knows the mapping for the destination IP address, 9.1.1.2, the broadcast ARP packet is transformed into a unicast packet destined for MAC z and forwarded onto ELAN 2. Similarly, if Client y ARPs for Server z, BCM 2 transforms the broadcast into a unicast packet that is forwarded directly to the LEC representing Server z. BBCM operates in an analogous manner when managing NetBIOS name queries.

BBCM is not only an integral component of the SuperVLAN support for large-scale LANE networks, but is also a generic bridging enhancement with an independent utility. BBCM can be activated independently for NetBIOS and/or IP. No support is provided for IPX. No BBCM equivalent exists for the BUS/BCM source route management (SRM) function. It is a global function and applies to all bridging ports.

10.2.3 Dynamic Protocol Filtering (DPF) for PVLANS

Dynamic Protocol Filtering (DPF) is another facet of the MSS Server's broadcast management capabilities. DPF makes the SuperELAN a SuperVLAN by dynamically partitioning the bridged network into several protocol-specific virtual LANs, or PVLANS. DPF then uses the PVLANS to limit the scope of broadcast/multicast frames that are normally forwarded over all active bridge ports.

DPF monitors the broadcast/multicast traffic received over each bridge port to learn the protocols and subnets being used on that segment. In Release 1.1 of MSS Server, the user controls the set of PVLANS to be managed by DPF via configuration. DPF support is provided for IP, IPX, and NetBIOS. A single NetBIOS PVLAN is supported, while multiple PVLANS may be configured for IP subnets and IPX networks. DPF manages the forwarding domain for each configured PVLAN, where the forwarding domain is the subset of bridge ports on which traffic for that PVLAN is being received. Transmission of broadcast/multicast frames for a particular PVLAN is limited to the forwarding domain for that PVLAN. Thus, the scope of broadcast/multicast packets for a particular protocol/subnet is reduced to those segments that are actually utilizing the protocol/subnet. More specifically, IP ARP traffic only travels over ELANs associated with that IP subnet, IPX broadcast/multicast traffic only travels over ELANs associated with that IPX network, and NetBIOS broadcast/multicast traffic only travels over ELANs using NetBIOS.

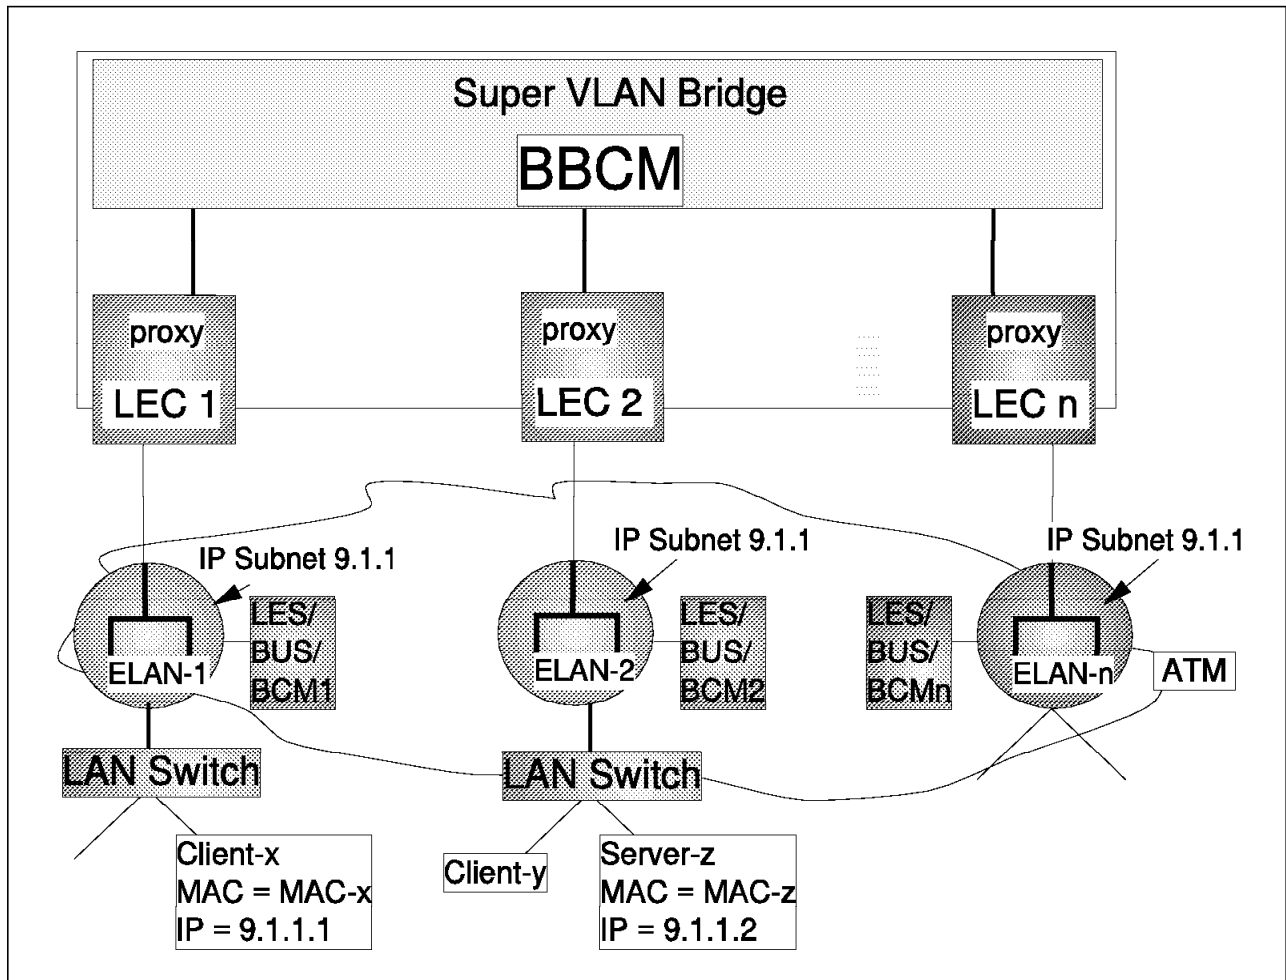


Figure 28. Bridging Broadcast Manager (BBCM)

DPF works in concert with BBCM to manage the propagation of broadcasts in large-scale LANE networks. If BBCM is successful in converting a broadcast into a unicast frame, then DPF is not needed. However, DPF can manage broadcasts in some cases where BBCM cannot. DPF operates at a much higher level of granularity than BBCM. Since DPF controls traffic on a per PVLAN basis, while BBCM controls broadcasts on a per station basis, DPF has a broader view of the network. Thus, in situations where BBCM does not have a mapping for the destination station, DPF can still limit the broadcast to the appropriate PVLAN. Additionally, in Release 1.1 of the MSS Server, DPF also supports a broader set of frame types; DPF is applied to IPX and NetBIOS frames with broadcast or multicast destination MAC addresses, and IP ARPs.

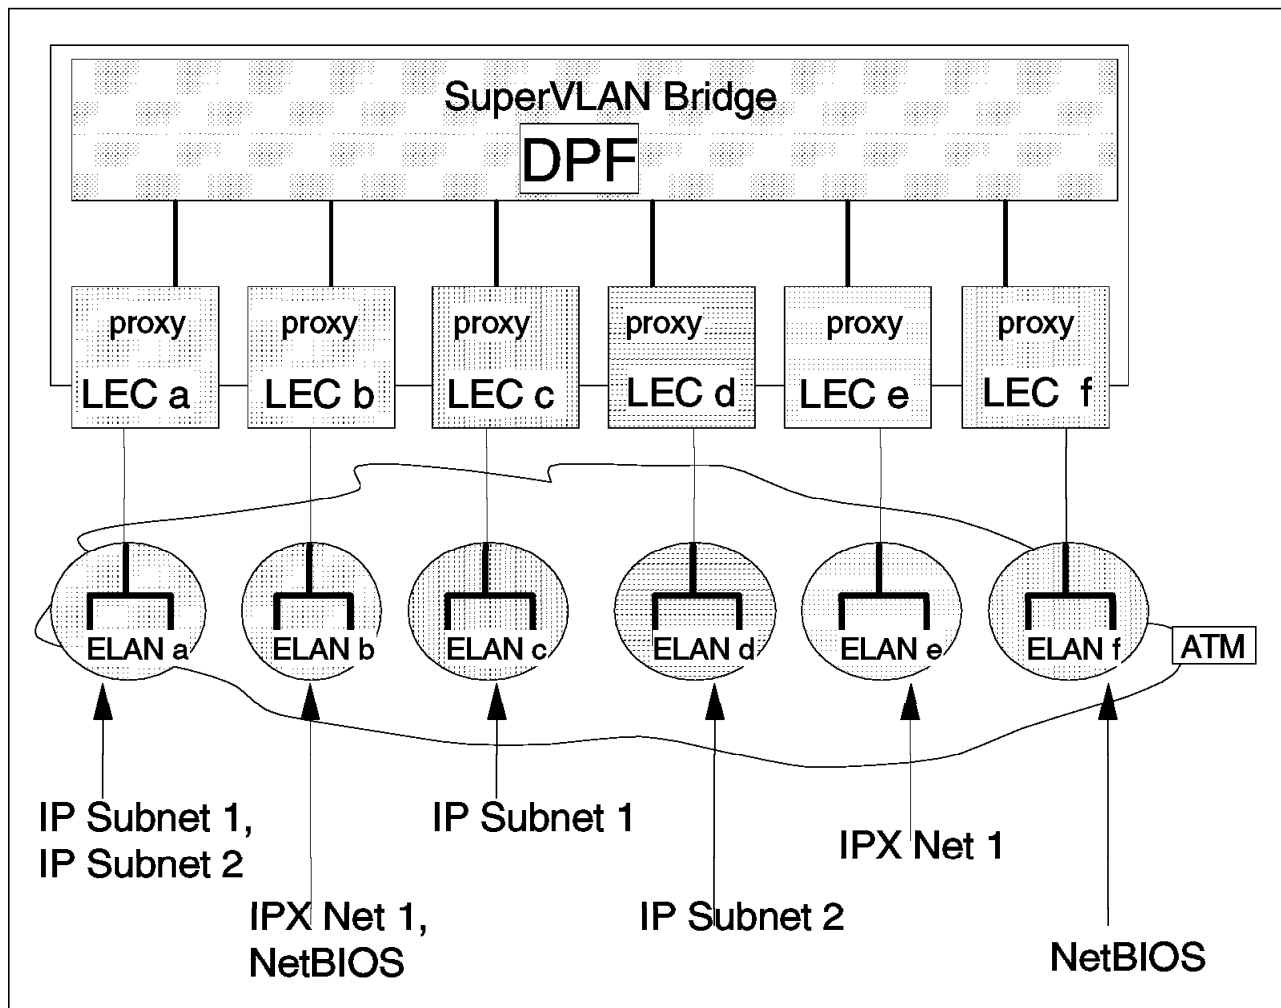


Figure 29. Dynamic Protocol Filtering

Figure 29 shows an example DPF configuration. In the example environment, multiple protocols/subnets are intermixed across a bridged environment. Assume that IP Subnet 1, IP Subnet 2, IPX Network 1, and NetBIOS are configured for DPF, and that MSS Server LECa and LECc have been identified as members of IP Subnet 1's forwarding domain based on traffic received on those ports. IP ARP traffic for destinations in IP subnet 1 is then restricted to ELAN a and ELAN c. Likewise, IPX and NetBIOS traffic is also restricted to a subset of the ELANs. Assuming that IP Subnet 3 has not been configured for DPF, any IP ARP traffic for Subnet 3 would be forwarded over all MSS Server bridge ports.

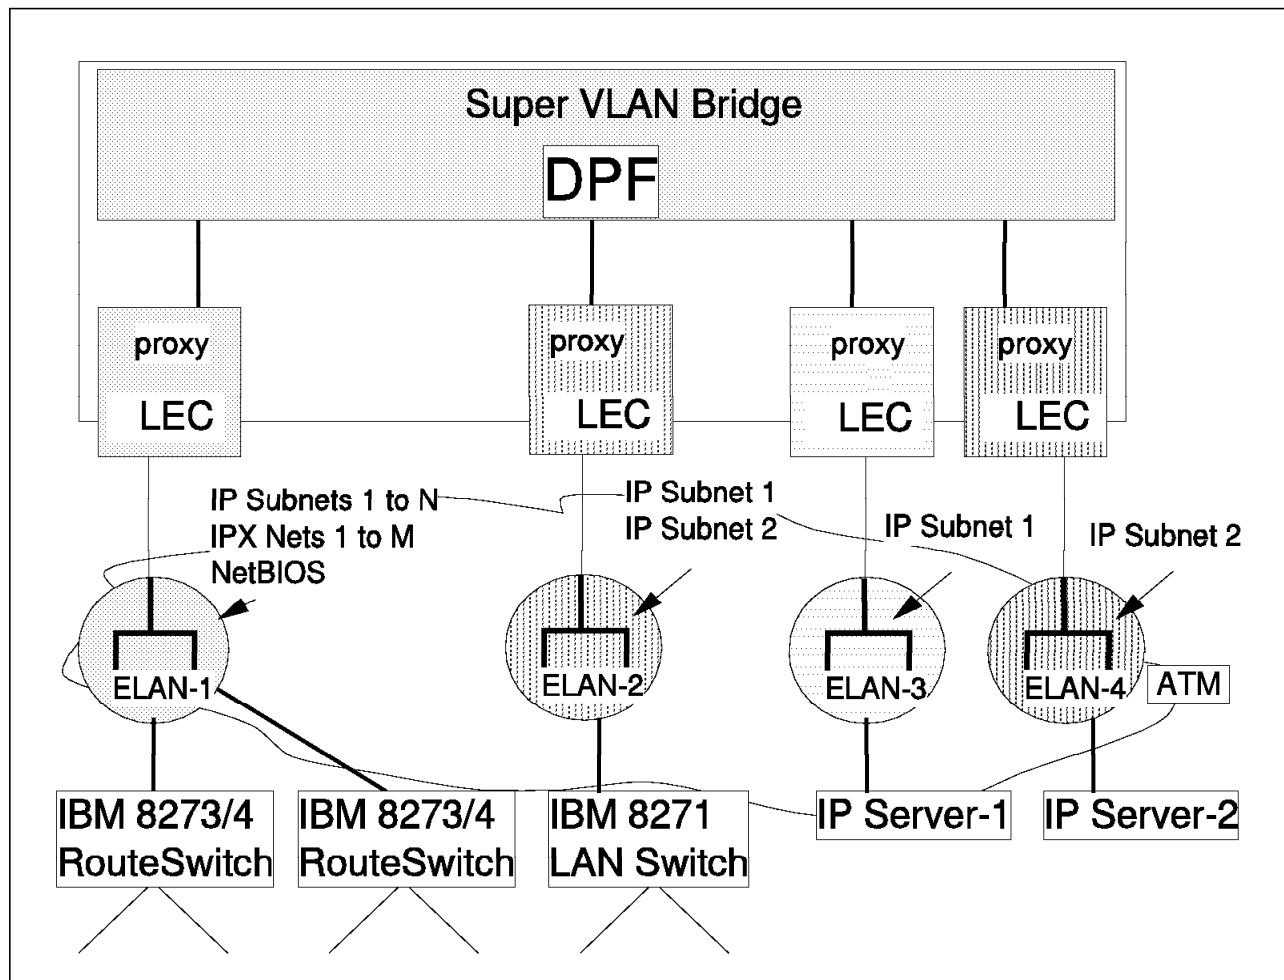


Figure 30. Integrated Use of VLANs

The SuperVLAN capabilities are useful both when ELANs are dedicated to a single PVLAN and when traffic for multiple PVLANS is carried on the same ELAN. There are several common situations where a single LEC carries traffic for multiple protocols/subnets. One situation arises when a multiprotocol host accesses the ATM network through a single LEC. Suppose an ELAN is composed of hosts that belong to the same IP subnet, but a few of these hosts also need access to an IPX server residing on another ELAN segment (which might be dedicated to IPX). A single LEC cannot join both an IP ELAN and an IPX ELAN. However, these hosts can benefit from the SuperVLAN capabilities that minimize broadcast traffic on the ELAN while still enabling direct communication with the IPX server. Similar situations occur when ATM bridges or LAN switches with port-based VLAN policies represent diverse sets of legacy LAN hosts. In these cases, the MSS Server is able to extend PVLAN function to products such as the IBM 8281 ATM Bridge and the IBM 8271/2 LAN Switches.

The SuperVLAN is also particularly useful when used in conjunction with the IBM 8273/4 RouteSwitch products, which contain powerful VLAN capabilities in their own right. Each RouteSwitch typically maps a group of VLANs to a single backbone ELAN that interconnects a collection of RouteSwitch devices as shown in Figure 30. This figure is intended to illustrate how the SuperVLAN can integrate a variety of VLAN capabilities into a cohesive network design. In this example, the RouteSwitches are supporting several VLANs, many of which are local to the collection of RouteSwitches (that is, all the members of these local

VLANs reside on legacy LAN segments behind a RouteSwitch). Since ELAN 1 is carrying the interRouteSwitch traffic, no other stations are attached to ELAN 1. Each RouteSwitch confines the traffic to the appropriate VLANs on the legacy LAN side, and the SuperVLAN ensures that traffic for local RouteSwitch VLANs is isolated to ELAN 1. The example also assumes that universal access is required to ATM-attached servers on IP Subnet 1 and IP Subnet 2. (That is, clients of these servers are located behind the IBM 8271 LAN Switch and every RouteSwitch.) Note that each of the servers, IP Server 1 and IP Server 2, resides on a dedicated ELAN. This design minimizes disturbances due to nuisance traffic, while preserving the capability to establish direct ATM connections with any of the LECs in the SuperVLAN.

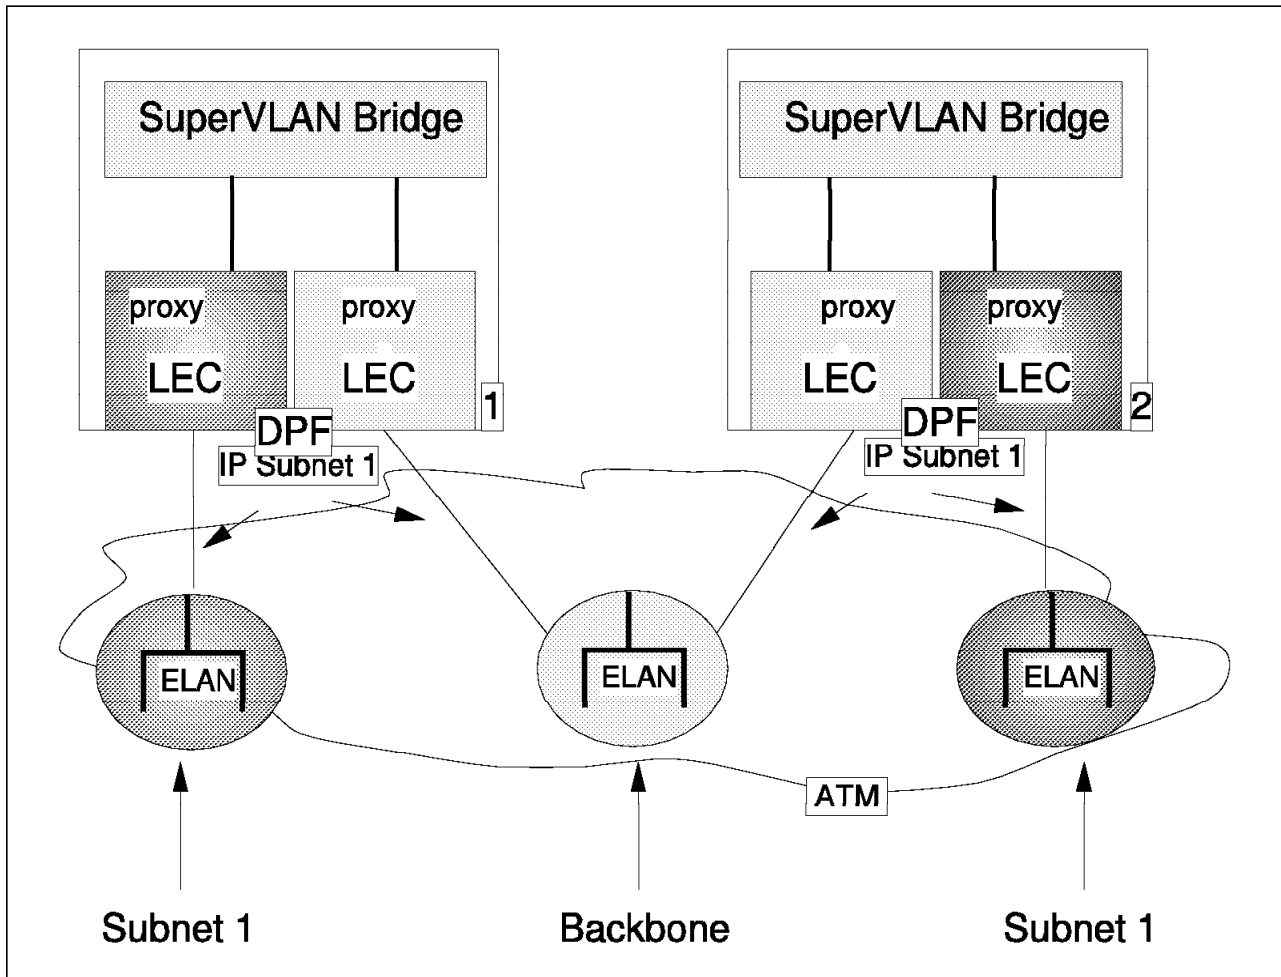


Figure 31. Partitioning a PVLAN across a Backbone Network

In the absence of traffic, a port's membership in a PVLAN forwarding domain will age out after a configurable period of time. The default aging times are 10,000 minutes for IP PVLANS, 10 minutes for IPX PVLANS, and 5000 minutes for NetBIOS. Alternatively, the aging times may be configured to infinity so that a port's membership in a PVLAN's forwarding domain never expires after being learned once.

Since membership in PVLAN forwarding domains is based on frame transmissions, it may be necessary to statically configure ports providing access to quiet stations, such as network printers or network monitoring equipment. Additionally, if a PVLAN is split across a backbone ELAN, as depicted in

Figure 31, then the PVLAN must be statically configured on the ports providing access to the backbone.

Like BBCM, DPF is a generic bridging enhancement. It can be used separately or in conjunction with the other SuperVLAN features.

10.2.4 IP Cut-Through Support for Zero-Hop Routing

DPF also includes support that enables IP stations to establish intersubnet data direct VCCs. This IP cut-through support will be described in the context of the example configuration depicted in Figure 32. The key to IP cut-through mode is that the IP address mask at certain stations is configured to be less restrictive than the actual subnet mask. In the example, the station with IP address 9.1.2.2 (which happens to be a server) is configured with a subnet mask of 255.255.0.0 although the actual subnet has a mask of 255.255.255.0.

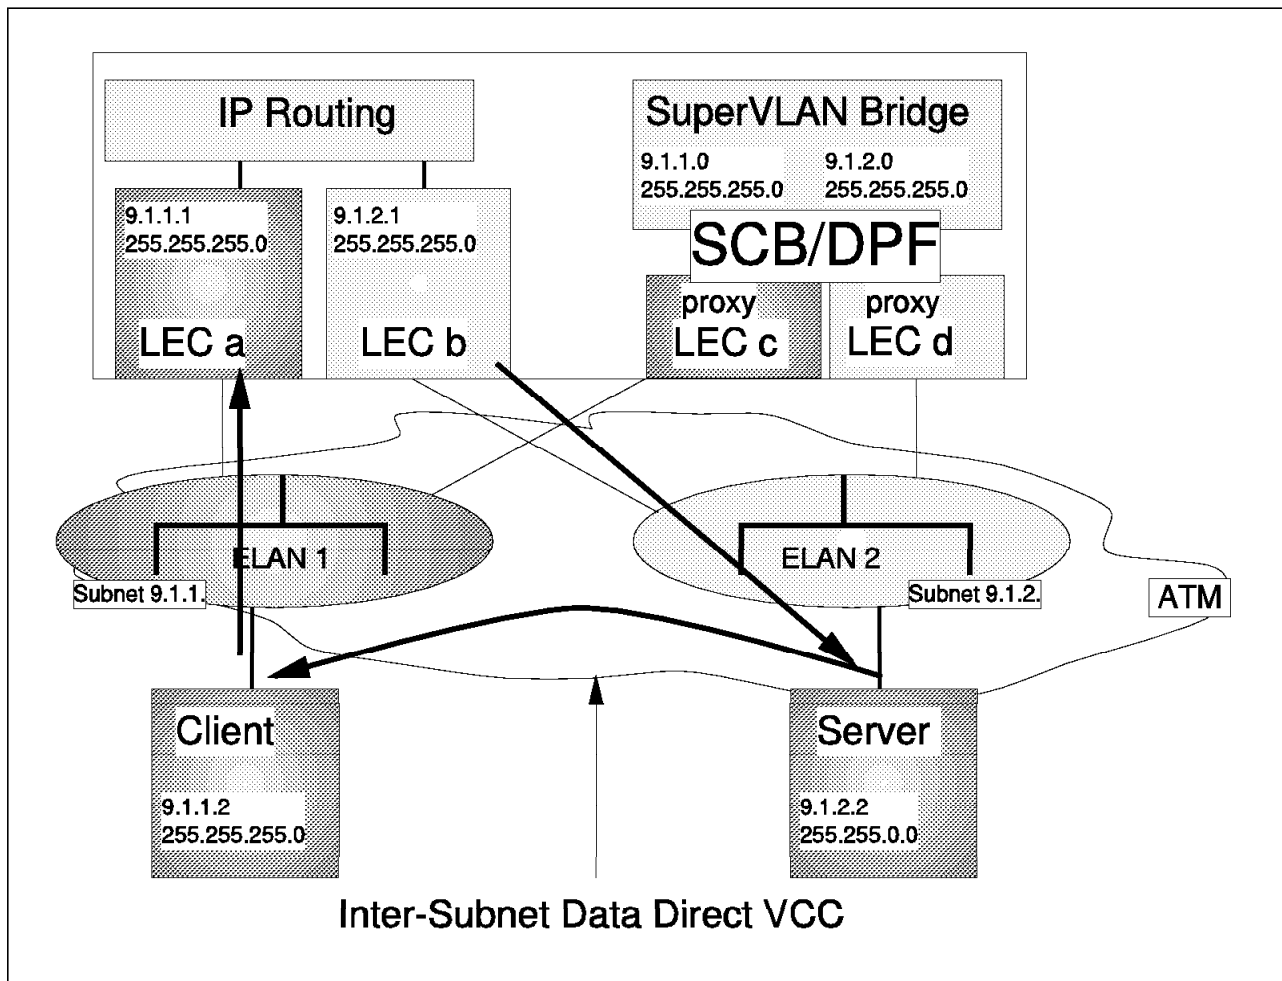


Figure 32. IP Cut-Through

First, let's look at what happens when the client with IP address 9.1.1.2 wants to communicate with the server at IP address 9.1.2.2. The client realizes that the server is not on its subnet, and sends packets destined for the server to its default gateway (which has an IP address of 9.1.1.1). The MSS server routes these packets to station 9.1.2.2. Thus, the client-to-server flow is routed through the MSS Server.

Now, let's examine the server-to-client flow. The server believes that the client is on the same subnet, so it transmits an IP ARP for 9.1.1.2. Normally, this IP ARP would not be forwarded by shortcut bridge LEC c because the source subnet (9.1.2) has not been learned on ELAN 1. However, in IP cut-through mode, the IP ARP is forwarded by shortcut bridge LEC c because the destination subnet (9.1.1) is known on ELAN 1. The client issues an IP ARP reply that is bridged to ELAN 2 through shortcut bridge LEC d. The server then sets up a data direct VCC to the client (if one does not already exist). Thus, the server-to-client flow is not routed. This is significant because the volume of traffic is often much greater in the server-to-client direction. It is also important to note that PVLAN broadcast domains are preserved in IP cut-through mode, since IP ARPs for destinations on subnet 9.1.2 are not forwarded onto ELAN 1 and, similarly, IP ARPs for destinations on subnet 9.1.1 are not forwarded onto ELAN 2.

To utilize IP cut-through mode, the system administrator must manipulate endstation IP subnet masks. The MSS Server provides both a bridged and a routed connection between IP stations, and individual stations choose their connection based on their subnet mask. The administrator has the flexibility to decide when a bridged connection is desirable by modifying the IP address mask of selected stations. This type of IP subnet mask manipulation may also be used in conventional bridged environments, but when used in a shortcut bridged environment, the bridged communications become direct communications over the ATM switch fabric.

— IP Cut-Through and IBM 8273/4 —

IP cut-through is compatible with ATM-attached hosts, with ATM bridges such as the IBM 8281, and with LAN switches such as the IBM 8271/2. However, care must be taken when using IP cut-through in conjunction with IBM 8273/4 RouteSwitch products. Currently, IP cut-through is not possible when the source and destination reside on different RouteSwitch VLANs (since the RouteSwitch VLAN policies prohibit direct interVLAN communications). IP cut-through transmission may be enabled and disabled on a subnet basis.

10.2.4.1 IP Cut-Through Benefits

As can be seen in Figure 32 on page 71 in IP cut-through mode, an asymmetrical connection between client and server results. The benefits are the following:

- Broadcast control

The IP broadcasts from the client stations connected to ELAN 1 end at MSS LE client A, and do not, due to DPF on LE client C, generate overhead on ELAN 2.

- Efficient routing

In most cases server-to-client traffic is higher than the traffic in the reverse direction. Because of the IP cut-through, the IP routing between server and client is very efficient.

A first alternative to this network construct is to discontinue the routing ports (LEC a and b) on MSS Server, use the same (large) subnet mask on all stations, and use BBCM on the MSS (LEC C and D) bridging ports for broadcast management. BBCM will avoid unnecessary broadcasting, while, due to their larger subnet mask, client and server stations can communicate directly in both directions.

A second alternative is to merge ELAN 1 and ELAN 2 into a single ELAN, discontinue the routing and bridging ports on MSS Server, use the same (large) subnet mask on all stations, and use BCM on the BUS of this ELAN for broadcast management. BCM will avoid unnecessary broadcasting, while client and server stations can communicate directly in both directions.

10.2.5 Configuring DPF and/or IP Cut-Through

Figure 33 depicts the configuration options available when configuring dynamic protocol filtering for IP.

Subnet	Mask	Name	Enabled	IP-Cut-From	IP-Cut-To	Age
172.16.0.0	255.255.255.0	SuperVL	enable	enable	disable	10000
172.16.1.0	255.255.255.0	SuperVL	enable	disable	enable	10000

IP Address: IP Mask: Generate -> IP Subnet:

Name: Age (in minutes):

☒ Enable this VLAN
☒ Enable IP-Cut-Through from this VLAN
☐ Enable IP-Cut-Through to this VLAN

Configure Individual Ports

Interfaces	Port Type
ET(1)	Auto-Detect
ET(3)	Auto-Detect

Reset Ports

Add Change Delete

Figure 33. DPF Configuration Options for IP

IP protocol VLANs (PVLANS) are defined by entering an IP address and a subnet mask. Per PVLAN you can configure if the VLAN is enabled and if IP cut-through to and/or from this VLAN is allowed. In addition, you specify the VLAN membership per bridging port.

The VLAN membership can be:

- Auto-detect (that is, dynamic, based on the source address of IP traffic received)

Note: Dynamic membership is subject to an aging mechanism. An aging timer of 0 sets the aging period to infinity.

- Inclusive (Port is always member of VLAN.)
- Exclusive (Port is never a member of the VLAN.)

The MSS Server makes forwarding decisions of IP ARP broadcast based on the configuration parameters entered in Figure 33. A port's membership in the forwarding domain of an IP subnet PVLAN is based on the *source* IP address of received IP ARP frames. IP ARP frames are handled in one or two modes, normal mode or cut-through mode. In normal mode, IP ARP frames are

forwarded to all ports (except the input port) in the forwarding domain of the source subnet PVLAN, and flooded to all ports if the source subnet is unknown.

IP cut-through is slightly more complex. In this mode, IP ARP frames are forwarded based on the *destination IP subnet*. The exact IP ARP forwarding rules are depicted in Figure 34.

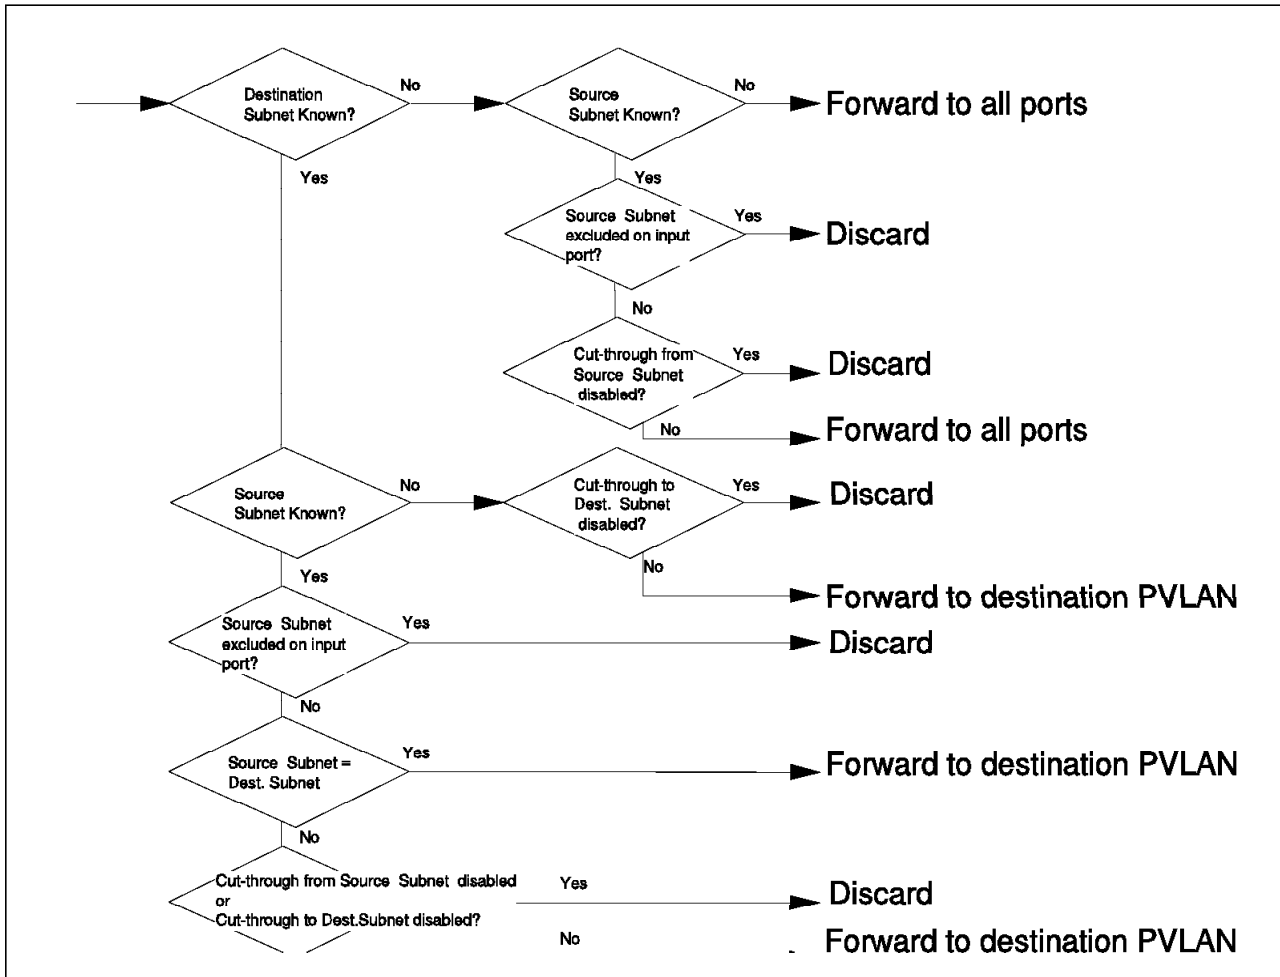


Figure 34. DPF IP Cut-Through Logic Flow

10.2.6 Configuring DPF for IPX

Figure 35 on page 75 depicts the configuration options available when configuring dynamic protocol filtering for IPX.

IPX Network Number	Name	Enabled	Age
1		enable	10

Network Number: Name:

Age (in minutes): ☒ Enable this VLAN

Configure Individual Ports

Interfaces	Port Type
ET(1)	Auto-Detect
ET(3)	Auto-Detect

Figure 35. DPF Configuration Options for IPX

IPX protocol VLANs (PVLANS) are defined by entering an IPX network number. Per PVLAN you can configure if the VLAN is enabled. In addition, you specify the VLAN membership per bridging port.

The VLAN membership can be:

- Auto-detect (that is, dynamic, based on the source address of IPX traffic)

Note: Dynamic membership is subject to an aging mechanism. An aging timer of 0 sets the aging period to infinity.

- Inclusive (Port is always member of VLAN.)
- Exclusive (Port is never a member of the VLAN.)

The MSS Server makes forwarding decisions of IPX broadcast traffic based on the configuration parameters entered in Figure 35 and the logic displayed in Figure 36 on page 76.

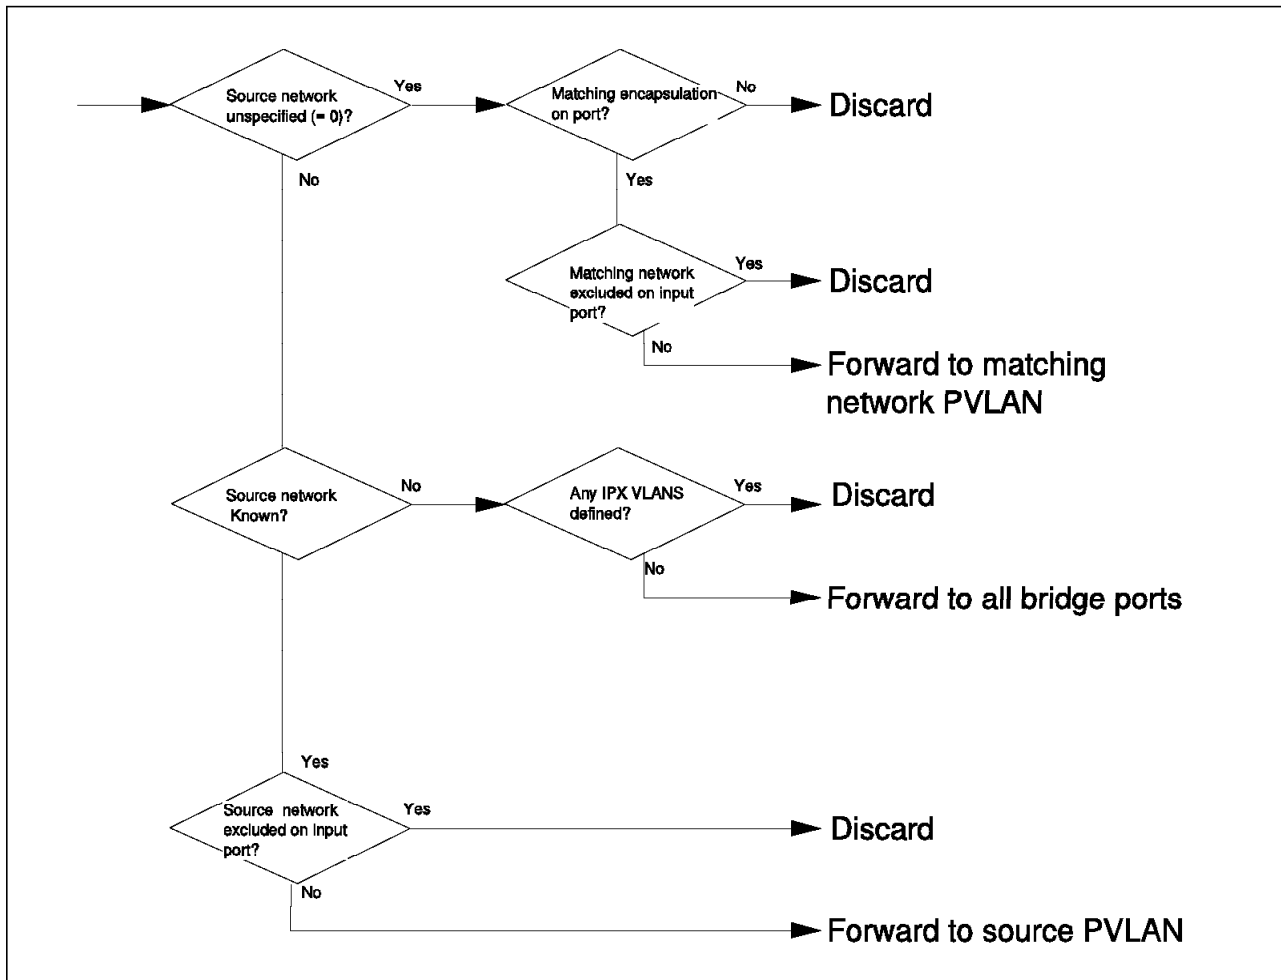


Figure 36. DPF IPX Logic Flow

Multiple IPX networks are permitted on a port as long as they use different encapsulation. IPX frames may also have an unspecified source network address (source network address field set to 0). When the source network is unspecified, DPF assumes that the network address is the unique IPX network on that port using the frame's encapsulation. If no IPX network on that port uses the frame encapsulation, then the frame is discarded.

10.2.7 Configuring DPF for NetBIOS

The PVLAN support for NetBIOS is simpler than that for IP or IPX. Since NetBIOS does not use subnets, there is only one NetBIOS forwarding domain. If a NetBIOS frame with a broadcast or multicast destination MAC address is received on a port, then that port is added to the NetBIOS forwarding domain. The membership is subject to an aging mechanism. An aging timer of 0 sets the aging period to infinity.

10.3 Multiple SuperVLAN Instances

Currently, only one bridge instance is supported per MSS Server. This is problematic with respect to supporting multiple independent SuperVLAN instances on a single MSS Server, which would be advantageous for creating virtual private networks in service provider environments. As a first step toward the desired goal, a SuperVLAN identifier has been associated with each

short-cut bridge port. This allows shortcut bridge ports to be grouped and partitioned into separate domains as illustrated in Figure 37 on page 77.




Figure 37. Multiple SuperVLAN Instances

Bridged communication is not allowed between different shortcut bridge domains. For example, frames received on a port with a given SuperVLAN ID are not forwarded on shortcut bridge ports with different SuperVLAN IDs. Conventional (that is, non-shortcut) bridge ports are considered to be members of all shortcut bridge domains.

Although this functionality should prove useful in many situations, dependencies on the single bridge instance introduce restrictions that may be prohibitive in some environments. First, only one spanning tree instance may be run. Ideally, it should be possible to run a spanning tree per shortcut bridge domain. However, at this time, the two options are: disable spanning tree, or run a single spanning tree instance over all shortcut bridge ports without regard to the domains. Another restriction is that MAC addresses must be unique across all shortcut bridge domains (since the domains share a common transparent bridge filtering database). Similarly, unique IP addresses are required if BBCM is enabled (since there is currently only one BBCM database per MSS Server).

10.3.1 Relationship between SuperVLANs and LNNI

IBM remains a strong supporter of the ATM Forum's LNNI Specification effort. As part of a comprehensive package designed to enable deployment of large-scale LANE networks, the SuperVLAN mechanisms will continue to add value in LNNI environments.

10.4 MSS and IBM 8274/3 Broadcast Control

The MSS Server and IBM 8274/3 both provide VLAN support for bandwidth management that is complementary and can be used together when designing a switched ATM/legacy LAN environment.

The MSS Server's broadcast management functions operate on ELANs, while, IBM 8274/3's broadcast management functions operate on legacy LANs. MSS operates in addition to the IBM 8274/3 broadcast management when using ATM-bridged (that is, using LAN emulation) connections between your IBM 8274/3s.

Note: When using trunking between IBM 8274/3's, data is unseen by MSS and not subject to its broadcast management functions.

10.4.1 Overview

The broadcast management functions that can be used in a switched network environment using the MSS Server and IBM 8274/3 are:

1. MSS Server

- Broadcast Management (BCM) (see 2 in Figure 38 on page 79)

BCM transforms IP, IPX, and NetBIOS broadcast frames into unicast frames. BCM is a BUS function that reduces traffic on ELANs for which the BUS function is performed. It is available on Ethernet and token-ring ELANs.

- Source Route Management (SRM) (see 2 in Figure 38 on page 79)

SRM transforms IP, IPX and NetBIOS explorer frames into specifically routed frames.

SRM is a BUS function that reduces traffic on ELANs for which the BUS function is performed. It is available for token-ring ELANs.

- Bridge Broadcast Management (BBCM) (see 1 in Figure 38 on page 79)

BBCM transforms IP and NetBIOS broadcast frames into unicast frames. BBCM is a bridging function that reduces traffic on the bridge ports that are part of a SuperVLAN. It is available for Ethernet and token-ring bridge ports (ELANs and RFC 1483 connections).

- Dynamic Protocol Filtering (DPF) (see 1 in Figure 38 on page 79)

DPF controls the bridging of IP, IPX and NetBIOS broadcast frames within a protocol VLAN. DPF is a bridging function that limits broadcasts to the bridge ports that are part of a PVLAN. It is available for Ethernet and token-ring bridge ports (ELANs and RFC 1483 connections).

Note: An MSS Server PVLAN is a subset of a SuperVLAN.

2. IBM 8274/3 (see 3 in Figure 38 on page 79)

The IBM 8274/3 VLAN support confines broadcasts within designated VLANs. IBM 8274/3 VLANs are based on pre-defined policies (rules):

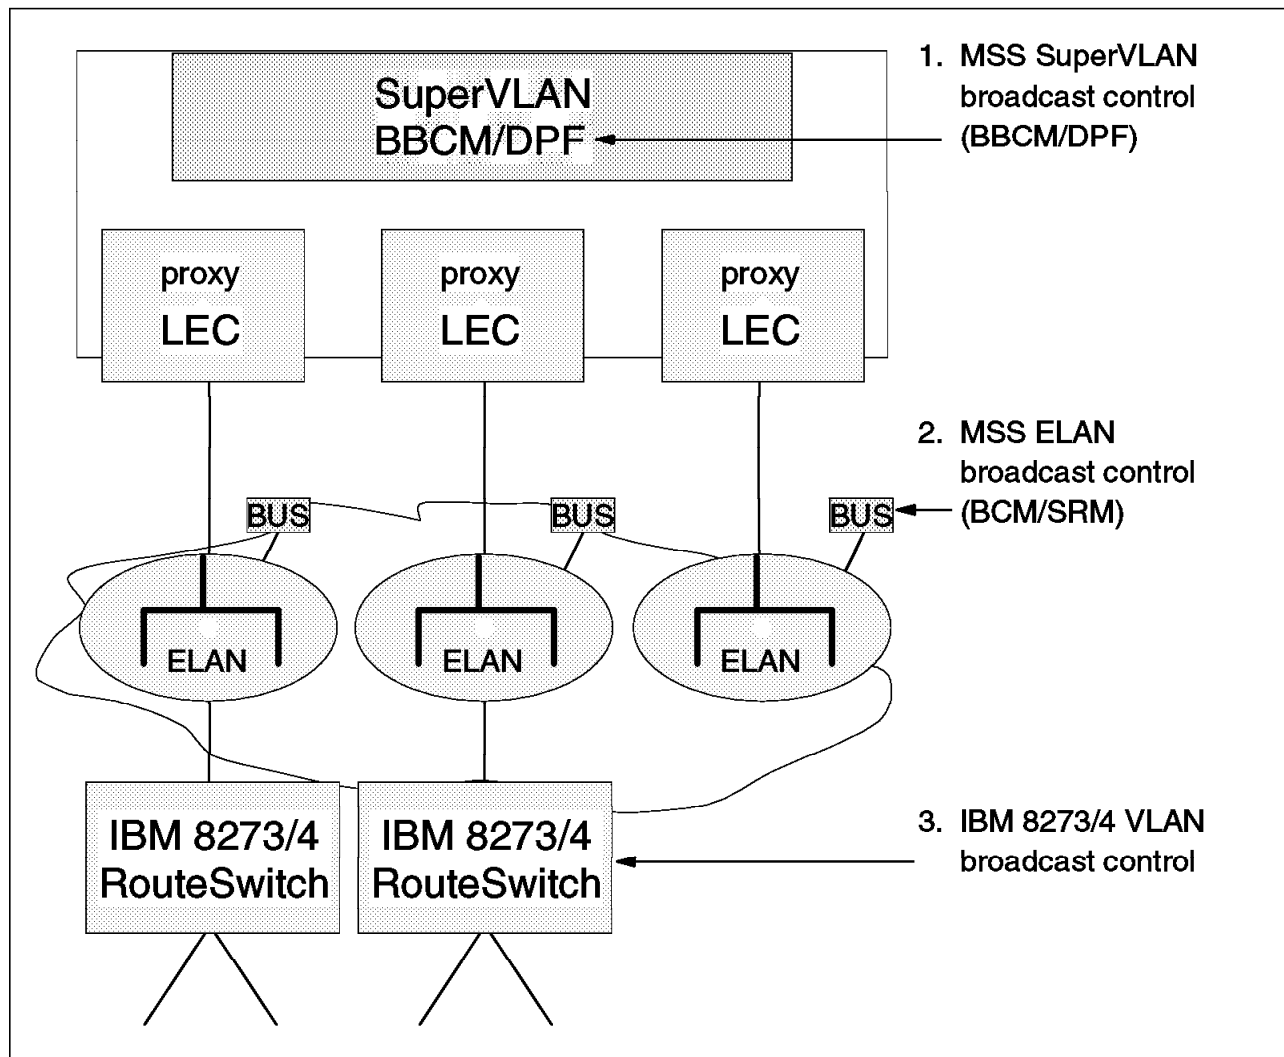


Figure 38. Different Levels of Broadcast Control

- Port rules
- MAC address rules
- Protocol rules (IP, IPX, DECnet, AppleTalk, Ethertype, SAP, SNAP)
- Network address rules (for example, IP subnet)
- User-defined rules

Multicast messages are confined to the legacy LANs that make up a VLAN.

10.4.2 Things You Should Know

When using the broadcast management functions in a switched network environment, be aware of the following:

1. MSS Server

- BCM versus BBCM (intra-ELAN versus inter-ELAN broadcast control)

BCM and BBCM perform functions that are similar. BBCM provides broadcast control *between* ELANs that make up a SuperELAN (that is, inter-ELAN broadcast control). BCM provides broadcast control *within* an ELAN (that is, intra-ELAN broadcast control). Using BCM on the BUSes of all the ELANs within the SuperELAN provides similar broadcast management as using BBCM.

The advantage of using BBCM is that your BUSES are offloaded. Maximum BUS performance can be reached by running in VCC slicing mode. The advantage of using BCM is that traffic between legacy LANs connected to the same ELAN is subject to the MSS broadcast control functions while BBCM's inter-ELAN broadcast control operates only on data sent between legacy LANs connected to bridged ELANs.

- Number of SuperELANs

MSS allows you to define multiple SuperELANs; however, MSS runs a single spanning tree protocol. Therefore, avoid multiple SuperELANs unless SuperELANs are isolated from each other (that is, no overlap of broadcast domains).

Note: A SuperELAN in this context is a set of bridge ports that are shortcut bridged together.

2. IBM 8274/3

- Parallel bridged connections

You can have multiple, parallel connections between IBM 8274/3s. However, as a single spanning tree is maintained per group, parallel connections within a group are blocked.

Note: Per group you can have multiple LE clients. LE clients assigned to different VLANs within the same group may be blocked due to having a single spanning tree.

- Quiet devices

Workstations that do not indicate (either by broadcast or unicast traffic) their presence need to be explicitly defined within the VLAN. A port rule should be used for this. Examples are IP and/or SNA servers and protocol analyzers.

Note: If you are using bridged (possibly via ELAN) connections, port rules may be required on any of your IBM 8274/3s.

- Moving hub-attached stations when using ATM LAN emulation

In the situation where you have two RouteSwitches (say X and Y) bridged together using ATM LAN emulation, connectivity problems may occur after moving a station from one RouteSwitch to another.

The problem occurs when:

- a. Two stations, A and B, are Ethernet-connected to the same RouteSwitch (X) and happily communicating. X has learned MAC/VLAN information about A and B.

Note: Station B is connected via an Ethernet hub.

- b. Station B is moved to Y. Note, the Ethernet hub is not disconnected and station B remains known on X.
- c. Station A and B will not be able to communicate until the cached entry for B ages out on RouteSwitch X.

Notes:

- a. The problem does not occur if station B is directly attached.
- b. This problem is resolved in RouteSwitch software program Release 3.0.

- Overlapping VLANs

The IBM 8274/3s assign VLAN membership to a workstation based on the contents of multicast or unicast messages. Unicast messages will only be interrogated when the source MAC address is unknown, multicast messages will always be interrogated.

Special care is required when using protocol VLANs if any of your workstations is a member part of multiple protocol VLANs (for example, IP, IPX, SNA etc.). Two cases should be distinguished:

- a. When using non-broadcast protocols (for example, SNA)

To establish a connection each SNA workstation sends a unicast message. If VLAN membership has already been assigned to the workstation (for example, due to a previous IP broadcast), this unicast message may be confined to a different protocol VLAN and, possibly, never arrive at its destination.

The problem can be avoided by starting the non-broadcast protocol first on any of your workstations, or, by making sure that your protocol VLANs completely overlap.

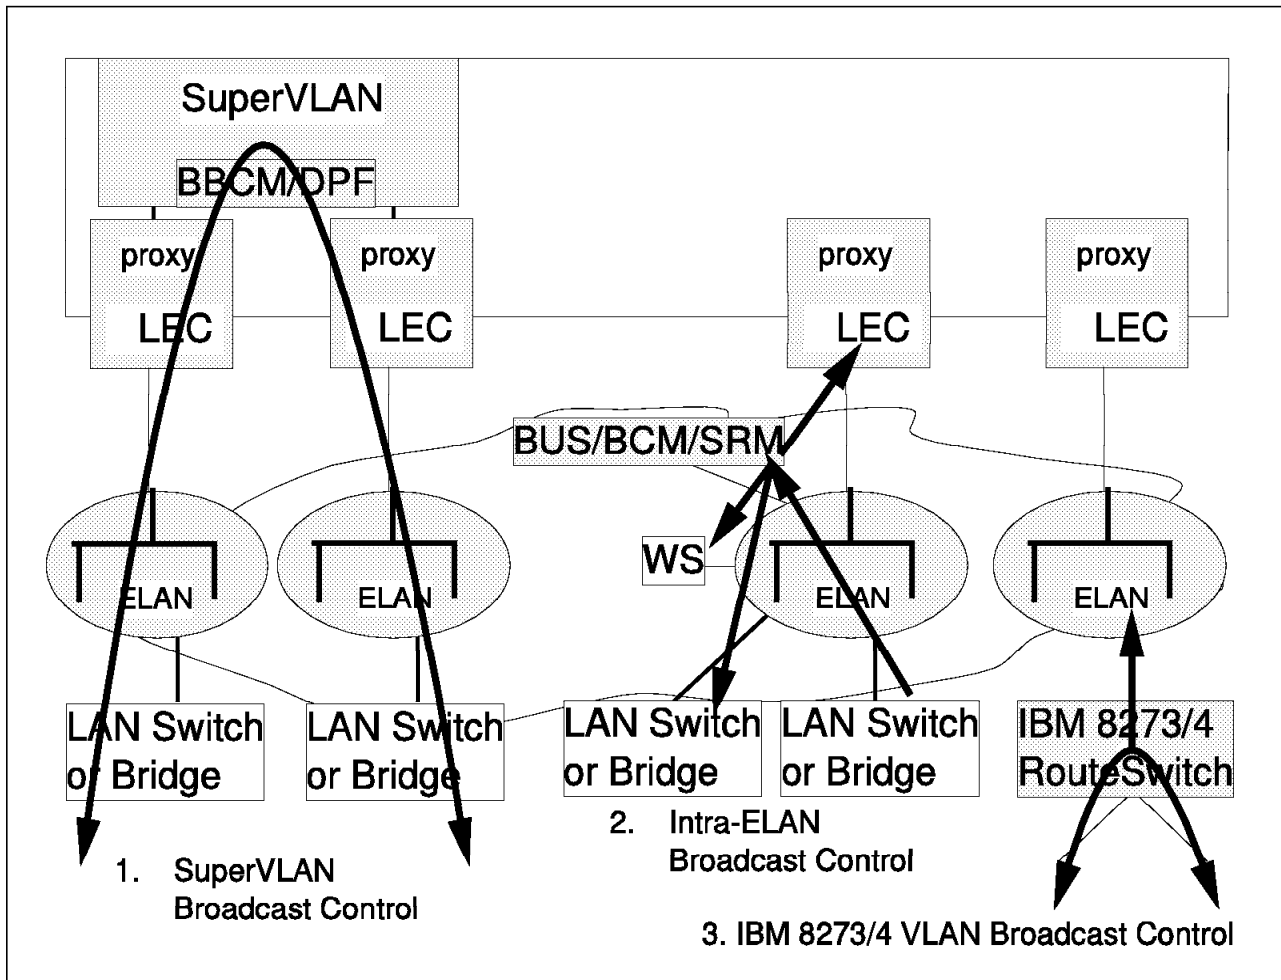


Figure 39. Intra-LAN, Intra-ELAN and Inter-ELAN Broadcast Management

b. When using broadcast protocols (for example, IP, IPX, etc.)

Due to timeouts and/or CAM overflows VLAN membership can be removed from an IBM 8274/3. VLAN membership will be re-learned on the first message sent by the workstation. Assuming a workstation has active connections for multiple protocols, VLAN membership will be regained for a single protocol and temporary connectivity problems may result. The protocol for which VLAN membership will be restored is unpredictable and depends on the first unicast data sent by the workstation.

Connectivity will be re-established after a multicast message has been sent. To avoid this problem make sure the protocol VLANs completely overlap.

- Bridging between groups on the same RouteSwitch

In the situation where you have a single 8274/3 on which you have defined two groups (such as X and Y) that are bridged together, connectivity problems may occur when temporarily disconnecting the bridged connection.

The problem occurs when:

- a. Two stations A and B are connected to a different group and happily communicating. As a result stations A and B are known in both groups.
- b. The bridged connection is temporarily disconnected. Because of the interrupt, stations A and B are only known within the group they physically connect to. Their presence on the bridged connection is no longer retained.
- c. Station A and B will not be able to communicate as their MAC/VLAN information is not re-learned on the bridged connection.

Note: A solution for an Ethernet bridged connection would be activating the MAC flooding feature that has become available in RouteSwitch program Release 2.1.4.

For details on the IBM 8274/3, see *IBM NWay's RouteSwitch Implementation Guide*, SG24-4581.

Figure 39 on page 82 displays the three types of broadcast control. BBCM/DPF control broadcast traffic across ELANs (inter-ELANs). BCM/SRM control broadcast traffic within ELANs (intra-ELANs). The IBM 8274/3 VLAN support controls broadcast traffic between legacy LANs.

10.4.3 An Example

Figure 40 on page 84 shows an example of an environment where both MSS and IBM 8274 broadcast management facilities are used. Depicted is a situation with three different buildings (1, 2 and 3), each having multiple LAN segments.

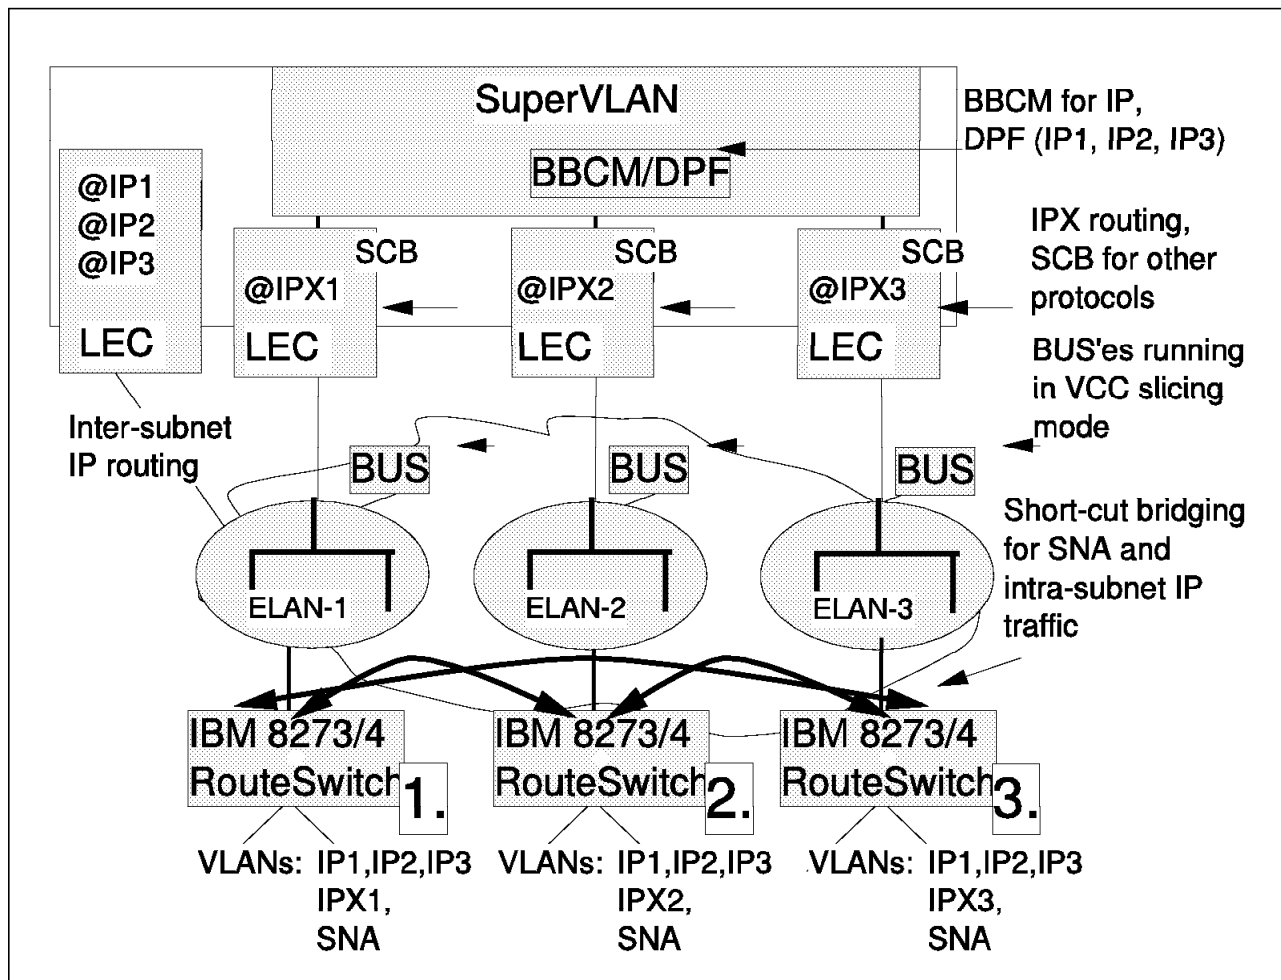


Figure 40. Routing, Switching and Broadcast Control

An IBM 8274 RouteSwitch is used for broadcast management between the legacy LANs within each building. To enable traffic between different buildings each IBM 8274 is connected to a separate ELAN. The MSS Server provides bridging, routing, and broadcast control functions between the different ELANs. The MSS Server provides IPX routing, SNA bridging, (intra-subnet) IP bridging for stations that are part of the same subnet, and (inter-subnet) IP routing for stations that are part of different IP subnets.

10.4.3.1 RouteSwitch Definitions

In each of the three buildings an IBM 8274 is used to establish IP, IPX, and SNA VLANs. To enable IP mobility of user stations between any of the buildings the same three IP subnet VLANs (IP1, IP2, and IP3) have been defined on all IBM 8274s. To enable IPX stations to connect to their local IPX server, different IPX VLANs (IPX1, IPX2, IPX3) have been defined on each of the RouteSwitches. The same SNA VLAN definitions are defined on all IBM 8274s.

Notes:

1. This setup provides mobility for IPX clients, but not for IPX servers.
2. The IBM 8274 LE client ports need to be explicitly defined in the IP and the SNA VLANs.

10.4.3.2 MSS Server Definitions

To enable MSS to route IPX between the different subnets (IPX1, IPX2, and IPX3) an IPX address has been defined on the SCB ports. The IPX addresses on the SCB ports prevent IPX switching between different buildings. IPX switching is unwanted in this situation as stations might connect to the wrong IPX server.

To enable MSS to route IP between the between different subnets (IP1, IP2, and IP3) an LE client has been defined on ELAN-1 on which three IP addresses have been defined, one within each subnet.

To enable MSS to bridge SNA and IP, shortcut bridging (SCB) ports have been defined on each of the ELANs. Using SCB the bridging will be very efficient as shortcut VCCs will be used between the Routeswitches.

BBCM has been activated for IP broadcast management on the SCB ports. DPF has been enabled for subnets IP1, IP2 and IP3. DPF is not required for IPX as the IPX data will be routed.

As no broadcast management functions are used on the BUSES, each of them can run in VCC slicing mode.

Chapter 11. Next Hop Resolution Protocol Support for IP

Next Hop Resolution Protocol (NHRP) support for IP is a significant feature of MSS Server Release 1.1. 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 may include multiple hops across the same subnetwork.

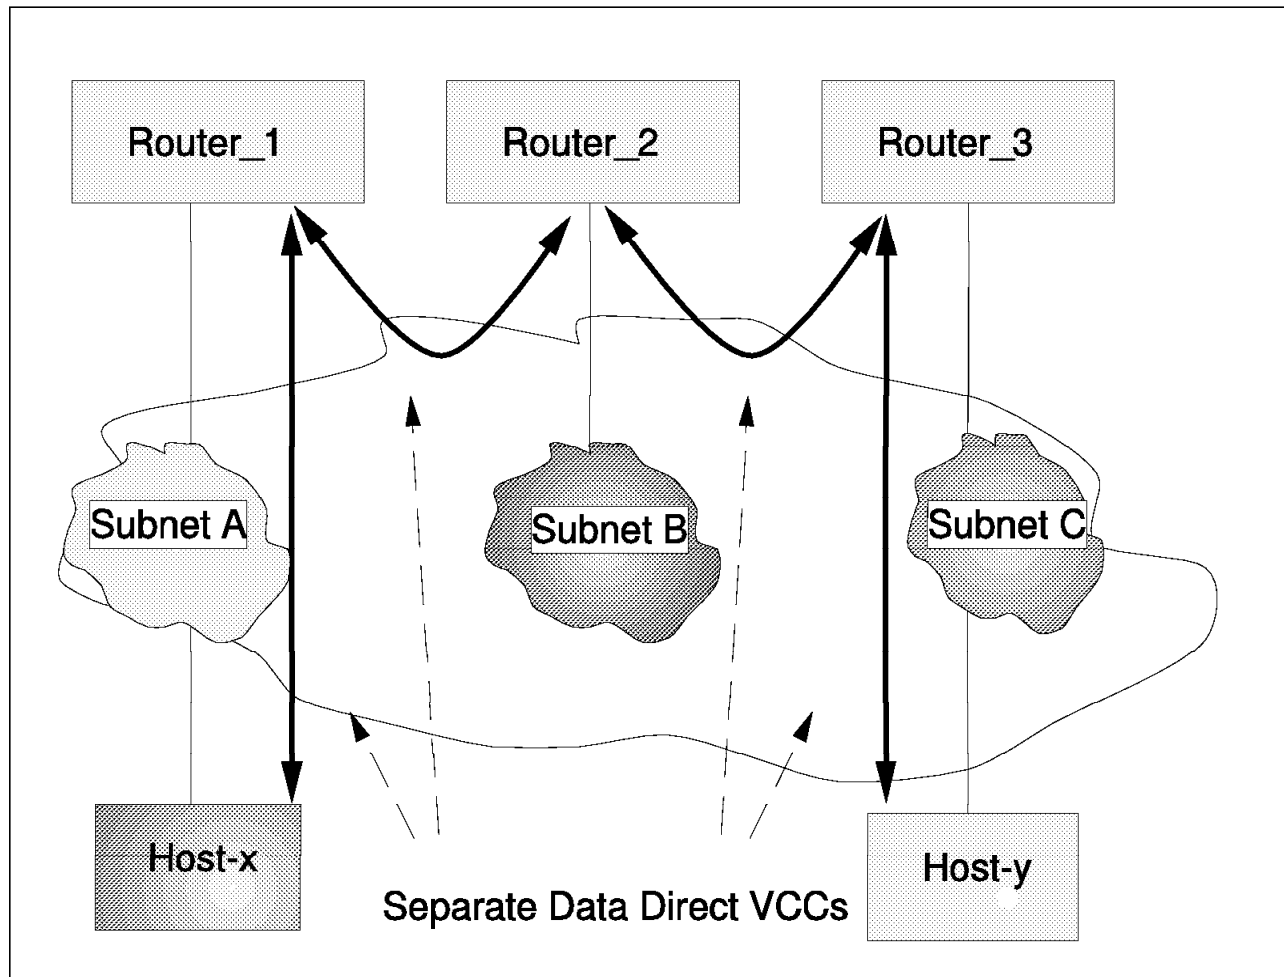


Figure 41. Conventional Routing over ATM

As an example, consider the network shown in Figure 41, 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 shortcut routes across ATM subnetworks. The shortcut 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. The MSS Server Release 1.1 implementation includes both NHS and NHC functions. The NHC allows the MSS Server to request shortcuts for stations that do not have their own NHC.

In the best case, NHRP enables zero-hop routing, where all the router hops are removed from the steady-state data path as illustrated in Figure 42. Zero-hop routing is possible when NHC software is installed on attached ATM hosts. To initiate establishment of a short-cut to IP Host y, the NHC in IP Host x issues a NHRP Resolution Request for Host y's IP address. As indicated above, NHRP Resolution Requests follow the routed path, so the request flows from IP Host x to NHS 1, from NHS 1 to NHS 2, and then on to NHS 3. NHS 3 responds to the request with the ATM address of IP Host y. When the NHRP Resolution Reply is returned to IP Host x, the ATM address is used to set up a shortcut VCC. Subsequent traffic for IP Host y is transmitted directly over this VCC, bypassing the three intermediate routers, which both lowers latency and increases system throughput.

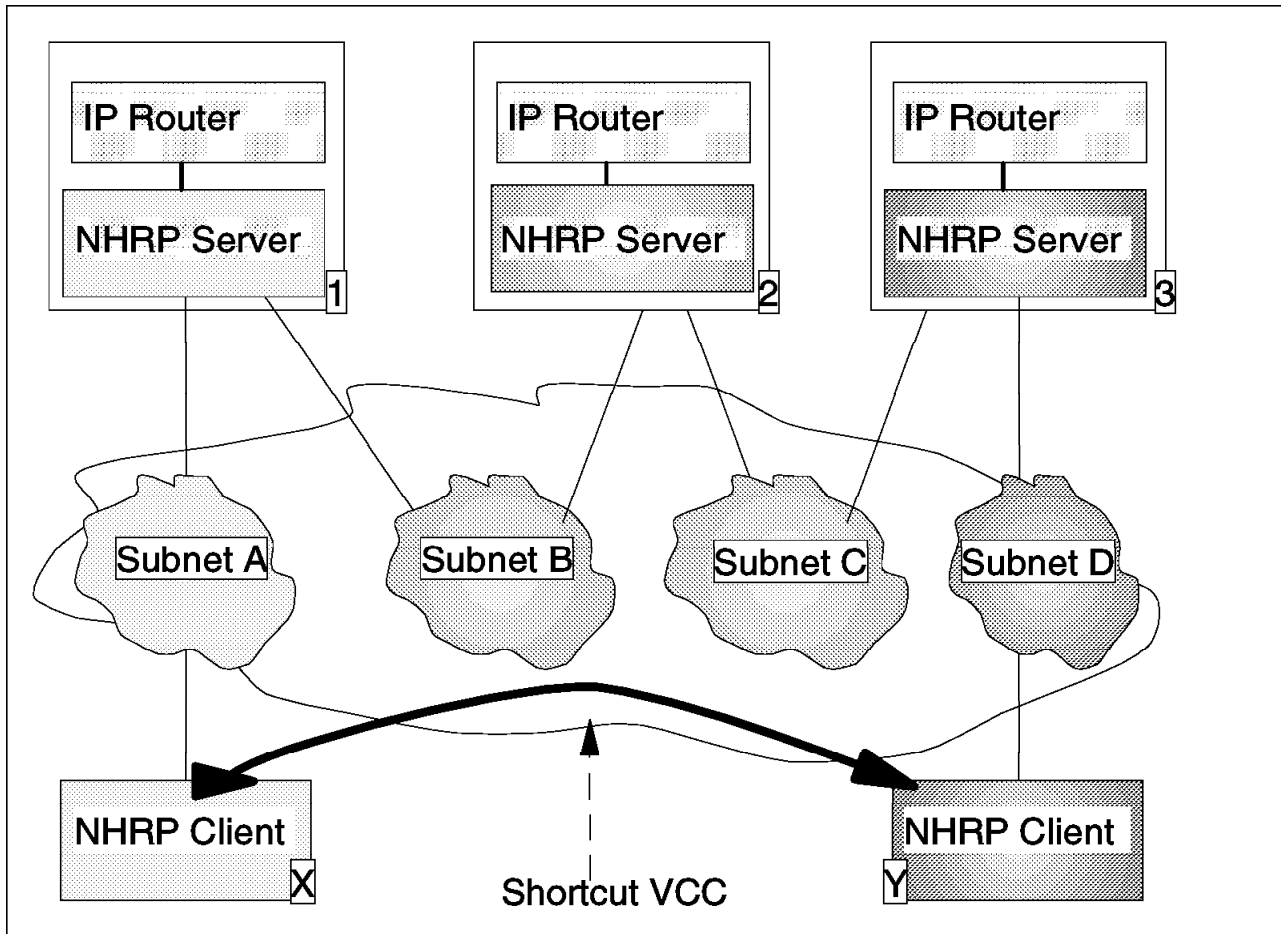


Figure 42. Zero-Hop Routing

In the preceding example, NHS 3 responds to the NHRP Resolution Request because NHS 3 serves IP Host y. A NHS serves destinations that reside on subnets local to the NHS (that is, subnets for which routing interfaces are defined). An NHS also serves destinations for which the routed path exits the ATM network at the NHS (for example, destinations accessible through a router

that is connected to the NHS via a non-ATM link). NHRP Resolution Requests are not forwarded to NHCs; when the NHS that serves the destination protocol address receives the request, the serving NHS responds on behalf of the destination. If the destination is on the ATM subnetwork, the serving NHS responds with the ATM address associated with the destination. An NHS may learn protocol address-to-ATM address mappings in multiple ways. NHCs register their ATM addresses with NHSs. Additionally, NHSs can use ARP procedures to learn protocol address-to-ATM address mappings, which enable shortcuts to ATM hosts without NHCs (for example, ATM hosts supporting Classical IP).

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 43 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.

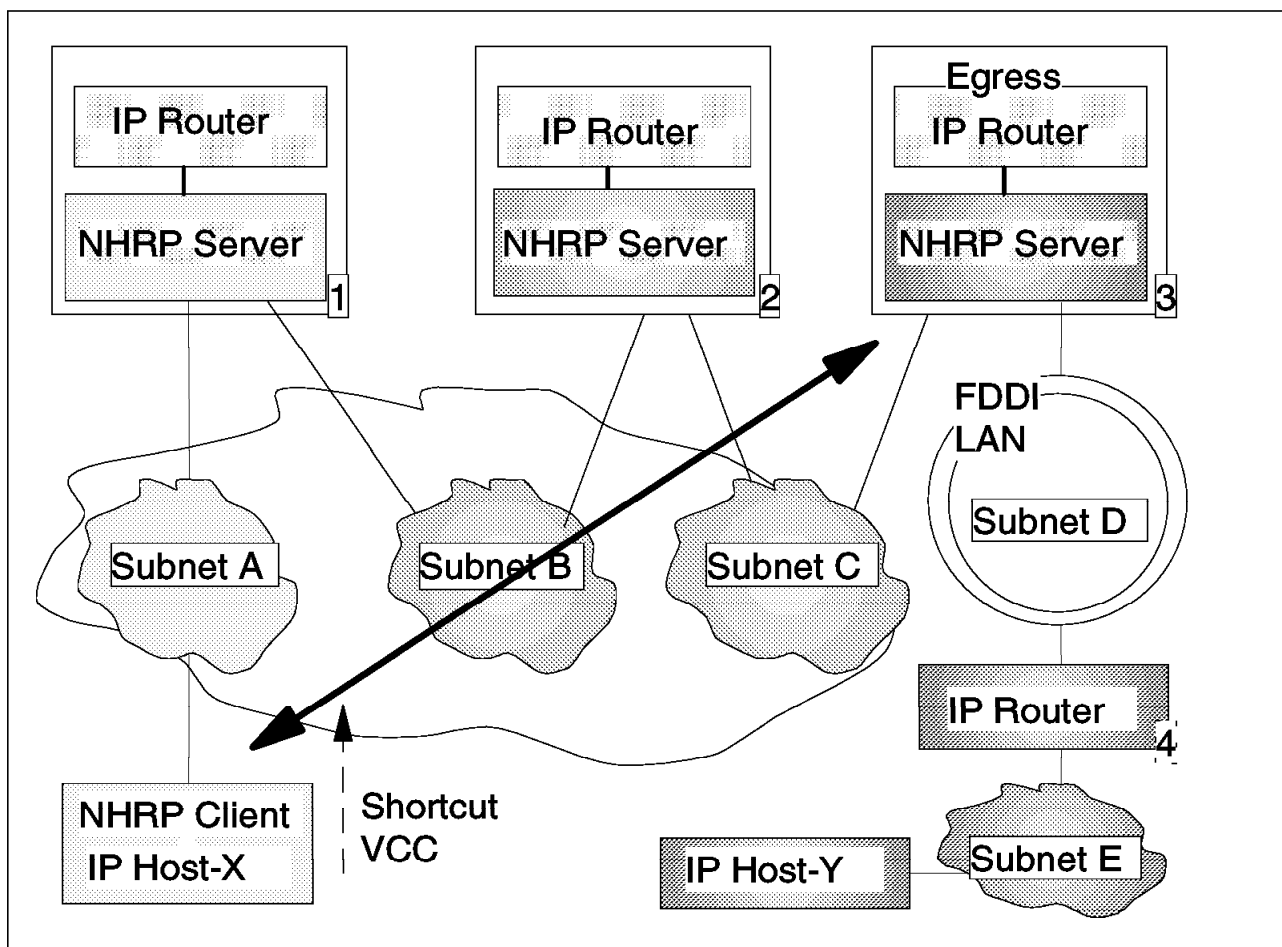


Figure 43. Shortcut to an Egress Router

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 shortcut 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 shortcut service. NHCs include

the holding time when registering their (protocol and ATM) addresses with an NHS. The MSS Server NHS also has a holding time configuration parameter that is used in responses issued on the behalf of non-registered destinations. (The default is 20 minutes.) In addition to the holding time safeguard, serving NHSs are responsible for sending NHRP Purge Requests if mappings they have provided become invalid prior to the expiration of the holding time. Upon receiving an NHRP Purge Request for a given mapping, a NHC must discontinue use of the associated shortcut. Mappings may become invalid due to deregistrations by NHCs, or as a result of routing protocol topology changes (for example, when the NHS is no longer serving as the egress point from the ATM subnetwork for a particular destination).

Since NHRP Resolution Requests follow the routed path, NHSs are coupled with routers, and routers without an NHS can discontinue NHRP reachability. In general, shortcuts are not possible when there is an intermediate router that does not support NHRP. For example, if there was not an NHS associated with Router 3 in Figure 43 on page 89, then the NHRP Resolution Request would not have been answered, and therefore, the shortcut VCC would not have been established. However, the MSS Server implementation of NHRP includes a configuration facility that enables shortcut VCCs to be established to routers that are not NHRP-capable. If the protocol address of a next-hop router is in the Exclude List, the MSS Server will not forward NHRP Resolution Requests to that router; instead, the MSS Server will respond to the requests with the ATM address of the next-hop router. In the context of the previous example, where there is no NHS at Router 3, including Router 3's IP address in the Exclude List at NHS 2 would allow a shortcut from IP Host x to Router 3.

As mentioned earlier, the MSS Server's NHRP implementation also includes support for hosts without NHCs. This support is called one-hop routing, since the ingress router (that is, MSS Server) establishes shortcut routes on behalf of hosts as depicted in Figure 44 on page 91. Note that the resulting routes between IP Host x and IP Host y are asymmetric. IP Host x sends intersubnet traffic to MSS Server 1, which then establishes a shortcut route to the destination, and similarly, IP Host y sends intersubnet traffic to MSS Server 2, which establishes shortcut routes on its behalf. Also note that neither host is required to support NHRP for the shortcut routes to be established. These shortcut routes are established by the MSS Server's NHC when the rate of traffic to a particular destination exceeds a configurable threshold. The default value for the Data Rate Threshold is 10 packets per second, which is useful in eliminating the overhead associated with creating shortcut VCCs for low volume or one-time traffic (such as SNMP traps).

11.1 ELAN Support

Support for ELANs is another important feature of the MSS NHRP implementation. This support has two parts. First, NHRP packets that are to follow the routed path are forwarded over ELANs as well as Classical IP subnets. Thus, if IP Subnet B in Figure 44 on page 91 were an ELAN, MSS Server 1 would forward NHRP Resolution Requests for IP Host y to MSS Server 2 over the ELAN VCCs (for example, over a data direct VCC between the two routers). 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. For interoperability reasons, such implementations should be connected to MSS Servers via LISs (which do employ RFC 1483 encapsulation), instead of ELANs.

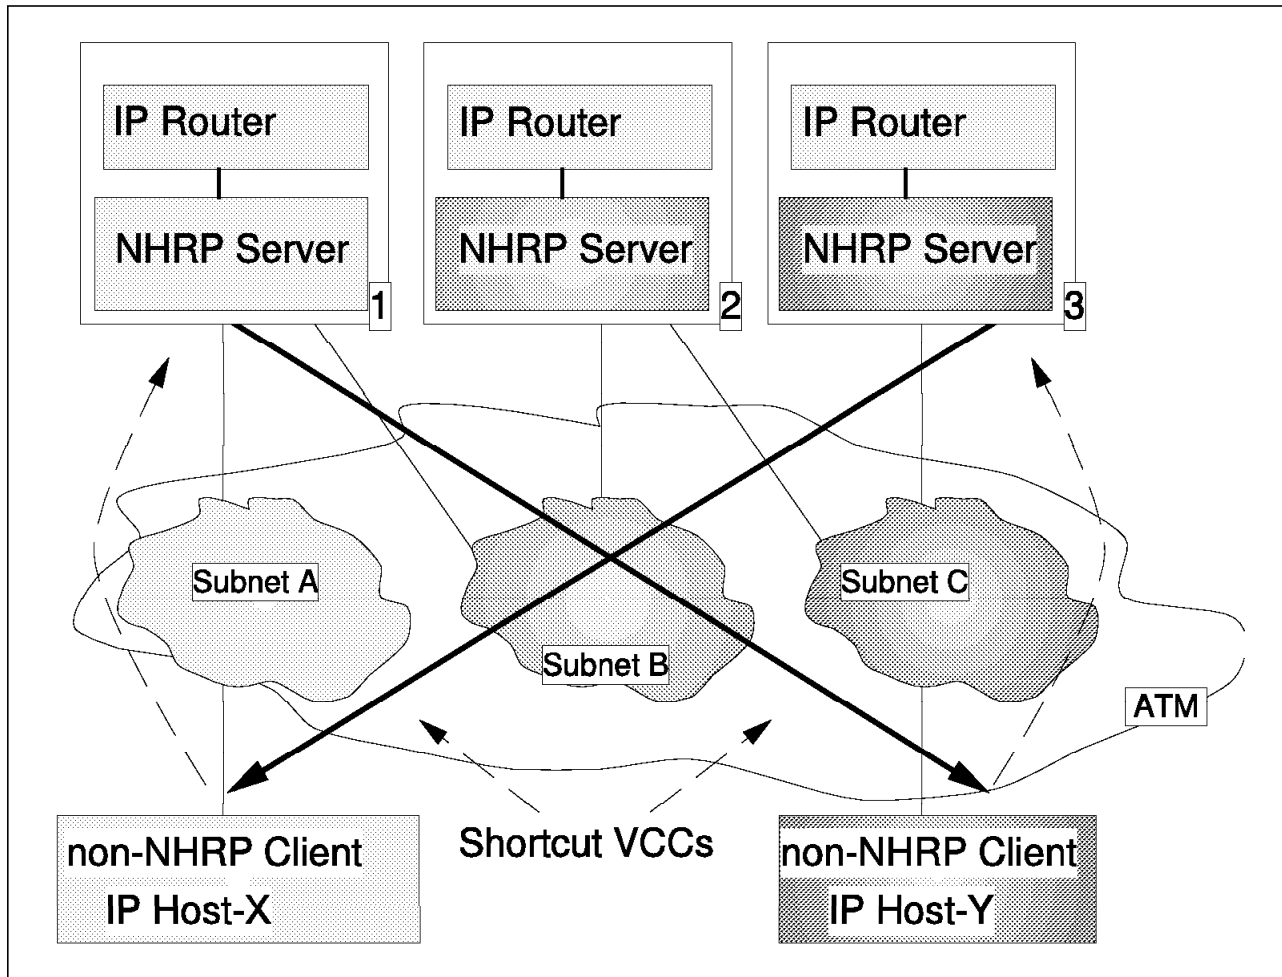


Figure 44. One-Hop Routing

The second part of the MSS Server's ELAN support is the capability to establish shortcuts to LAN emulation clients. This is an important value-add, since it extends the benefits of NHRP to the large installed base of LANE equipment, including LAN switches with ATM uplinks. Although the NHRP specification does not define support for LANE encapsulation, it does define a generic mechanism for vendor-private extensions. This mechanism is interoperable because NHRP implementations that don't understand the extensions simply ignore them. The MSS Server uses vendor-private extensions to support LANE shortcuts.

One-hop routing with LANE shortcuts is illustrated in Figure 45 on page 92. In this example, both source and destination IP stations reside on legacy LANs behind ATM-attached LAN switches. MSS Server 1 initiates a shortcut to IP Host y on behalf of IP Host x, and includes the LANE shortcut 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 shortcut 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.

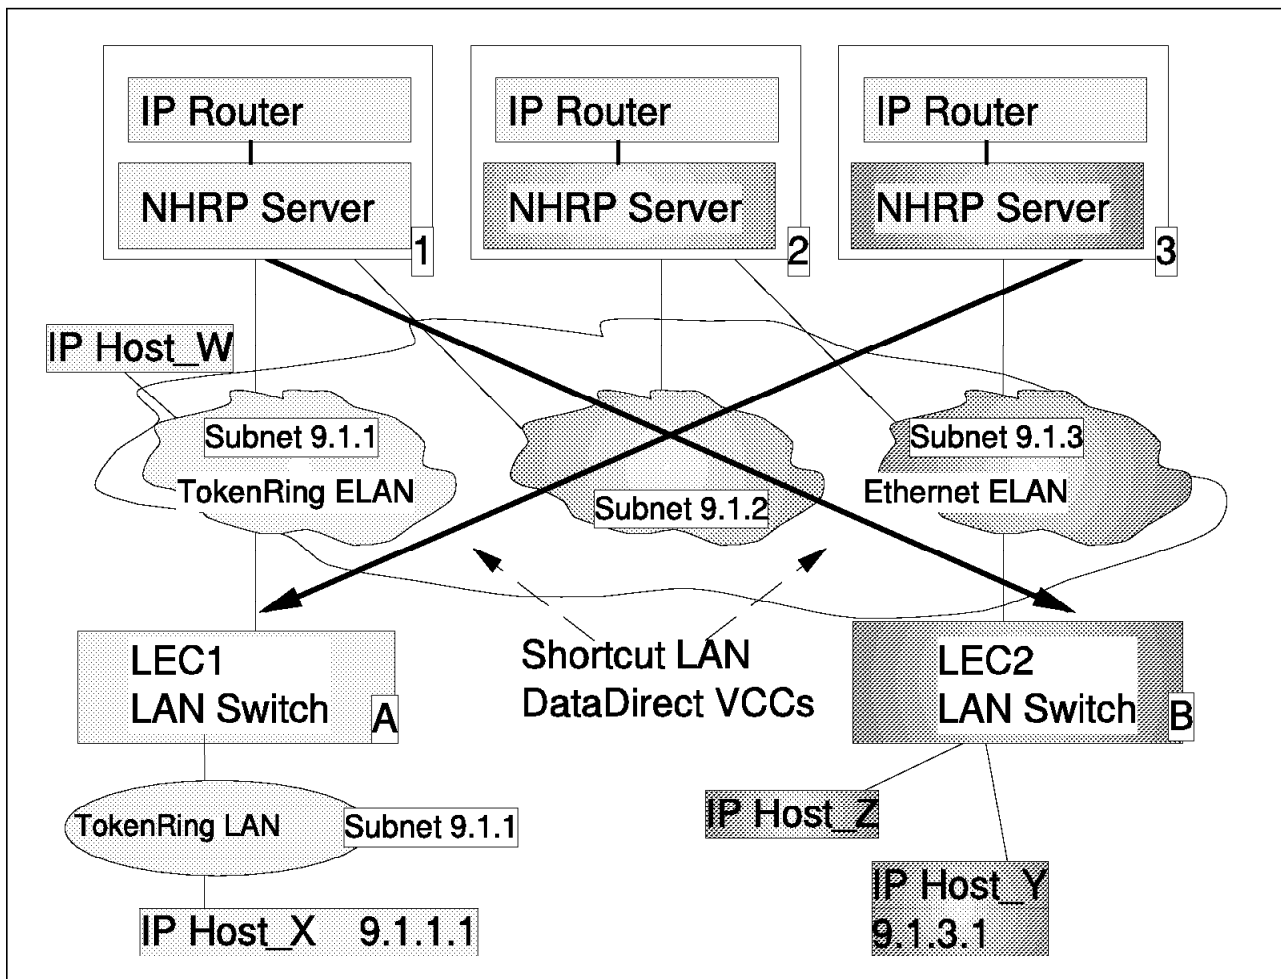


Figure 45. One-Hop Routing with LANE Shortcuts

MSS Server 2 establishes a shortcut for traffic to IP Host x in a similar manner. However, in this case, the LANE shortcut 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 shortcut routing mechanisms are equivalent for ATM-attached hosts with LANE interfaces, such as IP Host w.

LANE shortcut VCCs are established by a new net type included in Release 1.1 of the MSS Server. A LANE Shortcut Interface (LSI) is automatically created for each physical ATM interface when NHRP is enabled. The LSI is capable of establishing shortcut data direct VCCs to LAN emulation clients on token-ring and Ethernet/IEEE 802.3 ELANs. These LSI connections are used in a unidirectional fashion, with traffic transmitted by the MSS Server but not

received. Since a routing function is being performed, packets transmitted over a LANE (or RFC 1483) shortcut VCC will be fragmented as necessary to resolve MTU size mismatches. However, it should be noted that a single MTU size is associated with each LSI. This interface MTU is the minimum of the MTUs for all destination subnets accessed via the LSI (The MTU of the destination subnet is returned in the NHRP Resolution Reply.)

By default, the LSI uses the MAC address burned into the ATM interface hardware as the source MAC address of frames transmitted over LANE shortcut VCCs. Any other MSS LECs that are members of the target ELAN must have different MAC addresses (to ensure devices on the ELAN that learn MAC address VCC associations from received frames do not become confused).

Important

As a precaution, it is recommended that each MSS Server LEC be configured with a unique locally administered MAC address when NHRP LANE shortcuts are enabled.

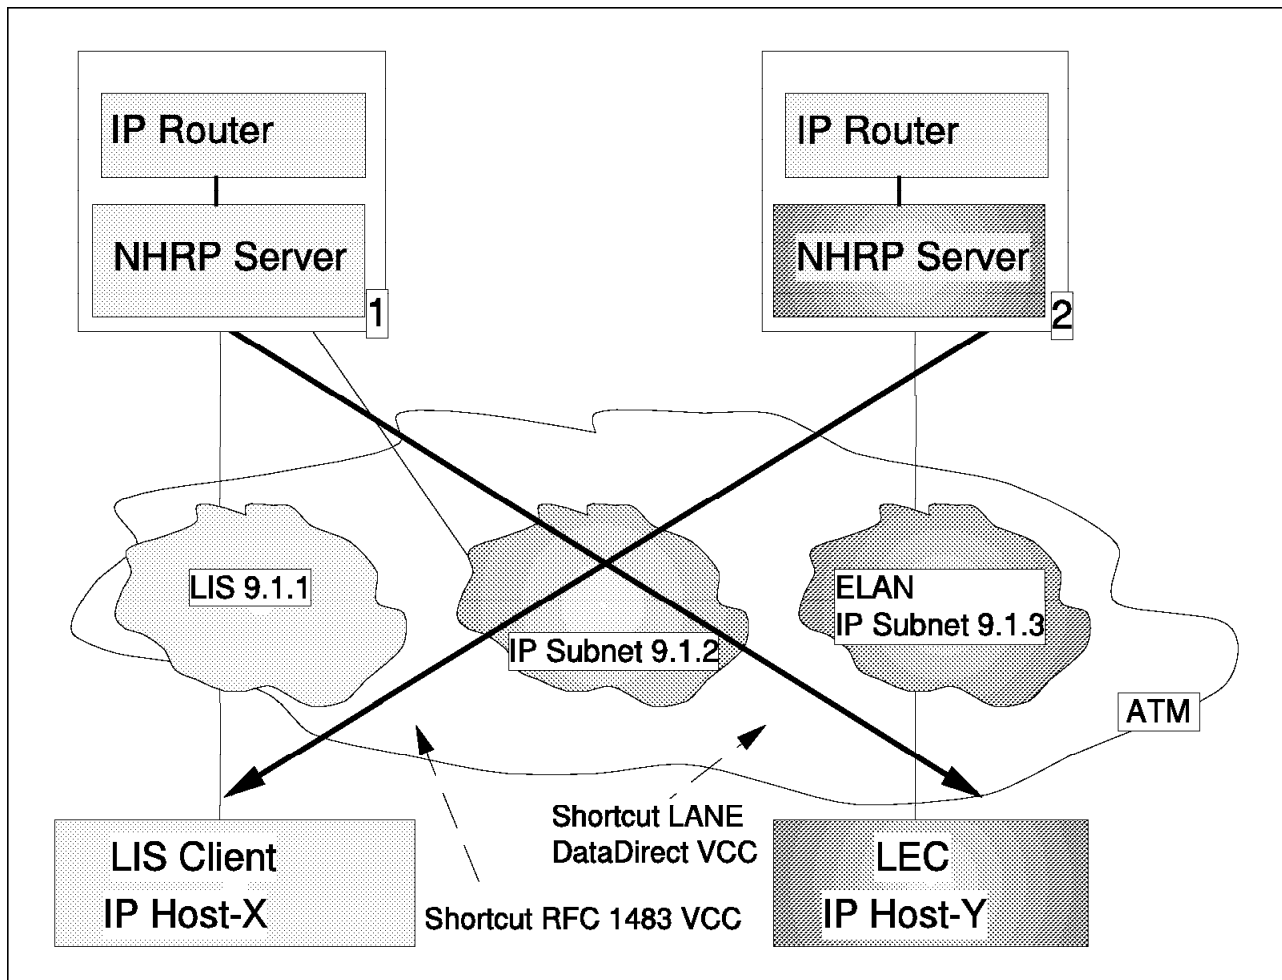


Figure 46. One-Hop Routing with Classical IP and LANE Hosts

NHRP utilizes the Classical IP client function to establish and accept short-cut RFC 1483 VCCs, which are generally bidirectional in nature. If multiple CIP clients are defined, NHRP selects the CIP client whose MTU size is most appropriate. If no CIP clients are defined, NHRP is still able to initiate shortcut

RFC 1483 VCCs, but establishment of shortcut RFC 1483 VCCs to the MSS Server is not possible.

One-hop routing is very flexible. Figure 45 on page 92 has already provided one example by depicting shortcut communications between hosts on token-ring and Ethernet LANs. Another example is provided in Figure 46 on page 93, which shows shortcut communications between Classical IP and LANE hosts. More generally, shortcuts may be established on behalf of hosts accessible via any of the supported interfaces: Classical IP LISs, token-ring ELANs, Ethernet/IEEE 802.3 ELANs, or an FDDI LAN. The type of shortcut established is based solely on the destination. RFC 1483 VCCs are established to Classical IP clients and destinations registered by NHCs, while LANE data direct VCCs are established to ELAN-resident devices.




Figure 47. Zero-Hop and One-Hop Routing Combined

Zero-hop and one-hop routing may also be combined as illustrated in Figure 47, where IP Host x has an NHC and IP Host y does not. In this case, MSS Server 2 establishes a shortcut VCC on behalf of IP Host y, while IP Host x is able to establish a shortcut VCC directly to IP Host y. Thus, MSS Server Release 1.1 can provide immediate benefits by supporting one-hop routing for all hosts, while simultaneously enabling the additional benefits available by zero-hop routing to be realized in an incremental and manageable manner as hosts are upgraded with NHCs.

The MSS Server implementation also includes a value-add extension designed to suppress the NHRP Domino Effect, which occurs when multiple NHCs along the routed path simultaneously attempt to set up a shortcut for the same traffic flow. (That is, before a shortcut has been established, each NHS/NHC pair along the routed path will see the same traffic, and this traffic may trigger all the NHCs to request a shortcut.) The NHRP domino effect introduces unnecessary signalling overhead since only the shortcut established by the NHC closest to the source will ultimately be utilized. The MSS Server limits the impact of the NHRP domino effect by cancelling any pending shortcut requests when an NHRP Resolution Request is received for the same destination. Reception of an NHRP Resolution Request cancels a pending shortcut request (by resetting Data Rate Threshold measurements) because reception of the request indicates that an upstream NHC is providing a shortcut for the traffic flow.

Security is one of the important functions provided by routers. In order to preserve router-based security capabilities in NHRP environments, protocol-layer access controls are applied when generating, forwarding, or responding to NHRP Resolution Requests. This allows shortcuts to be prohibited in situations where routed path communications is not allowed. Two modes are provided. In source and destination mode, shortcut requests received from hosts are only denied when there is a filter specifying the source/destination pair. Destination-only mode is more restrictive, as shortcut requests are denied if any filters apply to the destination protocol address. Regardless of the configuration setting, Destination-only mode is always applied to requests received from routers because routers may use shortcut VCCs for traffic originating from multiple sources.

The MSS Server Release 1.1 implementation of NHRP is based on Version 11 of the Internet Draft Specification, which has been submitted for approval and is expected to gain RFC status soon. Although it is unlikely that the specification will be modified again before achieving RFC status, if that should happen, the MSS Server implementation will be updated with any changes necessary to ensure standard compliance. In summary, NHRP minimizes the number of router hops that must be traversed in the steady-state data path, which improves network performance by reducing latency, while also reducing router bandwidth requirements. For example, one-hop routing can cut latency and routing bandwidth requirements by 50% in networks where two router hops are typical (which is common when routers are interconnected via a backbone subnet). Furthermore, the MSS Server implementation includes a number of value-add enhancements such as support for hosts and routers that are not NHRP-capable, extension of NHRP benefits to LAN emulation environments, suppression of the NHRP domino effect, and preservation of router-based security.

11.2 MSS 2.0 Zero-Hop Route Server (RouteSwitching)

RouteSwitching is a new function in MSS 2.0 that allows two LAN-attached devices on different subnets to communicate directly, without the need for traditional routing hops between them. This is a logical continuation of the MSS 1.1 one-hop routing capability for LAN attached devices and the zero-hop routing capability available to ATM-attached devices. Now, Ethernet and token-ring devices can have the performance and response characteristics available on layer-2 switched networks, with the IP routing robustness of layer-3 routed networks, without the additional complexity or expense of layer-3 switching in every workgroup LAN switch.

The terms *RouteSwitch* and *zero-hop route* have the same meaning in respect to MSS and LAN Emulation. These terms have different meanings for other products. IBM 8274 LAN Switch software is known as RouteSwitch software; zero-hop routes can also be achieved for ATMARP clients (with NHC function) in Classical IP over ATM environments, as discussed earlier in this chapter.

The zero-hop IP routing function is accomplished via a client/server protocol between an MSS RouteSwitching client (RSC) and a RouteSwitching server (RSS). The protocol used is an enhanced implementation of NHRP. The RSC is software in devices such as PCs and servers. The RSS is built into the microcode of MSS. LAN stations on Ethernet or token-ring with the RSC software are supported if they are attached to LAN switches that connect to an ATM network via ATM Forum-compliant LAN Emulation, and the LES/BUS and default IP gateway functions are supplied by MSS at Version 2.0.

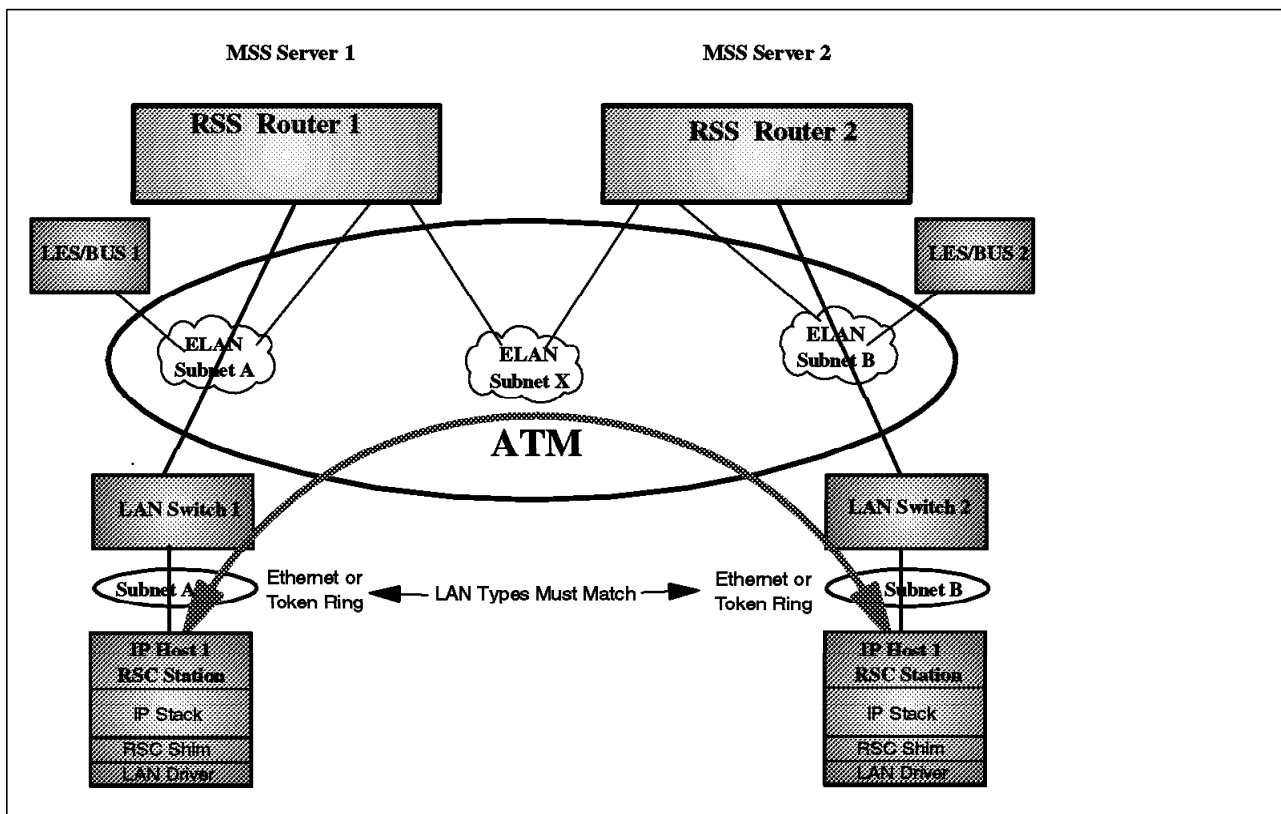


Figure 48. Zero-Hop Routing for LAN Attached Devices

11.2.1 Overview of RouteSwitching Clients

RSC software is specific to the layer-3 protocol (IP, IPX, etc.), LAN type (Ethernet or token-ring), LAN adapter type (NIC), and device operating system. The currently available RSC support is available for IP on Ethernet and token-ring PCs.

For Ethernet, the RSC software is a protocol driver between the station's network device driver and IP protocol stack. It will be made available as software compatible with multiple vendors' Ethernet adapters and IP software stacks. Below are the supported operating systems and network interfaces:

- Windows NT Workstation and Server 4.0 (NDIS 4.0)

- DOS, Windows 3.1, Windows for Workgroups, Windows 95 (16 and 32 bit ODI)
- Novell IntraNetWare V3.12 and later (32-bit server ODI)

For token-ring, the RSC software is embedded in the token-ring device driver, and will be available with specific types of token-ring adapters. Below are the supported operating systems and network interfaces:

- Windows NT Workstation and Server 4.0 (NDIS 4.0)
- OS/2 Warp V3 and later (NDIS 2.01)
- Novell IntraNetWare V3.12 and later (32 bit server ODI)

11.2.2 RouteSwitching Environments for Clients

The RouteSwitching technology offers flexibility so that stations can be migrated to RSC, while other stations continue operating with no change. This is due to MSS supplying traditional routing function while simultaneously supplying advanced functions such as NHRP and RSS. MSS is able to do this, since zero-hop routing and NHRP are compatible with traditional routing protocols. We now explore different scenarios for zero-hop routing.

11.2.2.1 Shortcut Routes

Each shortcut route is unidirectional, thus allowing flexibility when various devices have different capabilities. Thus an RSC can have a zero-hop route to a LAN device that has no RSC capability, while the non-RSC device may achieve a one-hop route to the RSC device (via an ATM-attached LAN switch and MSS 1.1 NHRP function).

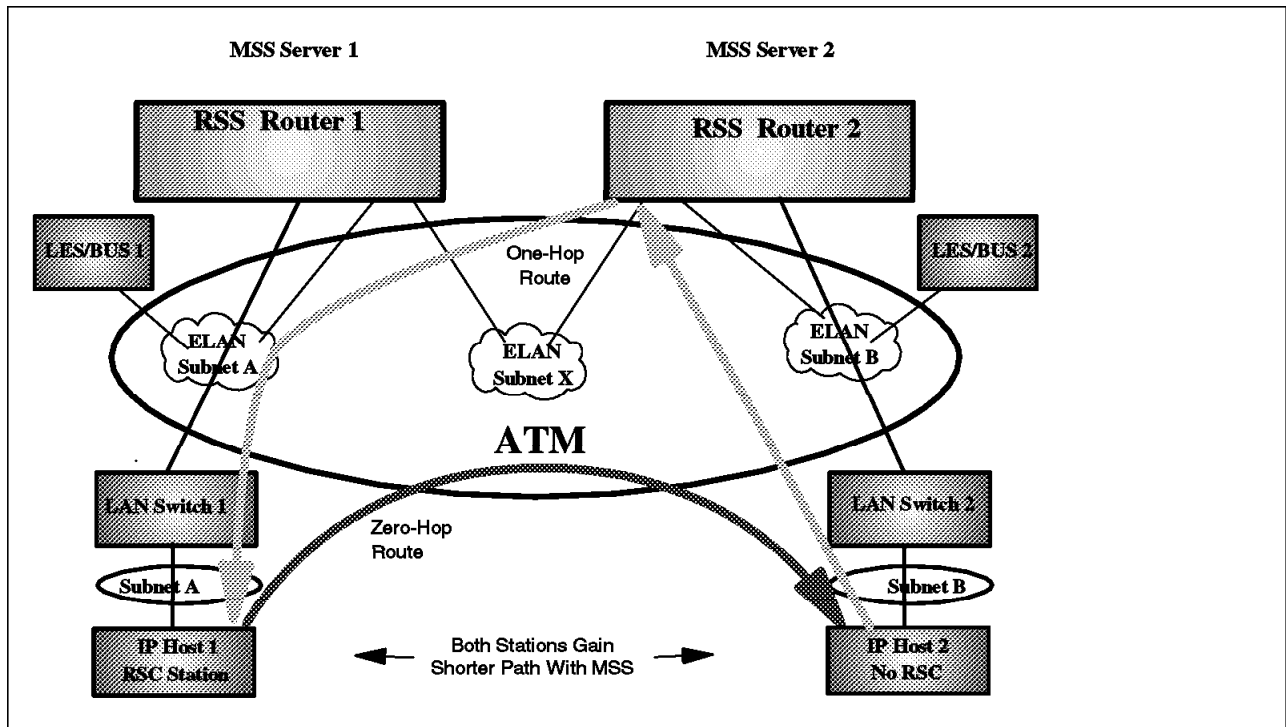


Figure 49. Zero-Hop Routing and One-Hop Routing between Two LAN Devices

11.2.2.2 Mixing of RSCs and Non-RSCs on Same Subnet

Likewise, there can be a mix of devices on the same subnet, some of which are RSCs, others are not. They all may benefit from the MSS RSS or NHRP function, whether obtaining zero-hop or one-hop shortcut routes to their destinations.

11.2.3 RouteSwitching Environments With MSS

11.2.3.1 Shortcut Routes with Non-RSS Routers in the Path

As in MSS 1.1 one-hop routing, if there are routers in the path that do not support RSS, the cut-through path can be set up from the RSC to the non-RSS router, still potentially offering a route with less hops. Note that the transit subnets must be emulated LANs and must match the LAN type of the RSC.

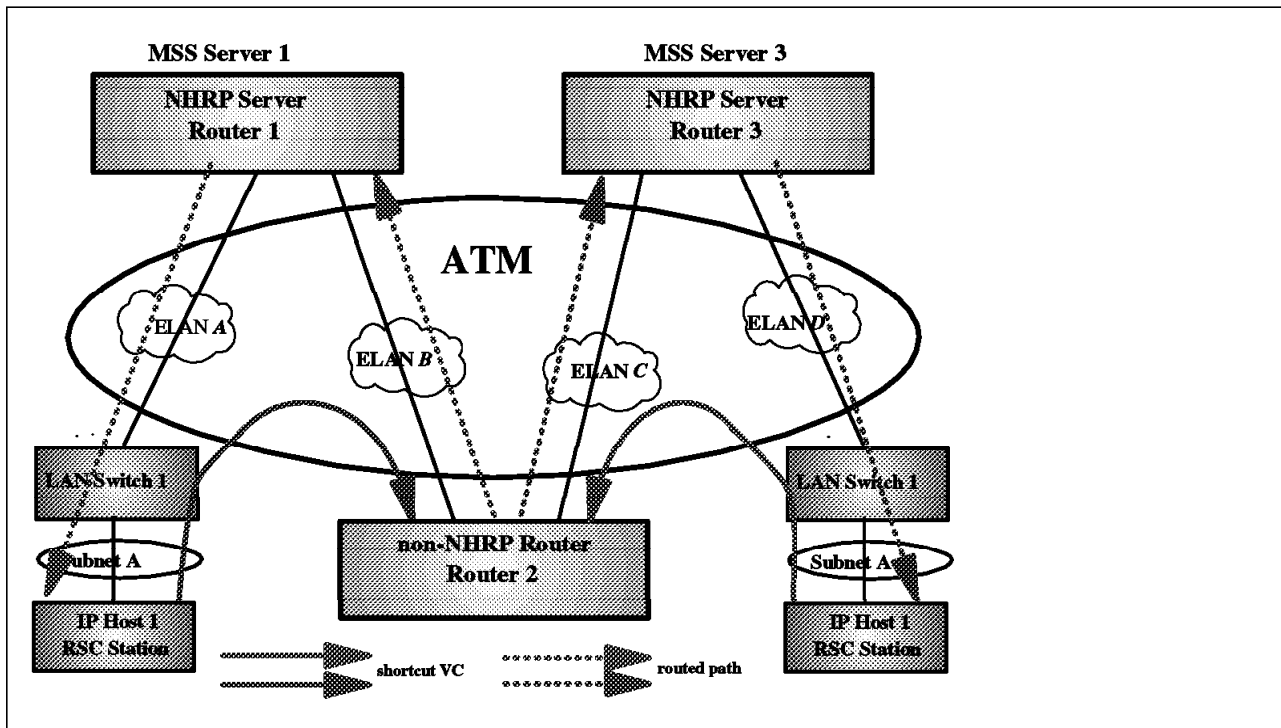


Figure 50. Reduced-Hop Routing with a Non-RSS Router in the Path

11.2.3.2 RSCs on Different LAN Types

To obtain a zero-hop route, the two RSCs must be on the same LAN type (Ethernet or token-ring). If they are not, a shortcut path (although not a zero-hop) may still be achieved as depicted in Figure 51 on page 99. This is the same path available to non-RSC clients, and is supported via MSS NHRP function.

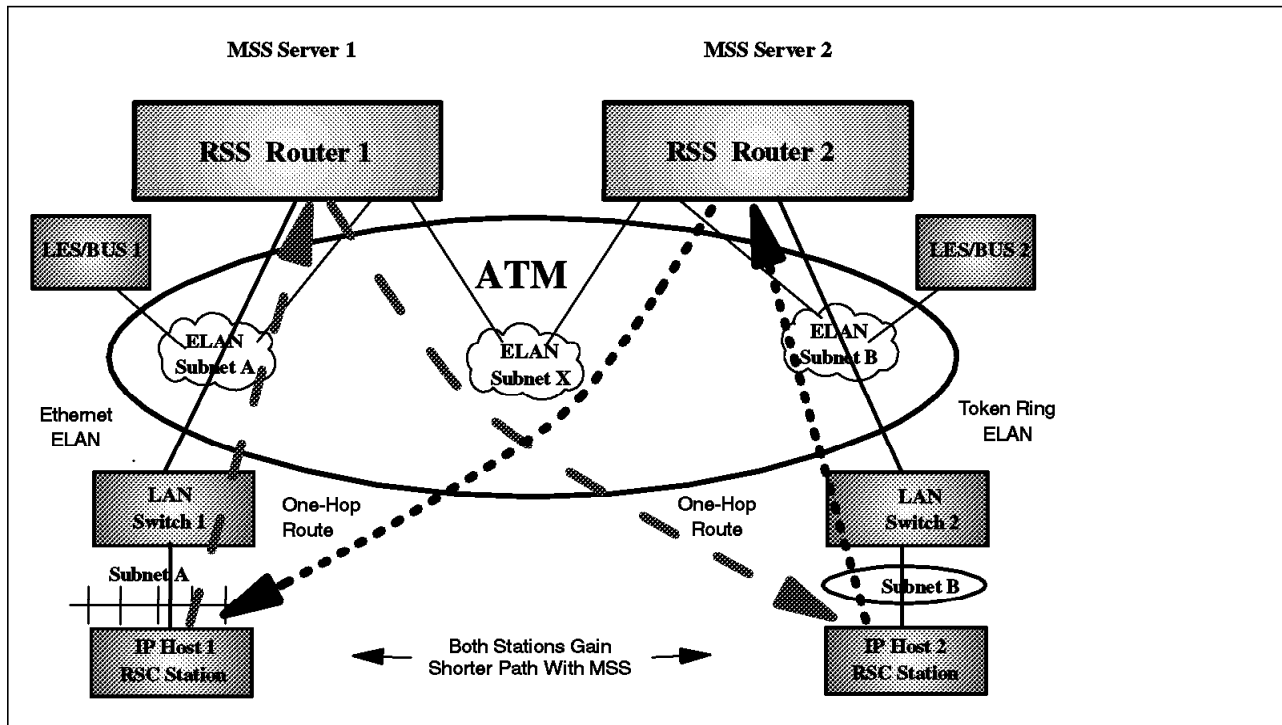


Figure 51. RSCs Achieving a Reduced Hop Route when LAN Types Differ

11.2.3.3 Super-ELANs and RouteSwitching

As described previously, a Super-ELAN can be designed to allow protocol-based VLANs. RouteSwitching can be used in a Super-ELAN (either token-ring or Ethernet).

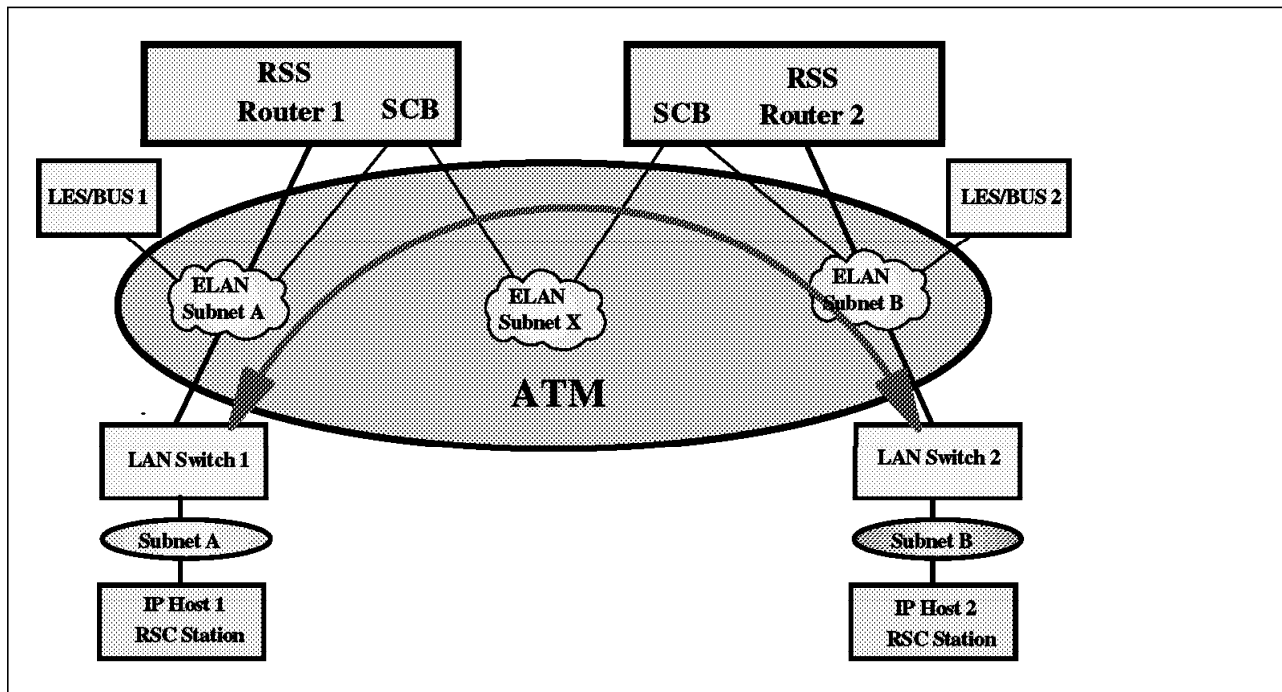


Figure 52. RSCs Zero-Hop Route in a Super ELAN Environment

11.2.3.4 Other Networking Considerations

The way that MSS achieves RouteSwitching differs based on the configuration of the network between the RSCs. The portion of the network with the zero-hop function must be completely Super-ELAN bridged, completely source-route bridged, or completely routed between ELANs.

If Ethernet ELANs are not connected by routing, then they must be connected by Super-ELAN SCBs. Transparent bridging between the ELANs is not supported.

The LAN stations (RSCs) must have the RSS set as their default router, as this is needed for the RSC logic to know a zero-hop route may be needed.

An RSS allows a maximum of 50 registrations per ELAN. In an Ethernet routed environment, this limits to the maximum number of shortcuts to 50 destinations from the same ingress (source) ELAN. In Super-ELANs, this is not limited, since the LES/BUS contains ATM-LAN address mapping for the destination.

The MTU size of the source must be less than or equal to the MTU size of the destination. This is because there are no intermediate routers to do IP fragmentation.

If layer-3 LAN switches are supporting RSCs and implementing policy-based VLANs, RouteSwitching shortcuts may not be established. An example is an RSC that belongs to a VLAN based on a MAC address. Understanding the VLAN configuration options for a LAN switch may allow for adjusting policies or new levels of microcode to overcome these situations.

Similar to the situation in Super-ELANs, there is a possibility that two LAN switches in different ELANs supporting RSCs might have the same LECID. Some LAN switches may look at the LECID of the LEC requesting the data direct VCC and not allow the VCC to be set up if it has the same LECID as itself. This can be avoided by using the MSS LES/BUS function, which allows LECID ranges to be configured per ELAN, and ensuring the ranges in various ELANs do not overlap.

11.2.4 RouteSwitching Flows

Below is a general flow for the setup of a zero-hop route between RSCs using Figure 53 on page 101 as the example. IP HostA wishes to communicate with IP HostB.

- During initialization, the RSC in a LAN device discovers the MAC address of the RSS via ARP. The RSS must be its default router. IP HostA discovers the RSS1's MAC address is MAC1 by ARPing for 9.1.1.9.
- While traffic destined for another subnet is being sent to the RSS1 (its default router), the client sends to RSS1 an NHRP Resolution Request to determine the layer 2 information of the destination. The Layer 2 information consists of the MAC address, Routing Information Field (RIF) for token-ring, and the ATM address of the egress LAN switch. IP HostA sends a request to RSS1 to determine the layer 2 information for IP HostB (9.1.3.1).
- The ingress RSS (RSS1) communicates with other RSSs along the routed path to the destination and obtains the layer 2 information from the egress RSS (the RSS on the destination LAN, RSS2) consisting of the MAC address (MACB), RIF from RSS2 to HostB (003,2,004), and the ATM address of LAN Switch2 (ATM2).

- The ingress RSS (RSS1) sends an NHRP Resolution Reply to the requesting client (IP HostA) with the destination layer 2 information (for IP HostB). This includes the MAC address (MACB) and the full RIF, which is a concatenation of RIF HostA-RSS1 (001,1,002), a virtual route descriptor (A.999.A), and RIF RSS2-HostB (003,2,004). RSS1 also registers the virtual route descriptor to ATM2 address binding with LES/BUS1 if needed (if not a Super-ELAN). The entire RIF is not actually used to source-route the packet, since LAN Switch 1 will check with LES/BUS1 for the next SRD (A,999), then set up a Data Direct VCC to LAN Switch 2.
- The RSC caches the information (MACB,001,1,002,A,999,A,003,2,004) and uses it to build data link headers for frames to HostB (9.1.3.1).
- Frames destined for IP HostB are then sent via the normal Layer-2 procedures, which will result in a shortcut data direct VCC between the LAN switches with ATM links (LECs). In this example, the LEC in LAN Switch1 issues an LE-ARP request for the next_RD (the virtual RD in the RIF, A,999), and receives the ATM address of LAN Switch2 (ATM2) in the LE-ARP response from LES/BUS1. LAN Switch1 then establishes a data direct VCC with LAN Switch2.
- This process is repeated by HostB in the opposite direction.

The flow for Ethernet ELANs is similar; it does not require any route descriptors and only the MAC address is needed.

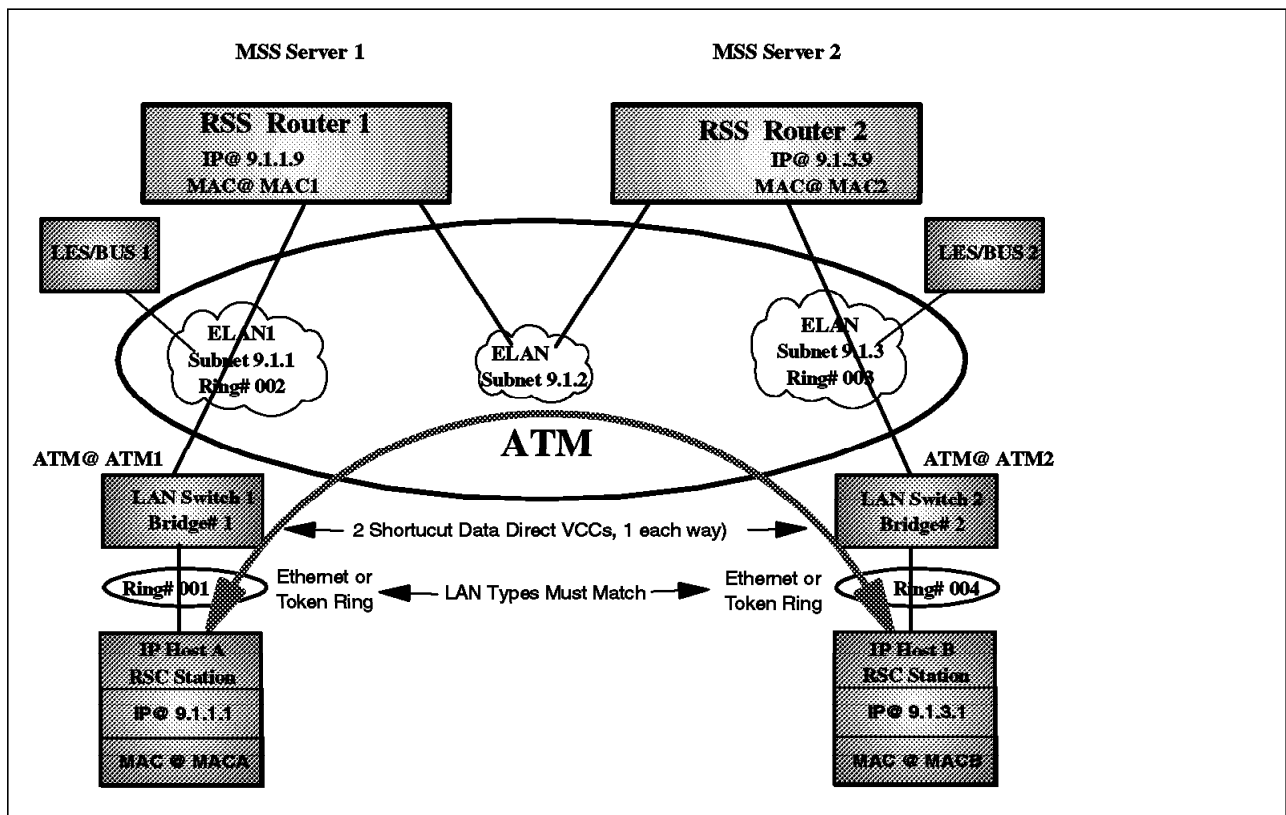


Figure 53. Network Diagram for Zero-Hop Flow Example

11.2.4.1 Additional Flow Considerations

Zero-hop routing uses the same techniques for refreshing shortcuts and purging invalid shortcuts as NHRP as described earlier in this chapter.

Chapter 12. Default IP Gateway Redundancy and LAN Emulation

Default IP gateways provide routing services for hosts that do not run routing topology protocols (for example, OSPF or RIP). These hosts are commonly configured with the IP address of their default gateway. Consequently, a backup default gateway must be accessible via the same IP address as the primary default gateway.

The MSS Server can provide the default gateway function for an ELAN when hosts are configured with the IP address of an MSS Server LE client. The MSS Server also provides support for redundant default IP gateways. Redundancy is important in preserving connectivity with other subnets in the event of a primary default gateway failure since, without a backup, hosts configured with the IP address of a failed default gateway will lose routing services, and therefore will not be able to communicate with stations on other subnets.

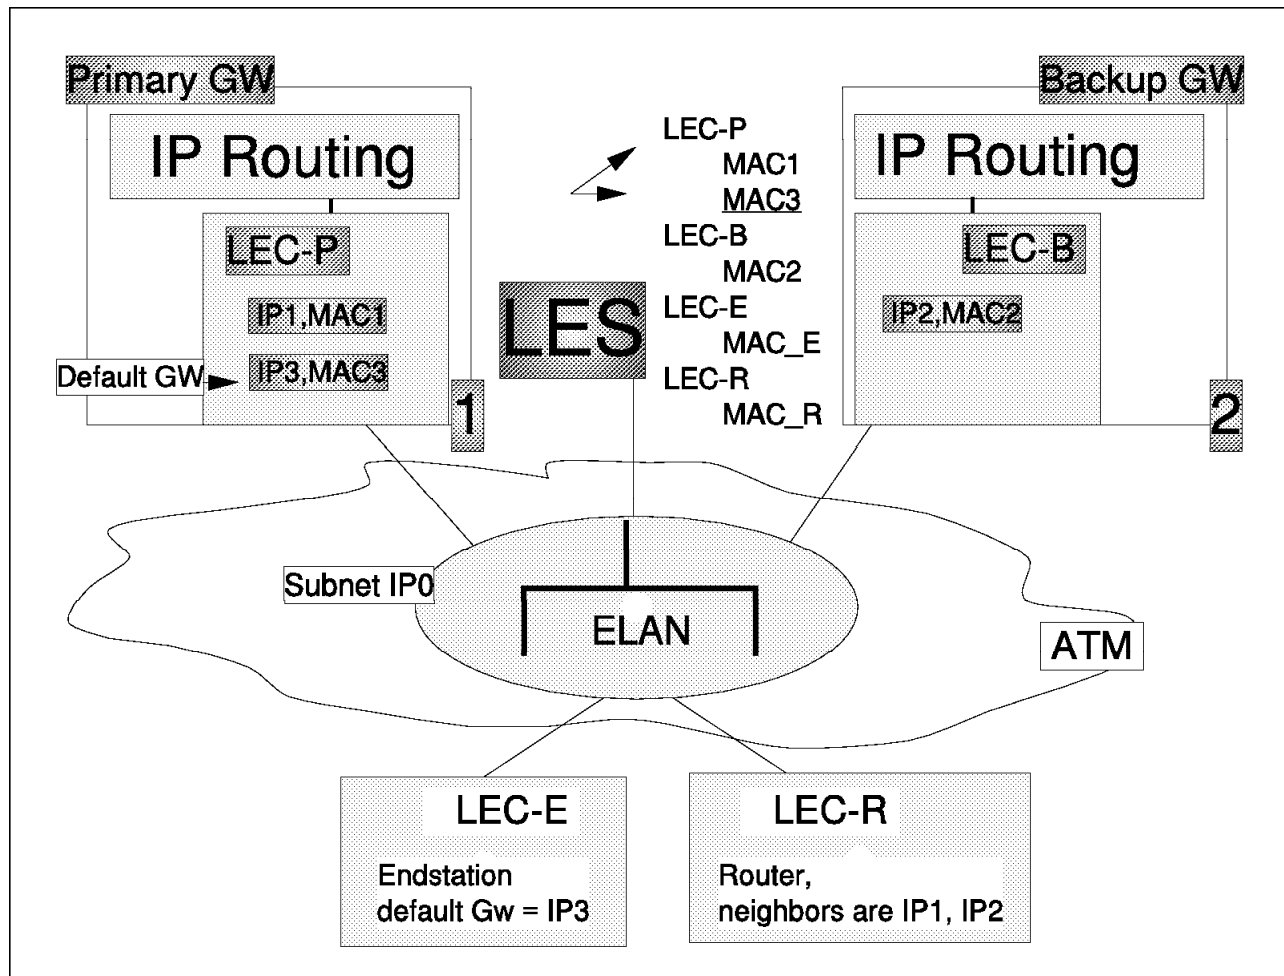


Figure 54. Default IP Gateway Redundancy - Normal Operation. The LES MAC address table contains two MAC addresses for the LE client that provides the default gateway (LEC-P).

In Release 1.0 of the MSS Server, support for default IP gateway redundancy was based on the LE service redundancy mechanisms. In the Release 1.0 approach, one MSS Server contains both the primary LES/BUS for the ELAN and the LEC associated with the primary default IP gateway. This LEC is configured to use the primary LES/BUS, while the LEC associated with the backup default IP

gateway is configured to use the backup LES/BUS. Both LECs are configured with the same IP address (that is, the default gateway IP address) and the same MAC address. As long as the primary LES/BUS is active, the backup LEC will not be allowed to join the ELAN, and default IP gateway services will be provided via the primary LEC. If the MSS Server containing the primaries should fail, the backup LES/BUS will become active and allow the backup LEC to join the ELAN. Then, default IP gateway services will be provided via the backup LEC. Since both LECs have the same IP and MAC addresses, hosts using default gateway services will not notice any difference.

Although the Release 1.0 solution has proved useful in a number of network designs, Release 1.1 offers an improved approach that provides several advantages. The Release 1.1 design will be described in the context of the example redundant default IP gateway configuration depicted in Figure 54 on page 103.

As in the Release 1.0 approach, the IP address of the default gateway, @IP3, is configured on both MSS Servers, and the same MAC address, MAC3, is associated with the default gateway IP address on both the primary and the backup. However, the Release 1.1 solution differs in that each MSS Server is also configured with a unique IP address that is associated with a normal routing interface. These routing interfaces are active on both the primary and the backup gateways, and routing topology protocols are associated with these interfaces. This allows other routers on the ELAN, such as router LEC-R in Figure 54 on page 103, to learn the backup gateway's connectivity while the primary is still active, which increases responsiveness in a backup situation. Additionally, it allows a router hop to be eliminated when an MSS Server hosting a backup gateway must route traffic to/through the ELAN for which backup gateway service is being provided.

Other advantages of the Release 1.1 approach include more intuitive configuration due to a straightforward user interface devoted to default IP gateway redundancy, and elimination of the requirement to co-locate the primary LES/BUS and the primary default IP gateway.

Although the LE clients supporting the primary and backup gateways are active simultaneously, only one MSS Server provides default gateway services at any given time. These LECs register the unique MAC address associated with their normal routing interface when joining the ELAN. Both LECs also attempt to register the MAC address, MAC 3, associated with the default gateway IP address. If this MAC address has already been registered, the LES will reject the registration request due to the duplication. In this case, the LEC will retry the registration later. The backup retries every 30 seconds, while the primary retries every 5 seconds.

Single LEC - Multiple MAC Addresses

The redundancy feature operates by virtue of the fact that a LEC can register multiple MAC addresses at the LES. With each MAC address a different IP address is associated.

Figure 55 on page 105 depicts the situation when the primary default gateway is inactive and the backup MSS Server has taken over. As can be seen at the LES, MAC addresses 2 and 3 have been registered for LEC-B.

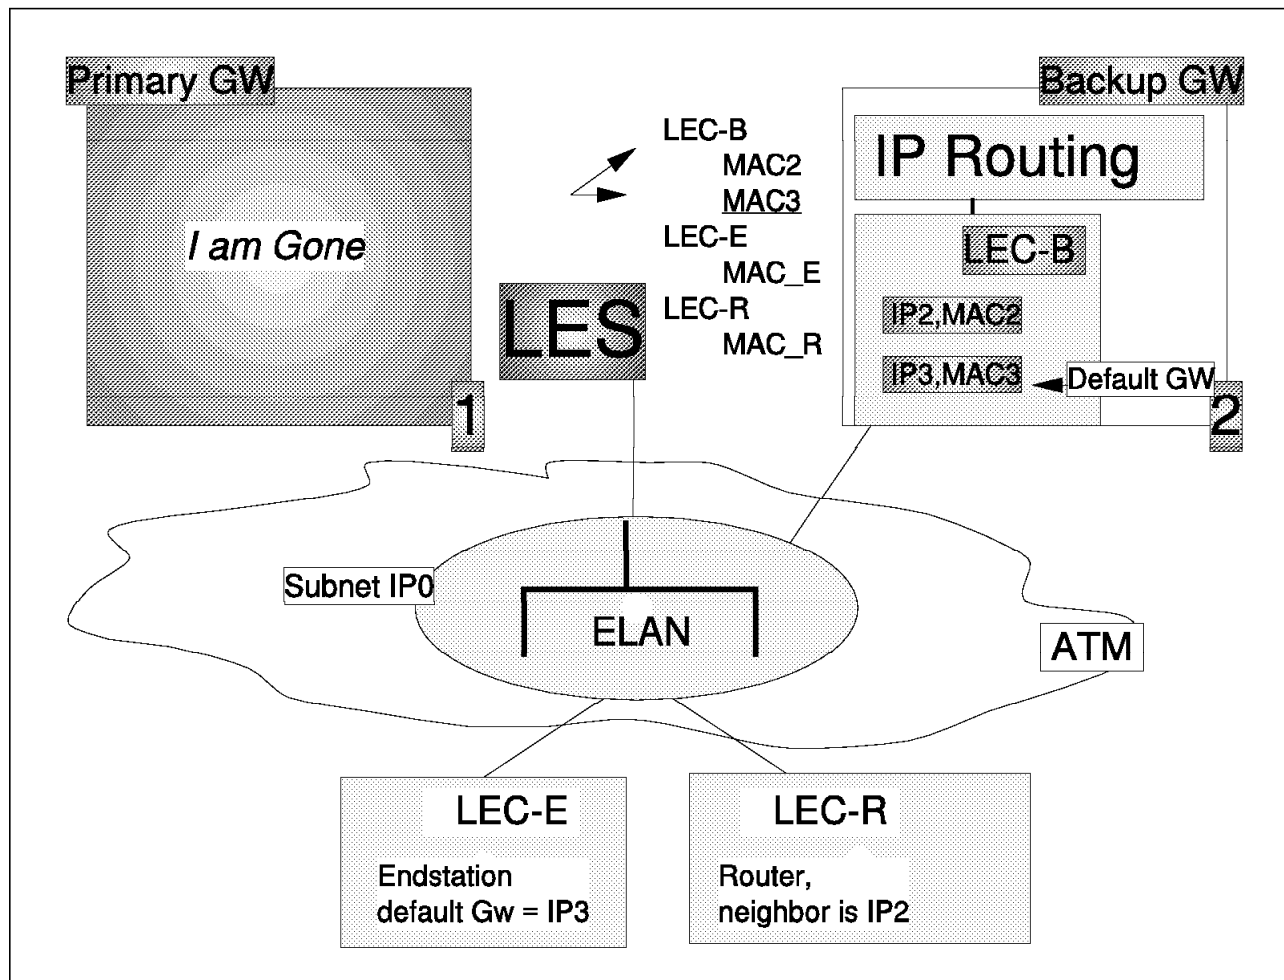


Figure 55. Default IP Gateway Redundancy - Backup. The LES MAC address table contains two MAC addresses for the LE client that provides the backup default gateway (LEC-B).

If the primary MSS Server becomes active again, it will join the ELAN and register MAC address 1. Registration of MAC address 3 will fail because the address has been registered by LEC-B. If the primary's registration fails, then the primary sends a message to the backup indicating its desire to become active as the default IP gateway for the ELAN. The backup responds to the message by de-registering the MAC address associated with the default gateway IP address. This allows the primary's next registration attempt to succeed.

The gateway that has successfully registered the MAC address provides default gateway services for the ELAN, and the other gateway ignores requests targeted to the default gateway IP address. Thus, if the primary is active, it will always provide the default gateway services, and if the primary should fail, the backup will take over within 30 seconds. From an endstation's perspective, the switch from primary to backup is seen only as a momentary outage, since the primary and backup gateways appear to be identical. The use of the same MAC address at both gateways is important in minimizing the disruption because it ensures that ARP entries cached at endstations remain valid.

From the MSS Server's perspective, the IP address associated with a default gateway is special in the sense that it has limited functionality. ARP responds to requests for this IP address with the MAC address used for the default gateway, and pings and Telnets to the address are also supported for network

management purposes, but the IP address is transparent to routing topology protocols, and the associated MAC address is never used when forwarding a routed frame onto the ELAN.

12.1 Configuration Issues

When defining redundant default IP gateways for ELANs be aware of the following:

Note: See 17.7, "Example - Redundancy" on page 178 for an example.

1. Make sure one MSS Server is defined as primary while the other is defined as the backup default gateway.
2. Make sure the IP address of the default gateway (IP3 in Figure 54 on page 103) is identical on both MSS Servers.
3. Make sure the MAC address of the default gateway (MAC3 in Figure 54 on page 103) is identical on both MSS Servers.
4. Make sure the IP addresses of the non-default gateway interfaces (IP1 and IP2 in Figure 54 on page 103) are different.
5. Make sure the MAC addresses of the non-default gateway interfaces (MAC1 and MAC2 in Figure 54 on page 103) are different.

There is no need to have the primary and backup MSS Server attached to the same ATM switch.

Default IP gateway redundancy can be combined with LES/BUS and LECS redundancy. Define the primary LES/BUS/LECS on the same node that provides the primary IP gateway function.

Attention

The MSS Release 1.1 that became generally available in March '97 does not support concurrent bridging or redundant default IP gateway for LAN emulation.

When we were completing this redbook a fix was tested that will be made available soon (see also page 181).

Chapter 13. Classical IP Redundancy and Multicast Features

This section describes mechanisms that improve Classical IP reliability and function. More specifically, the mechanisms address the following three issues:

1. The single point of failure represented by the ATMARP server for a logical IP subnet (LIS)
2. The single point of failure represented by the default gateway for the LIS
3. The lack of support for multicasts and broadcasts in a LIS

Classical IP is specified in RFC 1577, which does not provide for ATMARP server backup or distributed ATMARP servers. The role of the ATMARP server is to provide IP-to-ATM address mappings to requesting ATMARP clients within a LIS. If the ATMARP server fails or connectivity to it is lost, ATMARP clients in that LIS will no longer be able to resolve IP addresses to ATM addresses, and therefore will not be able to establish new VCCs. Furthermore, ATMARP clients will not be able to refresh existing mappings via the ATMARP server. Thus, failure of the ATMARP server can cripple communications on the associated LIS.

If a LIS spans locations, all ATMARP clients must communicate with a single ATMARP server even if it is in a different physical location. A loss of ATM connectivity will have the same effect as the ATMARP server failing. Thus an ATM link outage can affect ATMARP clients as much as ATMARP server failure, depending on the length of interruption.

The default gateway for a LIS provides the same function as the default IP gateway for an ELAN described in Chapter 12, "Default IP Gateway Redundancy and LAN Emulation" on page 103. Namely, the default gateway for a LIS provides routing services for hosts that do not run routing topology protocols (such as OSPF). Since hosts are commonly configured with the IP address of their default gateway, a default gateway failure will result in the loss of routing services for hosts on the LIS.

Both of these problems are significant because they impact the ability to deploy Classical IP equipment in networks carrying mission-critical traffic.

Although Release 1.0 of the MSS Server provided support for ATMARP server redundancy, limitations associated with the support restricted network design flexibility. Release 1.1 eliminates some of these restrictions. Limitations of the Release 1.0 ATMARP server redundancy support include: (1) both the primary and backup ATMARP server must be attached to the same ATM switch, (2) the user has no configuration control over which ATMARP server is the primary and which is the backup (that is, the first one to power up and register with the switch is the primary), and (3) the MSS Servers hosting the primary and backup ATMARP servers cannot perform IP routing. The first limitation is still present in Release 1.1, but the other two limitations have been removed. Furthermore, default gateway redundancy support for LISs is also provided in Release 1.1. MSS 2.0 adds the ability to have multiple ATMARP servers supporting one LIS and may be used in conjunction with redundant ATM ARP servers. This allows placement of ARP servers where appropriate.

Without the ability to support broadcasts and multicasts, Classical IP over ATM was limited in its usability. Functions such as IP routing protocols, DHCP and multicasting were either not available or made difficult to implement. MSS 2.0 supports a new set of protocols specified in RFC 2022, IP Multicast over ATM.

ATMARP clients supporting these protocols can now use IP functions requiring broadcasting or IP multicasting.

13.1 MSS 1.1

13.1.1 ATMARP Server Redundancy

Providing ATMARP server redundancy in a flexible manner is difficult because most ATMARP clients are configured with the ATM address of a single ATMARP server. Thus, to be widely applicable, a backup ATMARP server must be accessible via the same ATM address as the primary. This can be accomplished by connecting two MSS Servers, one with the primary ATMARP server and one with the backup, to the same 8260 ATM switch. By virtue of being connected to the same IBM 8260 ATM switch, the network prefix of the ATMARP servers' ATM addresses will be the same. Given the same network prefix, all that remains is for the backup to assume the primary's ESI and selector. Note that this approach also works with ATM switches other than the 8260, provided that both MSS Servers receive the same network prefix from the switch.

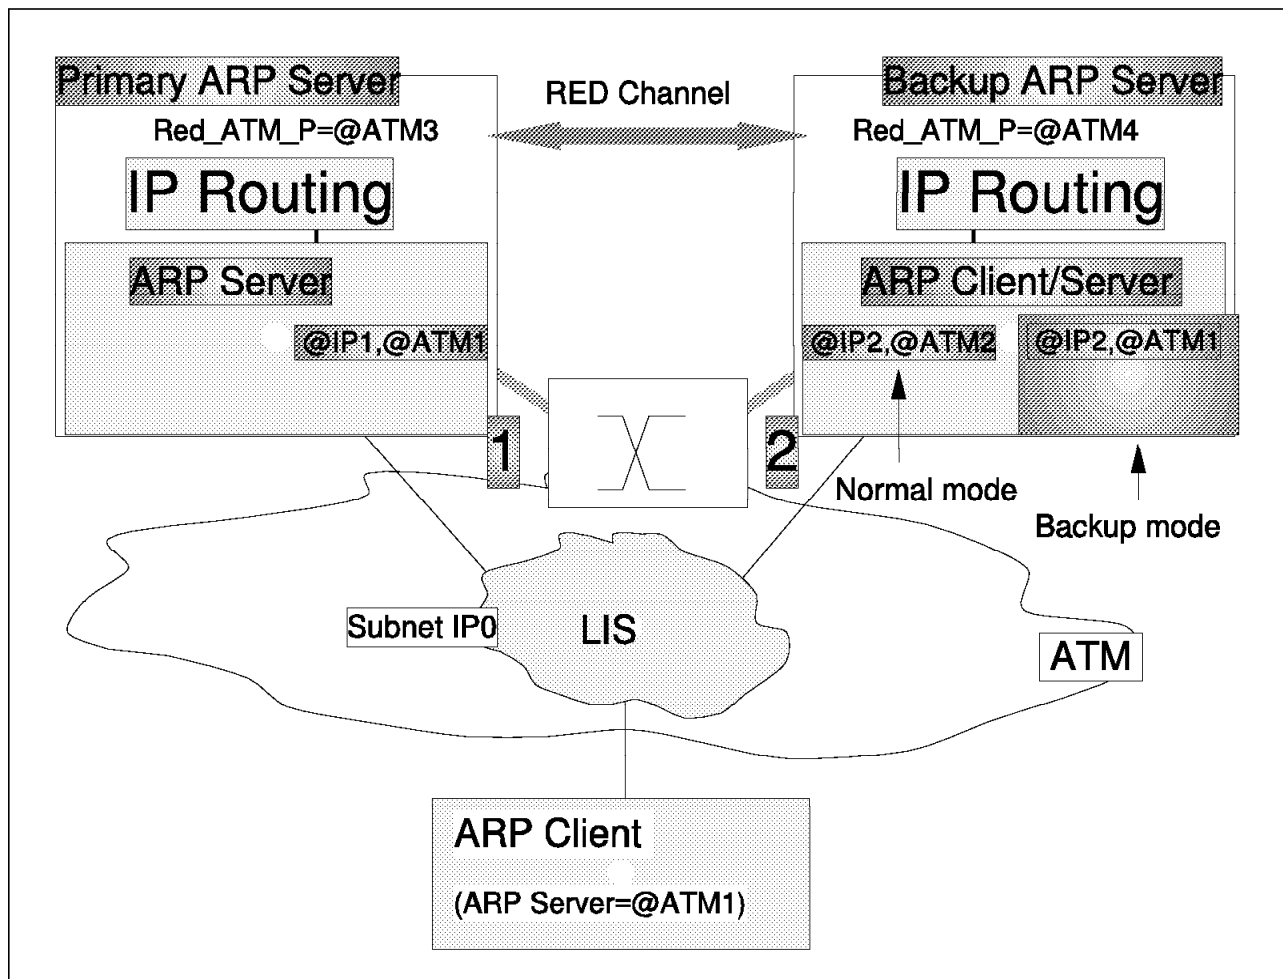


Figure 56. ATMARP Server Redundancy

During normal operation the primary MSS Server provides the ATMARP server function using @ATM1 and @IP1. On the backup MSS Server an ATMARP client will be active with a different ATM and IP address (depicted as @ATM2 and

@IP2 respectively in Figure 56). Using different ATM addresses and different IP addresses at the primary and backup enables both MSS Servers to simultaneously connect to the same LIS.

The ARP server redundancy has been made possible by the implementation of a mechanism for the backup ATMARP server to assume the primary's ATM address. The question is how does the backup know that the primary is not available to service address resolution requests for the LIS?

In Release 1.1, the user configures which MSS Server is to host the primary ATMARP server and which is to host the backup. The primary is then responsible for establishing a redundancy VCC, or RED channel, to the backup. The presence of this VCC indicates that the primary is serving the LIS. Conversely, the absence of the VCC indicates that the secondary should become active. So, when the RED channel is not present, the backup ATMARP server attempts to register the ATM address (@ATM1) of the primary ATMARP server with the ATM switch. After successfully registering this ATM address, the backup ATMARP server becomes active and services address resolution requests for the LIS. If the RED channel is established while the backup ATMARP server is actively servicing the LIS, the backup ATMARP server de-registers the primary ATMARP server's ATM address (@ATM1) so that the primary ATMARP server can register its ATM address and begin/resume servicing the LIS. In addition it re-registers its ATM address for normal operation (@ATM2).

Since the RED channel must be able to be established when the backup ATMARP server has registered the primary ATMARP server's ATM address, the ATM addresses used to establish the RED channel must have ESIs that are different than the ESI used for the primary ATMARP server's functional ATM address. The ESIs must be different from each other as well (so that they can both be registered with the switch).

13.1.2 ATMARP Server/Default IP Gateway Redundancy

The use of different ATM addresses and different IP addresses at the primary and backup enables both MSS Servers to simultaneously provide IP routing functions. In the failure scenario, the backup de-registers its ATM address and re-registers using the primary's ATM address. Therefore, when external ATMARP clients try to re-connect with their ATMARP server, the calls will be directed to the backup, which now functions as both an ATMARP client and the active ATMARP server for the LIS. The backup retains its original IP address, and does not assume the IP address of the primary. A third IP address is used to provide default gateway redundancy.

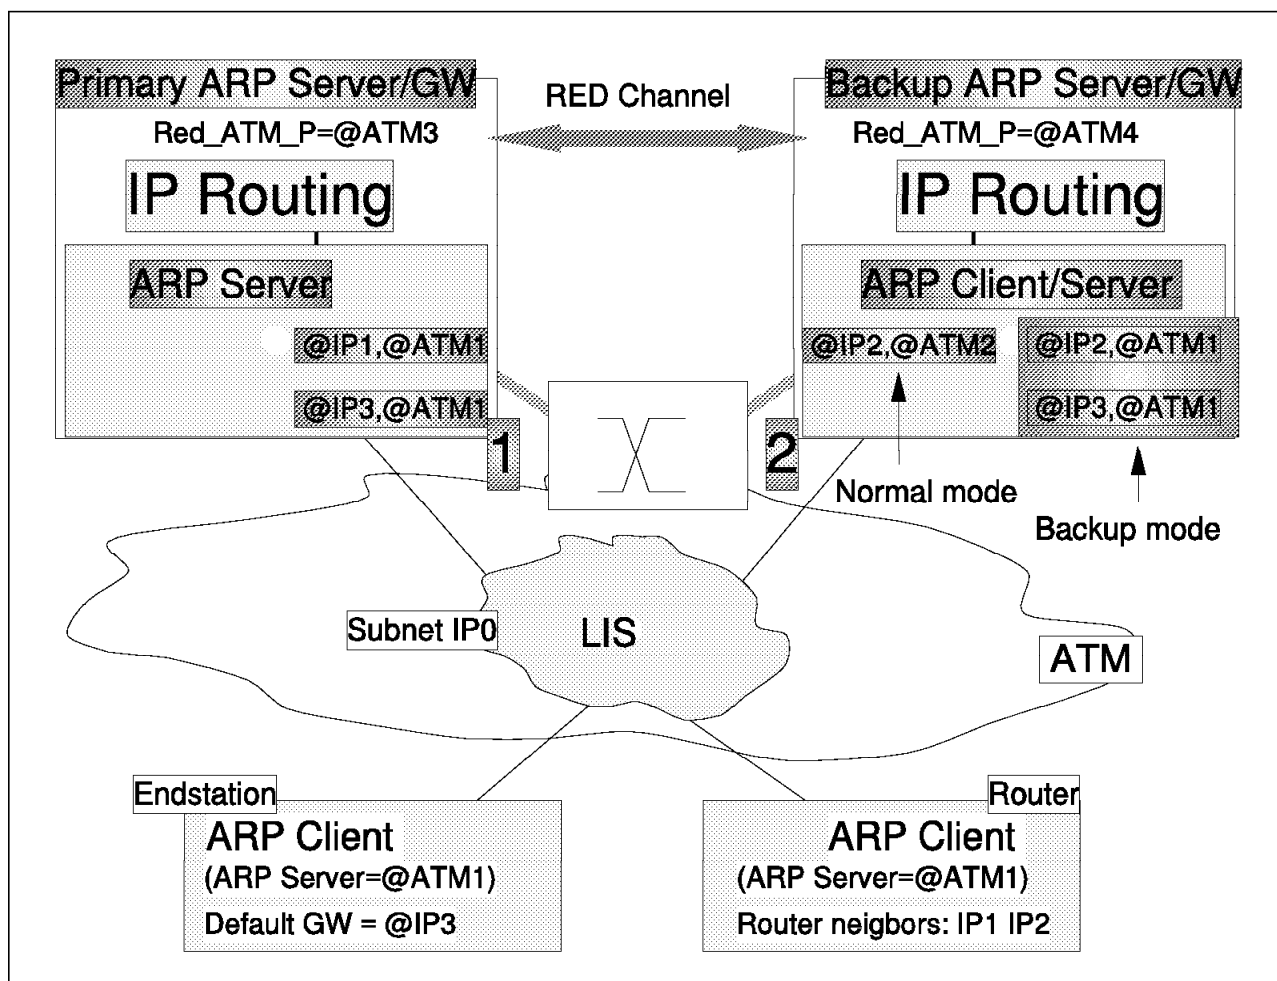


Figure 57. ATMARP Server/Default Gateway Redundancy

Default gateway redundancy support for LISs is provided as an optional extension of ATMARP server redundancy, and is similar to the default IP gateway redundancy support for ELANs that was described in Chapter 12, "Default IP Gateway Redundancy and LAN Emulation" on page 103.

As in the ELAN case, both the primary and the backup are configured with the same default gateway IP address, and use of the default gateway IP address is restricted. The active ATMARP server for the LIS responds to ARP requests for the default gateway IP address (with its own ATM address), and the ATMARP client associated with the active ATMARP server provides default gateway routing services for the LIS. When the primary ATMARP server is active, the backup does not respond to requests for the default gateway IP address. However, if the backup takes over as the active ATMARP server, it also assumes responsibility for the default gateway function.

13.1.3 IP Addresses

In addition to answering ARP requests for the default gateway IP address, the MSS Server containing the active default gateway also responds to ping and TELNET requests destined for the default gateway IP address. However, the default gateway IP address is not used for normal routing purposes; the unique IP addresses associated with the ATMARP clients are used instead. For example, if the default gateway IP address for LIS 9.1.1 is 9.1.1.3, then the

primary might use IP address 9.1.1.1 for normal routing purposes, and the backup might use IP address 9.1.1.2. With this design, the default gateway redundancy mechanisms are transparent to routing topology protocols, and the backup can provide non-default gateway routing functions for the LIS while the primary is active (for example, forwarding traffic to/through the LIS and exchanging topology information with neighboring routers on the LIS).

To summarize, when both the primary and backup ATMARP servers are operational and default gateway redundancy is enabled, the primary has two IP addresses on the LIS: an ATMARP client/server IP address (@IP1) and a default gateway IP address (@IP3), while the backup has a single ATMARP client IP address on the LIS (@IP2). If the primary fails, the primary will obviously cease to have any appearance on the LIS, and the backup will have two IP addresses on the LIS: its original ATMARP client IP address (@IP2) and the newly assumed default gateway IP address (@IP3). The backup will also assume the role of the active ATMARP server for the LIS by taking over the ATM address of primary ATMARP server. If the primary recovers and establishes the RED channel to the backup, the backup will de-register the primary's ATM address (@ATM1) and relinquish ATMARP server responsibility for the LIS. Furthermore, it will relinquish responsibility for responding to requests for the default gateway IP address (@IP3), register its client ATM address (@ATM2), and resume operation as an ATMARP client on the LIS with the primary as its ATMARP server.

13.1.4 ATM Addresses

There are several restrictions associated with the ATM addresses of ATMARP entities participating in the redundancy mechanisms (that is, client/server pairs serving as primary ATMARP server/default gateway and clients providing backup ATMARP server/default gateway functions). First, since the solution requires that two MSS Servers be given the same network prefix, then all ATM addresses containing this prefix that are registered by one of the MSS Servers must include an ESI that is not used by any component on the other MSS Server. (That is, no ESI should be used by both MSS Servers.) Additionally, the ATM address of each primary ATMARP client/server pair (@ATM1 and @ATM2) that uses the shared network prefix must include a unique ESI that is not used by any other component on the same MSS Server. (That is, the ESI of a primary ATMARP client/server pair must be unique across both MSS Servers.) Each ATMARP entity must also have a unique ATM address dedicated to its RED channel (@ATM3 and @ATM4). This address, like the ATM addresses of all ATMARP entities, must be unique among all components employing RFC 1483 encapsulation, and furthermore, if the address is for the RED channel of a primary ATMARP client/server pair, the ESI portion of the ATM address must not match the ESI portion of the client/server's functional ATM address (that is, the ATM address used for ATMARP client/server functions).

Configuration of the partner's RED channel ATM address is also required. The primary will place the call for the RED channel to the backup, and the backup will only accept redundancy VCC calls from the primary's RED ATM address. The establishment of the RED channel is all that is required for the primary to communicate its status to the backup. Currently, no data is exchanged over the RED channel.

13.1.5 Configuration Issues

When defining a redundant ARP server for Classical IP subnets be aware of the following:

Note: See 17.7, “Example - Redundancy” on page 178 for an example.

1. Make sure that the primary and backup MSS Servers are attached to the same ATM switch.
2. Identify your primary ATMARP server. On this node:
 - a. Define two ESIs, one for the ARP server (@ESI1) and one to establish the redundancy VCC (@ESI3).
 - b. Configure an ARP server. Use @ESI1 for its ATM address (@ATM1). Assign a unique IP address (@IP1).
 - c. Configure ARP server redundancy. Indicate the use of @ESI3 for the ATM address used by the primary node to establish the redundancy VCC.
 - d. Indicate this redundant ARP server is primary. Enter the RED channel's ATM address (@ATM4) of the backup node.

Note: This enables the primary node to establish the RED channel.

3. Identify your backup ATMARP server. On this node:
 - a. Define two ESIs, one for the ARP client (@ESI2) and one to establish the redundancy VCC (@ESI4).
 - b. Configure an ARP client. Use @ESI2 for the client's ATM address (@ATM2). Define @ATM1 as the ARP server ATM address for this client. Assign a unique IP address (@IP2).
 - c. Configure ARP server redundancy. Indicate the use of @ESI4 for the ATM address used by the backup node for the redundancy VCC.
 - d. Indicate this redundant ARP server is backup (not primary). Enter the RED channel's ATM address (@ATM3) of the primary node.

Note: This enables the backup node to verify the RED channel is established by the primary.

- e. Configure a redundant backup ARP server. Indicate the use of @ESI1 for the ATM address when operating as an ARP server.

When defining default IP gateway functions in addition to the ARP server backup facility remember to enable default IP gateway redundancy on the primary and backup MSS Servers. Assign the same default IP gateway address (@IP3) on both nodes.

13.2 MSS 2.0

13.2.1 Distributed ATM ARP Server

MSS 1.1 ATMARP Server redundancy offered the ability to have redundancy for a LIS. However, the MSS supplying redundancy had to be attached to the same ATM switch, and only one ATMARP server could be active. If the ATM LIS spanned sites, all ATMARP clients would have to communicate to the one MSS.

MSS 2.0 allows for multiple, or distributed ATMARP servers in one LIS; the servers communicate with each other to keep their ATMARP databases synchronized. This new function allows ATMARP (RFC 1577) clients to connect

to an ATMARP server that is nearer to them, while communication across a larger LIS and to other subnets is maintained. This function is possible due to a new protocol which allows the sharing of server databases, known as the Server Cache Synchronization Protocol (SCSP); it is an IETF Draft specification (SCSP-NBMA, available from <http://ds.internic.net/internet-drafts/> currently named draft-ietf-ion-scsp-atmarp-02.txt).

Along with this is an enhanced draft specification for ATMARP clients (http://ds.internic.net/internet-drafts, named draft-ietf-ion-ipatm-classic2-03) which allows the configuration of multiple ATMARP servers in a client. These clients, known as RFC 1577+ clients, gain ATMARP server redundancy by having a list of the distributed ATMARP servers for a LIS. If communications is lost to the current ATMARP server, the client will attempt connection to the next server in its preconfigured list. In a LIS with all RFC 1577+ clients, the distributed ATMARP servers can also supply redundancy without the limitation that they all be on the same ATM switch. Multiple ATMARP servers also allow for load balancing between MSSs that may be performing other functions. They can also allow for a larger LIS.

However, for existing ATMARP (RFC 1577) clients, this does not supply ATMARP server redundancy. These current clients can only define one ATMARP server for a LIS, and thus cannot take advantage of the multiple ATMARP servers. These clients still need the redundant ATMARP server function available in MSS 1.1 and 2.0.

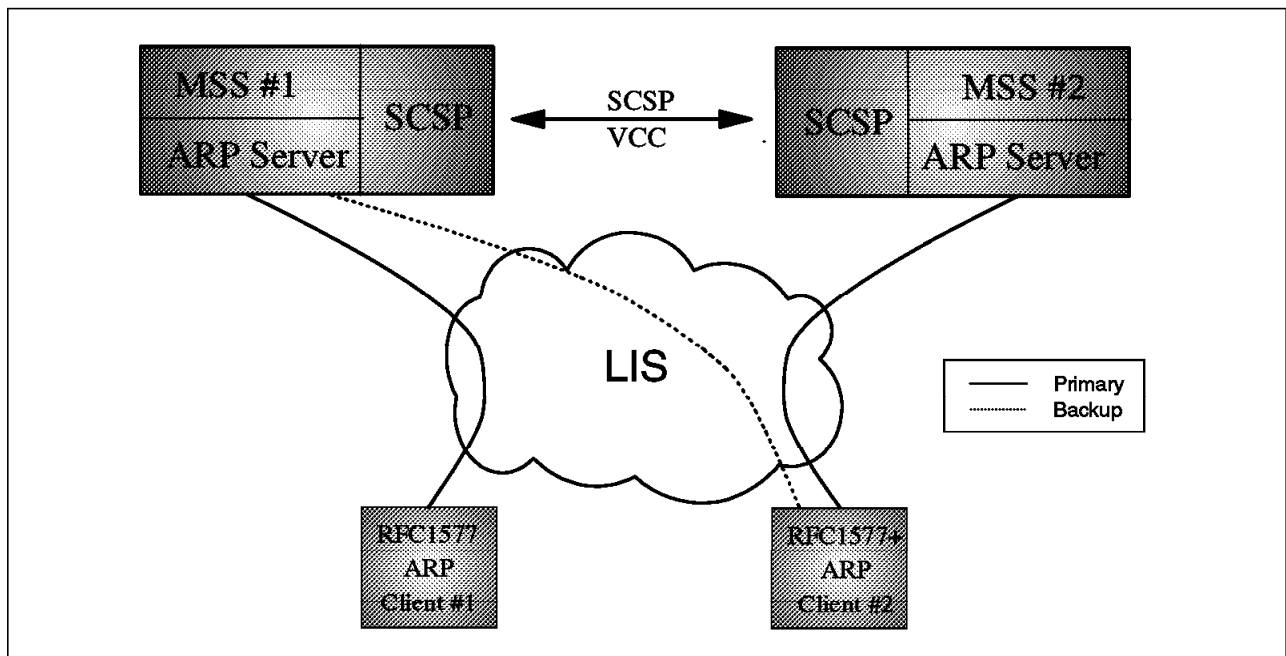


Figure 58. Example of Two ATMARP Servers, RFC 1577 and RFC 1577+ Clients

13.2.1.1 RFC 1577+ Client Considerations and Flows

As mentioned above, an RFC 1577+ must have locally defined the list of ATM addresses for the primary and backup ATMARP servers. Using Figure 58 as an example, client #2, is configured with the ATM address of MSS #2 as its primary ATMARP server. If client #1 wishes to communicate with client #2, it will request the ATM address associated with the IP address of client #2. Even though the

clients are using different ATMARP servers, MSS #1 knows the ARP entries of MSS #2 and supplies the information to client #1.

If client #2 loses communications with MSS #2 (either the VCC is terminated and/or a new call setup fails when attempting to request or refresh an ATMARP entry), it will use MSS#1 as its server with no loss of connectivity. If client #1 loses communications with MSS #1, it will eventually lose connectivity since it has no ability to have backup servers configured because it is an RFC 1577 ARP client (not an RFC 1577+ ARP client), which has no idea about the ARP server residing on MSS #2.

13.2.1.2 ATMARP Server Considerations and Flows

When there are multiple ATMARP servers in a LIS, they must use SCSP. Each server must have (at least one of) the ATM addresses of the other servers in the list configured. These servers are known as directly connected servers (DCSs) to the server that has them defined. This server will then set up a VCC with each of the configured ATMARP servers (DCSs) and exchange database updates (ARP cache changes) in order to attain duplicate databases between the servers. The SCSP function in MSS interacts with the ATMARP server function to receive and report cache changes.

The collection of ATMARP servers in a LIS is known as a server group (SG). Not every server needs to be configured with the address of every other server. Although this full mesh improves redundancy, it also increases the amount of overhead VCCs and traffic. This is only an issue with more than three ATMARP servers in an SG. An example of a configuration with three ATMARP servers is shown in Figure 59; MSS #1 has MSS #2 and MSS #3 as its DCSs. MSS #2s has MSS #1 and MSS #3 as its DCSs. MSS #3 has MSS #1 and MSS #2 as its DCSs. These MSSs can all be on different ATM switches anywhere in the network. A single break in connection between any two DCSs does not stop synchronization of the databases.

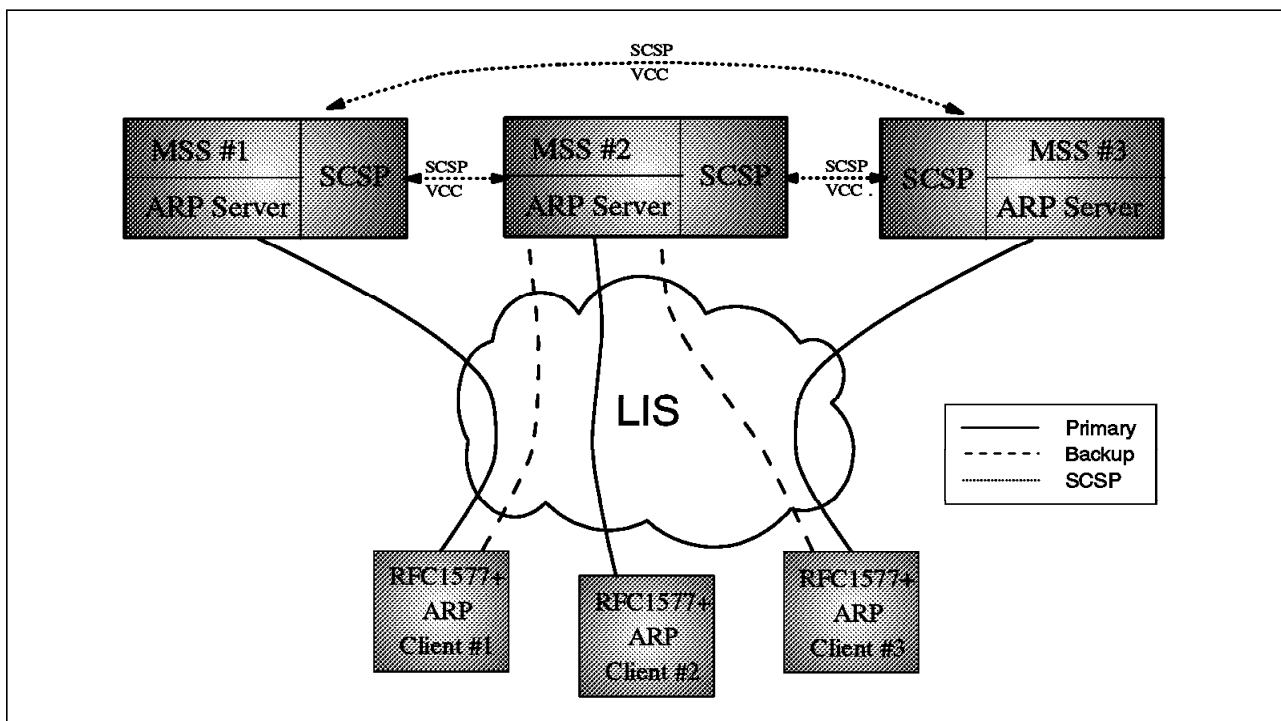


Figure 59. Example of a Server Group (SG) with Three ATMARP Servers in a Ring Topology

In Figure 59, if MSS #2 became inoperative, it would result in loss of connectivity for client #2 (which has only a single ATMARP server defined). If there were no DCS definitions between MSS #1 and MSS #3, and MSS #2 became inoperative, client #1 and client #3 would eventually lose connectivity to each other, since MSS #1 and MSS #3 would not be able to keep their ATMARP databases synchronized (and the ATMARP entries would age out of their tables).

When an MSS initializes, the SCSP function will establish VCCs to the configured DCSs, negotiate a role (master or slave), send ATMARP summaries, then start synchronizing databases. Then as changes occur on an ATMARP server, the changes would be sent by the SCSP function to its DCSs. These DCSs would then update their ATMARP databases. DCSs know their partners are still active with periodic HELLO messages.

13.2.1.3 Distributed and Redundant ATMARP Server Considerations

MSS V2.0 can support distributed ATMARP services for RFC 1577+ clients and support redundant ATMARP services for RFC 1577 clients simultaneously. To do this still requires that the redundant servers be attached to the same ATM switch, since the ATMARP server ATM address is configured into a client. Figure 60 shows a LIS with three ATMARP servers in an SG. MSS #1 is also a primary redundant server while MSS #2 is a backup redundant server.

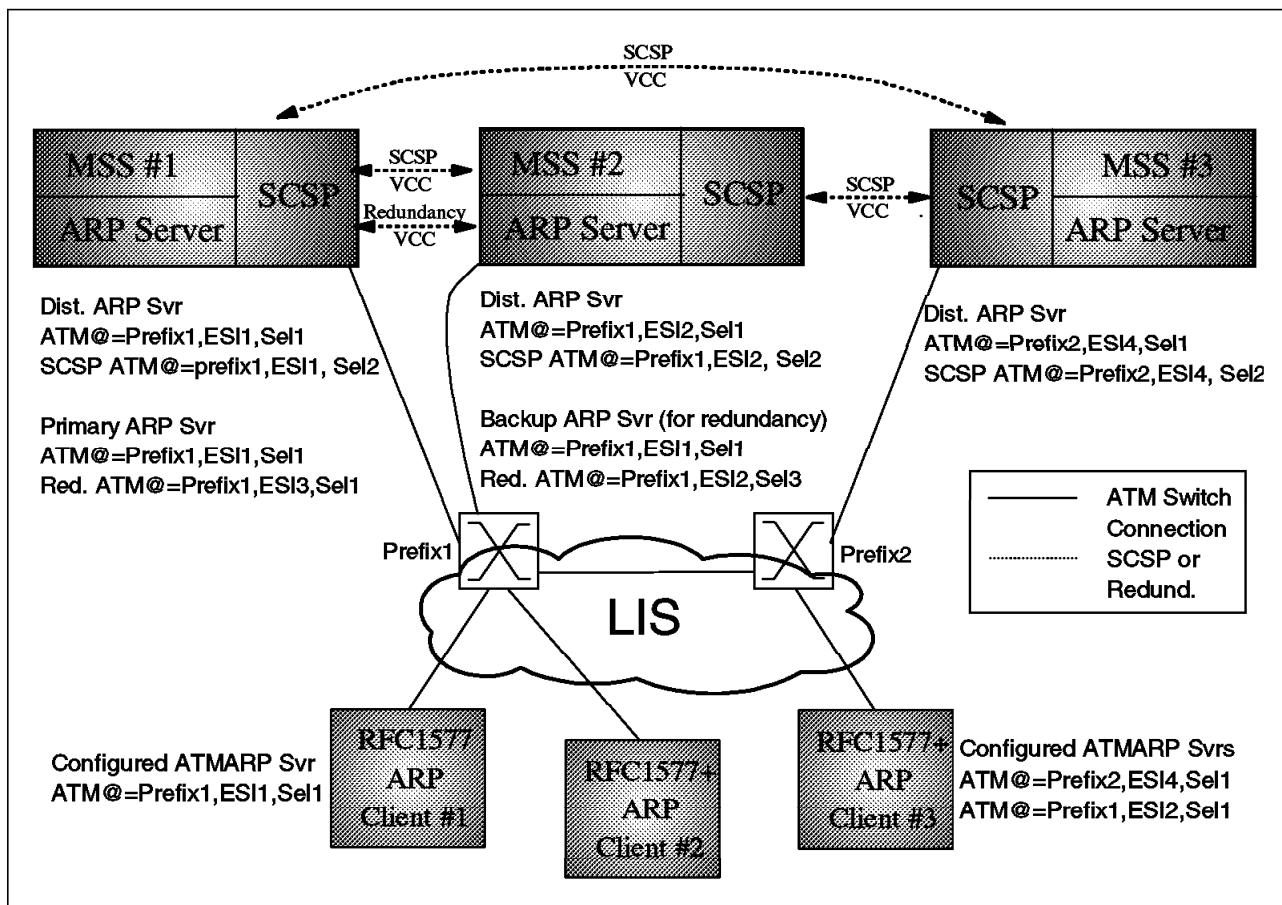


Figure 60. Example of a Server Group (SG) Providing ATMARP Server Redundancy

The redundancy mechanism is shown with the participating MSSs having the same ATM address (Prefix1, ESI1, Selector1) in addition to the ATM addresses

used for the RED (redundancy VCC) channel (Prefix1, ESI3, Selector1 in MSS #1; Prefix1, ESI2, Selector3 for MSS#2). If MSS #1 fails, the same steps occur as in MSS V1.1 ATMARP redundancy. If MSS #1 subsequently recovers, MSS #2 will deregister ATM address Prefix1, ESI1, Selector1 and MSS #1 will resume the ATMARP server function. This is also the same as the MSS 1.1 function.

13.2.1.4 Compatibility of ATMARP Functions between V1.1 and V2.0

MSS 2.0 is fully backward compatible with MSS V1.1 for ATMARP function. Thus in Figure 56 on page 108, the primary ATMARP server could be MSS 1.1 while the backup ATMARP server could be 2.0. However, in Figure 60 on page 115, all MSSs must be at V2.0, since each is a distributed ATMARP server running SCSP.

13.2.2 IP Multicast over ATM

Multicasting is an efficient method of sending information from one device to many devices simultaneously. The source sends a single copy of the IP datagrams, which are received by the destination devices. This is widely used in IP networks today for routing protocols (OSPF, RIP, BGP, EGP), DHCP, and audio/video distribution. MSS V2.0 supports IP Multicast over ATM as described below.

In shared media LANs, multicasts and broadcasts are implemented by sending packets to the MAC address FF.FF.FF.FF.FF.FF. However, IP multicast also allows for these multiple destinations to span subnets, which means that methods besides MAC layer broadcasting are needed. These methods are defined in RFC 1112.

With IP multicasting the source does not know who will receive the multicasts; it sends the packets to an IP group multicast address in the IP address space between 224.0.0.0 and 239.255.255.255. The IGMP protocol specified in RFC 1112 allows devices to dynamically join and leave multicast groups. IP routers and hosts may participate. IP routers use the Multicast Open Shortest Path First (MOSPF) protocol to advertise multicast groups across networks. IP multicasting assumes that the data link layer provides broadcast addressing and a connectionless transport.

ATM networks can provide connectionless transport and broadcasting when devices participate in LAN Emulation (LANE). However, for devices using Classical IP (RFC 1577/1755) there was no connectionless service or broadcast support until recently. IP Multicast over ATM (RFC 2022) provides this by defining these services for environments such as Classical IP.

RFC 2022 includes three basic services, all of which are included in MSS V2.0. The services, Multicast Address Resolution Server (MARS), MARS Client, and Multicast Server (MCS) are described below. MSS V2.0 supports all of these services.

The advantage of MARS is that it enables ATM devices to multicast efficiently. It allows real-time applications (such as video over IP) to operate with low delay, since a sender can directly use ATM and send data to multiple destinations simultaneously. Finally, MARS enables functions such as DHCP and routing protocols for ATMARP clients.

13.2.2.1 MARS Server and Client Overview

MARS is an extension to the ATMARP server, providing a mapping between an IP multicast address and the set of ATM addresses which are part of that multicast group. There is one MARS server for each LIS; thus MARS serves clients within a single LIS only. MSS V2.0 can provide MARS support for up to 32 LISs simultaneously (a MARS instance for each LIS).

The MARS client is an extension to the ATMARP client, providing the ability to query a MARS to retrieve the ATM addresses of members in a multicast group. The address of the MARS must be configured in the client. Since this standard is new to Classical IP, there are few devices implementing MARS client as of February, 1998.

A MARS client creates a point-to-point VC with a MARS to register itself as a member. The MARS will add the client to a point-to-multipoint VC known as the cluster control VC, which is used to distribute multicast group updates to all members. Clients then ask to join or leave specific multicast groups via messages to the MARS on the point-to-point VC (PtP VC). These updates are distributed to members via the cluster control VC (CCVC). Note that MARS does not get involved with the actual sending of multicasts. It only helps with the creation of the data path between source and destinations. The ATM MARS clients participating in an ATM level multicast and using the same MARS are known as a *MARS cluster*.

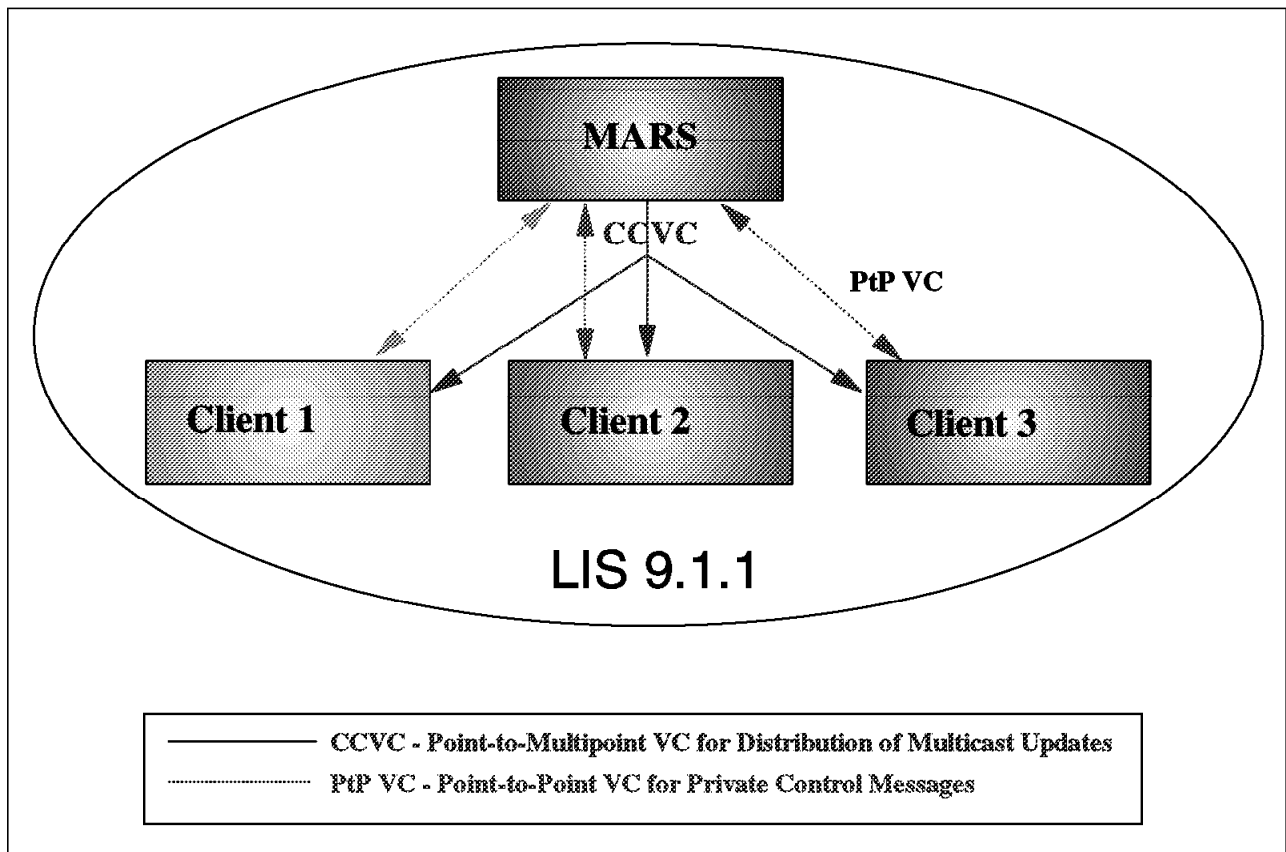


Figure 61. Example of a MARS Cluster and the Basic Connections

13.2.2.2 MARS Multicast Flow Options and the MCS

When a MARS client needs to transmit to a multicast group, it requests and receives from MARS the list of ATM addresses associated with the multicast group; this is known as the host map. The client will then request a point-to-multipoint VC via the ATM network which includes every MARS client in the multicast group. An example of this with client 1 and client 2 having established P2MP VCs is shown in Figure 62 on page 119. If every MARS client in a multicast group sends to the group, a full VC mesh would be achieved. This method of clients directly sending multicast data to the group has the following advantages:

- Optimal performance by utilizing the ATM switch fabric directly
- Low latency since cells arrive directly to the destinations, which is optimum for real-time applications such as video over IP.

However, there are considerations in using this method.

- More buffer overhead in the ATM switches due to more VCs
- More signaling overhead in the ATM switches due to more VCs
- More overhead when membership to a multicast group changes, since each client will need to modify its point-to-multipoint VC (PtMP VC)

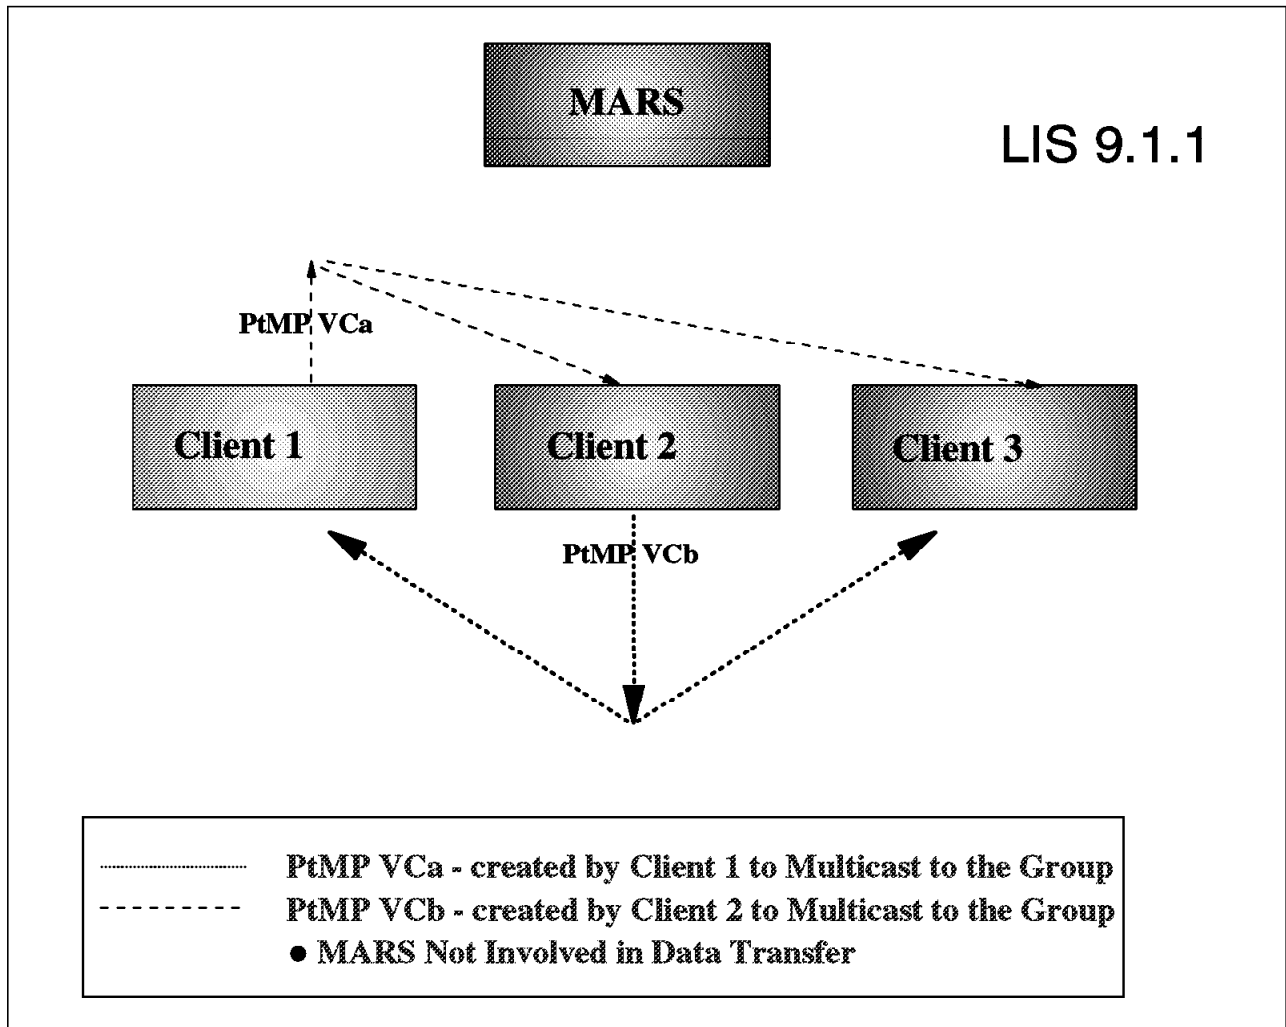


Figure 62. Example of a MARS Cluster with Three Clients. Clients 1 and 2 have PtMP connections to the multicast group.

There is another method of handling the point-to-multipoint (PtMP) connections for IP Multicasting. The MultiCast Server (MCS) is a proxy service that forwards multicast traffic on behalf of MARS clients. Instead of every client having their own point-to-multipoint VC, the MCS has a single PtMP VC for each multicast group. Clients send multicast traffic to the MCS via a point-to-point data VC, which the MCS then sends via the appropriate PtMP VC.

While the overhead is lower from an ATM perspective using MCS, performance will be lower, since the MCS is a focal point of traffic, must reassemble cells, and process and segment the IP multicast packets back into cells on a PtMP data VC. MCS processing includes inserting information in the packet so that a MARS client can ignore a multicast for which it was the source (since the sending client is also part of the PtMP data VC).

MARS allows for the benefits of both methods. VC meshes and MCS can be used together; each multicast group can be either served by MCS or VC meshes. A single MCS can serve multiple group addresses, but a group address can only be served by a single MCS. There can be multiple MCSs in a LIS, serving different multicast group addresses.

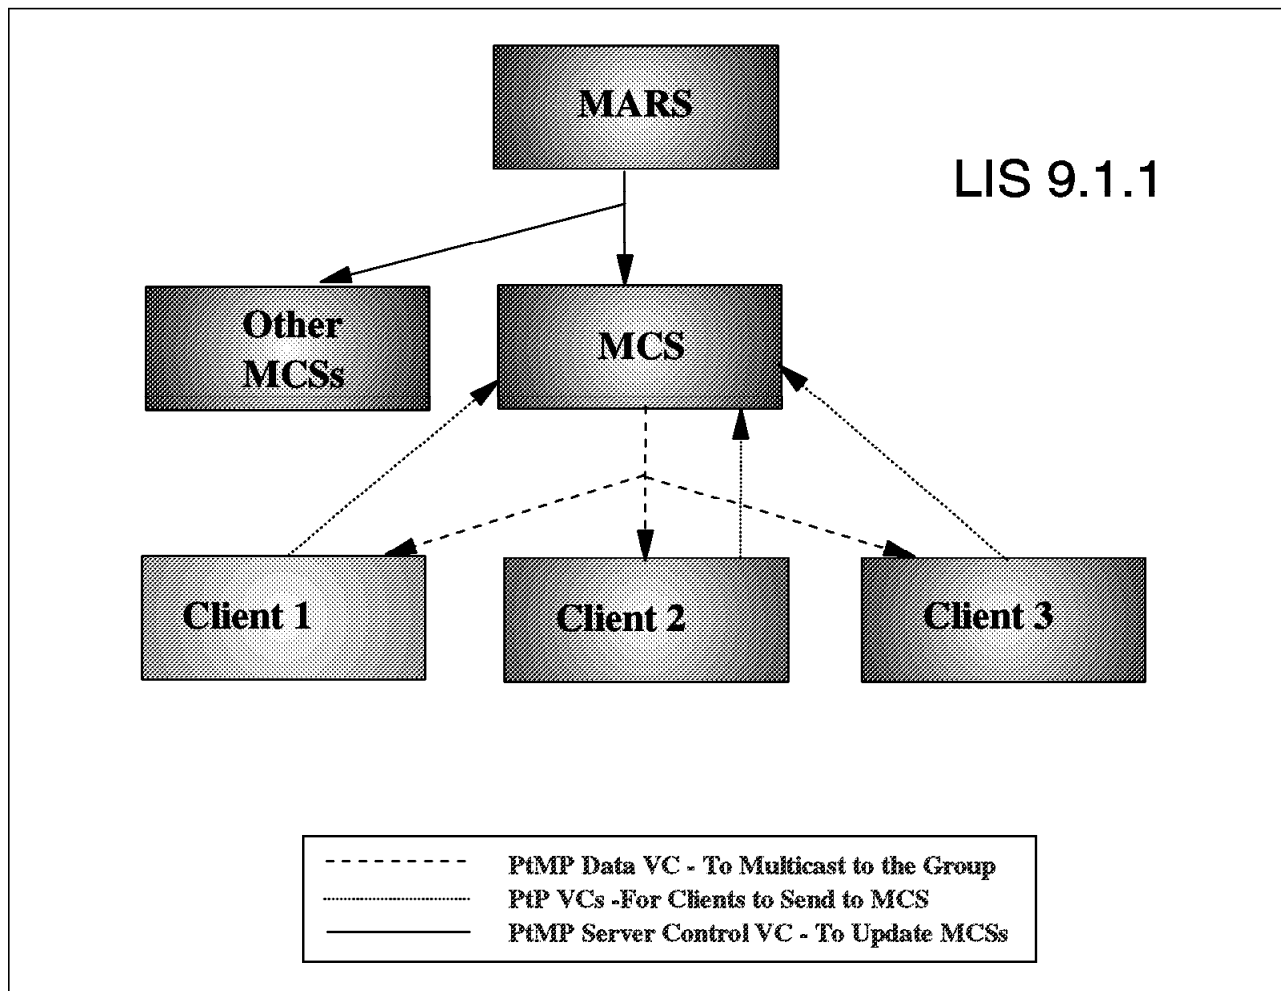


Figure 63. Example of a MARS Cluster with Clients Using MCS for a Single Multicast Group

Notice in the above diagram, that there is a VC between MARS and MCS called the server control VC (SCVC). This point-to-multipoint connection allows for MARS to distribute multicast group membership updates to MCS. Although in this example there is only one MCS, there could be multiple MCSs (supporting different multicast groups) in the same LIS. MARS would update all MCSs via the SCVC for group membership changes.

Using Figure 63, a summary of general operations follows:

- MARS clients register with MARS via PtP VCs.
- MARS adds the clients to the cluster control VC (CCVC).
- MARS adds the client ATM addresses to the host map for the appropriate groups, updates all clients via the CCVC, and updates all MCSs via the SCVC.
- When a client needs to send multicast traffic to a group, it queries the MARS via a PtP VC and receives either the list of ATM addresses of the MARS clients (if the group uses meshed VCs) or the address of the MCS.
- The client then sets up a PtMP data VC (for a meshed VC group) or a PtP data VC (to the MCS) to send multicasts for that group.

13.2.2.3 IP Multicasting across IP Networks

Along with the functions discussed, MSS V2.0 supports routing protocols used to allow IP multicasting across subnets and IP networks. These include Multicast OSPF (MOSPF) and the Distance Vector Multicast Routing Protocol (DVRMP). In a Classical IP network, MSS V2.0 can supply the services to support IP multicasting over ATM across IP subnets. Figure 64 shows an example of this, with ATMARP client #1 in subnet 9.1.1 (mask 255.255.255.0) sending to the multicast address 230.1.1.3. This multicast traffic is distributed to both subnet 9.1.1 and 9.1.2 via two PtMP VCs. The first is a mesh VC from client #1, and the second is a mesh VC from client #4 (MSS router ATMARP client on subnet 9.1.2).

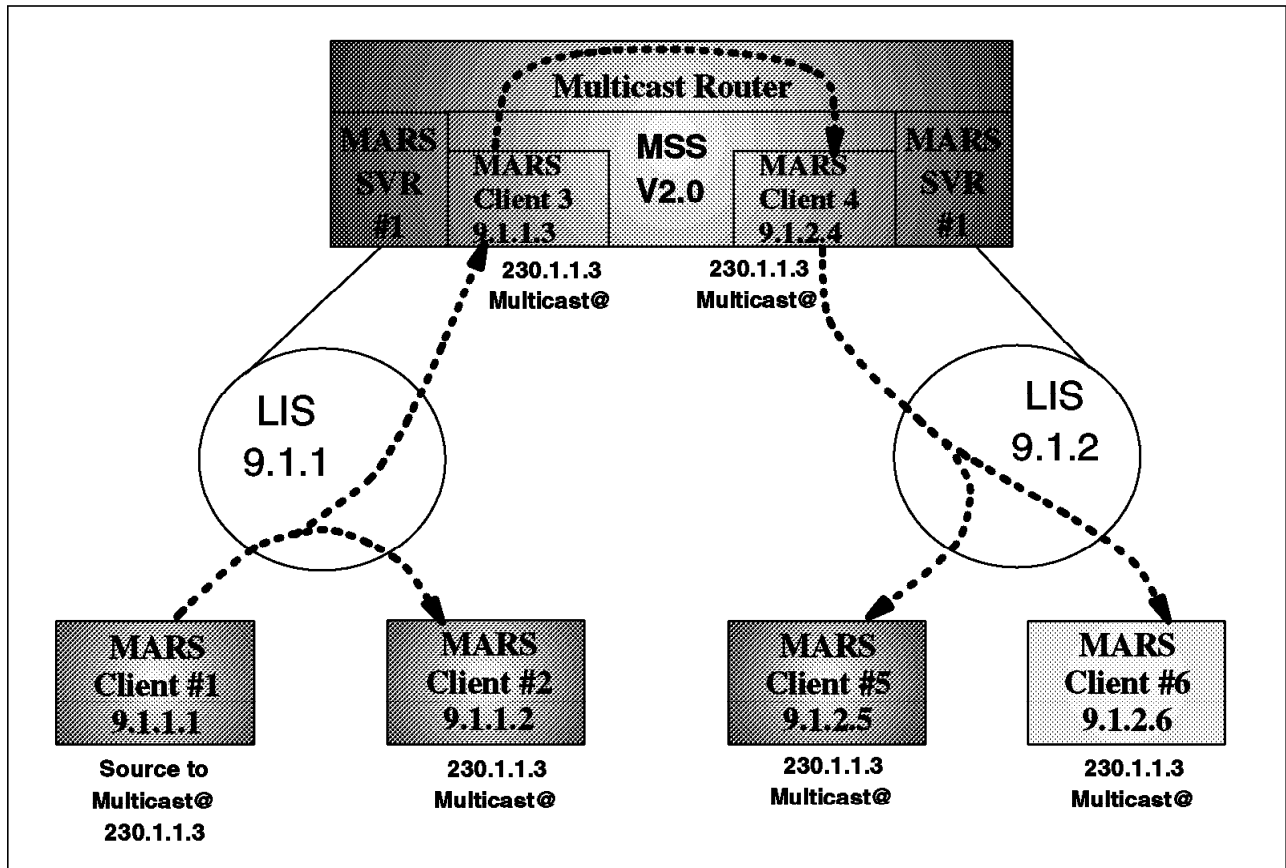


Figure 64. Example of an IP Multicast across Multiple Local IP Subnets Using MARS

13.2.2.4 Additional Considerations

MSS V2.0 allows configuration of MARS using the console or Web server. Configuration via the MSS configurator will be added later. In MSS V2.0, the MCS must be defined in the same MSS as the MARS.

Chapter 14. BUS Performance Enhancements

The potential for outstanding performance is one of the major benefits of ATM networks. Fast switching technology promotes high data transfer rates. Applications designed to run over legacy LANs access the ATM network via LAN Emulation. These applications can then utilize the large bandwidth of the ATM network. The data transfer speeds provided by LAN Emulation can be characterized by two performance measures: performance of data transfers over data direct VCCs and performance of data transfers over multicast send VCCs. Thus, BUS performance is one of the key measures for LAN Emulation and, by extension, ATM performance.

This section describes improvements to the MSS Server BUS that result in significant performance improvements.

LE clients send multicast frames and frames destined for unresolved LAN destinations over their multicast send VCC to the BUS. The performance of data transfers over multicast send VCCs is limited by the BUS's ability to process these frames.

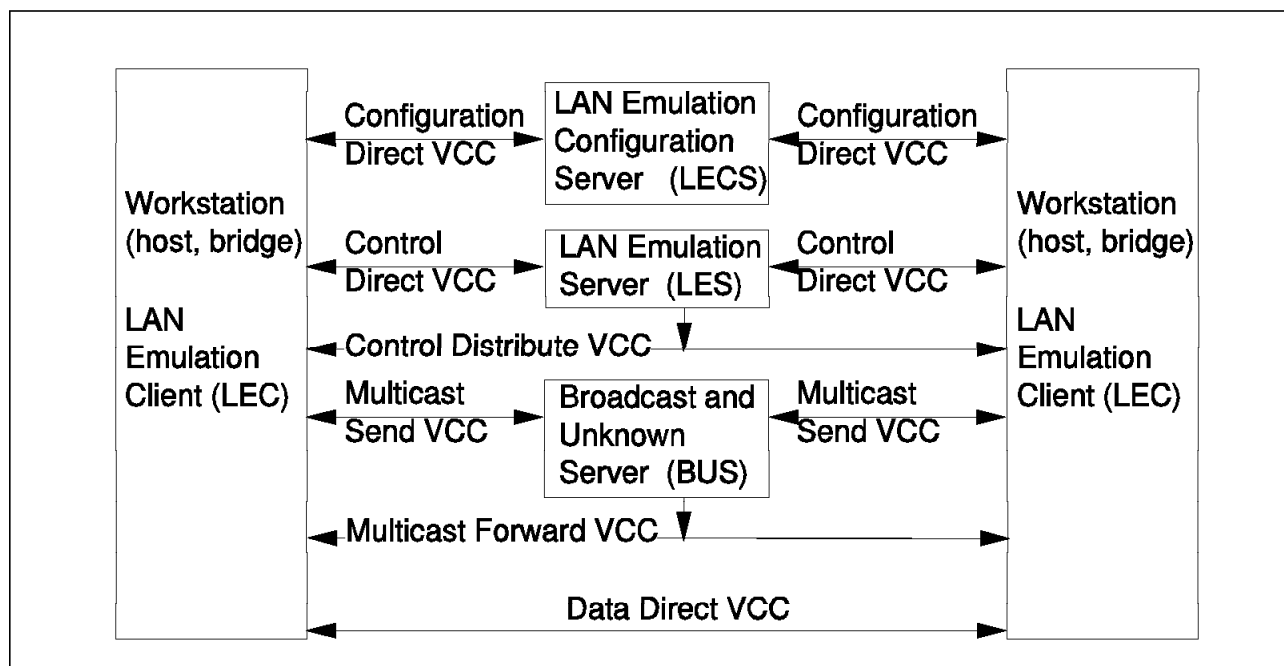


Figure 65. ATM Forum LAN Emulation VCCs

A major factor that influences the performance of the MSS Server BUS is its intelligent BUS feature. The MSS Server BUS may be configured to inspect incoming packets and direct them to either all LE clients or the LE clients configured as proxies. Additionally, BroadCast Manager (BCM) may be enabled to transform some broadcast frames into unicast frames. These functions decrease the utilization of the ATM network and result in fewer interruptions at user stations. In bridged networks, the effects are compounded because broadcast and unknown unicast frames are normally transmitted on many legacy segments.

In Release 1.1 of the MSS Server, options will be provided to trade BUS intelligence for BUS forwarding power. Although the BUS loses some function, its forwarding performance improves significantly. In addition to the current intelligent mode of operation, two additional BUS operational modes will be available to the user: adapter mode and VCC splicing mode. Adapter mode differs from the current BUS mode in that frames are not copied into system memory; instead, they are kept in the adapter's memory space.

Most intelligent BUS functions only require a few accesses to the frames and are still available, but BCM, which requires many accesses to the frame, is not available in adapter mode. VCC splicing mode is an enhancement that utilizes an ATM adapter feature that enables all frames received on one VCC to be immediately forwarded on another VCC. In VCC splicing mode, the forwarding function is implemented in hardware; the BUS and ATM device driver are not even notified that a frame was received. Thus, intelligent BUS functions and statistical information are unavailable in VCC splicing mode.

14.1 Benefits of the BUS Performance Enhancements

When operating in either of the two new modes, BUS forwarding performance improves significantly. The MSS Release 1.0 BUS yields a performance on the order of 67,000 64-octet pps. In either adapter or VCC splicing mode, performance in excess of 100,000 64-octet packets per second (pps) can be reached. Media speed is reached in VCC splicing mode when packets are longer than 1024 bytes, and in adapter mode when packets are 2048 bytes or larger.

VCC splicing offers one more advantage. Because no system bus or processor cycles are required its forwarding power scales linearly with the addition of a second ATM adapter. Thus, cumulative BUS forwarding rates in excess of 200,000 pps are possible when two ATM adapters are installed in an IBM 8210 MSS Server.

Additionally, due to increased processing speeds, the BUS buffer requirements are also decreased, freeing buffers for use by other system entities. In VCC splicing mode, the CPU is not interrupted to process BUS frames, and utilization of system bus bandwidth is reduced in both modes.

The BUS mode is configurable on a per ELAN basis.

14.2 Limitations of the BUS Performance Enhancements

As BUS forwarding capacity increases, BUS intelligence decreases. In each of the enhanced performance modes, a certain amount of intelligence is lost. When operating in adapter mode, BCM cannot operate. When operating in splicing mode, all intelligence is lost; the BUS merely serves as an agent to splice multicast send VCCs to the multicast forward VCC.

14.3 Details on BUS Modes

The following sections give some details on how the three BUS modes have been implemented.

14.3.1 System Mode

The BUS function introduced in MSS Release 1.0 operates as depicted in Figure 66. The ATM adapter performs segmentation and reassembly (SAR) functions and notifies the ATM device driver when a frame is received. The frame is copied from adapter memory into system memory and the BUS is notified of the frame's arrival. The BUS can then perform any of its intelligent functions on the copy of the frame held in system memory. When the BUS has finished processing the frame, it passes the frame to the ATM device driver. The device driver initiates a read operation to copy the frame back into adapter memory and, upon completion, signals the adapter to send the frame. This mode of BUS operation is referred to as system mode.




Figure 66. BUS - System Mode

14.3.2 Adapter Mode

Adapter mode eliminates the copy operation of the frame to and from system memory as depicted in Figure 67 on page 126. For the BUS to access the frame, it must access the memory on the ATM adapter. The need to read the whole frame into memory is largely avoided by using adapter mode.

When a frame is received by the BUS only a small amount of the frame is buffered into system memory. Since only a small amount of the frame is available to the BUS, BCM functions cannot be used in adapter mode. The other intelligent features of the BUS are still available. Such features include partitioning the unicast forwarding domain, BUS monitoring, and BUS statistics. For serviceability, packet tracing is allowed in adapter mode, but system

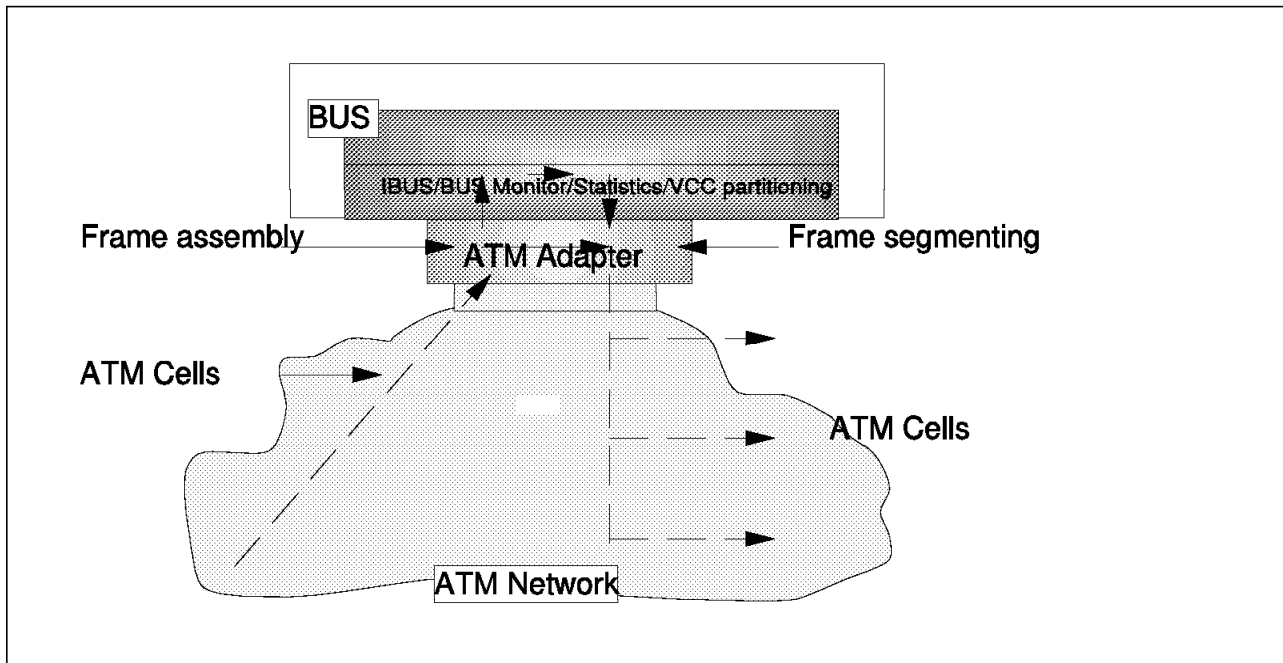


Figure 67. BUS - Adapter Mode

performance will be severely degraded if extensive tracing is done in adapter mode.

14.3.3 VCC Splicing Mode

VCC splicing mode builds on the adapter mode by completely removing the processor of the MSS Server from the data path. The ATM adapter performs SAR functions to build an incoming frame, and immediately transmits the frame on a pre-specified VCC. This is depicted in Figure 68.

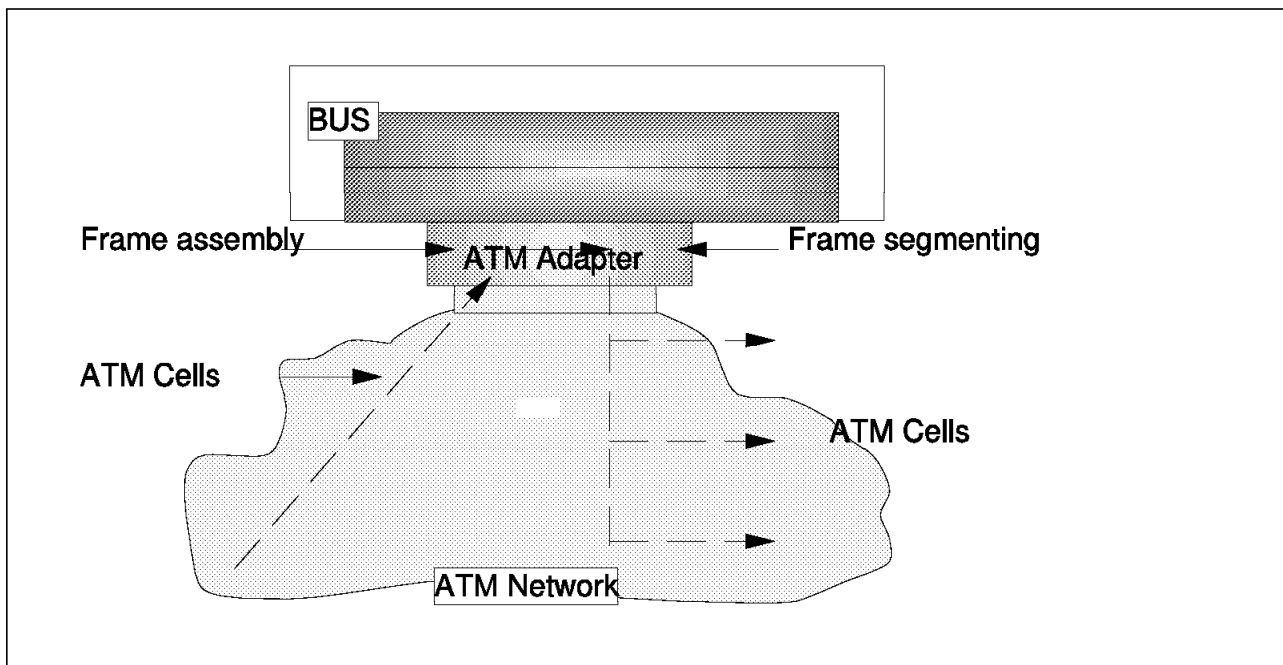


Figure 68. BUS - VCC Splicing

To perform VCC splicing, the BUS informs the ATM layer of the VCCs that should be spliced together. Additionally, all options for intelligent BUS processing are disabled; these are BCM, partitioning the unicast forwarding domain, and BUS monitor. For serviceability, packet tracing may be enabled in VCC splicing mode. However, if packet tracing is enabled, the BUS will dynamically switch to system mode operation. If packet tracing is subsequently disabled, the BUS will dynamically resume VCC splicing mode operation. Note when in this mode, the BUS and ATM interface statistical information does not reflect the data spliced between multicast VCCs.

Chapter 15. Miscellaneous Enhancements MSS 1.1

This section describes several additional enhancements included in MSS Server Release 1.1 that provide functional, performance, and capacity-related improvements.

15.1 Sharing of RFC 1483 PVCs

MSS Server Release 1.0 included two protocol entities that employed data encapsulation using the routed frame format specified in RFC 1483:

1. The native ATM support for IP (Classical IP)
2. The native ATM support for IPX routers

MSS Server Release 1.1 includes support for bridging using the bridged frame format specified in RFC 1483 (see Chapter 5, "RFC 1483 Bridging over ATM" on page 21).

In Release 1.0, the protocol entities employing RFC 1483 encapsulation could not share the same VCC. (That is, both IP and IPX frames could not be carried on a single VCC.) In Release 1.1, all of the protocol entities employing RFC 1483 encapsulation may share the same PVC. (That is, IP, IPX, and bridged format frames may all be multiplexed onto a single PVC.) PVC sharing is configured by using the same VPI/VCI when configuring PVCs at each of the protocol entities. As an example see Figure 12 on page 25.

Note that the traffic characteristics (for example, best-effort or reserved bandwidth, peak cell rate, etc.) of a shared PVC should be configured the same at each protocol entity. If the traffic characteristics are not the same, then the parameters of the first protocol entity to be initialized will be used for all protocols sharing the PVC.

The technique that enables multiple protocols to share an RFC 1483 VCC is called LLC encapsulation. Multiplexing via LLC encapsulation can be especially beneficial on connections between two routers or brouters (where a brouter is a device that can provide both bridging and routing services).

15.2 Multiple Internal LE Clients on the Same ELAN

In Release 1.0, the MSS Server supported only one internal LE client per ELAN. One reason for this is that the signalling implementation did not allow multiple leaves on the same point-to-multipoint VCC to be resident on the MSS Server. Since LE clients on the same ELAN need to be able to be added to the same control distribute and multicast forward VCCs, this restriction prohibited support for multiple internal LE clients that are members of the same ELAN. This restriction has been removed in Release 1.1, and the MSS Server is able to support multiple internal LE clients that are members of the same ELAN. This enhancement enables the MSS Server to function as a router for subnets that span multiple bridged ELAN segments (that is, subnets with members on multiple ELANs) by allowing both a bridging LE client and a routing LE client to be connected to the same ELAN segment as previously illustrated in Figure 24 on page 61.

15.2.1 ATM Interface VCC Binding

A feature called VCC binding has been added in Release 1.1 of the MSS Server. VCC binding improves performance when both ends of a point-to-point VCC reside on the same MSS Server ATM interface. Rather than transmit the frame to the ATM switch only to have it returned, VCC binding allows the frame to be transferred internally. In Release 1.1 of the MSS Server, the LE client has been modified to take advantage of VCC binding when its multicast send VCC is to a co-located LES/BUS and the BUS is not in VCC splicing mode. This is particularly useful in multiport bridging configurations where frames may have to be replicated on many ELANs. It should be noted that this application of VCC binding is limited to configurations where the BUS is in system mode. (That is, VCC binding is not used when the BUS is in adapter mode or VCC splicing mode.) For details on BUS modes, see Chapter 14, "BUS Performance Enhancements" on page 123.

15.3 Dynamic Allocation of LE_ARP Cache Entries

In Release 1.0, MSS Server LE clients pre-allocated the configured maximum number of LE_ARP cache entries. This was problematic because it was desirable to keep the maximum small in order to conserve storage, yet it was also desirable to have a relatively large maximum since the size of the LE_ARP cache limits the number of LAN destination mappings that can be maintained, as well as the number of data direct VCCs that may be established by an LE client. Thus, the size of an LE client's LE_ARP cache became an important performance/capacity-tuning parameter. (A cache that is too small can lead to address resolution and VCC thrashing, while a cache that is too large wastes memory and limits capacity.) In order to simplify this situation, LE clients no longer pre-allocate the configured maximum number of LE_ARP cache entries. Instead, in MSS Server Release 1.1, LE clients dynamically allocate LE_ARP cache entries as needed (in steps of 50 entries up to the configured maximum).

15.3.1 Recommendation

The LE_ARP parameter provides an upper boundary to the number of MAC addresses that can be cached per MSS internal LE client. Especially in a situation where remote proxy LE clients are used to connect to the MSS Server, the number of cached MAC addresses may be considerable as legacy LAN MAC addresses that are reachable via the proxy can be cached as well.

Note: In a token-ring environment legacy LAN MAC addresses are not cached. Instead the route descriptor (LAN segment number and bridge number) are cached. Because many stations may share the same route descriptor, it is less likely that the maximum will be reached.

Conforming to the LAN specifications, the current MSS release re-validates cached entries rather than, for example, deleting least recently used entries. As a result, new MAC destinations may not be accepted and connectivity problems might occur. Because memory is allocated based on the actual number of MAC addresses cached, it is advised to set the LE_ARP parameter to its maximum value (that is, 1024).

Note: A PTF has been scheduled for the MSS operational code to raise the default to 1024. At the time of writing it was unknown when this PTF would become available. It was also not known when this change would be propagated to the MSS configuration program.

If this maximum is not sufficient, redesign of your network (for example, using multiple LE clients and using smaller IP subnets) might be required. It can be expected that the next release of MSS will resolve this (potential) problem.

15.4 SRAM Size Increase

For historical reasons, the memory area where configuration records are stored is referred to as SRAM. In MSS Server Release 1.1, SRAM size has been increased to 256 KB. As a result of this change, the maximum number of LECS policy values and the maximum number of LES/BUS instances that can be configured have increased. The new upper bounds are provided in *MSS Release 1.1 Quick Guide*.

15.5 New ESI Selector LECS Assignment Policy

A new LECS assignment policy has been added in MSS Server Release 1.1: the ESI selector ELAN assignment policy. In Release 1.0, the MSS Server LECS supported an ATM address assignment policy. This policy was flexible in that it allowed variable-length values, but the variable-length value was limited to a prefix of the ATM address (that is, the policy was based on the first N bytes of the LE client's ATM address, where $1 \leq N \leq 20$), which prohibited the use of a policy that included only the ESI and the selector components of the LE client's ATM address. The addition of the ESI selector ELAN assignment policy rectifies this situation. The policy simplifies ELAN management when LE clients move from one ATM switch to another since the network prefix component of an LE client's ATM address is determined by the ATM switch that it is connected to, but moved LE clients retain their ESI and selector. The ESI-selector assignment policy also provides another significant benefit: ESI-selector policy value configuration records are 50% smaller than ATM address policy value configuration records. The smaller size allows the MSS Server LECS to support twice as many LE clients when ESI selector policies are used instead of ATM address policies. This is especially significant because the ATM address is the only ELAN assignment policy criteria that is guaranteed to be present in the LE_CONFIGURE_REQUEST.

15.5.1 Maximum Size of BCM's IPX Forwarding List

BCM IPX employs a cutoff mechanism designed to protect the MSS Server from performance degradation. In MSS Server Release 1.0, BCM IPX maintained at most 50 members in its forwarding list for a given ELAN. In MSS Release 1.1, the maximum size of the BCM IPX's forwarding list is now configurable on an ELAN basis with a range from 5 to 65,535, and a default of 50.

BCM's IPX forwarding list consists of all the dynamically discovered IPX routers or servers in the IPX network for the ELAN, plus any BCM IPX static targets currently active on the ELAN. For each broadcast frame managed by BCM IPX, the frame is individually transmitted to each of the members on the forwarding list. If the number of members in the forwarding list ever exceeds the configured maximum, BCM IPX automatically shuts itself down, allowing the BUS to forward IPX broadcasts unmanaged. Thus, this configurable parameter establishes a protective upper limit on the impact that the BCM IPX frame multiplying effect can have on MSS Server performance. Increasing the parameter from its default of 50 should be considered only when the MSS Server has excess capacity

available to handle this frame multiplying effect, and when broadcasts need to be kept to an absolute minimum at the expense of MSS Server performance.

15.5.2 Persistent SET Support for LE Service MIBs

In Release 1.0 of the MSS Server, SNMP SET requests to the LAN Emulation service MIBs were not saved in SRAM or non-volatile storage. Thus, while changes requested via SNMP SETs were applied to operational LE service components, the changes did not persist across a reboot.

In Release 1.1 of the MSS Server, configuration changes requested via SNMP SETs to the LE emulation service MIBs are saved in both SRAM and non-volatile storage thus backing up the SRAM. (The changes are saved to the non-volatile storage media that the box was booted from, either hard file or flash.) This means that changes made by the Nways Campus Manager product are not lost due to reboots, power outages, etc.

15.6 SR-TB Translational Bridging Support For IP

In SR-TB mode, both source route and transparent bridges are active, and frames are translated between the two bridges. SR-TB translation was supported for SNA and NetBIOS protocols in Release 1.0 of the MSS Server. SR-TB translation support for IP has been added in Release 1.1 of the MSS Server.

15.7 Association of Voice Files with Load Image

In MSS Server Release 1.0, the files used for the voice response functions are stored on the hard file and are not associated with a particular load image. In Release 1.1, the voice files are integrated into the load images, which simplifies upgrades. (That is, new voice files can now be easily associated with new load images.)

Chapter 16. Miscellaneous Enhancements MSS 2.0

This section describes several additional enhancements included in MSS Version 2.0 that provide operational and configuration improvements.

16.1 Dynamic Reconfiguration

Dynamic Reconfiguration (DR) allows for changes to the MSS Server without restarting it. DR is important for providing high availability, where a reboot of MSS during production is not desirable. In MSS 1.1, most parameters can not be changed (and the new changes become effective) without restarting MSS.

An important example of this is the ability to dynamically add bridging or routing interfaces. If the MSS was configured with spare interfaces, routing and/or bridging can be defined and immediately used on those interfaces. So by configuring spare interfaces an MSS can have new routing and bridging ports added dynamically.

A spare interface is a set of control blocks reserved for dynamic reconfiguration, and is easily defined via the MSS Version 2.0 configuration program, in the Navigation panel at NW;Network Device/System/General/ in the field labeled Spare Interfaces.

DR will be implemented in phases; this first phase includes some restrictions:

- A spare interface can not be activated if any interface has been deleted since the last reboot.
- A spare interface cannot be defined with a protocol or feature that is not currently active in the MSS.
- BGP and MARS can be configured dynamically (t 6) on a spare interface, but not dynamically activated.
- A spare interface can be configured but not dynamically activated if it exceeds the header/trailer size of all existing interfaces. An example is if an MSS has only Ethernet LEC interfaces, a token-ring LEC cannot be dynamically configured and activated.
- APPN can not be dynamically configured on a spare interface. However, APPN can be dynamically configured on a currently active interface (for example, one that is currently defined for IP routing).
- For dynamic reconfiguration IPX, IPX must be globally enabled and defined on at least one interface. IPX static routes, static services, and interface filters can be dynamically configured, but not dynamically activated. (Similar to BGP/MARS, it requires a reboot.)
- NetBIOS filters can be configured on spare interfaces, but not dynamically activated (similar to BGP/MARS).
- Bridging personality can not be changed by activating a spare interface. For example, if an MSS is configured for bridging and is performing transparent bridging on Ethernet LECs, a token-ring spare interface can not be dynamically configured and activated for bridging. This would require the bridge to now support SR-TB.

16.2 Dynamic Linking and Loading

Dynamic linking and loading allows selective loading of functions from operational code image into MSS memory. If a function is not configured, it is not loaded into memory. In Version 2.0, only the APPN function can be dynamically loaded.

In previous releases, MSS software was available as a single software package with a single load module. In MSS Version 2.0, the software package (feature #8707) consists of multiple load modules for MSS. The entire software package must be loaded into a bank (PCMCIA disk or flash) before it can be loaded into MSS memory. To update the software package, the entire software package must be reloaded into a bank.

16.3 Time Activated Reboot

The MSS can be automatically rebooted at a specified date and time to any bank and configuration. This is accomplished via the MSS console (either via telnet or serial port). Below is an example.

```
MSS NCC *t 6

MSS NCC Config>boot
Boot configuration
MSS NCC Boot config>timedload activate
```

Bank	Description	Date
BankA	IMAGE - AVAIL	01 Jan 1970 00:20
	CONFIG 1 - AVAIL	27 Jan 1970 23:52
	CONFIG 2 - AVAIL L osclab3 latest, no brg	01 Jan 1970 00:21
	CONFIG 3 - AVAIL *L mss 10/14 primary	02 Jan 1970 02:57
	CONFIG 4 - AVAIL	
BankB	IMAGE - ACTIVE	
	CONFIG 1 - ACTIVE *L v2 1st primary 12/4	01 Jan 1997 01:02
	CONFIG 2 - AVAIL osclab3 2nd latest	01 Jan 1970 00:41
	CONFIG 3 - AVAIL	27 Jan 1970 23:57
	CONFIG 4 - AVAIL	01 Jan 1970 03:04
BankF	IMAGE - AVAIL	02 Jan 1970 03:10
	CONFIG 1 - AVAIL	01 Jan 1970 00:26
	CONFIG 2 - AVAIL L osclab3 latest, no brg	02 Jan 1970 02:50
	CONFIG 3 - AVAIL *L mss 10/14 primary	02 Jan 1970 03:36
	CONFIG 4 - AVAIL v2 1st primary 12/4	01 Jan 1997 01:01

```
* - Last Used Config      L - Config File is Locked

Time Activated Load Processing...

Select the bank to use: (A, B, F): (A) b
Do you want to put load modules into the bank? (Yes, No, Quit): (Yes) no

Do you want to put a configuration into the bank? (Yes, No, Quit): (Yes) no

Select the configuration to use: (1, 2, 3, 4): (1) 2

Time of day to load the router (YYYYMMDDHHMM): ()? 199812312359
The load timer has been activated.
MSS NCC Boot config>
```


To turn off Timed Activated Reboot, the following console command is used.

```
MSS NCC Boot config>timedload deactivate  
Deactivate Load Timer Processing...  
  
Do you want to deactivate the load timer? (Yes, No, Quit): (No) yes  
The load timer has been deactivated.
```

Chapter 17. Configuration Examples MSS Version 1.1

After being physically installed, the MSS Server must be configured to meet the individual requirements of your network. The configuration examples included in this section will help you to simplify the process of getting your MSS Server operational. The purpose of the configuration examples is to get you started, and show you how to use the MSS Release 1.1 performance improvement and redundancy features.

MSS Release 1.1

All configuration examples assume MSS Release 1.1 and the use of the configuration program. Make sure that your MSS Server and MSS Configuration Program are at the right software level.

The configuration examples are ready for use, although, depending on your configuration, some additional configuration may be required. They can be downloaded (see <ftp://www.redbooks.ibm.com/redbooks/SG242115/49xx.cdb>) from the Internet. All configuration examples are kept in a single configuration database, which can be loaded into the MSS Configuration Program and be transferred to the MSS using the communication option.

The following configuration examples have been included:

- Basic Connectivity (see 17.3, "Example - Basic Connectivity" on page 149)
A basic configuration example to get the MSS operational. Provides services for three different network types, including IP routing between them. The example allows you to build up a network, including ATM edge devices or end systems, without additional customization.
- Shortcut Bridging (SCB) (see 17.4, "Example - Shortcut Bridging" on page 154)
MSS Server performance should be considered as a critical point for large networks. SCB allows you to distribute LES/BUS functions over several MSS Servers, resulting in SuperVLANs that scale up much further than ELANs controlled by a single MSS Server. This configuration example gets you started when you want to use shortcut bridging.
- IP Cut-Through (see 17.5, "Example - IP Cut-Through" on page 164)
IP cut-through provides broadcast reduction for large IP networks while avoiding extra routing hops. This configuration example gets you started when you want to use IP cut-through.
- Backward Connectivity (see 17.6, "Example - Backward Connectivity" on page 171)
This example shows how to interconnect ATM networks to existing networks. It shows support for routing to an FDDI ring and bridging using RFC 1483 encapsulation.
- Redundancy (see 17.7, "Example - Redundancy" on page 178)
Mission critical applications require the elimination of single points of failure. This example shows you how to provide default IP gateway backup when using Classical IP and LAN emulation, LECS, and LES/BUS functions.

Depending on the configuration example you need one or two MSS Servers. Two MSS Servers are required for the shortcut bridging and the redundancy example, but the other examples require a single MSS Server.

Prior to discussing each example, we describe the steps that are required to get the MSS from its shipping box to being a running part of your network.

17.1 Preparation

The steps required to use the configuration examples in this publication are:

1 Prepare your ATM switch.

Note: All commands entered assume an IBM 8260 ATM switch, but other ATM switches can be used as well.

To connect an IBM 8210 MSS Server you need an ATM switch that provides a 155 Mbps interface with UNI 3.0 or UNI 3.1.

To use the IBM 8260 ATM switching capabilities make sure an ATM control point and switch module (A_CPSW) has been installed. For the 8 MB A_CPSW it is recommended to use code level 2.1.0 or higher. For the 16 MB A_CPSW use code level 2.5.1 or higher.

Verify and, if necessary, upgrade the microcode on your ATM switch. Obtain your current code level (see Flash EEPROM version) with:

```
8260ATM>show device
```

The A_CPSW operational code can be downloaded from the Internet at <http://www.networking.ibm.com/826/826fix.html>.

2 Prepare you MSS Server.

Make sure your MSS Server is running MSS Release 1.1. The operational code for the stand-alone IBM 8210 can be obtained from <http://www.networking.ibm.com/nes/nesswitc.htm#8210v11>. The operational code for the MSS Server module can be obtained from <http://www.networking.ibm.com/826/826fix.html>.

3 Connect the MSS Server to the ATM network.

For details see 17.1.1, "Connecting MSS Server to the ATM Network" on page 139.

4 Configure your ATM switches to enable dynamic access to the LECS.

For details see 17.1.2, "Allow Dynamic Access to the LECS" on page 140.

Note: This step must be repeated for every ATM switch in your network.

5 Prepare your configuration workstation.

To activate, and possibly change the configuration examples, you need a configuration station running MSS Configuration Program Release 1.1. The configuration program package is available for download from the Internet at <http://www.networking.ibm.com/nes/nesswitc.htm#8210v11>. It consists of a stand-alone software package for AIX, Windows or OS/2.

After the configuration program has been upgraded retrieve the configuration examples at <ftp://www.redbooks.ibm.com/redbooks/SG242115/49xx.cdb>. The file provided at that location is a single configuration program database that contains each of the configuration examples. The examples are named according to the headings in this publication. We recommend storing the configuration database in the same directory as the configuration program.

From the configuration station you need a serial connection to the MSS Server. As an alternative, for customers that have a running ATM network, existing IP connectivity between the configuration station and the MSS Server can be used.

6 Establish IP connectivity between the MSS Server and your configuration station.

For details see 17.1.3, “Establishing IP Connectivity” on page 141.

7 Configure proper SNMP access on the MSS Server.

In addition to the IP connectivity the configuration program must be authorized to access the MSS Server using SNMP.

For details see 17.1.4, “Configuring SNMP” on page 143.

8 Select, update, and activate a configuration example.

Follow the instructions in 17.2, “Selecting and Transferring a Sample Configuration” on page 144.

17.1.1 Connecting MSS Server to the ATM Network

The procedure to connect a stand-alone IBM 8210 MSS Server to your ATM network is slightly different from installing an MSS Server module for the IBM 8260. We outline both:

- IBM 8210 MSS Server

Attach the fiber pair to port 1 (lower slot) of the IBM 8210. Connect the other end of the fiber pair to the ATM switch. Enable UNI on the ATM switch for the port to which the MSS Server connects, using:

```
8260ATM>set port 3.2 enable uni
```

Power on the MSS Server. Verify proper wiring and RX-TX matching by checking the link state. Look at the LEDs beneath the port or enter:

```
8260ATM>show port 3.2
```

Note: Replace port 3.2 with the port number you are using.

- MSS Server module for IBM 8260

The IBM 8260 might be powered down. Insert the MSS module and power up the IBM 8260. Enable UNI by entering the following command at the ATM switch:

```
8260ATM>set port 13.1 enable uni
```

Verify proper attachment by entering:

```
8260ATM>show port 13.1
```

Note: Replace port 13.1 with the port number you are using.

17.1.2 Allow Dynamic Access to the LECS

LE clients use a number of ways to connect to their LAN emulation configuration server (LECS). The order in which connectivity is tried, is:

1. Using a hard-coded LECS address
2. Using ILMI (interim local management interface)
3. Using the LECS WKA (well-known address)
4. Using the well-known PVC

Using ILMI and/or the LECS WKA requires configuration effort on *every* ATM switch to which LE clients connect.

For ILMI access the ATM address of the primary and backup, if any, LECSes must be defined. LE clients connect to the LECS in the order in which you define their addresses. An LE client will only connect to a backup LECS if it is unable to connect to the primary LECS.

LE clients use the WKA if they are unable to connect to the LECS using ILMI. To enable the use of the LECS WKA only a single LECS ATM address can be specified. If defined, the ATM switch maps the LECS WKA address to this pre-configured LECS address.

The LECS ATM address(es) are 20-bytes long and a concatenation of:

1. 13-byte network prefix
2. 6-byte ESI (end system identifier)
3. 1-byte SEL (selector)

The ESIs used in the configuration examples are 40.00.82.10.00.00 for the primary (or only) LECS, and 40.00.82.10.00.02 for the backup LECS. The SEL is always 00.

The network prefix must be obtained from the ATM switch to which your MSS Server connects. The prefix can be found with the following 8260 CPSW command:

```
8260ATM>show device
```

Among other things, your display will show something similar to:

A-CPSW

ATM address: **39.09.85.11.11.11.11.11.11.11.01.01.40.00.00.82.60.A1.00**

The first 13 bytes of this address, 39.09.85.11.11.11.11.11.11.11.01.01. in our example, shows the ATM network prefix for the ATM switch on which the command was entered.

Obtain the network prefix for your primary LECS and your backup LECS. They will be different if your MSS Servers connect to different ATM switches:

Assuming your network prefixes are:

Primary LECS: **pp.pp.pp.pp.pp.pp.pp.pp.pp.pp.pp.pp.**

Backup LECS: **bb.bb.bb.bb.bb.bb.bb.bb.bb.bb.bb.bb.bb.**

Define the ILMI LECS addresses for the primary and backup LECS using:

Note: The second command may be omitted if you do not have a backup LECS.

For the LE clients that do not support learning the LECS via ILMI you need to define the WKA-to-real-LECS mapping. Use the following command:

Verify the correct setting and save the changes by entering:

Entries 2 and 3 will be used if the attached devices retrieve the LECS address(es) using ILMI. LECS redundancy is only possible if LE clients use ILMI. The order in which you define the addresses defines the order in which LECSes are contacted.

If LE clients try to establish a connection to the LECS WKA address, the ATM address (WKA active) specified in the first entry will be substituted. No LECS redundancy is available when using the WKA address.

Note: Most LE clients will only use the WKA if connection establishment using ILMI-learned addresses fails.

17.1.3 Establishing IP Connectivity

In the remainder of this section we assume that the communication option of the configuration program will be used to activate the configuration examples. The configuration option requires IP connectivity between your configuration station and the MSS Server. This can be done either out-of-band or in-band. Out-of-band IP access requires a point-to-point SLIP connection. In-band IP access uses ATM connectivity.

The SLIP connection will in general be used to activate the first configuration while the in-band (ATM) connectivity will be used for further configuration. In-band connectivity can be realized using the example discussed in 17.3, “Example - Basic Connectivity” on page 149.

17.1.3.1 Using SLIP

To enable the use of a SLIP connection attach your configuration station to the MSS Server via either the EIA 232 service port (remote, or local via null modem cable) or the PCMCIA modem. Note, that only one of these connections can be active at a time.

The default connection parameters are: 19.2 kbps, no parity, 8 data bits, 1 stop bit. The default IP addresses are 10.1.1.2 for the MSS Server and 10.1.1.3 for the configuration node station.

If you have OS/2 running on the configuration node, the SLIP connection is established by entering:

```
mode com1: 19200,n,8
start /min slip -com1 -ifconfig 10.1.1.3 10.1.1.2 -mtu 1500 -speed 19200
```

17.1.3.2 In-Band Connection

In-band IP connections might be used for MSS Servers that are already operational. After loading the basic connectivity example, see 17.3, "Example - Basic Connectivity" on page 149, where the following interfaces for in-band management are defined:

- Ethernet (Emulated LAN) at IP address 172.16.0.1
- ATM (Classical IP client) at IP address 172.17.0.1
- Token-ring (Emulated LAN) at IP address 172.18.0.1

Any of these interfaces can be used for in-band management purposes.

17.1.3.3 Using XModem

If neither SLIP nor in-band configuration is possible, a serial ASCII terminal connection can be used as an alternative. Either the EIA 232 service port or the PCMCIA modem can be used to connect your terminal (emulator).

When using a terminal emulator program, XModem file transfer can be used to retrieve a communication example from the configuration station. Make sure that the communication example you want to load is available on the terminal emulator workstation in *exported* form. See 17.2.4.1, "Export a Configuration" on page 146 for how to produce an exported file.

If for any reason you are unable to use the configuration program (and use IP or XModem to transfer configuration files), you have to configure the MSS Server using the Web browser facilities or the command line interface. For details, see *Understanding and Using the MSS Server, SG24-4915*.

Use the Configuration Program

Due to its ease-of-use, its parameter consistency checking, its help files, and its configuration file management we recommend the use of the Configuration Program.

17.1.4 Configuring SNMP

The configuration program provides a communication option (see 17.2.4, “Downloading a Configuration” on page 146) to load configuration files into your MSS Servers from your configuration station. The communication option requires IP connectivity to the MSS Server and SNMP write access.

Granting SNMP write access on the MSS Server requires that you define the following on the MSS Server:

1. A community name

The configuration station should use this community name when activating a configuration file.

Note: In our examples we use community name admin.

2. Privileges of the community

The community should at least have write access. To enable retrieval of a configuration we have added read access as well.

3. IP host(s) that are part of the community

For example, the IP address used by the configuration program when communicating with the MSS Server and your network management stations.

The SNMP configuration should be entered using the command line interface or using the Web browser facilities.

Assuming the line command is used, Telnet to the MSS Server using any of its IP addresses. In case of a SLIP connection use 10.1.1.2.

```
telnet 10.1.1.2
```

Log in to the MSS (note that the initial configuration does not require authorization), enter SNMP configuration, define the community name admin, and give read-write access to the configuration station at 10.1.1.3 (the SLIP address). To do so enter the following commands.

```
*talk 5
+protocol SNMP
SNMP>add community admin
Community added successfully
SNMP Config>set community access write_read_trap admin
Access set successfully
SNMP>add address admin
IP Address [0.0.0.0]? 10.1.1.3
IP Mask [255.255.255.255]? 255.255.255.0
Address added successfully
SNMP>exit
+CTRL-P
*logout
```

Notes:

1. For the IP mask we entered 255.255.255.0 to avoid traps being sent to the configuration station.
2. It is not necessary to reload the MSS Server after adding the SNMP community.

17.2 Selecting and Transferring a Sample Configuration

The configuration program is key to the management of the your MSS Server configuration files. From a single workstation you are able to generate configuration files, download them to the MSS Server, or retrieve active configurations. Reloading your MSS Servers can be time-initiated or immediate. The following sections briefly describe the actions that are required to get started with the configuration examples.

17.2.1 Terminology

When using the configuration examples provided some user-specific customization is required. When discussing the recommended customization we refer to menus of the configuration program. We do that in a way that guides you from the configuration program Navigation Window (NW) to any of the detailed configuration menus. For example if a parameter has to be changed in a menu that can be found in the Router/System subtree of the Navigation Window, and the menu opens by clicking on General, and the parameter to be changed is called Location, we will call this:

NW;Router/System/General;Location.

Figure 69 depicts the configuration screen referred to.

System General

System name
MSS_A

Location
ITSO_Raleigh

Contact
Heiner Heng

☐ System console

Max packet buffers 0 Packet size 0

Inactivity timer 0

Help URL

Figure 69. Terminology

For the IP addresses and subnet masks we will use the following notation: 172.16.0.1/16. The first four bytes denote the IP address. The number behind the dash shows the number of bits set to 1 in the subnet mask. For example:

172.16.0.1/16 - IP address 172.16.0.1, subnet mask 255.255.0.0

172.16.1.1/24 - IP address 172.16.1.1, subnet mask 255.255.255.0

17.2.2 Opening a Configuration

Start the configuration program by clicking on the **MSS Release 1.1** icon or by entering `cfg` at the command line, whichever is appropriate for your system. After some greeting remarks, the Navigation Window (NW) will appear with an overview of your configuration options.

Enter menu `NW;Configure/Open_Configuration...` to select the database file downloaded from <ftp://www.redbooks.ibm.com/redbooks/SG242115/49xx.cdb>. The examples discussed in the following sections will appear in the Configurations window. Select one by clicking on it. Review the contents by clicking on **Configuration Summary**. Clicking on **OK** will bring up the opening confirmation dialog.

Note: Pressing F1 at any time will raise a help window with advice on the topic you are currently working on.

17.2.3 Optional Customization

The configuration examples are designed to run with minor customization for most installations. The basic connectivity example is especially designed for immediate use. However, you might want to change the system's name, location and contact. This can be done by using the `NW;Router/System` subtree.

Customization is mandatory for the examples that involve more than one MSS Server as these contain ATM addresses that are installation-dependent. The required customization steps are included in the description of these examples.

17.2.3.1 Changing the System Information

The `NW;Router/System/General` menu shows some options for the MSS Server. It is recommended you enter the information that is specific to your environment. Although the values have no operational consequences they might simplify the management of the MSS Server. Be careful with activating the System console button as this enables console-login and requires you to define an administrative user (see 17.2.3.3, "Adding Users" on page 146) before the console can be used.

17.2.3.2 Changing the Time of Day

MSS is able to retrieve the current time from a time server using the protocol specified in RFC 868.

Note: This differs from the network time protocol (NTP) and simple network time protocol (SNTP) specified in RFC 1305 and RFC 1769, respectively.

If you do not know if you have an RFC 868 time server set the polling interval to 0 (that is, no polling). Otherwise, specify the time server, time zone and polling interval in the `NW;Router/System/Time_of_Day` menu.

17.2.3.3 Adding Users

In our configuration examples we have defined an administrative user `admin` using password `ibm8210`. You may wish to change the password or to define additional users. This is done using the `NW;Router/System/Add_User` menu. Different levels of user privileges are available. Monitor access allows the user to display system information. Operational access allows the user to display system information and to enter temporary configuration changes, but does not allow making permanent configuration changes. Administrative users have access to all command line facilities.

17.2.3.4 Saving Current Configuration

For those installations that already have an active MSS Server we recommend you save the current configuration before activating a new configuration file. Retrieve your current router configuration first and store it in a backup database file, for example, `backup.cdb`.

To retrieve the configuration enter the `Configure/Communications/Single_Router` menu. Select **retrieve configuration** and enter one of the IP addresses of the MSS Server (for example, its SLIP address `10.1.1.2`). Save the configuration file using the **Configure/Save_configuration_as...** menu. Enter `backup` as your database file name and click on **OK**.

17.2.3.5 Saving Changed Configuration

We recommend you save the example configuration files that have been changed into a separate database file, for example, `customized.cdb`. This can be done by entering the `Configure/Save_configuration_as..` menu. Enter `customized` as your database file name and click on **OK**. For subsequent open and save operations, this database will appear in the selection window for the databases.

17.2.4 Downloading a Configuration

After the necessary changes have been applied the configuration file can be downloaded to your MSS Server. This function can be accessed from the **NW;Configure/Communications/Single_Router** menu. Enter one of the IP addresses of the MSS Server. We have preconfigured MSS Server's SLIP address as `10.1.12`.

The SNMP community is set to `admin`. This community name allows you to transfer new configuration files, as discussed in 17.1.4, "Configuring SNMP" on page 143. Enter the `Send_Configuration_to_router` button to send a configuration. To activate the configuration file select the **Restart_router** button as well. Clicking on **OK** will start the actual send process.

17.2.4.1 Export a Configuration

As discussed in 17.1.3.3, "Using XModem" on page 142, an alternative to using the configuration option is to export a configuration example, and use XModem to transfer the configuration file to the MSS Server. This procedure requires that the exported files reside on a serial line attached ASCII terminal emulator station.

After a configuration example has been opened (see 17.2.2, "Opening a Configuration" on page 145), use the **Configure/Create_router_configuration** menu to create an exported configuration file. Enter, for example, `export.cfg` and click on **OK**.

17.2.5 Loading a Configuration Using XModem

To upload the exported file you can use the following procedure:

- 1 Restart the MSS Server. Interrupt the boot sequence by pressing Ctrl+C until you are prompted for the MSS supervisory password (mss).

```
+-----Supervisory Password-----+
| Please Enter Supervisory Password |
|                                   |
|          ***                     |
+-----+-----+-----+-----+
```

- 2** Select option 4 in System Management Services.

```
Select one:                               System Management Services

1. Manage Configuration
2. Boot Sequence Selection
3. Select Device to Test
4. Utilities

-----
Enter   -   Esc=Quit  -   F1=Help   -   F3=Reboot  -   F9=Start OS -
```

- 3** Select option 12 in System Management Utilities.

```
Select one:                               System Management Utilities
1. Set Supervisory Password
2. Enable Unattended Start Mode
3. Disable Unattended Start Mode
4. Remove Supervisory Password
5. Update System Firmware
6. Display Error Log
7. View or Set Vital Product Data
8. Copy Remote Files
9. Remote Initial Program Load Setup
10. Manipulate Dead Man Timer
11. Display Event Log
12. Change Management

-----
Enter   -   Esc=Quit -   F1=Help -
```

- 4 Select option **XMODEM Software** from MSS Software Control.

Note: Specify you want a CONFIG file loaded.

```
Select one:                                MSS Software Control
Describe Software
Control Rebooting of Router
Control Dumping of Router
Copy Software
Erase Software
List Software
Set Boot Information
TFTP Software
XMODEM Software

-----
Enter   -   Esc=Quit -   F1=Help  -
```

- 5 Indicate to which bank and config file you want the new configuration example loaded. During our testing we used Bank A and Config 1.

Important

Copy the configuration examples to a bank to which the MSS R1.1 operational code has been copied before.

BANK A	BANK B	BANK F
IMAGE - AVAIL	IMAGE - PENDING	IMAGE - AVAIL
CONFIG 1 - NONE	CONFIG 1 - PENDING	CONFIG 1 - AVAIL
CONFIG 2 - NONE	CONFIG 2 - AVAIL	CONFIG 2 - AVAIL
CONFIG 3 - NONE	CONFIG 3 - AVAIL	CONFIG 3 - AVAIL
CONFIG 4 - NONE	CONFIG 4 - AVAIL	CONFIG 4 - AVAIL

TFTP Software
XMODEM Software

Enter - Esc=Quit - F1=Help -

- 6 Start your XModem file transfer when requested. After the file is loaded select **List Software** from MSS Software Control to verify the result of the file transfer. The config file you have just loaded should have status AVAIL.

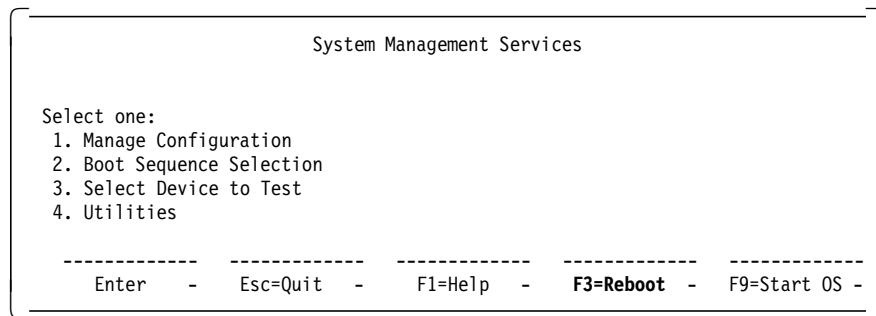
BANK A	BANK B	BANK F
IMAGE - AVAIL	IMAGE - PENDING	IMAGE - AVAIL
CONFIG 1 - AVAIL	CONFIG 1 - PENDING	CONFIG 1 - AVAIL
CONFIG 2 - NONE	CONFIG 2 - AVAIL	CONFIG 2 - AVAIL
CONFIG 3 - NONE	CONFIG 3 - AVAIL	CONFIG 3 - AVAIL
CONFIG 4 - NONE	CONFIG 4 - AVAIL	CONFIG 4 - AVAIL

- 7 Press Enter and select option **Set Boot Information** from MSS Software Control to make sure MSS will be started with the proper configuration file.

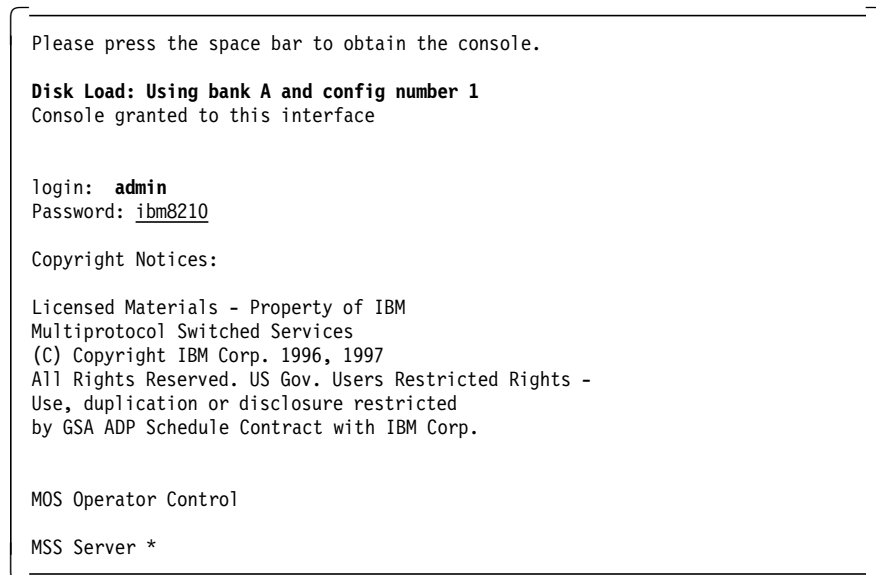
Select one:	MSS Software Control
Describe Software	
Control Rebooting of Router	
Control Dumping of Router	
Copy Software	
Erase Software	
List Software	
Set Boot Information	
TFTP Software	
XMODEM Software	

Enter - Esc=Quit - F1=Help -

- 8 Select the proper bank and config file (in our example, Bank A and Config 1). Set the duration parameter to either ONCE or PERMANENT.
- 9 Press ESC twice and the following screen appears:



- 10 Press F3 to reboot your machine. When prompted, enter YES to restart MSS.
- 11 The following display shows from which bank/config file MSS has been started. Enter the proper user ID and password and verify the configuration.



Note: All our configuration examples assume user ID admin and password ibm8210.

17.3 Example - Basic Connectivity

Today's ATM networks use two different approaches to support IP-based applications: Classical IP over ATM (RFC 1577) and, Ethernet or token-ring LAN Emulation. The basic connectivity configuration example contains the required server functions for a Classical IP subnet, an Ethernet ELAN, and a token-ring ELAN. This configuration can be used to get your MSS running and as a starting point for further configuration.

17.3.1 Configuration Details

Figure 70 on page 150 shows the basic connectivity example.

A single ESI (40.00.82.10.00.00) is used to address the ATM server and client functions on the MSS Server. Each of the functional elements is selected by a different selector.

The LECS/LES security interface allows authentication of LE clients joining an ELAN. See NW;Router/Devices/LAN_Emulation/LECS/LECS_LES Security.

- LES/BUS pairs - Selectors 02 and 03

We have defined two ELANs in this example: one token-ring (ELAN_T) and one Ethernet (ELAN_E) emulated LAN. The LES/BUS pair defined for the token-ring ELAN is using selector 02, and the Ethernet LES/BUS pair uses selector 03 (see NW;Router/Devices/LAN_Emulation/ELANs/ELANs).

The BUSES are running in adapter mode (see NW;Router/Devices/LAN_Emulation/ELANs/ELANs;local-LES-BUS/General-2).

- LE Clients - Selectors 04 and 05

The MSS Server will join the ELANs using two internal LE clients: one token-ring LE client using selector 04, and one Ethernet LE client using selector 05. These clients are able to communicate with other clients in the same emulated LAN. By adding an IP address to each LE client we have enabled IP routing on the ELAN interfaces (see NW;Router/Protocols/IP/Interfaces). No bridging has been defined in this example.

Administrative user admin has been defined using password ibm8210. We recommend you change the password to prevent unauthorized access to your MSS Server.

17.3.2 Logical IP Layout

Figure 71 shows the IP structure of this configuration example.




Figure 71. Basic Connectivity - IP Connectivity

Three Class B subnets are defined with MSS providing IP routing between each. The addresses are taken from a pool of unassigned networks (see RFC 1597, Address Allocation for Private Internets). The addresses have to be changed if you plan to connect to public IP facilities. Class B networks have been chosen to enable the setup of large flat networks.

For management purposes, the IP addresses 172.16.0.3, 172.17.0.3, 172.18.0.3, and 10.1.1.3 have been defined with write access within SNMP community admin (see NW;Router/System/SNMP_Config/Communities/Details;Addresses).

17.3.3 Verifying the Basic Configuration

The next section shows how you can verify proper operation of your just loaded configuration. Note that some of the information displayed is configuration-dependent, for example the ATM addresses.

1 Log in to the MSS Server.

login: **admin**
Password: ibm8210

Copyright Notices:

Licensed Materials - Property of IBM
Multiprotocol Switched Services
(C) Copyright IBM Corp. 1996, 1997
All Rights Reserved. US Gov. Users Restricted Rights -
Use, duplication or disclosure restricted
by GSA ADP Schedule Contract with IBM Corp.

MOS Operator Control

2 Enter GWCON and display the active configuration.

MSS Server ***talk 5**

MSS Server **+configuration**

IBM 8210 Nways Multiprotocol Switching Server
Host name: MSS Server
Version: 8210-MSS Feature 8706 V1 R1.1 PTF 0 RPQ 0 cc1_16e

Num	Name	Protocol
0	IP	DOD-IP
3	ARP	Address Resolution
11	SNMP	Simple Network Management Protocol
29	NHRP	Next Hop Routing Protocol

Num	Name	Feature
2	MCF	MAC Filtering
6	QOS	Quality of Service

3 Networks:

Net	Interface	MAC/Data-Link	Hardware	State
0	ATM/0	ATM	CHARM ATM	Up
1	TKR/0	Token-Ring/802.5	CHARM ATM	Up
2	Eth/0	Ethernet/IEEE 802.3	CHARM ATM	Up

3 Display the ATM addresses being used by the MSS Server.

```

MSS Server +network 0
ATM Console
MSS Server ATM+interface 0
ATM Interface Console
MSS Server ATM Interface+list address

```

Network Prefix	ATM Address	ESI	SEL
39.09.85.11.11.11.11.11.11.11.11.01.01.40.00.82.10.00.00.04			
39.09.85.11.11.11.11.11.11.11.11.01.01.40.00.82.10.00.00.05			
39.09.85.11.11.11.11.11.11.11.11.01.01.40.00.82.10.00.00.10			
39.09.85.11.11.11.11.11.11.11.11.01.01.40.00.82.10.00.00.03			
39.09.85.11.11.11.11.11.11.11.11.01.01.40.00.82.10.00.00.02			
39.09.85.11.11.11.11.11.11.11.11.01.01.40.00.82.10.00.00.01			
39.09.85.11.11.11.11.11.11.11.11.01.01.40.00.82.10.00.00.00			

```

MSS Server ATM Interface+exit

```

4 Display the status of the Ethernet and token-ring LES/BUS pairs.

```

MSS Server ATM+le-s

LE-Services Console
MSS Server LE-SERVICES+li all

ELAN Type (E=Ethernet/802.3, T=Token Ring/802.5)
Interface #
LES-BUS State (UP=Up, ID=Idle, ND=Net Down, ER=Error/Down,
               **=Other; Work with specific LES-BUS to see actual state)
ELAN Name
LES ATM Addr
-----
E 0 UP ELAN_E 39098511111111111111111111111111010140008210000003
T 0 UP ELAN_T 39098511111111111111111111111111010140008210000002

```

5 Display the status of the token-ring ELAN.

Note: In addition to the MSS internal LE client, we have connected one external LE client.

```

MSS Server LE-SERVICES+work ELAN_T
LE-Services Console for an existing LES-BUS Pair
MSS Server EXISTING LES-BUS 'ELAN_T'+da li all lec

Number of LEC's to display: 2

LEC-LES and LEC-BUS State (UP=Up, ID=Idle, --. --.
**=Other; Show specific LEC to see actual)
v v
LEC ID State #ATM #Reg #Lrnd
LEC Primary ATM Address Proxy ID LES BUS Adrs MACs MACs
-----
39098511111111111111111111111111010140008210000005 N 0001 UP UP 1 1 0
39098511111111111111111111111111010302007777000081 N 0002 UP UP 1 1 0

MSS Server EXISTING LES-BUS 'ELAN_T'+exit

```

6 Display the status of the Ethernet ELAN.

Note: In addition to the MSS internal LE client, we have connected one external LE client.

```

MSS Server LE-SERVICES+work ELAN_E
LE-Services Console for an existing LES-BUS Pair
MSS Server EXISTING LES-BUS 'ELAN_E'+da li all lec

Number of LEC's to display: 2

```

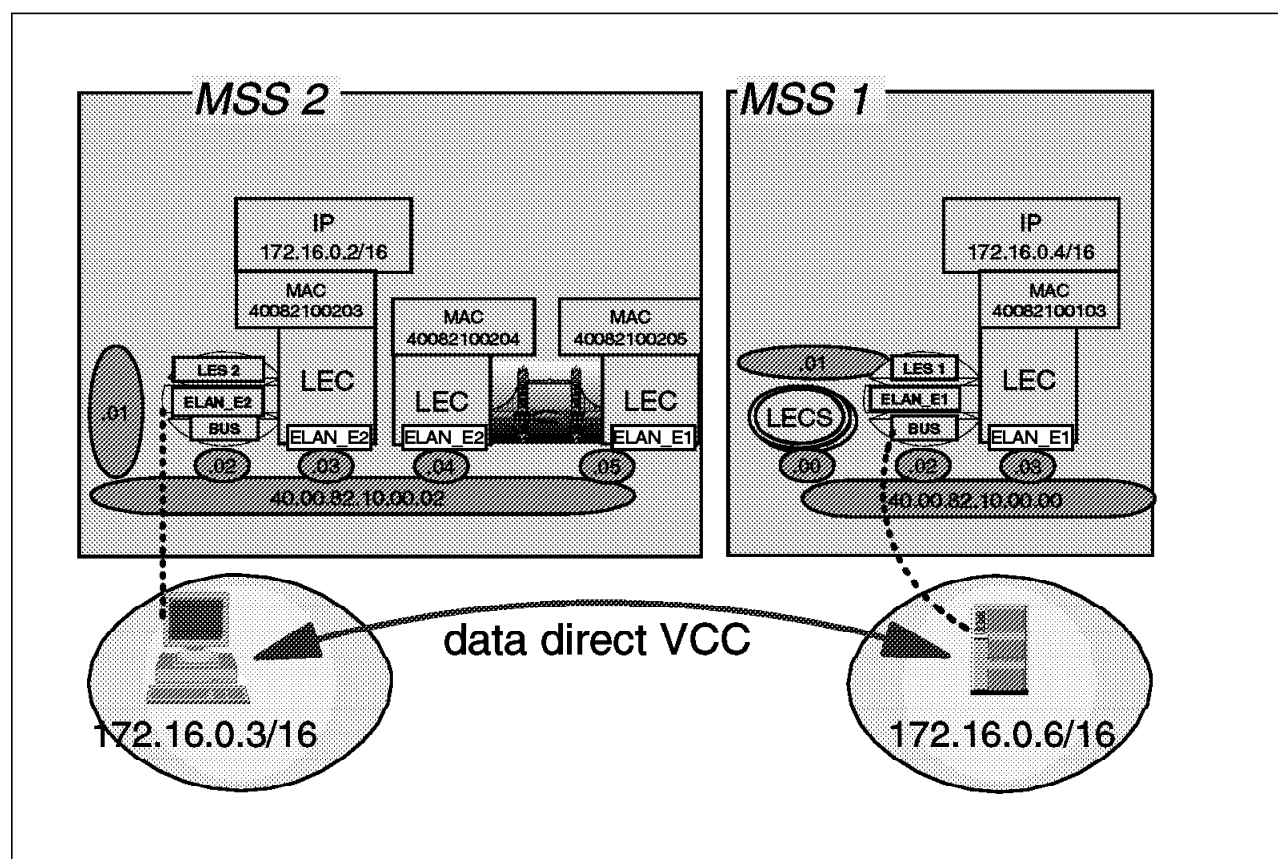
```

LEC-LES and LEC-BUS State (UP=Up, ID=Idle, --. --.
**=Other; Show specific LEC to see actual)
v v

```


Note: This example requires two distinct MSS Servers to be installed. In case you have only one MSS, we recommend you configure a single ELAN and avoid the complexity of having multiple ELANs.

Figure 72 depicts the shortcut bridging configuration example.



The configuration details for the MSS Servers are detailed in the next sections.

17.4.1.1 MSS Server 1

MSS Server 1 uses a single ESI (40.00.82.10.00.00) to address its ATM server and client functions. The selectors used by the various functional elements are:

- LECS - Selector 00

Using a LECS adds flexibility, redundancy, and security to your ELANs. We have defined a LECS to control Ethernet ELANs ELAN_E1 and ELAN_E2 (see NW;Router/Devices/LAN_Emulation/General).

LECS policies have been added to direct LE clients to the proper LES based on the ELAN name provided. For a definition of the enabled policies, see NW;Router/Devices/LAN_Emulation/LECS_Assignment_Policies. The policy values for ELAN_E1 are defined in:

NW;Router/Devices/LAN_Emulation/ELANs/ELANs;local-LES-policy-Values

(ELAN name and type). The policy values for ELAN_E2 are defined in:

NW;Router/Devices/LAN_Emulation/ELANs/ELANs;Remote-LES/ELAN Name.

Reconfiguration Required

The (remote) LES ATM address for ELAN_E2 is currently defined as 39.09.85.11.11.11.11.11.11.01.01.40.00.82.10.00.02.02. As this ATM address is implementation-dependent, reconfiguration of MSS Server 1 is required.

- LECS/LES Security Interface - Selector 01

The LECS/LES security interface allows authentication of LE clients joining an ELAN (see NW;Router/Devices/LAN_Emulation/LECS/LECS_LES Security).

- LES/BUS pair - Selector 02

This MSS Server provides the LES/BUS pair for ELAN_E1 using selector 02.

The backup LES/BUS for this ELAN runs on MSS Server 2. A redundancy VCC between the primary and backup LES/BUS prevents both LES/BUSes being active at the same time. Parameter settings can be reviewed in NW;Router/Devices/LAN_Emulation/ELANs/ELANs.

The LES will assign LEC IDs ranging from 1 to 1000. The BUS is running in adapter mode (see NW;Router/Devices/LAN_Emulation/ELANs/ELANs;local-LES-BUS/General-2).

- LE Client - Selector 03

MSS Server 1 contains a single internal LE client for ELAN_E1 (see NW;Router/Devices/LAN_Emulation/LEC_Interfaces/LEC_Interfaces).

By adding an IP address to the LE client we have enabled IP routing on the ELAN interface (see NW;Router/Protocols/IP/Interfaces).

Note: The purpose of this LE client is to verify the SCB connectivity and provide IP access for management purposes.

17.4.1.2 MSS Server 2

MSS Server 2 uses a single ESI (40.00.82.10.00.02) to address its ATM server and client functions. The selectors used by the various functional elements are:

- LES/BUS pair - Selector 02

MSS Server 2 provides the LES/BUS for ELAN_E2, using selector 02. Parameter settings can be reviewed in

NW;Router/Devices/LAN_Emulation/ELANs/ELANs.

The LES will assign LEC IDs ranging from 1001 to 2000. The BUS is running in adapter mode (see NW;Router/Devices/LAN_Emulation/ELANS/ELANS;local-LES-BUS/General-2).

- Shortcut Bridging LE clients - Selectors 04 and 05

Several LE clients are defined on MSS Server 2. For details, see NW;Router/Devices/LAN_Emulation/LEC_Interfaces/LEC_Interfaces.

Shortcut bridging between ELANS ELAN_E1 and ELAN_E2 is activated by enabling SCB on the bridging LE clients connected to ELAN_E1 and ELAN_E2 (see NW;Router/Bridging/Interfaces;Configure/Super_ELAN).

- Routing LE client - Selector 03

To enable IP routing, we have added an IP address to the routing LE client connected to ELAN_E2. For details, see NW;Router/Protocols/IP/Interfaces.

Note: The purpose of this LE client is to verify the SCB connectivity and provide IP access for management purposes.

17.4.1.3 Common Definitions

Both MSS Servers are configured to be SNMP manageable from IP host 172.16.0.3. On both MSS Servers we have defined default IP gateway 172.16.0.1.

Administrative user admin has been defined using password ibm8210. We recommend you change the password to prevent unauthorized access to your MSS Server.

17.4.2 Required Customization

Because MSS Server 1 uses LECS definitions for a remote LES, the ATM address of the LES/BUS for ELAN_E2 has to be updated to your local ATM configuration, before the example configuration can be used.

It is essential to find out what the ATM network prefix is of the ATM switches to which your MSS Server 2 connects. Refer to 17.1.2, "Allow Dynamic Access to the LECS" on page 140 for a description of how to obtain the ATM network prefix.

Open the configuration file for MSS Server 1 and select **NW;Router/Devices/LAN_Emulation/LECS/ELANS/ELANS;Remote_LES** for ELAN_E2. Enter the ATM prefix of MSS Server 2 followed by 40.00.82.10.00.02.02. Save the configuration in your customized.cdb database (see 17.2.3.5, "Saving Changed Configuration" on page 146).

No customization is required for MSS Server 2 as long as the LECS definitions on the adjacent ATM switch are properly set. For details see 17.1.2, "Allow Dynamic Access to the LECS" on page 140.

17.4.3 Logical IP Layout

This example contains a single IP network mapped onto two ELANS. Unicast traffic will flow directly between LE clients once data direct VCCs have been established. The SCB function enables the establishment of inter-ELAN VCCs. IP (and NetBIOS) broadcast traffic is controlled using BBCM on the shortcut bridging ports on MSS Server 1.




Figure 73. Shortcut Bridging - IP Connectivity

Figure 73 shows the IP addressing structure used for the example configuration. A single Class B subnet is used. The address is taken from a pool of unassigned networks (see *RFC 1597, Address Allocation for Private Internets*). The addresses have to be changed if you plan to connect to public IP facilities. A Class B network has been chosen to enable the setup of a large flat network.

On both MSS Servers we have added a default gateway statement to a router with IP address 172.16.0.1. If your default router has a different address, change the MSS definitions accordingly (see `NW;Router/Protocols/IP/Default_Gateway`). An alternative is to enable a dynamic routing protocol.

For management purposes, the IP addresses 172.16.0.3 and 10.1.1.3 have been defined with write access within SNMP community admin (see `NW;Router/System/SNMP_Config/Communities/Details/Addresses`).

17.4.4 Verifying Shortcut Bridging

The next section shows how you can verify proper operation of your just loaded configuration. Note, that some of the information displayed is configuration-dependent, for example, the ATM addresses.

The displays are obtained from MSS Servers 1 and 2 and external devices.

17.4.4.1 MSS Server 1

- 1 Log in to MSS Server 1.


```

MSS Server 1 ATM+le-s

LE-Services Console
MSS Server 1 LE-SERVICES+li all

ELAN Type (E=Ethernet/802.3, T=Token Ring/802.5)
Interface #
LES-BUS State (UP=Up, ID=Idle, ND=Net Down, ER=Error/Down,
**=Other; Work with specific LES-BUS to see actual state)
ELAN Name LES ATM Addr
-----
E 0 UP ELAN_E1 39098511111111111111010140008210000002

```

5 Display the status of ELAN_E1.

Note: The display shows an LE client on MSS 1, an LE client from MSS 2 and an external LE client

```

MSS Server 1 LE-SERVICES+work ELAN_E1
LE-Services Console for an existing LES-BUS Pair
MSS Server 1 EXISTING LES-BUS 'ELAN_E1'+da li all lec

```

Number of LEC's to display: 3

```

LEC-LES and LEC-BUS State (UP=Up, ID=Idle, --. --.
**=Other; Show specific LEC to see actual) v v
LEC State #ATM #Reg #Lrnd
Primary ATM Address Proxy ID LES BUS Adrs MACs MACs
-----
39098511111111111111010140008210000003 N 0001 UP UP 1 1 0
39098511111111111111010140008210000205 Y 0002 UP UP 1 1 0
39098511111111111111010302007770000081 N 0003 UP UP 1 1 0

```

```

MSS Server 1 EXISTING LES-BUS 'ELAN_E1'+exit
MSS Server 1 LE-SERVICES+exit
MSS Server 1 ATM+exit

```

6 Display the IP interfaces.

```

MSS Server 1 +pro ip
MSS Server 1 IP>int

```

```

Interface IP Address(es) Mask(s)
Eth/0 172.16.0.4 255.255.0.0

```

7 Ping to IP address 172.16.0.2 on MSS 2.

```

MSS Server 1 IP>ping 172.16.0.2

```

```

PING 172.16.0.4 -> 172.16.0.2: 56 data bytes, ttl=64, every 1 sec.
56 data bytes from 172.16.0.2: icmp_seq=0. ttl=255. time=40. ms
56 data bytes from 172.16.0.2: icmp_seq=1. ttl=255. time=0. ms

```

```

----172.16.0.2 PING Statistics----
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max = 0/20/40 ms

```

8 Ping to IP address 172.16.0.6 on the external LE client connected to ELAN_E1.

```

MSS Server 1 IP>ping 172.16.0.6

```

```

PING 172.16.0.4 -> 172.16.0.6: 56 data bytes, ttl=64, every 1 sec.
56 data bytes from 172.16.0.6: icmp_seq=0. ttl=255. time=0. ms
56 data bytes from 172.16.0.6: icmp_seq=1. ttl=255. time=0. ms

```

```

----172.16.0.6 PING Statistics----
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max = 0/0/0 ms

```

- 9 Ping to IP address 172.16.0.3 on the external LE client connected to ELAN_E2.

MSS Server 1 IP>**ping 172.16.0.3**

```

PING 172.16.0.4 -> 172.16.0.3: 56 data bytes, ttl=64, every 1 sec.
56 data bytes from 172.16.0.3: icmp_seq=0. ttl=255. time=0. ms
56 data bytes from 172.16.0.3: icmp_seq=1. ttl=255. time=0. ms

```

```

----172.16.0.3 PING Statistics----
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max = 0/0/0 ms

```

17.4.4.2 MSS Server 2

- 1 Log in to MSS Server 2.

```

login: admin
Password: ibm8210

```

Copyright Notices:

```

Licensed Materials - Property of IBM
Multiprotocol Switched Services
(C) Copyright IBM Corp. 1996, 1997
All Rights Reserved. US Gov. Users Restricted Rights -
Use, duplication or disclosure restricted
by GSA ADP Schedule Contract with IBM Corp.

```

MOS Operator Control

- 2 Enter GWCON and display the active configuration.

MSS Server 2 ***talk 5**

MSS Server 2 **+configuration**

```

IBM 8210 Nways Multiprotocol Switching Server
Host name: MSS Server 2
Version: 8210-MSS Feature 8706 V1 R1.1 PTF 0 RPQ 0    cc1_16e

```

```

Num Name  Protocol
0  IP      DOD-IP
3  ARP     Address Resolution
11 SNMP    Simple Network Management Protocol
29 NHRP    Next Hop Routing Protocol

```

```

Num Name  Feature
2  MCF     MAC Filtering
6  QOS     Quality of Service

```

```

4 Networks:
Net Interface  MAC/Data-Link          Hardware      State
0  ATM/0       ATM                    CHARM ATM    Up
1  Eth/0       Ethernet/IEEE 802.3    CHARM ATM    Up

```



```

Bridge type:          STB
Number of ports:      2
STP Participation:    IEEE802.1d

```

Port	Interface	State	MAC Address	Modes	Maximum MSDU	Segment	Flags
1	Eth/1	Up	40-00-82-10-02-04	T	1520		RD
2	Eth/2	Up	40-00-82-10-02-05	T	1520		RD

Flags: RE = IBMRT PC behavior Enabled, RD = IBMRT PC behavior Disabled

```

SR bridge number:    1
SR virtual segment:  001
Adaptive segment:    000

```

7 Display the status of the superELAN.

```

MSS Server 2 ASRT>1i super-elan
Number of Super ELAN bridge cache entries: 0
ATM Address                      Age Port Transid    Lecid CtrlFrame
-----
390985111111111111111010302007777000081  0  2          0  0000      N/A
390985111111111111111010302008888000081 15  1          0  0000      N/A
390985111111111111111010140008210000003 135 2          0  0000      N/A
390985111111111111111010140008210000203 15  1          0  0000      N/A
MSS Server 2 ASRT>exit

```

8 Display the IP interfaces.

```

MSS Server 2 +pro ip
MSS Server 2 IP>int

```

Interface	IP Address(es)	Mask(s)
Eth/0	172.16.0.4	255.255.0.0

9 Ping to IP address 172.16.0.4 on MSS 1.

```

MSS Server 2 IP>ping 172.16.0.4

```

```

PING 172.16.0.2 -> 172.16.0.4: 56 data bytes, ttl=64, every 1 sec.
56 data bytes from 172.16.0.4: icmp_seq=0. ttl=255. time=40. ms
56 data bytes from 172.16.0.4: icmp_seq=1. ttl=255. time=0. ms

```

```

----172.16.0.4 PING Statistics----
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max = 0/20/40 ms

```

10 Ping to IP address 172.16.0.6 on the external LE client connected to ELAN_E2.

```

MSS Server 2 IP>ping 172.16.0.6

```

```

PING 172.16.0.2 -> 172.16.0.6: 56 data bytes, ttl=64, every 1 sec.
56 data bytes from 172.16.0.6: icmp_seq=0. ttl=255. time=0. ms
56 data bytes from 172.16.0.6: icmp_seq=1. ttl=255. time=0. ms

```

```

----172.16.0.6 PING Statistics----
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max = 0/0/0 ms

```

11 Ping to IP address 172.16.0.3 on an external LE client connected to ELAN_E2.

MSS Server 2 IP>ping 172.16.0.3

```
PING 172.16.0.2 -> 172.16.0.3: 56 data bytes, ttl=64, every 1 sec.  
56 data bytes from 172.16.0.3: icmp_seq=0. ttl=255. time=0. ms  
56 data bytes from 172.16.0.3: icmp_seq=1. ttl=255. time=0. ms
```

```
----172.16.0.3 PING Statistics----  
2 packets transmitted, 2 packets received, 0% packet loss  
round-trip min/avg/max = 0/0/0 ms
```

17.4.4.3 Enhanced Mode

The final test we did was pinging between the two external LE clients. After activating the ping, we restarted MSS Server 2. During the restart we noticed no loss of pings, indicating that the shortcut VCC remained active.

Enhanced Mode

During ELAN configuration of the external LE clients make sure that the enhanced mode parameter is set to YES. If this parameter is set to NO the data direct VCCs between the LE clients will be dropped if the LES becomes unavailable.

17.5 Example - IP Cut-Through

A flat network design is desirable as it maximizes peer-to-peer communication without any intermediate routing. However, many reasons are possible for which partitioned (subnetted) networks will continue to exist. The most important are:

- To avoid broadcast traffic flooding the whole network
- Inability (for example, lack of control) to implement a single addressing structure on all user stations

The MSS IP cut-through facility provides an elegant solution for preventing unnecessary router hops while minimizing broadcasts and avoiding having to change the IP addressing structure on your client stations. For details, see 10.2.4, "IP Cut-Through Support for Zero-Hop Routing" on page 71.

This configuration example shows the design of a flat network with IP cut-through between two Ethernet ELANs.

17.5.1 Configuration Details

Figure 74 on page 165 shows the configuration details of the IP cut-through configuration example.




Figure 74. IP Cut-Through - Configuration Detail

A single ESI (40.00.82.10.00.00) is used to address the ATM server and client functions on the MSS Server. Each of the functional elements is selected by a different selector.

The selectors used are:

- LECS - Selector 00

Using a LECS adds flexibility, redundancy, and security to your ELANs. We have defined a LECS to control Ethernet ELANs ELAN_E1 and ELAN_E2 (see NW;Router/Devices/LAN_Emulation/General).

LECS policies have been added to direct LE clients to the proper LES based on the ELAN name provided or the LAN type. The ELAN name policy has highest priority. When no proper ELAN name has been defined Ethernet (and unspecified) LE clients will be directed to the LES for ELAN_E2. Token-ring LE clients will not be able to join any of the ELANs. For a definition of the enabled policies, see NW;Router/Devices/LAN_Emulation/LECS_Assignment_Policies. The policy value for ELAN_E1 and ELAN_E2 are defined in: NW;Router/Devices/LAN_Emulation/ELANs/ELANs;local-LES-policy-Values/ELAN Name.

- LECS/LES Security Interface - Selector 01

The LECS/LES security interface allows authentication of LE clients joining an ELAN (see NW;Router/Devices/LAN_Emulation/LECS/LECS_LES Security).

- LES/BUS pairs - Selector 02 and 03

We have defined two Ethernet ELANs in this example, ELAN_E1 and ELAN_E2. The LES/BUS pair for ELAN_E1 uses selector 02. The LES/BUS pair for ELAN_E2 uses selector 03. Parameter settings can be reviewed in NW;Router/Devices/LAN_Emulation/ELANs/ELANs.

The LES for ELAN_E1 will assign LEC IDs ranging from 1 to 1000, while the LES for ELAN_E2 assigns LEC IDs ranging from 1001 to 2000. The BUSes are running in adapter mode. For details on the LES/BUS definitions, see NW;Router/Devices/LAN_Emulation/ELANs/ELANs;local-LES-BUS/General-2.

- Shortcut Bridging LE clients - Selectors 04 and 05

Several LE clients are defined on the MSS Server. For details, see NW;Router/Devices/LAN_Emulation/LEC_Interfaces/LEC_Interfaces. The bridging LE client connected to ELAN_E1 uses selector 04, while the bridging LE client connected to ELAN_E2 selector 05.

By shortcut bridging the LE clients connected to ELAN_E1 and ELAN_E2, a SuperELAN containing both ELANs is created (see NW;Router/Bridging/Interfaces;Configure/Super_ELAN). Dynamic protocol filtering (DPF) is activated on the bridge ports for IP subnet 172.16.0.0/24 using SuperVLAN1 and for IP subnet 172.16.1.0/24 using SuperVLAN2 (see NW;Router/Bridging/Filtering/Dynamic_Protocol_Filtering/IP).

Note: The definitions we have entered are displayed in Figure 33 on page 73.

- Routing LE clients - Selectors 06 and 07

Two LE clients have been defined for routing purposes. The routing LE client connected to ELAN_E1 uses selector 06, while the routing LE client connected to ELAN_E2 uses selector 07. IP addresses have been added to enable IP routing on the ELAN interfaces. For details, see NW;Router/Protocols/IP/Interfaces.

Administrative user admin has been defined using password ibm8210. We recommend you change the password to prevent unauthorized access to your MSS Server.

17.5.2 Logical IP Layout

The configuration example depicted in Figure 75 on page 167 shows an IP network in which we have defined two SuperVLANs: SuperVLAN1 containing subnet 172.16.0.0/24 and SuperVLAN2 containing subnet 172.16.1.0/24. The dynamic protocol filtering definitions (see also Figure 33 on page 73) allow IP cut-through *from* subnet 172.16.0.0/24 *to* subnet 172.16.1.0.

Note: The NetBIOS PVLAN capabilities have been enabled as well. This PVLAN is named SuperVLAN3.




Figure 75. IP Cut-Through - IP Connectivity

Class B subnet address 172.16.0.0 is taken from a pool of unassigned networks (see RFC 1597, *Address Allocation for Private Internets*). The addresses have to be changed if you plan to connect to public IP facilities. A Class B network has been chosen to enable the setup of a large flat network.

All IP stations using subnet mask 255.255.0.0 can set up data direct VCCs; however, stations that use subnetting (in the example using IP subnet mask 255.255.255.0) require the routing functions of MSS. For example, data flowing from station 172.16.0.3/16 is sent over a data direct VCC, while data sent by station 172.16.1.3/24 requires routing on MSS.

On MSS we have added a default gateway statement to a router with IP address 172.16.0.1. If your default router has a different address, change the MSS definitions accordingly (see NW;Router/Protocols/IP/Default_Gateway). An alternative is to enable a dynamic routing protocol.

For management purposes, the IP addresses 172.16.0.3 and 10.1.1.3 have been defined with write access within SNMP community admin (see NW;Router/System/SNMP_Config/Communities/Details;Addresses).

17.5.3 Verifying IP Cut-Through

The next section shows how you can verify proper operation of your just loaded configuration. Note, that some of the information displayed is configuration-dependent, for example the ATM addresses.

- 1 Log in to the MSS Server.

login: **admin**
Password: ibm8210

Copyright Notices:

Licensed Materials - Property of IBM
Multiprotocol Switched Services
(C) Copyright IBM Corp. 1996, 1997
All Rights Reserved. US Gov. Users Restricted Rights -
Use, duplication or disclosure restricted
by GSA ADP Schedule Contract with IBM Corp.

MOS Operator Control

2 Enter GWCON and display the active configuration.

MSS Server ***talk 5**

MSS Server **+configuration**

IBM 8210 Nways Multiprotocol Switching Server
Host name: MSS Server
Version: 8210-MSS Feature 8706 V1 R1.1 PTF 0 RPQ 0 cc1_16e

Num	Name	Protocol
0	IP	DOD-IP
3	ARP	Address Resolution
11	SNMP	Simple Network Management Protocol
29	NHRP	Next Hop Routing Protocol

Num	Name	Feature
2	MCF	MAC Filtering
6	QOS	Quality of Service

5 Networks:

Net	Interface	MAC/Data-Link	Hardware	State
0	ATM/0	ATM	CHARM ATM	Up
1	Eth/0	Ethernet/IEEE 802.3	CHARM ATM	Up
2	Eth/1	Ethernet/IEEE 802.3	CHARM ATM	Up
3	Eth/2	Ethernet/IEEE 802.3	CHARM ATM	Up
4	Eth/3	Ethernet/IEEE 802.3	CHARM ATM	Up

3 Display the ATM addresses being used by the MSS Server.

MSS Server **+network 0**

ATM Console

MSS Server ATM**+interface 0**

ATM Interface Console

MSS Server ATM Interface**+list address**

Network Prefix	ATM Address	ESI	SEL
39.09.85.11.11.11.11.11.11.11.01.01.40.00.82.10.00.00.04			
39.09.85.11.11.11.11.11.11.11.01.01.40.00.82.10.00.00.06			
39.09.85.11.11.11.11.11.11.11.01.01.40.00.82.10.00.00.05			
39.09.85.11.11.11.11.11.11.11.01.01.40.00.82.10.00.00.07			
39.09.85.11.11.11.11.11.11.11.01.01.40.00.82.10.00.00.02			

```
39.09.85.11.11.11.11.11.11.11.01.01.40.00.82.10.00.00.03
39.09.85.11.11.11.11.11.11.11.01.01.40.00.82.10.00.00.01
39.09.85.11.11.11.11.11.11.11.01.01.40.00.82.10.00.00.00
MSS Server ATM Interface+exit
```

4 Display the status of the Ethernet LES/BUS pairs.

MSS Server ATM+1e-s

LE-Services Console

MSS Server LE-SERVICES+li all

ELAN Type (E=Ethernet/802.3, T=Token Ring/802.5)

```

Interface #
| LES-BUS State (UP=Up, ID=Idle, ND=Net Down, ER=Error/Down,
|               **=Other; Work with specific LES-BUS to see actual state)
| ELAN Name      LES ATM Addr
-----
E 0 UP ELAN_E1   3909851111111111111111111010140008210000002
E 0 UP ELAN_E2   39098511111111111111111111111010140008210000003

```

5 Display the status of ELAN_E1.

Note: In addition to the two MSS internal LE clients, we have connected one external LE client.

MSS Server LE-SERVICES+work ELAN E1

LE-Services Console for an existing LES-BUS Pair

MSS Server EXISTING LES-BUS 'ELAN E1'+da li all lec

Number of LEC's to display: 3

LEC-LES and LEC-BUS State (UP=Up, ID=Idle, --. --.							
**=Other; Show specific LEC to see actual) v v							
LEC Primary ATM Address	Proxy	LEC ID	State LES BUS	#ATM Adrs	#Reg MACs	#Lrnd MACs	
-----	-	----	-- --	----	-----	-----	
3909851111111111111110101400082100000006	N	0001	UP UP	1	1	0	
3909851111111111111110101400082100000004	Y	0002	UP UP	1	1	0	
39098511111111111111101010302008880000081	N	0003	UP UP	1	1	0	

```
MSS Server EXISTING LES-BUS 'ELAN E1'+exit
```

6 Display the status of ELAN_E2.

Note: In addition to the two MSS internal LE clients, we have connected one external LE client.

MSS Server LE-SERVICES+work ELAN_E2

LE-Services Console for an existing LES-BUS Pair

MSS Server EXISTING LES-BUS 'ELAN E2'+da li all lec

Number of LEC's to display: 3

LEC-LES and LEC-BUS State (UP=Up, ID=Idle, --, --, **=Other; Show specific LEC to see actual)						v	v		#ATM	#Reg	#Lrnd
						LEC	State				
LEC Primary ATM Address	Proxy	ID	LES	BUS				Adrs	MACs	MACs	
390985111111111111111111010140008210000007	N	03E9	UP	UP				1	1	0	
390985111111111111111111010140008210000005	Y	03EA	UP	UP				1	1	0	
39098511111111111111111101010302007777000081	N	03EB	UP	UP				1	1	0	

```

MSS Server EXISTING LES-BUS 'ELAN E2'+exit

```

MSS Server LE-SERVICES+exit

MSS Server ATM+exit

7 Display the IP interfaces.

MSS Server **+pro ip**
MSS Server IP>**int**

Interface	IP Address(es)	Mask(s)
Eth/1	172.16.0.2	255.255.255.0
Eth/3	172.16.1.1	255.255.255.0

8 Ping to IP address 172.16.0.3.

Note: We pinged to the external LE client connected to the ELAN_E1.

MSS Server IP>**ping 172.16.0.3**

PING 172.16.0.2 -> 172.16.0.3: 56 data bytes, ttl=64, every 1 sec.
56 data bytes from 172.16.0.3: icmp_seq=0. ttl=255. time=40. ms
56 data bytes from 172.16.0.3: icmp_seq=1. ttl=255. time=0. ms

----172.16.0.3 PING Statistics----
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max = 0/20/40 ms

9 Ping to IP address 172.16.1.3.

Note: We pinged to the external LE client connected to the ELAN_E2.

MSS Server IP>**ping 172.16.1.3**

PING 172.16.1.1 -> 172.16.1.3: 56 data bytes, ttl=64, every 1 sec.
56 data bytes from 172.16.1.3: icmp_seq=0. ttl=255. time=40. ms
56 data bytes from 172.16.1.3: icmp_seq=1. ttl=255. time=0. ms

----172.16.1.3 PING Statistics----
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max = 0/20/40 ms

10 Show the asymmetrical connection between IP address 172.16.1.3 and 172.16.0.3.

a. Display the route taken from IP address 172.16.1.3 to 172.16.0.3.

Note: This display was taken on an OS/2 workstation attached to ELAN_E2.

c:\tracerte **172.16.0.3**

0	172.16.1.1	(172.16.1.1)	0 ms	0 ms	0 ms
1	172.16.1.1	(172.16.1.1)	0 ms	0 ms	0 ms
2	172.16.0.3	(172.16.0.3)	32 ms	0 ms	0 ms

b. Display the route taken from IP address 172.16.0.3 to 172.16.1.3.

Note: This display was taken on an OS/2 workstation attached to ELAN_E1.

c:\tracerte **172.16.1.3**

0	172.16.1.3	(172.16.1.3)	0 ms	0 ms	0 ms
---	------------	--------------	------	------	------

As can be seen, the traffic from 172.16.0.3/16 to 172.16.1.3/24 takes a shortcut.

17.6 Example - Backward Connectivity

MSS Server Release 1.1 provides advanced functionality for SuperELANs and protocol-based VLANs. It also introduces support for internetworking with devices that are not able to take part in emulated LANs or CIP networks. Routing to existing FDDI networks is supported by the introduction of an FDDI adapter for the IBM 8210 MSS Server (see 4.1, "FDDI Interface" on page 11). The support for the RFC 1483 bridged frame format enables bridging over predefined PVCs (see Chapter 5, "RFC 1483 Bridging over ATM" on page 21).

In this example we depict how to define an FDDI attachment and an RFC 1483 bridging port. In addition, by using an emulated LAN, we demonstrate how to bridge between the RFC 1483 and the ELAN bridge port, and how to (IP) route between the FDDI LAN and the emulated LAN.

Note: IPX and AppleTalk routing can be enabled in a similar way by adding IPX and AppleTalk addresses to the ELAN and FDDI interface.

To verify the IP connectivity we connect an IBM 8274 RouteSwitch at the other end of the RFC 1483 PVC. (For definitions, see 17.6.3, "IBM 8274 Definitions" on page 173.)

17.6.1 Configuration Details

Figure 76 on page 172 shows the details of the backward connectivity example. A single ESI (40.00.82.10.00.00) is used to address the ATM server and client functions on the MSS Server. Each of the functional elements is selected by a different selector.

The selectors used are:

- LECS - Selector 00

Using a LECS adds flexibility, redundancy, and security to your ELANs. We have defined a LECS to control Ethernet ELAN ELAN_E (see NW;Router/Devices/LAN_Emulation/General).

LECS policies have been added to direct LE clients to the proper LES based on the network type (Ethernet or unspecified) or ELAN name (ELAN_E). For a definition of the enabled policies, see NW;Router/Devices/LAN_Emulation/LECS_Assignment_Policies. The policy values are defined in NW;Router/Devices/LAN_Emulation/ELANs/ELANs;local-LES-policy-Values (ELAN name and type).

- LECS/LES Security Interface - Selector 01

The LECS/LES security interface allows authentication of LE clients joining an ELAN (see NW;Router/Devices/LAN_Emulation/LECS/LECS_LES Security).

- LES/BUS pair - Selector 02

We have defined a single Ethernet ELAN in this example, ELAN_E. The LES/BUS for ELAN_E uses selector 02. Parameter settings can be reviewed in NW;Router/Devices/LAN_Emulation/ELANs/ELANs. The BUS is running in adapter mode (see NW;Router/Devices/LAN_Emulation/ELANs/ELANs;local-LES-BUS/General-2).

- Bridging LE client - Selector 03




Figure 76. Backward Connectivity - Configuration Detail

Several LE clients are defined on the MSS Server. For details, see `NW;Router/Devices/LAN_Emulation/LEC_Interfaces/LEC_Interfaces`.

The bridging LE client connects to ELAN_E using selector 03. It enables bridging between the RFC 1484 bridged connection and ELAN_E. BBCM (IP and NetBIOS) has been enabled for this interface.

- Routing LE client - Selector 04

The routing LE client connects to ELAN_E using selector 04. An IP address has been added to enable IP routing on the ELAN (`NW;Router/Protocols/IP/Interfaces`). As a result IP connectivity exists between the FDDI LAN and ELAN_E.

- RFC 1483 Bridge Port (Runtime selector)

On the first ATM attachment we have defined a PVC (VPI 0, VCI 69) for the transport of RFC 1483 bridged frames (see `NW;Router/Bridging/Interface/Configure`). All frames received on this interface are bridged using transparent bridging (see `NW;Router/Bridging/ATM-Virtual-Connections/General`). BBCM (IP and NetBIOS) has been enabled for this interface.

Note: The first selector that will be assigned at run time is X' C8'.

In addition we have included the following definitions:

- FDDI interface

We have enabled an FDDI interface using the default parameters (see NW;Router/Devices/Interfaces). On the FDDI interface we have defined an IP address to enable IP routing to the other IP ports on the MSS Server (see NW;Router/Protocols/IP/Interfaces).

Note: As we did not have an FDDI adapter we were not able to verify our FDDI definitions.

- Bridging Definitions

Transparent bridging has been enabled (see NW;Router/Bridging/General). The ELAN and the RFC 1483 PVC have been defined as transparent bridging ports. For the RFC 1483 PVC, see NW;Router/Bridging/ATM_Virtual_Connections.

- Spanning Tree Protocol

The spanning tree protocol has been enabled. The bridge address for the transparent bridging function is 400082100001. This address must be set as MSS uses the MAC address of the lowest bridge number for its bridge address and no MAC address can be defined for an RFC 1483 bridge port. For definitions, see NW;Router/Bridging/Spanning-Tree-Protocol/General.

Administrative user admin has been defined using password ibm8210. We recommend you change the password to prevent unauthorized access to your MSS Server.

17.6.2 IBM 8260 PVC Definitions

To define a PVC on the IBM 8260 we entered:

```
8260ATM1> Set PVC 12.1 1 16.1 1 channel 0.69 0.69
```

We were using port 16.1 for the MSS Server and port 12.1 for IBM 8274 and used VPI/VCI 0.69 on both ends. The status of PVCs can be displayed with:

```
8260ATM1> show pvc all
```

Local end point				! Remote end point !						
Port	id	type	Vpi/Vci	! Port	Vpi/Vci	HNb!	role	!QOS!	Status	
12.01	1	PTP-PVC	0/69	!16.01	0/69	1!	Primary	! BE!	Active	
16.01	1001	PTP-PVC	0/69	!12.01	0/69	1!	Secondary	! BE!	Active	

17.6.3 IBM 8274 Definitions

The following steps are required to define an RFC 1483 point-to-point bridging connection on the IBM 8274.

Note: Our example assumes group 3, internal IP address 172.16.0.5, and VPI/VCI 0/69 on the IBM 8274.

- Create a new group (CRGP).

Enter no to FR routing, CIP, route/IP, enable IPX, and config interfaces.

- Create the 1483 service (CAS) for group 3/1. Set the following parameters:

```

1 = rfc1483
2 = 6
10 = 2
4 = 3
5 = 69 (the vci number, the only allowable VPI is 0 !)

```

- **VAS** to add the service.
- **ADDVP 3 6/9** to add an Ethernet port to group 3.
- Select **optimized** as the bridge mode. Take defaults for the remaining parameters.
- Enter VI to verify that the RFC 1483 port (in our case 3/1/1) and the Ethernet port (6/9 in our case) are in group 3.
- Enter STC 3 and change the spanning tree priority to 10.
- Save, to save the configuration.

In addition, we configured an IP router arm (IP address 172.16.0.5) in group 3 to verify connectivity.

Note: For details on the IBM 8273 and 8274, see *IBM Nways RouteSwitch Implementation Guide*, SG24-4881.

17.6.4 Logical IP Layout

Figure 77 on page 175 shows the IP structure of this configuration example.

Two Class B subnets are defined with MSS providing IP routing between them. On the FDDI attachment we use subnet 172.17.0.0/16 while on the bridged ELAN/RFC 1483 environment we use 172.16.0.0/16. The addresses are taken from a pool of unassigned networks (see *RFC 1597, Address Allocation for Private Internets*). The addresses have to be changed if you plan to connect to public IP facilities. Class B networks have been chosen to enable the setup of large flat networks.

One of the MSS internal LE clients connected to ELAN_E uses IP address 172.16.0.1/16. An external LE client on the ELAN uses IP address 172.16.0.3. On the IBM 8274 to which we have defined the RFC 1483 connection we have configured IP address 172.16.0.5.

For management purposes, IP hosts 172.16.0.3, 172.17.0.3, and 10.1.1.3 have been defined with write access within SNMP community admin (see NW;Router/System/SNMP_Config/Communities/Details;Addresses).

17.6.5 Verifying Backward Connectivity

The next section shows how you can verify proper operation of your just loaded configuration. Note, that some of the information displayed is configuration dependent, for example the ATM addresses.

1 Log in to the MSS Server.

```

login: admin
Password: ibm8210

```

Copyright Notices:

```

Licensed Materials - Property of IBM
Multiprotocol Switched Services
(C) Copyright IBM Corp. 1996, 1997
All Rights Reserved. US Gov. Users Restricted Rights -

```

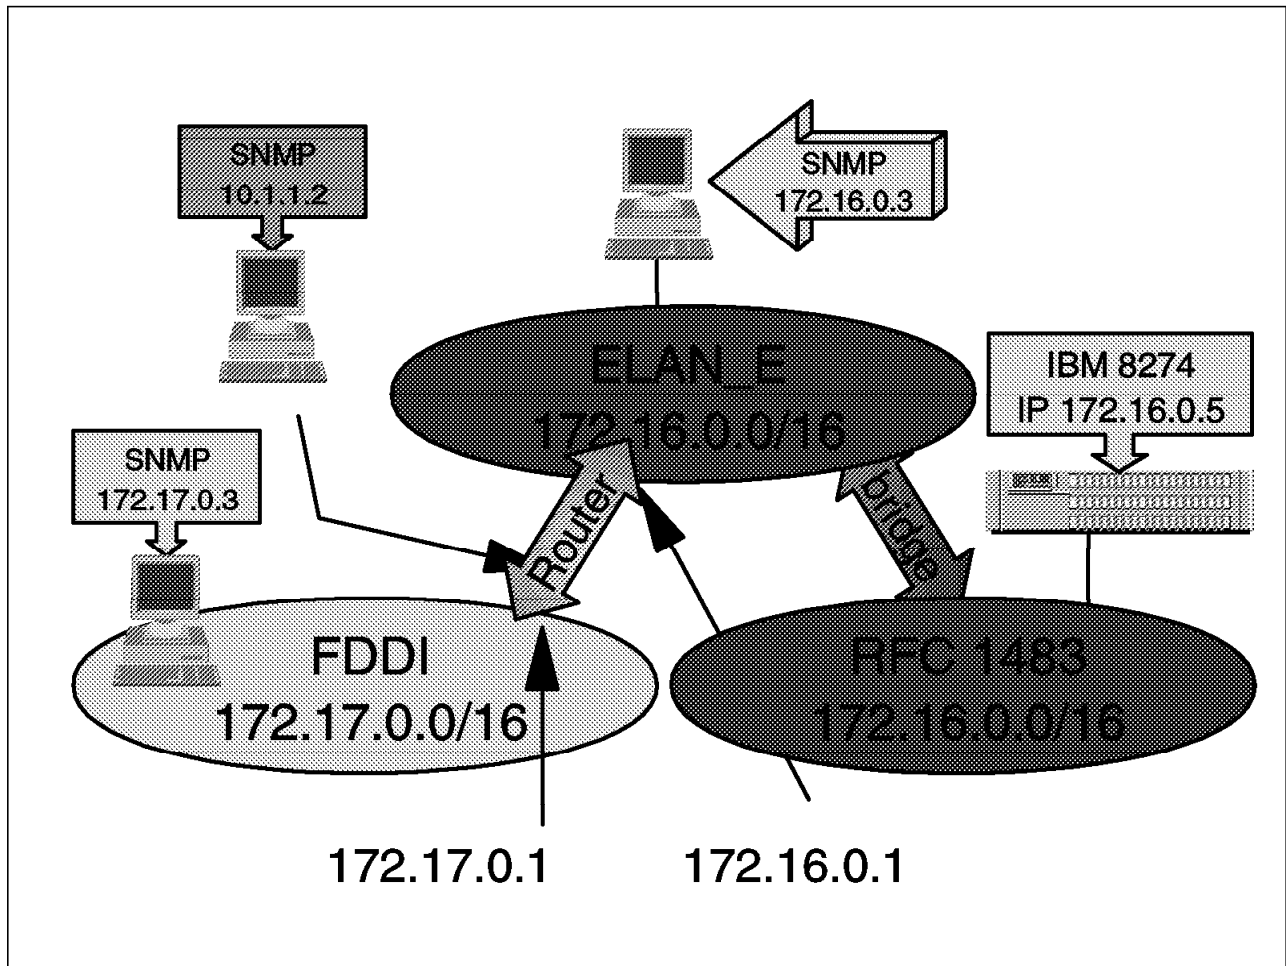



Figure 77. Backward Connectivity - IP Connectivity

Use, duplication or disclosure restricted
by GSA ADP Schedule Contract with IBM Corp.

MOS Operator Control

2 Enter GWCON and display the active configuration.

MSS Server ***talk 5**

MSS Server **+configuration**

IBM 8210 Nways Multiprotocol Switching Server
Host name: MSS Server
Version: 8210-MSS Feature 8706 V1 R1.1 PTF 0 RPQ 0 cc1_16e

Num	Name	Protocol
0	IP	DOD-IP
3	ARP	Address Resolution
11	SNMP	Simple Network Management Protocol
29	NHRP	Next Hop Routing Protocol

Num	Name	Feature
2	MCF	MAC Filtering
6	QOS	Quality of Service

4 Networks:				
Net	Interface	MAC/Data-Link	Hardware	State
0	ATM/0	ATM	CHARM ATM	Up
1	Eth/0	Ethernet/IEEE 802.3	CHARM ATM	Up


```

MSS Server +pro arp
MSS Server ARP>dis 0
Active Channel List : Net 0
New Channel List : Net 0
PVC Channel List : Net 0
      P/S FLAGS LIST  VPI/VCI      FwdPcr      FwdScr      MaxSDUsz Control P2P
1) P   80   03   0/69   155000000 155000000      9188      F      T
      Tgt Addr:
      Client Address (owner): Port no. 1
MSS Server ARP>exit

```

7 Display the status of the bridge.

```

MSS Server +pro asrt
MSS Server ASRT>li bridge
Bridge ID (prio/add):      32768/40-00-82-10-00-01
Bridge state:              Enabled
UB-Encapsulation:          Disabled
Bridge type:               SR-TB
Number of ports:           2
STP Participation:         IEEE802.1d on TB ports only

Port  Interface      State  MAC Address      Modes  Maximum  Segment  Flags
1    AT/0:0:69        Up     00-00-00-00-00-00  T      9234      RD
2    Eth/0            Up     40-00-82-10-00-03  T      1520      RD

Flags:  RE = IBMRT PC behavior Enabled,  RD = IBMRT PC behavior Disabled

SR bridge number:         1
SR virtual segment:       101
Adaptive segment:         EE1
MSS Server ASRT>exit

```

8 Display the IP interfaces.

```

MSS Server +pro ip
MSS Server IP>int

Interface  IP Address(es)  Mask(s)
Eth/1      172.16.0.1      255.255.0.0
FDDI/0     172.17.0.1      255.255.255.0

```

9 Verify RFC 1483 connectivity by pinging to an IP router arm on the IBM 8274 (IP address 172.16.0.5).

```

MSS Server IP>ping 172.16.0.5

PING 172.16.0.1 -> 172.16.0.5: 56 data bytes, ttl=64, every 1 sec.
56 data bytes from 172.16.0.5: icmp_seq=0. ttl=32. time=0. ms
56 data bytes from 172.16.0.5: icmp_seq=1. ttl=32. time=0. ms

----172.16.0.5 PING Statistics----
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max = 0/0/0 ms

```

10 Verify ELAN connectivity by pinging to the external LE client connected to ELAN_E (IP address 172.16.0.3).

```

MSS Server IP>ping 172.16.0.3

PING 172.16.0.2 -> 172.16.0.3: 56 data bytes, ttl=64, every 1 sec.
56 data bytes from 172.16.0.3: icmp_seq=0. ttl=255. time=40. ms
56 data bytes from 172.16.0.3: icmp_seq=1. ttl=255. time=0. ms

```

```
----172.16.0.3 PING Statistics----  
2 packets transmitted, 2 packets received, 0% packet loss  
round-trip min/avg/max = 0/20/40 ms
```

- 11 To verify connectivity between the IBM 8274 and the Ethernet LE client, we pinged from (OS/2) IP address 172.16.0.3 to IBM 8274 IP address 172.16.0.5.

Note: This display was taken from an OS/2 station (172.16.0.3) connected to ELAN_E.

```
c:\ping 172.16.0.5  
PING 172.16.0.5: 56 data bytes  
64 bytes from 172.16.0.5: icmp_seq=0. time=0. ms  
64 bytes from 172.16.0.5: icmp_seq=1. time=0. ms
```

```
----172.16.0.5 PING Statistics----  
2 packets transmitted, 2 packets received, 0% packet loss  
round-trip (ms) min/avg/max = 0/0/0
```

17.7 Example - Redundancy

The MSS Server redundancy features have been enhanced in Release 1.1. New functions are:

1. Redundant default IP gateway support for LAN emulation (see Chapter 12, "Default IP Gateway Redundancy and LAN Emulation" on page 103)

Two MSS Servers can provide mutual backup for a default ELAN IP gateway. During normal operation both MSS Servers have an active connection to the IP subnet. A pre-defined primary/backup relation is defined to make sure only one MSS Server is actively providing the gateway function. The backup station takes over responsibility when the primary fails. The primary station regains responsibility when reactivated.

The default IP gateway functions for ELAN do not require that both MSS Servers are connected to the same ATM switch, see Chapter 12, "Default IP Gateway Redundancy and LAN Emulation" on page 103.

2. Redundant default IP gateway support for Classical IP (see Chapter 13, "Classical IP Redundancy and Multicast Features" on page 107)

Two MSS Servers can provide mutual backup for a default Classical IP gateway. This function is equivalent to the default gateway for ELAN. Both MSS Servers have an active connection to the IP subnet at the same time. A pre-defined primary/backup relation is defined to make sure only one MSS Server is active. The backup station takes over responsibility when the primary fails. The primary station regains responsibility when reactivated.

The default IP gateway functions for Classical IP requires that both MSS Servers are connected to the same ATM switch. The default IP gateway function operates in conjunction with the ARP server function (that is, the MSS Server that provides the default gateway is ARP server for the same IP subnet).

3. Redundant LES/BUS support

The LES/BUS redundancy feature in R1.1 offers two options:

- a. Single active LES/BUS

Two LES/BUS pairs are defined for mutual backup using a pre-defined primary/backup relation. A redundancy VCC is used to make sure that

the backup LES only accepts LE clients joining the ELAN when the primary is not available.

Note: This is the same LES/BUS redundancy feature as introduced in MSS R1.0.

b. Concurrent active LES/BUSes

Two LES/BUS pairs are defined for mutual backup; however, both LES/BUSes can be active at the same time. The operation of the LES/BUSes is not synchronized and no redundancy VCC is used.

Notes:

- 1) Special care is required to make sure connectivity exists between LE clients controlled by different LESes (that is, joining a different ELAN).
- 2) This mode of redundant LES/BUS operation is the default.

The configuration depicted in Figure 78 on page 180 provides an example of all of the aforementioned redundancy features. By using LECSes on both MSS Servers redundant LECS support has also been implemented.

Let's summarize the redundancy features that are used in this example:

1. Default IP gateway for LAN Emulation

MSS Server 1 provides the primary default IP gateway for LAN emulation IP subnet 172.16.0.0/16 using an LE client connected to Ethernet ELAN ELAN_GW. For normal IP routing this LE client registers MAC address 40.00.82.10.01.06 while using IP address 172.16.0.4. For the default gateway functions it registers MAC address 40.00.82.10.00.00 to enable the use of IP address 172.16.0.1.

MSS Server 2 provides the backup IP gateway using its LE client connected to Ethernet ELAN ELAN_GW. During normal operation the LE client registers MAC address 40.00.82.10.02.06 while using IP address 172.16.0.2. Registering MAC address 40.00.82.10.00.00 will fail until the default IP gateway on MSS Server 1 becomes inactive. When in backup mode both MAC addresses will be registered: one for IP address 172.16.0.2, and one, using the newly acquired MAC address, for IP address 172.16.0.1.

2. Default IP gateway for Classical IP

MSS Server 1 provides the backup IP gateway for LIS 172.17.0.0/16. During normal operation it connects to the LIS as a LIS client using IP address 172.17.0.4. In backup mode, it provides the ARP server (using IP address 172.17.0.4), and the default IP gateway (IP address 172.17.0.1). The operation of both ARP servers is synchronized using a redundancy VCC.

MSS Server 2 provides the primary default IP gateway for logical IP subnet (LIS) 172.17.0.0/16 using IP address 172.17.0.1. MSS Server 2 is also the ARP server (IP address 172.17.0.2) for the logical IP subnet.

3. Redundant LES/BUS support (using a redundancy VCC)

MSS Servers 1 and 2 use the R1.0-compliant LES/BUS redundancy feature to provide LES/BUS redundancy for ELAN_GW. MSS Server 1 is defined as the primary LES/BUS. The operation of both LES/BUSes is synchronized using a redundancy VCC.

4. Redundant LES/BUS support (no redundancy VCC)

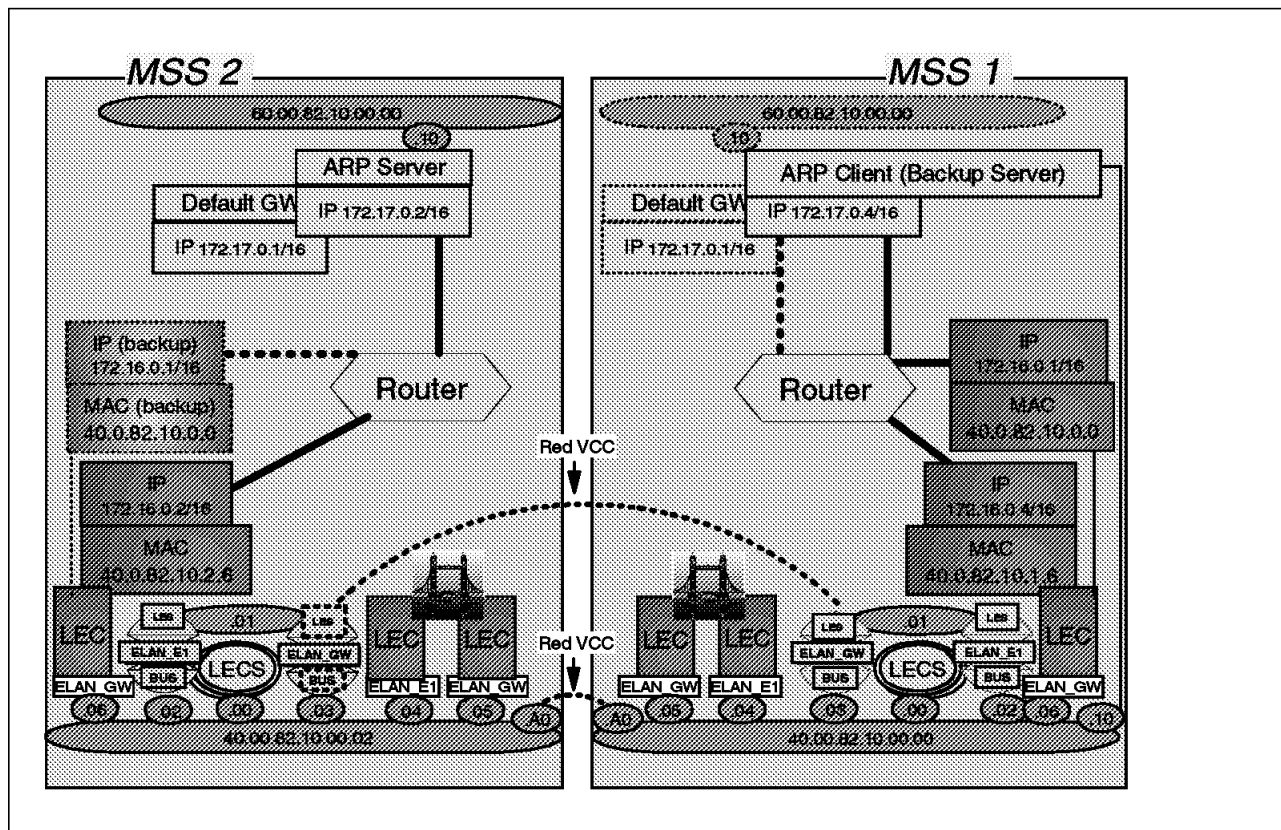


Figure 78. Redundancy - Configuration Detail

We are using the new redundancy features introduced in MSS R1.1, to provide LES/BUS redundancy for ELAN_E1.

On MSS Server 1 we have defined ELAN_E1 with a local LES/BUS as the primary LES/BUS. The LES/BUS for ELAN_E1 on MSS Server 2 is defined as its backup. On MSS Server 2 we have entered similar definitions: ELAN_E1 with a local LES/BUS as the primary LES/BUS, and the LES/BUS for ELAN_E1 on MSS Server 1 as its backup.

No redundancy VCC is used between the LES/BUSes for ELAN_E1 and both LES/BUSes are active at the same time. Therefore, LE clients can connect to either LES/BUS at the same time.

In this setup LE clients can connect to:

- a. ELAN_GW
- b. ELAN_E1 controlled by MSS Server 1
- c. ELAN_E1 controlled by MSS Server 2

To enable switched connectivity between any of the LE clients we have defined two shortcut bridging ports, one on ELAN_GW and ELAN_E1, on both MSS Servers.

5. Redundant LECS support

To enable the use of LES/BUS redundancy the use of a LECS is required. LECS redundancy is introduced by defining a LECS on both MSS Servers.

The LECS definitions on MSS Servers 1 and 2 are similar. LE clients will be assigned to an ELAN based on the ELAN name or ELAN type provided during registration. LE clients using ELAN_GW will be directed to the proper LES for ELAN_GW, and LE clients using ELAN_E1 will be directed to the local LES

for this ELAN. Ethernet LE clients using an undefined ELAN name will be directed to ELAN_E1.

Because of the SCB ports, switched connectivity exists between LE clients connected to ELAN_E1 and LE clients connected to ELAN_GW.

Check PTF

The MSS Server 1.1 code that was made available in March does not support bridging and default IP gateway for LAN emulation, concurrently. During testing we used an early version of the MSS code which solves this problem. Do not implement this scenario unless you have installed the updated MSS code.

Full redundancy can be achieved when your LE clients obtain the LECS ATM address from their adjacent ATM switch. Based on the definitions on your ATM switch (see 17.1.2, "Allow Dynamic Access to the LECS" on page 140) either the LECS on MSS Server 1 or MSS Server 2 will be tried first.

17.7.1 LES/BUS Service Distribution

To provide the redundant default ELAN IP gateway we have configured a more complex variant in which each MSS Server is connected to two, shortcut bridged, ELANs. External LE clients connect to ELAN_E1 (controlled by MSS Server 1 and 2 at the same time) while the MSS internal LE client providing the default IP gateway function is active on ELAN_GW.

A more simple variant of providing a redundant default IP gateway is using a single ELAN (for example, ELAN_GW) and connecting all MSS internal and external LE clients to the same ELAN.

The advantage of our more complex configuration is that the load on LESEs and BUSes for ELAN_E1 is shared while using shortcut data direct VCCs between LE clients and the IP gateway.

Shortcut Bridging

Shortcut bridging is currently only supported for Ethernet ELANs. If you want to set up a default IP gateway scenario for token-ring ELANs, use the simple variant and connect all your LE clients to a single non-bridged ELAN.

The ELAN that the LE clients connect to depends on the LECS it is using. Because of the Classical IP redundancy both MSS Servers have to connect to the same ATM switch.

17.7.2 Configuration Details

Figure 78 on page 180 depicts the redundancy example. The definitions entered on each of the MSS Servers are examined in the following sections.

17.7.2.1 MSS Server 1

MSS Server 1 uses ESI 40.00.82.10.00.00 for its ELAN server and client functions, its LIS client (during normal operation), and to establish the Classical IP redundancy VCC. The selectors used by these functions are:

- LECS - Selector 00

Using a LECS adds flexibility, redundancy, and security to your ELANs. We have defined a LECS to control Ethernet ELANs ELAN_E1 and ELAN_GW (see NW;Router/Devices/LAN_Emulation/General).

LECS policies have been added to direct LE clients to the proper LES based on their ELAN name or ELAN type. Ethernet LE clients that provide an undefined name are connected to ELAN_E1. For a definition of the enabled policies, see NW;Router/Devices/LAN_Emulation/LECS_Assignment_Policies. The policy values for ELAN_E1 and ELAN_GW are defined in: NW;Router/Devices/LAN_Emulation/ELANs/ELANs;local-LES-policy-Values (ELAN name and type).

- LECS/LES Security Interface - Selector 01

The LECS/LES security interface allows authentication of LE clients joining an ELAN (see NW;Router/Devices/LAN_Emulation/LECS/LECS_LES Security).

- LES/BUS pair - Selector 02 and 03

A local LES/BUS (selector 02) is defined for ELAN_E1. The parameter settings can be reviewed in NW;Router/Devices/LAN_Emulation/ELANs/ELANs. The ATM address of the backup LES on MSS Server 2 has been specified in NW;Router/Devices/LAN_Emulation/LECS/ELANS/ELANS;Remote_LES).

Reconfiguration Required

The (remote) LES ATM address for ELAN_E1 is currently defined as 39.09.85.11.11.11.11.11.11.01.01.40.00.82.10.00.02.02. As this ATM address is implementation-dependent, reconfiguration of MSS Server 1 is required.

A second local LES/BUS pair is defined using selector 03. This LES/BUS operates as the primary LES/BUS for ELAN_GW. A redundancy VCC to the backup LES/BUS on MSS Server 2 has been defined to avoid both LES/BUSes being active at the same time. The parameter settings can be reviewed in NW;Router/Devices/LAN_Emulation/ELANs/ELANs. The ATM address of the backup LES on MSS Server 2 has been specified in NW;Router/Devices/LAN_Emulation/LECS/ELANS/ELANS;Remote_LES) and NW;Router/Devices/LAN_Emulation/LECS/ELANS/ELANS;Local_LES_BUS/Redundancy).

Reconfiguration Required

The (remote) LES ATM address for ELAN_GW is currently defined as 39.09.85.11.11.11.11.11.11.11.01.01.40.00.82.10.00.02.03. As this ATM address is implementation-dependent, reconfiguration of MSS Server 1 is required.

Note: Change the address at both places.

The LES for ELAN_E1 assigns LEC IDs ranging from 1 to 1000. The LES for ELAN_GW assigns LEC IDs ranging from 1001 to 2000. The BUSes are

running in adapter mode (see
NW;Router/Devices/LAN_Emulation/ELANs/ELANs;local-LES-BUS/General-2).

- Shortcut Bridging LE clients - Selectors 04 and 05

Several LE clients are defined on MSS Server 1. For details, see
NW;Router/Devices/LAN_Emulation/LEC_Interfaces/LEC_Interfaces.

Shortcut bridging between ELAN_E1 and ELAN_GW is activated by enabling
SCB (NW;Router/Bridging/Interfaces;Configure/Super_ELAN) on the bridging
LE clients.

The bridging function enables LE clients connected to ELAN_E1 to
communicate with the default IP gateway that has been defined on
ELAN_GW. Shortcut bridging bypasses the bridging function once a data
direct VCC has been established. BBCM (IP and NetBIOS) has been enabled
for these interfaces.

To avoid isolation of ELANs it is important that the internal LE client for
ELAN_E1 is controlled by the local LES/BUS pair. Therefore, we use a
hard-coded LES address for this LE client (see
NW;Router/Devices/LAN_Emulation/LEC_Interfaces/LEC_Interfaces;
Servers/LE_Server_ATM_Address).

Reconfiguration Required

The (local) LES ATM address for ELAN_E1 is currently defined as
39.09.85.11.11.11.11.11.11.01.01.40.00.82.10.00.00.02. As this ATM
address is implementation dependent, reconfiguration of MSS Server 1 is
required.

- Default GW (routing) LE client - Selector 06

An LE client has been defined using ELAN name ELAN_GW. An IP address
(172.16.0.4/16) has been added to enable IP connectivity for IP network
172.16.0.0. For details, see NW;Router/Protocols/IP/Interfaces. This LE client
is using MAC address 40.00.82.10.01.06.

This LE client is defined as the primary default IP gateway using IP address
172.16.0.1/16 and MAC address 40.00.82.10.00.00 (see
NW;Router/Protocols/IP/Redundant_Gateway_Interface).

- LIS Client - Selector 10

A LIS client (172.17.0.4/16) is defined using the (remote) ARP server on MSS
2. The ARP server's ATM address has been defined in
NW;Router/Devices/Protocols/Classical_IP_over_ATM/Classical_IP_over_ATM;
ARP_Server.

Reconfiguration Required

The ARP server's ATM address is currently defined as
39.09.85.11.11.11.11.11.11.01.01.60.00.82.10.00.00.10. As this ATM
address is implementation-dependent, reconfiguration of MSS Server 1 is
required.

The LIS client has also been configured as a backup default IP gateway
(using IP address 172.17.0.1). For more details, see
NW;Router/Devices/Protocols/Classical_IP_over_ATM/Classical_IP_over_ATM;
Red_Backup.

- Classical IP Gateway Redundancy VCC - Selector A0

The LIS client provides backup for the ARP server. To synchronize their activity, a redundancy VCC is established using a pre-configured ATM address. For more details, see NW;Router/Devices/Protocols/Classical_IP_over_ATM/Classical_IP_over_ATM;Redun_Server.

Reconfiguration Required

The remote redundancy VCC ATM address is defined as 39.09.85.11.11.11.11.11.11.01.01.40.00.82.10.00.02.A0. As this ATM address is implementation-dependent, reconfiguration of MSS Server 1 is required.

MSS Server 1 uses ESI 60.00.82.10.00.00 for its backup ARP server and default Classical IP gateway function. The selector used is 10. For details, see NW;Router/Devices/Protocols/Classical_IP_over_ATM/Classical_IP_over_ATM;Redun_Backup.

17.7.2.2 MSS Server 2

MSS Server 2 uses ESI 40.00.82.10.00.02 for its ELAN server and client functions and to establish the Classical IP redundancy VCC. The selectors used by these functions are:

- LECS - Selector 00

Using a LECS adds flexibility, redundancy, and security to your ELANs. We have defined a LECS to control Ethernet ELANs ELAN_E1 and ELAN_GW (see NW;Router/Devices/LAN_Emulation/General).

LECS policies have been added to direct LE clients to the proper LES based on their ELAN name or ELAN type. Ethernet LE clients that provide an undefined name are connected to ELAN_E1. For a definition of the enabled policies, see NW;Router/Devices/LAN_Emulation/LECS_Assignment_Policies. The policy values for ELAN_E1 are ELAN_GW are defined in: NW;Router/Devices/LAN_Emulation/ELANs/ELANs;local-LES-policy-Values (ELAN name and type).

- LECS/LES Security Interface - Selector 01

The LECS/LES security interface allows authentication of LE clients joining an ELAN (see NW;Router/Devices/LAN_Emulation/LECS/LECS_LES Security).

- LES/BUS pairs - Selector 02 and 03

A local LES/BUS (selector 02) is defined for ELAN_E1. Parameter settings can be reviewed in NW;Router/Devices/LAN_Emulation/ELANs/ELANs. The ATM address of the backup LES on MSS Server 1 has been specified in NW;Router/Devices/LAN_Emulation/LECS/ELANS/ELANS;Remote_LES) and NW;Router/Devices/LAN_Emulation/LECS/ELANS/ELANS;Local_LES_BUS/Redundancy).

Reconfiguration Required

The (remote) LES ATM address for ELAN_E1 is currently defined as 39.09.85.11.11.11.11.11.11.01.01.40.00.82.10.00.00.02. As this ATM address is implementation-dependent, reconfiguration of MSS Server 2 is required.

Note: Change the address at both places.

A second local LES/BUS pair is defined using selector 03. This LES/BUS operates as the backup LES/BUS for ELAN_GW. A redundancy VCC is used to the backup LES/BUS on MSS Server 1, to avoid both LES/BUSes being active at the same time. The parameter settings can be reviewed in NW;Router/Devices/LAN_Emulation/ELANS/ELANS. The ATM address of the primary LES on MSS Server 1 has been specified in NW;Router/Devices/LAN_Emulation/LECS/ELANS/ELANS;Remote_LES) and NW;Router/Devices/LAN_Emulation/LECS/ELANS/ELANS;Local_LES_BUS/Redundancy).

Reconfiguration Required

The (remote) LES ATM address for ELAN_GW is currently defined as 39.09.85.11.11.11.11.11.11.01.01.40.00.82.10.00.00.03. As this ATM address is implementation-dependent, reconfiguration of MSS Server 2 is required.

Note: Change the address at both places.

The LES for ELAN_GW assigns LEC IDs ranging from 1001 to 2000. The LES for ELAN_E1 assigns LEC IDs ranging from 2001 to 3000. The BUSes are running in adapter mode; see NW;Router/Devices/LAN_Emulation/ELANS/ELANS;local-LES-BUS/General-2.

- Shortcut Bridging LE clients - Selectors 04 and 05

Two LE-clients are defined on MSS Server 2. For details, see NW;Router/Devices/LAN_Emulation/LEC_Interfaces/LEC_Interfaces.

Shortcut bridging between ELAN_E2 and ELAN_GW is activated by enabling SCB (NW;Router/Bridging/Interfaces/Configure/Super_ELAN) on the bridging LE clients.

The bridging function enables LE clients connected to ELAN_E1 to communicate with the default IP gateway that has been defined on ELAN_GW. Shortcut bridging bypasses the bridging function once a data direct VCC has been established. BBCM (IP and NetBIOS) has been enabled for these interfaces.

To avoid isolation of ELANS it is important that the internal LE client for ELAN_E1 is controlled by the local LES/BUS pair. Therefore, we use a hard-coded LES address for this LE client; see NW;Router/Devices/LAN_Emulation/LEC_Interfaces/LEC_Interfaces;Servers/LE_Server_ATM_Address).

Reconfiguration Required

The (local) LES ATM address for ELAN_E1 is currently defined as 39.09.85.11.11.11.11.11.11.01.01.40.00.82.10.00.02.02. As this ATM address is implementation-dependent, reconfiguration of MSS Server 2 is required.

- Default GW (routing) LE client - Selector 06

A LE client has been defined using ELAN name ELAN_GW. An IP address (172.16.0.2/16) has been added to enable IP connectivity for IP network 172.16.0.0. For details see NW;Router/Protocols/IP/Interfaces. This LE client is using MAC address 40.00.82.10.02.06.

This LE client is defined as the backup default IP gateway using IP address 172.16.0.1/16 and MAC address 40.00.82.10.00.00.
(NW;Router/Protocols/IP/Redundant_Gateway_Interface).

When joining the ELAN, the LE client will try to register MAC addresses 40.00.82.10.02.00 and 40.00.82.10.00.00. Registering 40.00.82.10.00.00 will fail as long as the primary IP gateway (on MSS Server 1) is active (see NW;Router/Devices/LAN_Emulation/LECS/ELANS/ELANS; Local_LES_BUS/Redundancy).

- ARP Server - Selector 10

A redundant ARP client (172.17.0.2/16) has been defined for LIS 172.17.0.0/16. For details, see NW;Router/Devices/Protocols/Classical_IP_over_ATM/Classical_IP_over_ATM.

The ARP server has been configured as the primary default IP gateway (using IP address 172.17.0.1). For more details, see NW;Router/Devices/Protocols/Classical_IP_over_ATM/Classical_IP_over_ATM; Red_Backup.

- Classical IP Gateway Redundancy VCC - Selector A0

The redundant ARP server signals its presence by establishing a redundancy VCC to a preconfigured ATM address (see NW;Router/Devices/Protocols/Classical_IP_over_ATM/Classical_IP_over_ATM; Red_Server).

Reconfiguration Required

The remote redundancy VCC ATM address is defined as 39.09.85.11.11.11.11.11.11.01.01.40.00.82.10.00.00.A0. As this ATM address is implementation-dependent, reconfiguration of MSS Server 2 is required.

MSS Server 2 uses ESI 60.00.82.10.00.00 for its ARP server and default IP gateway function. The selector used is 10.

17.7.2.3 Common Definitions

Both MSS Servers are configured to be SNMP manageable from IP host 172.16.0.3.

Administrative user admin has been defined using password ibm8210. We recommend you change the password to prevent unauthorized access to your MSS Server.

17.7.3 Customization

Before you can use the example configuration, you have to adjust ATM addresses according to your local ATM configuration. Refer to 17.1.2, "Allow Dynamic Access to the LECS" on page 140 for a description of how to obtain the ATM prefixes of your MSS Servers.

Important

If you are using the Classical IP default gateway function, attach both MSS Servers to the same ATM switch. In this case the ATM network prefix is the same for both MSS Servers.

17.7.3.1 MSS Server 1

The following definitions need to be reconfigured:

1. The remote LES ATM addresses in the LECS/ELAN definitions for ELAN_GW and ELAN_E1.

MSS Server 1 provides the primary LES/BUS for ELAN_E1. No redundancy VCC is being used. To specify the ATM address of the backup LES on MSS Server 2 for ELAN_E1 open the configuration file for MSS Server 1 and select **NW;Router/Devices/LAN_Emulation/LECS/ELANS/ELANS;Remote_LES** for ELAN_E1. Enter the ATM prefix of MSS Server 2 followed by 40.00.82.10.00.02.02.

2. The ATM address of the partner LES/BUS for the redundancy VCC of ELAN_GW.

MSS Server 1 also provides the primary LES/BUS for ELAN_GW. A redundancy VCC is being used to avoid primary and backup LESes from being active at the same time. To specify the ATM address of the backup LES on MSS Server 2 for ELAN_GW open the configuration file for MSS Server 1 and select

NW;Router/Devices/LAN_Emulation/LECS/ELANS/ELANS;Remote_LES for ELAN_GW. Enter the ATM prefix of MSS Server 2 followed by 40.00.82.10.00.02.03. Enter the same ATM address in **NW;Router/Devices/LAN_Emulation/LECS/ELANS/ELANS;Local_LES_BUS/Redundancy**.

3. The ATM address of the ARP server of the LIS client.

The LIS client uses an ARP server running on MSS Server 2. To enter the proper ATM address, select

NW;Router/Devices/Protocols/Classical_IP_over_ATM/Classical_IP_over_ATM;ARP_Server Enter the ATM prefix of MSS Server 2 followed by 60.00.82.10.00.00.10.

4. The ATM address of the partner node for the redundancy VCC of the Classical IP default gateway.

To use the backup ARP server function a redundancy VCC has to be enabled between primary and backup node. Select

NW;Router/Devices/Protocols/Classical_IP_over_ATM/Classical_IP_over_ATM;Redun_Server. and enter the ATM prefix of MSS Server 2 followed by 40.00.82.10.00.02.A0.

5. The LES ATM addresses for the internal LE client for ELAN_E1.

MSS Server 1 provides an internal LE client connected to ELAN_E1 with a hard-coded LES address pointing to the local LES. To specify the ATM address of the LES on MSS Server 1 for LE client ELAN_E1 open the configuration file for MSS Server 1 and select

NW;Router/Devices/LAN_Emulation/LEC_Interfaces/LEC_Interfaces;Servers/LE_Server_ATM_Address for LE client ELAN_E1. Enter the ATM prefix of MSS Server 1 followed by 40.00.82.10.00.00.02.

Save the configuration in your customized.cdb database (see 17.2.3.5, "Saving Changed Configuration" on page 146) and download it to MSS Server 1.

17.7.3.2 MSS Server 2

The following information has to be entered before this configuration can be used:

1. The remote LES ATM addresses in the LECS/ELAN definitions for ELAN_GW and ELAN_E1.

MSS Server 2 provides the primary LES/BUS for ELAN_E1. No redundancy VCC is being used. To specify the ATM address of the backup LES on MSS Server 1 for ELAN_E1 open the configuration file for MSS Server 2 and select **NW;Router/Devices/LAN_Emulation/LECS/ELANS/ELANS;Remote_LES** for ELAN_E1. Enter the ATM prefix of MSS Server 1 followed by 40.00.82.10.00.00.02.

2. The ATM address of the partner LES/BUS.

MSS Server 2 provides the backup LES/BUS for ELAN_GW. A redundancy VCC is being used to avoid primary and backup LESes from being active at the same time. To specify the ATM address of the primary LES on MSS Server 1 for ELAN_GW open the configuration file for MSS Server 2 and select **NW;Router/Devices/LAN_Emulation/LECS/ELANS/ELANS;Remote_LES** for ELAN_GW. Enter the ATM prefix of MSS Server 1 followed by 40.00.82.10.00.00.03. Enter the same ATM address in **NW;Router/Devices/LAN_Emulation/LECS/ELANS/ELANS;Local_LES_BUS/Redundancy** for the redundancy VCC of ELAN_GW.

3. The ATM address of the partner node for the redundancy VCC of the Classical IP default gateway.

The ARP server is primary with a backup defined on MSS Server 1. To enable the redundancy VCC to be enabled between the primary and backup node select

NW;Router/Devices/Protocols/Classical_IP_over_ATM/Classical_IP_over_ATM;Redun_Server. Enter the ATM prefix of MSS Server 1 followed by 40.00.82.10.00.00.A0.

4. The LES ATM addresses for the internal LE client for ELAN_E1.

MSS Server 2 provides an internal LE client connected to ELAN_E1 with a hard-coded LES address pointing to the local LES. To specify the ATM address of the LES on MSS Server 2 for LE client ELAN_E1: open the configuration file for MSS Server 1 and select

NW;Router/Devices/LAN_Emulation/LEC_Interfaces/LEC_Interfaces;Servers/LE_Server_ATM_Address for LE client ELAN_E1. Enter the ATM prefix of MSS Server 2 followed by 40.00.82.10.00.02.02.

Save the configuration in your customized.cdb database (see 17.2.3.5, "Saving Changed Configuration" on page 146) and download it to MSS Server 2.

17.7.4 Logical IP Layout

This example contains a single IP network mapped onto all three ELANs. (ELAN_E1 controlled by MSS 1 is different from ELAN_E1 controlled by MSS 2.) Unicast traffic will flow directly between LE clients once data direct VCCs have been established. The SCB function enables the establishment of inter-ELAN VCCs. Broadcasts are managed using the BBCM and BCM functions.

Figure 79 on page 189 shows the IP addressing structure used for the example configuration. A single class B subnet is used. The address is taken from a pool of unassigned networks (see *RFC 1597, Address Allocation for Private*




Figure 79. Redundancy - IP Connectivity

Internets). The addresses have to be changed if you plan to connect to public IP facilities. A class B network has been chosen to enable the setup of a large flat network.

On the MSS Server we have added a default gateway statement to a router with IP address 172.16.0.1. If your default router has a different address change the MSS definitions accordingly (see NW;Router/Protocols/IP/Default_Gateway). An alternative is to enable a dynamic routing protocol.

For management purposes, the IP addresses 172.16.0.3, 172.17.0.3, and 10.1.1.3 have been defined with write access within SNMP community admin (see NW;Router/System/SNMP_Config/Communities/Details;Addresses).

17.7.5 Verifying Redundancy

The next section shows how you can verify proper operation of your just loaded configuration.

Note: Some of the information displayed is configuration-dependent, for example, the ATM addresses.

The displays are obtained from MSS Servers 1 and 2 and external devices.

17.7.5.1 MSS Server 1

- 1 Log in to MSS Server 1.

login: **admin**
Password: ibm8210

Copyright Notices:

Licensed Materials - Property of IBM
Multiprotocol Switched Services
(C) Copyright IBM Corp. 1996, 1997
All Rights Reserved. US Gov. Users Restricted Rights -
Use, duplication or disclosure restricted
by GSA ADP Schedule Contract with IBM Corp.

MOS Operator Control

2 Enter GWCON and display the active configuration.

MSS Server 1 ***talk 5**

MSS Server 1 **+configuration**

IBM 8210 Nways Multiprotocol Switching Server
Host name: MSS Server 1
Version: 8210-MSS Feature 8706 V1 R1.1 PTF 0 RPQ 0 **cc1_17b test-load**

Num	Name	Protocol
0	IP	DOD-IP
3	ARP	Address Resolution
11	SNMP	Simple Network Management Protocol
29	NHRP	Next Hop Routing Protocol

Num	Name	Feature
2	MCF	MAC Filtering
6	QOS	Quality of Service

4 Networks:

Net	Interface	MAC/Data-Link	Hardware	State
0	ATM/0	ATM	CHARM ATM	Up
1	Eth/0	Ethernet/IEEE 802.3	CHARM ATM	Up
2	Eth/1	Ethernet/IEEE 802.3	CHARM ATM	Up
3	Eth/2	Ethernet/IEEE 802.3	CHARM ATM	Up

Note: Make sure you are using the MSS code that supports bridging and the redundant default IP gateway for LAN emulation, concurrently.

3 Display the ATM addresses being used by MSS Server 1.

MSS Server 1 **ATM+le-s**

LE-Services Console
MSS Server 1 **LE-SERVICES+li all**

Network Prefix	ATM Address	ESI	SEL
39.09.85.11.11.11.11.11.11.11.01.01.40.00.82.10.00.00.04			
39.09.85.11.11.11.11.11.11.11.01.01.40.00.82.10.00.00.05			
39.09.85.11.11.11.11.11.11.11.01.01.40.00.82.10.00.00.06			

4 Display the status of the Ethernet LES/BUS pairs.

5 Display the status of the LES/BUS for ELAN_E1.

```

-LES-BUS Options-
BUS Mode:                               Adapter
Security (LECS Validation of Joins):    Disabled

```

```
-Broadcast Manager Configuration-
IP BCM: Disabled
IPX BCM: Disabled
NetBIOS BCM: Disabled
BCM IP Cache Aging Time: 5
BCM IPX Cache Aging Time: 3
BCM NetBIOS Cache Aging Time: 15
BCM IPX Maximum Forwarding List: 50
No BCM IPX Static Entries defined
```

```
MSS Server 1 EXISTING LES-BUS 'ELAN_E1'+exit
MSS Server 1 LE-SERVICES+exit
```

```

-Current Configuration-
  LES-BUS Enabled/Disabled:      Enabled
  ATM Device number:              0
  End System Identifier (ESI):    40.00.82.10.00.00
  Selector Byte:                  0x03
  ELAN Type:      (S2)            Ethernet
  Max Frame Size: (S3)            1516
  Control Timeout: (S4)           120
  Max Frame Age:  (S5)            1

```

```

LECID Range Minimum:          1001
LECID Range Maximum:          2000
Validate Best Effort Peak Cell Rate (PCR): No
Control Distribute VCC Traffic Type: Best Effort VCC
Control Distribute VCC PCR in Kbps: 155000
Control Direct VCC Max Reserved Bandwidth: 0
Multicast Forward VCC Traffic Type: Best Effort VCC
Multicast Forward VCC PCR in Kbps: 155000
Multicast Send VCC MAX Reserved Bandwidth: 0

-LES-BUS Options-
BUS Mode:                      Adapter
Security (LECS Validation of Joins): Enabled
Partition LE_ARP_REQUEST Forwarding Domain: Yes
LE_ARP RESPONSE Destination:  One client
Partition Unicast Frame Domain: Yes
Redundancy:                    Enabled
Redundancy Role:               Primary LES-BUS
ATM address of Backup LES-BUS: 390985111111111111111010140008210000203
ATM address trace filter value: 0000000000000000000000000000000000000000
ATM address trace filter mask: FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF

-BUS Monitor Configuration-
Monitor Host Usage of BUS:      Disabled
# Top Hosts to Record:         10
# Seconds in each sample interval: 10
# Minutes between sample intervals: 30
Frame sampling rate:           1 out of 10

-Broadcast Manager Configuration-
IP BCM:                        Disabled
IPX BCM:                       Disabled
NetBIOS BCM:                   Disabled
BCM IP Cache Aging Time:       5
BCM IPX Cache Aging Time:      3
BCM NetBIOS Cache Aging Time:  15
BCM IPX Maximum Forwarding List: 50
No BCM IPX Static Entries defined

```

8 Display the LE clients that have joined ELAN_GW.

```

MSS Server 1 EXISTING LES-BUS 'ELAN_GW'+da 1i all lec
Number of LEC's to display: 4

```

```

      LEC-LES and LEC-BUS State (UP=Up, ID=Idle, --. --.
      **=Other; Show specific LEC to see actual)      v  v
      LEC      State #ATM #Reg #Lrnd
      Primary ATM Address Proxy ID LES BUS Adrs MACs MACs
-----
390985111111111111111111111010140008210000006 N 03E9 UP UP 1 2 0
390985111111111111111111111010140008210000005 Y 03EA UP UP 1 1 0
3909851111111111111111111110101400082100000206 N 03EB UP UP 1 1 0
3909851111111111111111111110101400082100000205 Y 03EC UP UP 1 1 0

```

9 Display the MAC addresses registered on ELAN_GW.

```

MSS Server 1 EXISTING LES-BUS 'ELAN_GW'+da 1i all reg
Number of Registered MAC's to display: 5

```

```

Registered
MAC Address   Registering ATM Address   Type   LEC   # BCM
-----
400082100000 390985111111111111111111111010140008210000006 R 03E9 0
400082100105 390985111111111111111111111010140008210000005 R 03EA 0
400082100106 390985111111111111111111111010140008210000006 R 03E9 0
400082100205 3909851111111111111111111110101400082100000205 R 03EC 0
400082100206 3909851111111111111111111110101400082100000206 R 03EB 0
MSS Server 1 EXISTING LES-BUS 'ELAN_GW'+exit
MSS Server 1 LE-SERVICES+exit
MSS Server 1 ATM+exit

```

Note: The MAC address for the ELAN IP gateway is registered by the active gateway (MSS 1 in our display).

10 Display the bridging status.

```
MSS Server 1 +pro asrt
MSS Server 1 ASRT>li bridge
Bridge ID (prio/add):      16384/40-00-82-10-01-04
Bridge state:              Enabled
UB-Encapsulation:         Disabled
Bridge type:               STB
Number of ports:           2
STP Participation:         IEEE802.1d
```

Port	Interface	State	MAC Address	Modes	Maximum MSDU	Segment	Flags
1	Eth/0	Up	40-00-82-10-01-04	T	1520		RD
2	Eth/1	Up	40-00-82-10-01-05	T	1520		RD

Flags: RE = IBMRT PC behavior Enabled, RD = IBMRT PC behavior Disabled

```
SR bridge number:      1
SR virtual segment:    001
Adaptive segment:      000
```

11 Verify that the shortcut bridging (SuperELAN) is active.

```
MSS Server 1 ASRT>li super
Number of Super ELAN bridge cache entries: 0
ATM Address                      Age Port Transid      Lecid CtrlFrame
-----
3909851111111111111111010140008210000006 180    2          0 0000      N/A
```

MSS Server 1 ASRT>exit

12 Display the status of the Classical IP connections.

```
MSS Server 1 +pro arp
MSS Server 1 ARP>di 0 ip
Active Channel List : Net 0
P/S  FLAGS LIST  VPI/VCI  FwdPcr  FwdScr  MaxSDUz Control P2P
1) S   80   01   0/892   155000000 155000000 9188    T    T
    Tgt Addr: 39.09.85.11.11.11.11.11.11.11.01.01.60.00.82.10.00.00.10
    Client Address (owner): 172.17.0.4
    Target Protocol Addresses: 172.17.0.2
2) S   00   01   0/323   155000000 155000000 9188    F    T
    Tgt Addr: 39.09.85.11.11.11.11.11.11.11.01.01.40.00.00.60.00.01.01
    Client Address (owner): 172.17.0.4
    Target Protocol Addresses: 172.17.0.3
New Channel List : Net 0
PVC Channel List : Net 0
```

13 Display the IP interfaces.

```
MSS Server 1 +pro ip
MSS Server 1 IP>int
```

Interface	IP Address(es)	Mask(s)
ATM/0	172.17.0.4	255.255.0.0
Eth/2	172.16.0.4	255.255.0.0

14 Display the status of the redundant gateways.

```
MSS Server 1 IP>red
Redundant Default IP Gateways for each interface:
inf 0 172.17.0.1 255.255.0.0 backup standby
inf 3 172.16.0.1 255.255.0.0 40.00.82.10.00.00 primary active
```

15 Verify the connectivity to a CIP-attached station (172.17.0.3).

```

MSS Server 1 IP>ping 172.17.0.3
PING 172.17.0.4 -> 172.17.0.3: 56 data bytes, ttl=64, every 1 sec.
56 data bytes from 172.17.0.3: icmp_seq=0. ttl=255. time=0. ms
56 data bytes from 172.17.0.3: icmp_seq=1. ttl=255. time=0. ms

----172.17.0.3 PING Statistics----
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max = 0/0/0 ms

```

16 Verify the connectivity to an ELAN-attached station (172.16.0.3).

```

MSS Server 1 IP>ping 172.16.0.3
PING 172.16.0.4 -> 172.16.0.3: 56 data bytes, ttl=64, every 1 sec.
56 data bytes from 172.16.0.3: icmp_seq=0. ttl=255. time=0. ms
56 data bytes from 172.16.0.3: icmp_seq=1. ttl=255. time=0. ms

----172.16.0.3 PING Statistics----
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max = 0/0/0 ms

```

17.7.5.2 MSS Server 2

1 Log in to MSS Server 2.

```

login: admin
Password: ibm8210

```

Copyright Notices:

```

Licensed Materials - Property of IBM
Multiprotocol Switched Services
(C) Copyright IBM Corp. 1996, 1997
All Rights Reserved. US Gov. Users Restricted Rights -
Use, duplication or disclosure restricted
by GSA ADP Schedule Contract with IBM Corp.

```

MOS Operator Control

2 Enter GWCON and display the active configuration.

MSS Server 2 ***talk 5**

MSS Server 2 **+configuration**

```

IBM 8210 Nways Multiprotocol Switching Server
Host name: MSS Server 2
Version: 8210-MSS Feature 8706 V1 R1.1 PTF 0 RPQ 0   cc1_17b test-load

```

```

Num Name  Protocol
0  IP      DOD-IP
3  ARP     Address Resolution
11 SNMP   Simple Network Management Protocol
29 NHRP   Next Hop Routing Protocol

```

```

Num Name  Feature
2  MCF     MAC Filtering
6  QOS     Quality of Service

```

4 Networks:

Net	Interface	MAC/Data-Link	Hardware	State
0	ATM/0	ATM	CHARM ATM	Up
1	Eth/0	Ethernet/IEEE 802.3	CHARM ATM	Up
2	Eth/1	Ethernet/IEEE 802.3	CHARM ATM	Up
3	Eth/2	Ethernet/IEEE 802.3	CHARM ATM	Up

Note: Make sure you are using the MSS code that supports bridging and redundant default IP gateway for LAN emulation, concurrently.

3 Display the ATM addresses being used by MSS Server 2.

```
MSS Server 2 +network 0
ATM Console
MSS Server 2 ATM+interface 0
ATM Interface Console
MSS Server 2 ATM Interface+list address
```

Network Prefix	ATM Address	ESI	SEL
39.09.85.11.11.11.11.11.11.11.01.01.40.00.82.10.00.02.05			
39.09.85.11.11.11.11.11.11.11.01.01.40.00.82.10.00.02.06			
39.09.85.11.11.11.11.11.11.11.01.01.60.00.82.10.00.00.10			
39.09.85.11.11.11.11.11.11.11.01.01.40.00.82.10.00.02.04			
39.09.85.11.11.11.11.11.11.11.01.01.40.00.82.10.00.02.A0			
39.09.85.11.11.11.11.11.11.11.01.01.40.00.82.10.00.02.02			
39.09.85.11.11.11.11.11.11.11.01.01.40.00.82.10.00.02.03			
39.09.85.11.11.11.11.11.11.11.01.01.40.00.82.10.00.02.01			
39.09.85.11.11.11.11.11.11.11.01.01.40.00.82.10.00.02.00			

```
MSS Server 2 ATM Interface+exit
```

4 Display the status of the Ethernet LES/BUS pairs.

```
MSS Server 2 ATM+le-s
LE-Services Console
MSS Server 2 LE-SERVICES+li all
ELAN Type (E=Ethernet/802.3, T=Token Ring/802.5)
Interface #
LES-BUS State (UP=Up, ID=Idle, ND=Net Down, ER=Error/Down,
**=Other; Work with specific LES-BUS to see actual state)
ELAN Name LES ATM Addr
-----
E 0 UP ELAN_E1 390985111111111111111111010140008210000202
E 0 UP ELAN_GW 390985111111111111111111010140008210000203
```

5 Display the status of the LES/BUS for ELAN_E1.

```
MSS Server 2 LE-SERVICES+work ELAN_E1
LE-Services Console for an existing LES-BUS Pair
ELAN Name: ELAN_E1
ELAN Type: Ethernet
ATM Device number: 0
# of Proxy LEC's: 1
# of Non-Proxy LEC's: 1
LES ATM Address: 390985111111111111111111010140008210000202

-Status-
LES-BUS State: OPERATIONAL
Redundancy VCC State: IDLE
Major Reason LES-BUS was last Down: none
Minor Reason LES-BUS was last Down: none
LES-BUS State last changed at: 00.00.04.66 (System Up Time)
LES-LEC Status Table changed at: 00.10.38.11 (System Up Time)
BUS-LEC Status Table changed at: 00.10.38.21 (System Up Time)
UNI Version: 3.0
IP BCM: INACTIVE
IPX BCM: INACTIVE
NetBIOS BCM: INACTIVE

-Current Configuration-
LES-BUS Enabled/Disabled: Enabled
```



```
MSS Server 2 EXISTING LES-BUS 'ELAN_E1'+exit
MSS Server 2 LE-SERVICES+exit
```

8 Display the status of the LES/BUS for ELAN_GW.

```

MSS Server 2 LE-SERVICES+work ELAN_GW
LE-Services Console for an existing LES-BUS Pair
ELAN Name:                ELAN_GW
  ELAN Type:                Ethernet
  ATM Device number:        0
  # of Proxy LEC's:         0
  # of Non-Proxy LEC's:     0
  LES ATM Address:          390985111111111111111010140008210000203

```

```
-Status-
LES-BUS State: OPERATIONAL
Redundancy VCC State: ESTABLISHED
Major Reason LES-BUS was last Down: none
Minor Reason LES-BUS was last Down: none
LES-BUS State last changed at: 00.00.04.66 (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)
UNI Version: 3.0
IP BCM: INACTIVE
IPX BCM: INACTIVE
NetBIOS BCM: INACTIVE
```

```

-Current Configuration-
LES-BUS Enabled/Disabled:      Enabled
ATM Device number:              0
End System Identifier (ESI):    40.00.82.10.00.02
Selector Byte:                  0x03
ELAN Type:                      (S2) Ethernet
Max Frame Size:                 (S3) 1516
Control Timeout:                (S4) 120
Max Frame Age:                  (S5) 1
LECID Range Minimum:            1001
LECID Range Maximum:            2000
Validate Best Effort Peak Cell Rate (PCR): No
Control Distribute VCC Traffic Type: Best Effort VCC
Control Distribute VCC PCR in Kbps: 155000
Control Direct VCC Max Reserved Bandwidth: 0
Multicast Forward VCC Traffic Type: Best Effort VCC
Multicast Forward VCC PCR in Kbps: 155000
Multicast Send VCC MAX Reserved Bandwidth: 0

```

```
-LES-BUS Options-  
BUS Mode: Adapter  
Security (LECS Validation of Joins): Enabled  
Partition LE_ARP_REQUEST Forwarding Domain: Yes  
LE_ARP_RESPONSE Destination: One client  
Partition Unicast Frame Domain: Yes  
Redundancy: Enabled  
Redundancy Role: Backup LES-BUS  
ATM address trace filter value: 000000000000000000000000000000000000  
ATM address trace filter mask: FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
```

```

-BUS Monitor Configuration-
Monitor Host Usage of BUS:      Disabled
# Top Hosts to Record:         10
# Seconds in each sample interval: 10
# Minutes between sample intervals: 30
Frame sampling rate:            1 out of 10

```

```
-Broadcast Manager Configuration-
IP BCM: Disabled
IPX BCM: Disabled
NetBIOS BCM: Disabled
BCM IP Cache Aging Time: 5
BCM IPX Cache Aging Time: 3
BCM NetBIOS Cache Aging Time: 15
BCM IPX Maximum Forwarding List: 50
No BCM IPX Static Entries defined
```


Note: Because of the redundancy VCC, the LES for ELAN_GW on MSS 2 will not accept any LE clients.

9 Display the LE clients that have joined ELAN_GW.

```
MSS Server 2 EXISTING LES-BUS 'ELAN_GW'+da 1i all lec
Entry was not found
MSS Server 1 EXISTING LES-BUS 'ELAN_GW'+exit
MSS Server 1 LE-SERVICES+exit
```

10 Display the bridging status.

```
MSS Server 2 ASRT>1i bridge
Bridge ID (prio/add):      16384/40-00-82-10-02-04
Bridge state:              Enabled
UB-Encapsulation:         Disabled
Bridge type:               STB
Number of ports:           2
STP Participation:         IEEE802.1d
```

Port	Interface	State	MAC Address	Modes	Maximum MSDU	Segment	Flags
1	Eth/0	Up	40-00-82-10-02-04	T	1520		RD
2	Eth/1	Up	40-00-82-10-02-05	T	1520		RD

Flags: RE = IBMRT PC behavior Enabled, RD = IBMRT PC behavior Disabled

```
SR bridge number:      1
SR virtual segment:    001
Adaptive segment:      000
```

11 Display the status of the shortcut (SuperELAN) bridging ports.

```
MSS Server 2 ASRT>1i super
Number of Super ELAN bridge cache entries: 0
ATM Address              Age Port Transid      Lecid CtrlFrame
-----
3909851111111111111110103020077770000081  45   1         0 0000      N/A
3909851111111111111110101400082100000006 120   2         0 0000      N/A
```

MSS Server 2 ASRT>exit

12 Display the status of the Classical IP connections.

```
MSS Server 2 +pro arp
MSS Server 2 ARP>di 0 ip
Active Channel List : Net 0
P/S  FLAGS LIST  VPI/VCI  FwdPcr  FwdScr  MaxSDUsz Control P2P
1) S   00   01  0/779  154999984 154999984 9188      F    T
   Tgt Addr: 39.09.85.11.11.11.11.11.11.11.11.11.11.11.11.11.01.01.40.00.82.10.00.00.10
   Client Address (owner): 172.17.0.2
   Target Protocol Addresses: 172.17.0.4
2) S   00   01  0/931  155000000 155000000 9188      F    T
   Tgt Addr: 39.09.85.11.11.11.11.11.11.11.11.11.11.11.11.11.01.01.40.00.00.60.00.01.01
   Client Address (owner): 172.17.0.2
   Target Protocol Addresses: 172.17.0.3
3) S   00   01  0/44   155000000 155000000 9188      F    T
   Tgt Addr: 39.09.85.11.11.11.11.11.11.11.11.11.11.11.11.11.01.01.40.00.00.60.00.01.01
   Client Address (owner): 172.17.0.2
   Target Protocol Addresses:
New Channel List : Net 0
PVC Channel List : Net 0
MSS Server 2 ARP>exit
```

Note: Because of prior IP traffic, VCCs have been established to the ARP server/client on MSS 2.

13 Display the IP interfaces.

```
MSS Server 2 +pro ip
MSS Server 2 IP>int
```

Interface	IP Address(es)	Mask(s)
ATM/0	172.17.0.2	255.255.0.0
Eth/2	172.16.0.2	255.255.0.0

14 Display the status of the redundant gateways.

```
MSS Server 2 IP>red
Redundant Default IP Gateways for each interface:
  inf 0 172.17.0.1      255.255.0.0      primary active
  inf 3 172.16.0.1      255.255.0.0      40.00.82.10.00.00 backup standby
```

15 Verify the connectivity to the CIP-attached station (172.17.0.3).

```
MSS Server 2 IP>ping 172.17.0.3
PING 172.17.0.2 -> 172.17.0.3: 56 data bytes, ttl=64, every 1 sec.
56 data bytes from 172.17.0.3: icmp_seq=0. ttl=255. time=0. ms
56 data bytes from 172.17.0.3: icmp_seq=1. ttl=255. time=0. ms

----172.17.0.3 PING Statistics----
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max = 0/0/0 ms
```

16 Verify the connectivity to the ELAN-attached station (172.16.0.3).

```
MSS Server 2 IP>ping 172.16.0.3
PING 172.16.0.2 -> 172.16.0.3: 56 data bytes, ttl=64, every 1 sec.
56 data bytes from 172.16.0.3: icmp_seq=0. ttl=255. time=40. ms
56 data bytes from 172.16.0.3: icmp_seq=1. ttl=255. time=0. ms

----172.16.0.3 PING Statistics----
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max = 0/20/40 ms
```

17.7.5.3 External Equipment

The default gateway IP addresses on MSS can be used for routing, to ping to MSS, and to Telnet into MSS. Traffic generated by MSS itself does not, in general, use these addresses as can be seen from the following displays:

- Traceroute from (OS/2) ELAN-attached station to the CIP-attached station.

```
c:\tracerte 172.17.0.3
0  172.16.0.4 (172.16.0.4)  0 ms  0 ms  0 ms
1  172.16.0.4 (172.16.0.4)  0 ms  32 ms  0 ms
2  172.17.0.3 (172.17.0.3)  0 ms  0 ms  0 ms
```

Note: Although the default GW is 172.16.0.1, MSS 1 responds to the traceroute message using IP address 172.16.0.4.

- Traceroute from the (RS/6000) CIP-attached station to an ELAN-attached station.

```
traceroute 172.16.0.3
traceroute to 172.16.0.3 (172.16.0.3), 30 hops max, 40 bytes packets
1  172.17.0.2 (172.17.0.2)  66 ms  22 ms  3 ms
2  172.16.0.3 (172.16.0.3)  4 ms  4 ms  4 ms
```

Note: Although the default GW is 172.17.0.1, MSS 2 responds to the traceroute message using IP address 172.17.0.2.

Chapter 18. Configuration Examples For MSS Version 2.0

This chapter provides configuration examples for the new or enhanced functions in MSS Version 2.0. Where appropriate, examples from the previous MSS V1.1 chapter will be used or referenced. Where possible, examples are shown using the MSS configuration program. MSS V2.0 is generally compatible with V1.1 (except for Super-ELAN) and is a superset of MSS V1.1 functions. We recommend the reading of V1.1 configuration examples, even if V2.0 is being implemented.

As in the previous chapter, the configuration examples included in this section will help you to simplify the process of getting your MSS Server operational or adding new functions for enhanced redundancy and performance

MSS Version 2.0

All configuration examples in this chapter assume MSS Version 2.0 unless otherwise noted. Make sure that your MSS Server and MSS Configuration Program are at the right software level.

The MSS configuration files are provided to detail how each function is defined. Although they can be used and modified for production, we recommend using them as a template and creating a new file. This can be done by starting two copies of the MSS configuration program and working with both the example and new configurations. This is especially useful when using a migrated configuration from V1.1. The files may be downloaded (see <ftp://www.redbooks.ibm.com/redbooks/SG242115/49xx.cdb>) from the Internet. The V2.0 configuration examples in this chapter are grouped in a configuration database, named MSSV2 in the same Internet Web page.

The most important part of migrating to MSS V2.0 is ensuring that all previous definitions (and functions) remain operational. The MSS configuration program must be used to initially migrate an MSS V1.1 configuration file. This can be done by opening the MSS V1.1 CDB using the V2.0 configuration program and loading the old file as shown below.

1. Start the MSS 2.0 configuration program.
2. Click the **Configure** action item (top left of navigation window).
3. Click **Open**, click on the database name and click on the configuration to be migrated.
4. Click **OK**. A window will pop up stating the configuration is an older version. Click **OK**.
5. After the config loads, click the **Configure** action item again.
6. Click **Save as** and save the newly converted file as a new name in a new database (CDB file-type).

The MSS and 8260 must be prepared and be at appropriate code levels in order to implement MSS V2.0 as shown in these configuration examples. The steps to perform this are outlined in 17.1, "Preparation" on page 138. Note that running MSS V2.0 requires additional memory in existing MSSs, and new software, both of which must be ordered for a fee from IBM. Instructions for the upgrade from MSS V1.1 to V2.0 are included with the delivered software, and also available

from the IBM Web site <http://www.networking.ibm.com/nes/nesswitc.htm>. To obtain the code or instructions from the Web however, requires a user ID and password which are included with the delivered V2.0 software. There are also a user ID and a password on ATM Corner for IBM internal use only.

The following configuration examples have been included:

- RouteSwitching (see 18.1, "Example - Zero-Hop IP Routing (RouteSwitching)")

RouteSwitching provides broadcast reduction for large IP networks while avoiding extra routing hops. Zero-hop routing, or RouteSwitching allows for no router hops across subnets. This configuration example gets you started when you want to use RouteSwitching.

- Distributed ATMARP Server and Redundancy (see 18.3, "Example - Distributed ATMARP Servers in a LIS" on page 213)

Mission-critical applications require the elimination of single points of failure. This example shows you how to provide distributed and redundant ATMARP servers to support current and newer ATM ARP clients.

- Multicasting in Classical IP over ATM (see 18.2, "Example - IP Multicasting in a LIS Using MARS" on page 208)

To support Classical IP with needed functions such as router protocols, DHCP and IP Multicast requires MARS. This configuration example shows how to use MSS to supply the needed functions.

- APPN over ATM Using LAN Emulation (see 18.4, "APPN Configuration" on page 218)

Networks with SNA and IP are being tightly integrated, yet the requirements for performance and throughput are increasing. This configuration example details supporting APPN HPR on an ATM network.

- Bridging Using RFC 1483 and SVCs (see 18.6, "Example - RFC 1483 SVC Bridging" on page 220)

When migrating from previous ATM networks, or when interoperating with frame relay networks, it may be required to support RFC 1483 bridge links. The new MSS capability of using SVCs in addition to PVCs for RFC 1483 bridging allows for more flexibility and potentially easier configuration across multiple ATM switches.

Depending on the configuration example you need one or two MSS Servers. Two MSS servers are required for the Super-ELAN shortcut bridging and Classical IP redundancy examples; other examples require one MSS Server.

18.1 Example - Zero-Hop IP Routing (RouteSwitching)

This example highlights the RouteSwitching capability from a LAN-based workstation, allowing broadcast protection and communications between subnets without a router hop in the network. It differs from the previous MSS 1.1 method of zero-hop routing since Super-ELANs are not required, and no reconfiguration of subnet masks (to make them span subnets) is needed. For details and client requirements, see 11.2, "MSS 2.0 Zero-Hop Route Server (RouteSwitching)" on page 95. Note that this example requires the RouteSwitch drivers available from IBM on LAN stations which will be gaining zero-hop routes.

18.1.1 Configuration Details

To differentiate the new 2.0 RouteSwitching capability, the example configuration will be two ELANs, each with one subnet, with only IP routing used to connect them (no Super-ELAN). This is shown in Figure 80.

Since this configuration is for example, there are IP no routing protocols or default gateway defined.

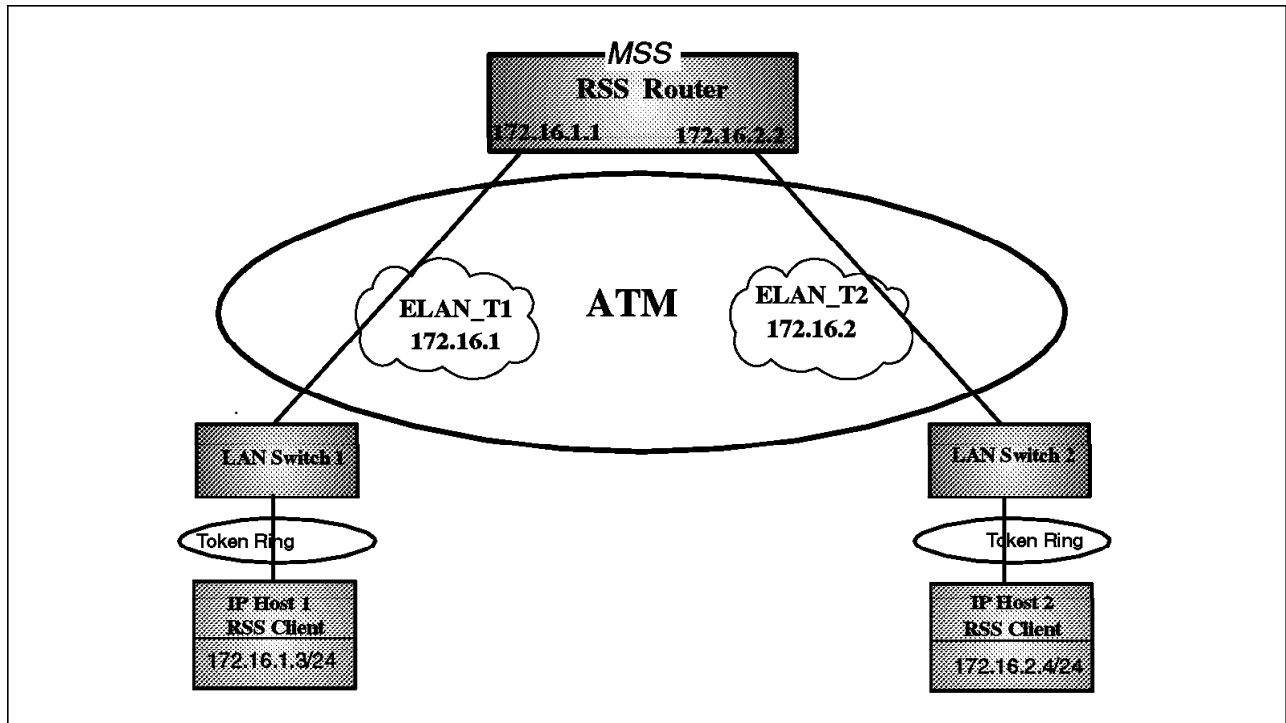


Figure 80. Zero-Hop Routing IP Connectivity

Figure 81 on page 204 shows the configuration details of the zero-hop RouteSwitching server example.

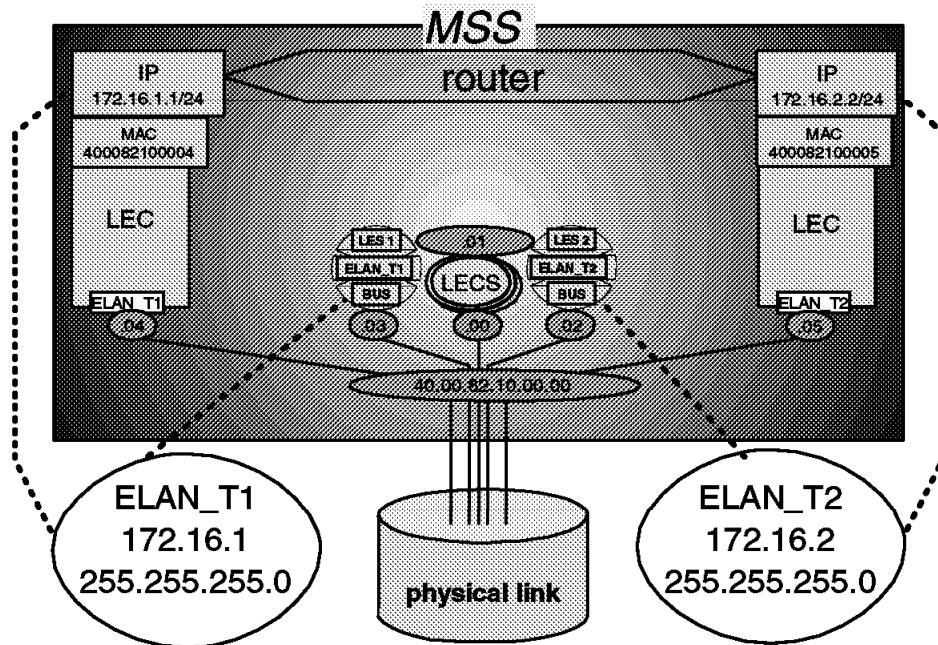


Figure 81. Zero-Hop Routing MSS Configuration

The following configuration is in database MSSV2.CDB, file ZHOP1. To review these definitions, we recommend starting the MSS Version 2.0 configuration program, opening the configuration file, and viewing the definitions as described below.

A single ESI (40.00.82.10.00.00) is used to address all functions on the MSS Server. Each of the functional elements is defined with a different selector byte.

The selectors used are:

- LECS - Selector 00

LECS is used to assign the two token-ring LAN switches to the appropriate ELAN (ELAN_T1 or ELAN_T2) using ELAN Name. The LECS policies are defined in NW;Network Device/LAN_Emulation/LECS/LECS Assignment Policies (which is the path in the configuration program navigation window). As defined, a LAN switch that is not configured to use LECS and give the ELAN name as a value will not be able to join an ELAN.

- LECS/LES Security Interface - Selector 01

LECS/LES Security ensures that clients do not join a LES without first meeting the policies as defined in the LECS for that ELAN. The settings can be reviewed in NW;Network Device/LAN_Emulation/ELANs/ELANs;Local-LES-Policy-Values/ELAN Name.

- LES/BUS pairs - Selector 02 and 03

We have defined two token-ring ELANs, ELAN_T1 and ELAN_T2. The LES/BUS uses Selector 02 for ELAN_T1 and 03 for ELAN_T2. Parameters can be reviewed in NW;Network Device/LAN_Emulation/ELANs/ELANs;Local-LESBUS/General-1.

The LES for ELAN_T1 will assign LEC IDs ranging from 1 to 1000. The LES for ELAN_T2 will assign LEC IDs ranging from 1001 to 2000. This ensures that if a cut-through path includes a "paranoid LEC", that there is no chance

that the two LECs have the same LECID. These parameters can be reviewed in NW;Network Device/LAN_Emulation/ELANs/ELANs;Local-LESBUS/General-2.

- LE Clients - Selector 04 and 05

Two LECs are defined for IP routing, one on each ELAN. The routing client on ELAN_T1 uses selector 4 and the client on ELAN_T2 uses selector 5. The parameters can be reviewed in NW;Network Device/LAN_Emulation/ELANs/LEC Interfaces/LEC Interfaces. An IP address has been defined on each to enable IP routing, with 172.16.1.1 for LEC interface #1 on ELAN_T1, and 172.2.2 for LEC interface #2 on ELAN_T2. The parameters can be reviewed in NW;Network Device/Protocols/IP/Interfaces.

Administrative user admin has been added using password ibm8210.

The configuration now needs to be defined for RouteSwitching.

Since zero-hop routing is based on NHRP, NHRP must be enabled as seen in NW;Network Device/Protocols/IP/NHRP/General. The defaults here are good to start with, although it is recommended to enable the NHRP check box. Since each interface defaults to using the box-level defaults, they don't need to be changed. These can be seen in NW;Network Device/Protocols/IP/NHRP/Interfaces.

There are some design considerations when enabling NHRP with routers having multiple paths (or back doors) to ATM-based IP subnets. If a shortcut path is achieved where information used in route selection is lost, a routing loop could occur. This is possible where the back door path exists with routers using multiple routing protocols with different metrics. For these (less likely) scenarios, MSS can be configured to disallow shortcut paths to specific routers or subnets. These can be configured in NW;Network Device/Protocols/IP/NHRP/Disallow Rtr-Rtr Short Cuts.

With NHRP configured, RouteSwitching can be enabled by checking the Route-Switching check box. Select the proper end-to-end mode of the network (routing or bridging). See 11.2.3.4, "Other Networking Considerations" on page 100 for an explanation of this issue. For token-ring ELANs with source route bridging, a range of segment numbers are needed for virtual route descriptors. These are used for registering zero-hop routes with the LES/BUS associated with the ingress RSS. The default value is 3840 to 4095 (range of 255). These numbers are in decimal, so a hex/decimal calculator is useful. Thus the default reserves ring numbers from 0xF00 to 0xFFFF. These values can be reviewed in NW;Network Device/Protocols/IP/NHRP/Route-Switching.

18.1.2 Verifying Zero-Hop Routing

The next section shows how we verified proper operation of the configuration. Note that some of the information displayed is configuration-dependent.

1. Log in to the MSS Server.

Login:

Password: ibm8210

Copyright Notices:

Licensed Materials - Property of IBM
Multiprotocol Switched Services

(C) Copyright IBM Corp. 1996, 1997
 All Rights Reserved. US Gov. Users Restricted Rights -
 Use, duplication or disclosure restricted
 by GSA ADP Schedule Contract with IBM Corp.

MOS Operator Control

2. Enter GWCON and display the active configuration.

MSS #1 Test ***talk 5**

MSS #1 Test **+configuration**

IBM 8210 Nways Multiprotocol Switching Server

Host name: MSS #1 Test

Version: 8210-MSS Feature 8707 V2.0 Mod 0 PTF 0 RPQ 0 cc2_18c

Num	Name	Protocol
0	IP	DOD-IP
3	ARP	Address Resolution
11	SNMP	Simple Network Management Protocol
29	NHRP	Next Hop Routing Protocol

Num	Name	Feature
2	MCF	MAC Filtering
6	QOS	Quality of Service

4 Networks:

Net	Interface	MAC/Data-Link	Hardware	State
0	ATM/0	ATM	ATM	Up
1	TKR/0	Token-Ring/802.5	ATM	Up
2	TKR/1	Token-Ring/802.5	ATM	Up
3	NHRPL/0	NHRP LANE Shortcut	ATM	Up

MSS #1 Test +

3. Display the ATM addresses being used by the MSS Server.

MSS #1 Test + **network 0**

ATM Console

MSS #1 Test ATM**+interface**

ATM Interface Console

MSS #1 Test ATM Interface**+list address**

ATM Address	Name
3911FF2299999900000000CCAA002035993200C9	
3911FF2299999900000000CCAA40008210000004	LEC 1 'ELAN_T1'
3911FF2299999900000000CCAA40008210000005	LEC 2 'ELAN_T2'
3911FF2299999900000000CCAA40008210000002	LES/BUS 'ELAN_T1'
3911FF2299999900000000CCAA40008210000003	LES/BUS 'ELAN_T2'
3911FF2299999900000000CCAA40008210000001	
3911FF2299999900000000CCAA40008210000000	
3911FF2299999900000000CCAA002035993200C8	

4. Display the status of the Ethernet LES/BUS pairs.

MSS #1 Test ATM Interface**+exit**

MSS #1 Test ATM**+le-s**

MSS #1 Test LE-SERVICES**+list**

ELAN Type (E=Ethernet/802.3, T=Token Ring/802.5)

Interface #	LES-BUS State	LES ATM Addr
	(UP=Up, ID=Idle, ND=Net Down, ER=Error/Down, **=Other; Work with specific LES-BUS to see actual state)	
	ELAN Name	

5. Display the status of ELAN_T1.

```
MSS #1 Test ATM+work ELAN_T1
MSS #1 Test LE-SERVICES+work ELAN_T1
```

Number of LEC's to display: 2

6. Display the status of ELAN_T2.

Number of LEC's to display: 2

7. Display the IP interfaces and routes.

MSS #1 Test IP>dump

```
Routing table size: 768 nets (49152 bytes), 3 nets known
                   0 nets hidden, 0 nets deleted, 1 nets inactive
                   0 routes used internally, 764 routes free
```

8. Ping from one LAN client (172.16.1.3) to another LAN client (172.16.2.4).

Token-ring drivers (shims) are available for Windows 16-bit ODI and Ethernet at <http://www.networking.ibm.com/nes/nesswite.htm#8210rsct>. For other RouteSwitch client drivers and documentation (PCI NICs): <ftp://ftp.networking.ibm.com/pub/products/lanprods/switch/nestoken.html> and look under PCI NICs.

The following commands were done from a Windows95 station with address 172.16.1.3.

```
c:\windows>tracert 172.16.2.4
Tracing route to 172.16.2.4 over a maximum of 30 hops
```

1	1 ms	1 ms	1 ms	172.16.1.1
2	4 ms	3 ms	3 ms	172.16.2.4

Trace complete.

```
c:\windows>FTP 172.16.2.4
Connected to 172.16.2.4
220 phlp75 FTP server (IBM OS/2 TCP/IP FTP Version 1.2) ready.
User (172.16.2.4:(none)):dmindel
331 Password required for dmindel.
Password:
230 User dmindel logged in.
ftp>bin
200 Type set to I.
ftp>cd e:\ftp
250 CWD command successful.
ftp>get test.fil
200 PORT command successful.
150 Opening BINARY mode data connection for test.fil (1773230 bytes).
1773230 bytes received in 8.23 seconds (215.46 Kbytes/sec)
ftp>quit
221 Quit command received. Goodbye.
c:\windows>tracert 172.16.2.4
```

Tracing route to 172.16.2.4 over a maximum of 30 hops

1	1 ms	1 ms	1 ms	172.16.2.4
---	------	------	------	------------

Trace complete.

Ensure that you have Windows installation CD or disks when upgrading the PCI token-ring drivers. We performed the driver upgrade by removing protocol bindings and current device driver, installing the new driver, then adding the original protocol bindings to the new driver.

18.2 Example - IP Multicasting in a LIS Using MARS

The ability to support multicasts and broadcasts in Classical IP subnets will maximize its usability. The following example is in configuration file MARS MSS#1, which uses the ZHOP1 configuration as a base.

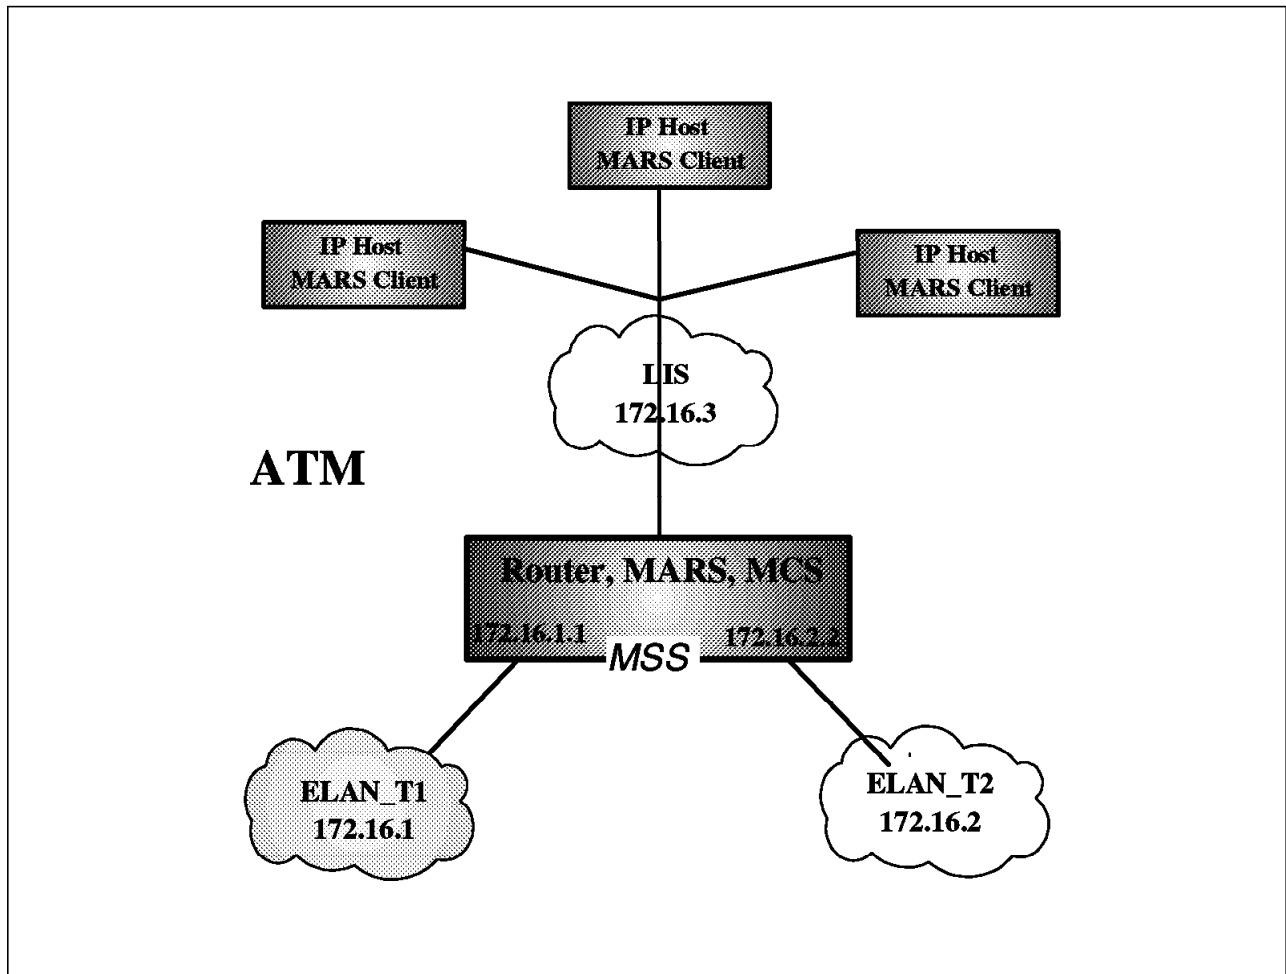


Figure 82. MARS/MCS and IP Routing Connectivity

Figure 82 Shows IP routing capability and highlights broadcasting and multicasting in a Classical IP subnet.

In this example, there is an additional ESI defined. MSS supports multiple ESIs on the same ATM interface. Figure 83 on page 210 details the MSS configuration including the previous example's ELAN and LEC definitions. Note that there is a single physical link, although it is represented in the diagram as two pipes.

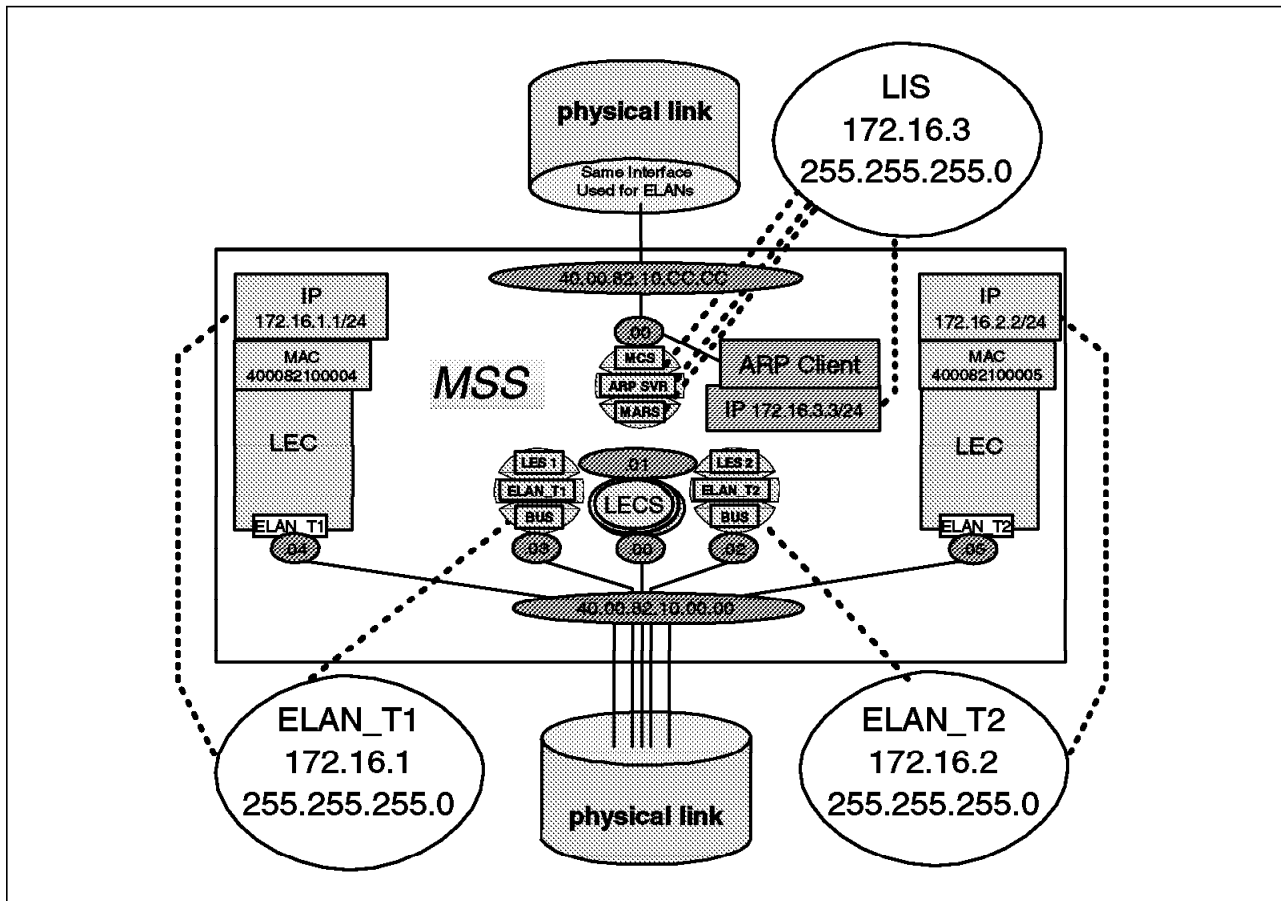


Figure 83. MSS Configuration with MARS/MCS and IP Routing

18.2.1 Setting Up Classical IP via the Configurator

We will be using a different ESI (40008210CCCC). This is defined in configuration MSSV2.cdb, MARS MSS #1. The parameters can be viewed in NW;Network Device/Interfaces:Interface 0, Configure.

We need to add an ATMARPs server and client function with an associated IP address. We define a new IP address (172.16.3.3). The parameters can be viewed in NW;Network Device/Protocols/IP/Interfaces:Interface 0, IP Addresses.

We then configure Classical IP with this MSS as the ATMARPs server. This can be viewed in NW;Network Device/Protocols/IP/Classical IP/Classical IP: Configure. Under the ARP Server tab, we select the **Client is also an ARP Server** box. Under the Client Address tab, we select the **40008210cccc** address with a selector of '00'.

Although this configuration is not using routing protocols; those that will may require additional definition. If the network uses OSPF, and IP Multicasting across subnets is desired, OSPF would be configured. In this case an additional configuration item (adding MOSPF) would be needed. This is done in NW;Network Device/Protocols/IP/OSPF/Multicast Forwarding.

18.2.2 Configuring MARS

Both the MARS and MCS definitions can currently only be done via console interface. This will be added to a future update to the MSS configurator.

1. Configure MARS.

```
t 6

Gateway user configuration
MSS #1 Test Config>p mars
MARS Configuration
MSS #1 Test MARS config>list all

MARS Server Configuration List:
    No MARS Servers have been added

MSS #1 Test MARS config>add mars-server private

Interface Number of MARS Server being created [0]? 0 .
Private NSAP Address: Specify unique 40 digits
ATM Address []? 39.11.ff.22.99.99.00.00.00.00.00.cc.aa.40.00.82.10.cc.cc.00
Validate PCR for best effort VCCs? [No]: no
Maximum Reserved Bandwidth for incoming VCCs (Kbps) [0]? 0
Use Best Effort Service for Control VCCs? [Yes]: yes
Peak Cell Rate of outbound control VCCs (Kbps) [0]? 0
Sustained Cell Rate of outbound control VCCs (Kbps) [0]? 0
Max SDU size (bytes) [9188]? 9188
Private Control VC Aging Timer (minutes) [20]? 20
Cluster Redirect Timer (seconds) [60]? 60
MARS Server Instance has been created

To enable MARS MCS support, select: change mcs-support
```

The interface chosen is that of Classical IP. The Private NSAP is the full ATM address of the MSS ATMARP Server. Since there is no QOS currently available, all bandwidth parameters have been set to '0'.

2. Configure MCS.

```
MSS #1 Test MARS config>change mcs

MARS Server Configuration List:
Mars Instance  Int.  ATM Address/Sub Address
[1]            [0]   39.11.FF.22.99.99.99.00.00.00.00.CC.AA.40.00.82.10.CC.CC.00

Enter MARS Server Instance Number to be changed
(Default of 0 will change nothing) [0]? 1
Enable MARS MCS Support? [No]: yes
Support Broadcast address? [No]: yes
Support Multicast address Range? [No]: yes
IP address of any device on the LIS [0.0.0.0]? 172.16.1.1
IP Subnet mask for the LIS [255.255.0.0]? 255.255.255.0
Starting group address range [224.0.0.0]? 224.0.0.0
Ending group address range [239.255.255.255]? 239.255.255.255
Send all multicast/broadcast traffic over CCVC? [No]: no
Use Best Effort Service for Data VCCs? [Yes]: yes
Peak Cell Rate of outbound Data VCCs (Kbps) [0]? 0
Sustained Cell Rate of outbound Data VCCs (Kbps) [0]? 0
Data VC Aging Timer (minutes) [20]? 20
MARS MCS Support has been changed.
```

The support for a broadcast address allows MSS to handle IP broadcasts that do not use an IP multicast (class D) address, such as RIP Version 1. The Multicast address range parameter allows MCS to handle IP multicasting. The parameters regarding IP address and subnet anchor the MCS to the proper IP subnet for broadcasts. Sending all multicast traffic over the ClusterControl VC (CCVC) instead of creating a PtMP VC for every

multicast group saves ATM and devices call setup resources. However, it may actually increase traffic, since every multicast will be sent to every ATMARP client, even if destined for only a few nodes.

3. Configure MARS client on MSS.

These commands add MARS capability to the ATMARP client by defining the MARS client and the ATM address of the MARS.

```
MSS #1 Test MARS config>ex

MSS #1 Test Config>p arp

ARP user configuration

MSS #1 Test ARP config>add multicast
Interface Number [0]? 0
Client IP Address [0.0.0.0]? 172.16.3.3
Configure multicast support? [No]: yes
Configure broadcast support? [No]: yes
Redirect Timer value (seconds) [90]? 90
Register Timer value (seconds) [10]? 10
Response Timer value (seconds) [5]? 5
Maximum number of consecutive missed Redirects [4]? 4
Maximum number of consecutive failed Register attempts [3]? 3
Maximum number of consecutive Response timer failures [3]? 3
Multicast VC Aging Timer (minutes) [20]? 20

MSS #1 Test ARP config>add mars private
Local Client IP Address [0.0.0.0]? 172.16.3.3
Private NSAP Address: Specify 40 digits
ATM Address []? 39.11.ff.22.99.99.99.00.00.00.00.cc.aa.40.00.82.10.cc.cc.00
Primary server? [Yes]: yes

MSS #1 Test ARP config>
```

18.2.3 Verifying MARS And MCS

The next section shows how we verified operation of the configuration. Note that some of the information displayed is configuration-dependent. At the time of testing, there were no external ATMARP clients with MARS support available, thus the verification is of configuration only. For details on monitoring MARS, see the IBM manual *MSS Configuring Protocols And Features* (SC30-3819), Chapter 27.

1. Log in to the MSS Server.

Login:
Password: ibm8210

Copyright Notices:

Licensed Materials - Property of IBM
Multiprotocol Switched Services
(C) Copyright IBM Corp. 1996, 1997
All Rights Reserved. US Gov. Users Restricted Rights -
Use, duplication or disclosure restricted
by GSA ADP Schedule Contract with IBM Corp.

MOS Operator Control

2. Enter GWCON and display the Global Statistics for MARS.

MSS #1 Test ***talk 5**

MSS #1 Test **+p mars**

MSS #1 Test MARS>**global info**

```
Global Status Flags: 0x0000      Global Error Flags: 0x0000
Global # Queue Overflows:          0
Global # Invalid Destinations:     0
MARS Server # Input Packet Overflows: 0
MARS Server # Queued Input  Packets: 0
MARS Server # Queued Output Packets: 0
MARS Client # Input Packet Overflows: 0
MARS Client # Queued Input  Packets: 0
MARS Client # Queued Output Packets: 0
MCS # Input Packet Overflows:      0
MCS # Queued Input  Packets:       0
MCS # Queued Output Packets:       0
```

3. Display definition and status for a specific MARS instance.

MSS #1 Test MARS>**instance**

Network interface [0]? **0**

```
Instance      ATM Address / Subaddress
[ 0]          39.11.FF.22.99.99.99.00.00.00.00.CC.AA.40.00.82.10.CC.CC.00
```

4. Display status of MARS or multicast groups.

MSS #1 Test MARS>**group server**

Network interface [0]?

```
Instance      ATM Address / Subaddress
[ 0]          39.11.FF.22.99.99.99.00.00.00.00.CC.AA.40.00.82.10.CC.CC.00
```

Instance number [0]?

For protocol 0x0800:

There are currently no groups.

The group command can also be used with the parameter clusters, which will list all multicast groups being served by MARS.

18.3 Example - Distributed ATMARP Servers in a LIS

This example shows two MSSs in a Classical IP environment supporting distributed and redundant ATMARP services. This allows RFC 1577+ clients access to a primary and backup ARP server regardless of location. For RFC 1577 clients, this allows for backup via a redundant ARP server that is attached to the same ATM switch as the primary.

18.3.1 Distributed and Redundant ATMARP Server Definition

The configuration files for this are Distributed ARP MSS#1 and Distributed ARP MSS#2 in MSS Configurator database mssv2.cdb. The configuration for MSS1 will be described.

The MSS 1 configuration includes the previous ELAN definitions, while MSS2 configuration contains only the distributed and the redundant ARP server definitions. Examples of redundant IP gateway and LES/BUS have been shown previously. Figure 84 on page 214 details the MSS configurations for this.

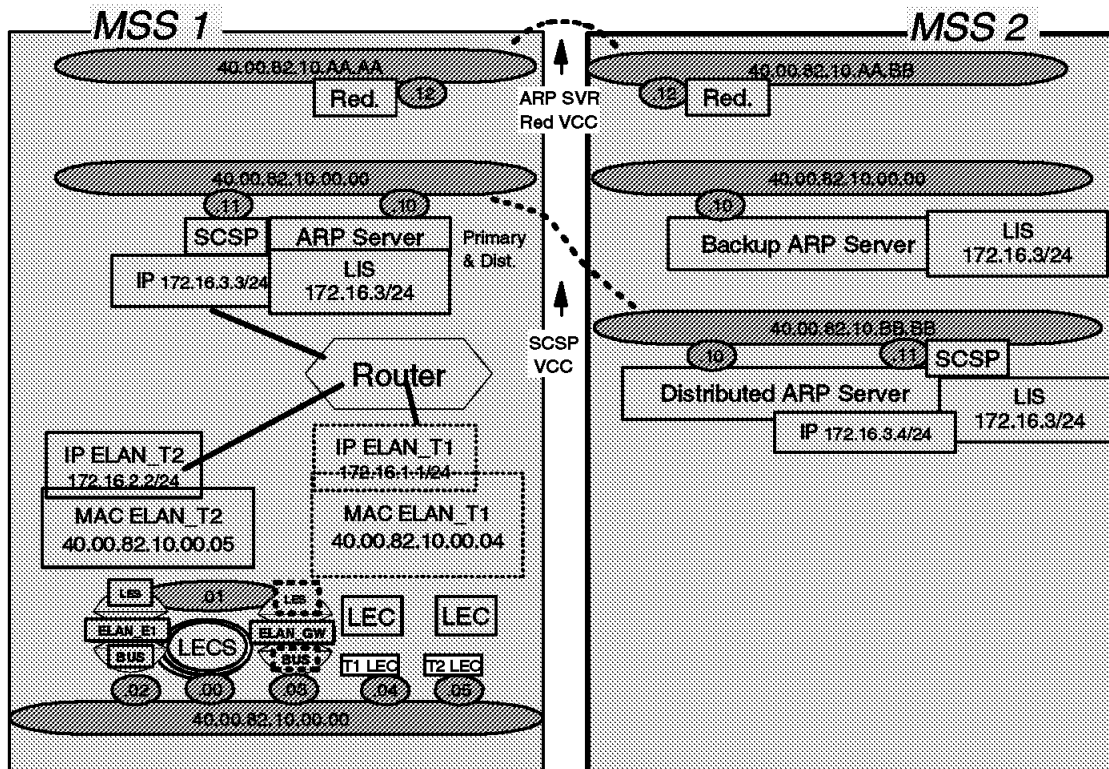


Figure 84. MSS Distributed and Redundant ATMARP Server Details

In Figure 85 on page 215, note that MSS #2 has two ATM interfaces, one attached to ATM switch 39...CC.AA and one attached to 39...CC.BB. This allows for RFC 1577+ clients attached to ATM switch 39...CC.BB to operate if ATM switch 39...CC.AA or MSS #1 becomes inoperative. This allows for RFC 1577 clients to operate if MSS #1 becomes inoperative. This figure shows the physical connections and address mapping between the two ARP servers.

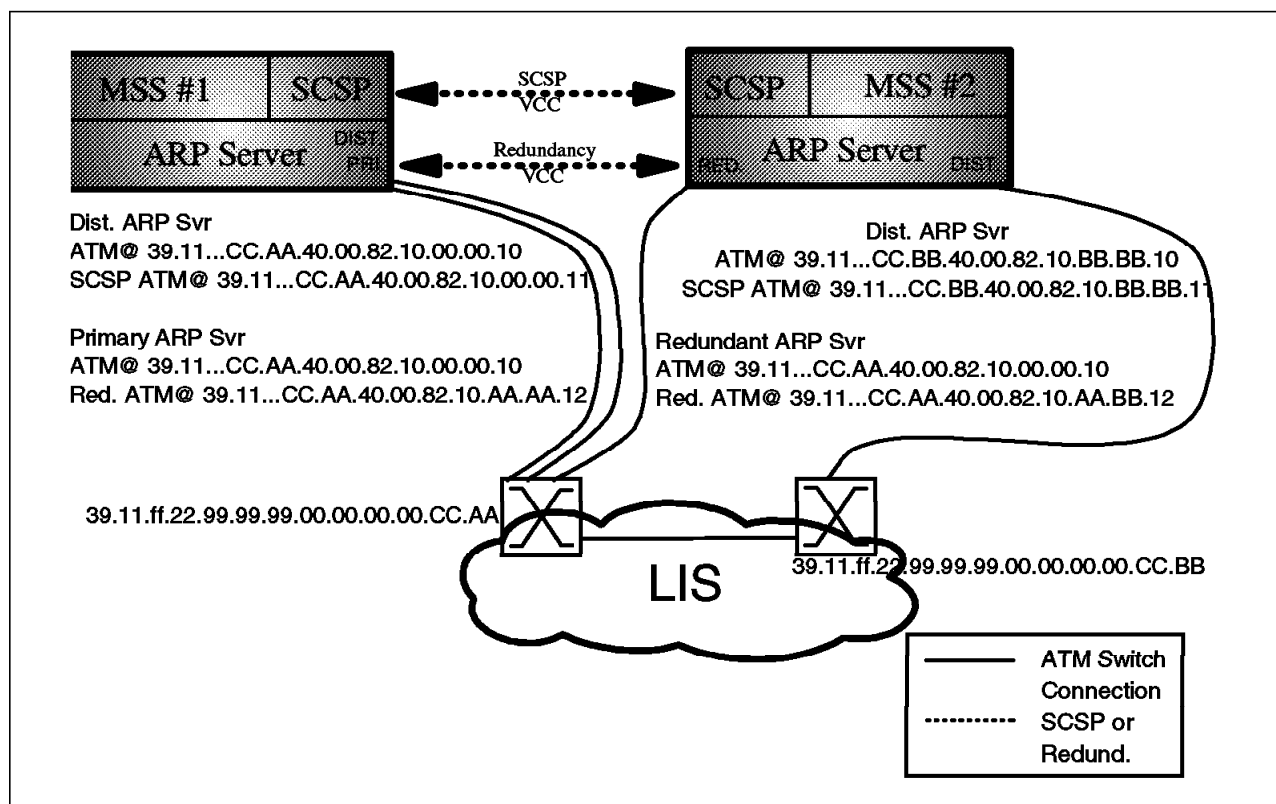


Figure 85. MSS Distributed and Redundant ATMARP Server ATM Addressing and Connections

The following configuration is in database MSSV2.CDB, file Distributed ARP MSS #1:

1. Adding an ESI

The ESI of 40.00.82.10.AA.AA is added for use with the redundancy VCC. This can be viewed in the NW;Network Device/Interface, Interface 0 Configure button/ESI tab. The ARP Server and Distributed ARP Server function is turned on. This can be viewed in NW;Network Device/Protocols/IP/Classical IP/Classical IP, Number 0 Configure button/ARP Server tab.

2. Setting ESI for ARP server and client.

The address and selector for the ARP Server is 40.00.82.10.00.00.10. This can be viewed in NW;Network Device/Protocols/IP/Classical IP/Classical IP, Number 0 Configure button/Client Address tab.

3. Activating distributed ATMARP server.

4. Defining the server group for distributed ARP server.

The Default Server Group of 1 has been chosen. This can be viewed in NW;Network Device/Protocols/IP/Classical IP/Classical IP, Number 0 Configure button/SCSP Parameter tab.

5. Setting the ESI for SCSP.

The ESI 40.00.82.10.00.00.11 is chosen for SCSP use. This can be viewed in NW;Network Device/Protocols/IP/Classical IP/Classical IP, Number 0 Configure button/SCSP Address tab.

6. Setting the local ESI and backup ARP server's ATM address for ARP server redundancy VCC.

The ESI of 40.00.82.10.AA.AA.12 has been chosen for the redundancy VCC. The ATM address of the backup ARP server's redundancy VCC is also defined. The definitions can be viewed in NW;Network Device/Protocols/IP/Classical IP/Classical IP, Number 0 Configure button/Redun Server tab.

Note

On some PCs, the MSS Configurator panel may not be tall enough to show the SCSP tabs. Just increase the height by dragging down the bottom of the panel and click on to the double arrow tab on the bottom-right of the window.

18.3.2 Verifying Distributed and Redundant ATMARP Server

The configuration supplied was tested on two MSSs, each with a single ATM interface. Since supporting the redundant ATMARP server and distributed ATMARP server requires two interfaces on an MSS, this verification did not produce meaningful results. The displays shown will have different results with a complete MSS configuration.

1. Display the IP addresses on MSS #1.

MSS #1 Test*T 5

MSS #1 Test ATM+p ip

MSS #1 Test IP>interf

Interface	IP Address(es)	Mask(s)
ATM/0	172.16.3.3	255.255.255.0
TKR/0	172.16.1.1	255.255.255.0
TKR/1	172.16.2.2	255.255.255.0

2. Display the RFC 1577 (Net 0) and LEC IP information.

MSS #1 Test +p arp

MSS #1 Test ARP>protocol

Network	Proto	(num)	AS	Protocol	Address(es)
0 ATM/0	IP	(0)	0800	172.16.3.3	
1 TKR/0	IP	(0)	0800	172.16.1.1	
1 TKR/0	SR	(16)	0000	40:00:82:10:00:04	
2 TKR/1	IP	(0)	0800	172.16.2.2	
2 TKR/1	SR	(16)	0000	40:00:82:10:00:05	

3. Display the status of the Classical IP connections.

MSS #1 Test +pro arp

MSS #1 Test ARP>di 0 ip

4. Display the ATMARP server redundancy state.

MSS #1 Test ARP>redundancy-state

Network number [0]0 Protocol [IP]? ip

If: 0 Prot: IP Clients configured with Redundancy

 Addr: 172.16.3.3 Pri/Secy: Pri Real Esi: Up Red. Esi: Up Red. Chnl: Down
 FLAGS: Real Client: C8 Red. Client: C8 RedFlags: C0
 Red. Channel: 0/0

```

Red. Channel: Source ATM address
39.11.FF.22.99.99.99.00.00.00.00.CC.BB.40.00.82.10.AA.AA.12
Red. Channel: Target ATM address
39.11.FF.22.99.99.99.00.00.00.00.CC.AA.40.00.82.10.AA.BB.12
Redundancy Status: Active

```

The redundancy display above should show Red. Chnl: Up when operating properly. Note the display of fully configured local and remote (partner) ATM addresses used for the redundancy VCC. Ensure these addresses represent the ESI and ATM network-part address of the ATM switch the MSSs are attached to.

5. Display the status of this distributed ATMRARP server.

```
MSS #1 Test ARP>ex
```

```
MSS #1 Test +p scsp
```

```

MSS #1 Test SCSP>stat
Network number [0]? 0
Server Group ID [0]? 1
DCS ID (hex) [0]? 1
DCS with that ID not found, listing all DCS's.

```

```
-----
DCS ID: Unknown
```

```

HFSM State: Down_Inop   DCS Hello Interval(sec): 0   DCS Dead Factor: 0
CAFSM State: Down Master/Slave: S   CA seq: 34C863D6   CSUS seq: 34C863D6
Cache Summary List sent?: yes   Cache Summary List ACKed?: yes
Cache Request List Size: 0   Cache ReTransmit List Size: 0   Age(sec): 9723
ATM Addr: 39.11.FF.22.99.99.99.00.00.00.00.CC.BB.40.00.82.10.BB.BB.10
VPI: 0   VCI: 0   Missed Hello Msgs: 0   RID doesn't match LSID: 0
Short Messages: 0   Sequence Mismatches: 0

```

6. Display the information about this distributed ARP server.

```

MSS #1 Test SCSP>list ser
Network number [0]?
      SGID Protocol      LSID DCSs ATM Addr
      1   ATMRARP  AC100303   1  39.11.FF.22.99.99.99.00.00.00.00.CC.BB.
                                40.00.82.10.00.00.11

```

This display maps the ATM address of the distributed ARP server with the LSID (also used as the DCSID, which is the IP address in hex format).

7. List the directly connected distributed ATMRARP servers.

```

MSS #1 Test SCSP>list dcs
Network number [0]? 0
Server Group ID [0]? 1
      DCS Id      HFSM State  CAFSM State  M/S  CRL Len  ReTran Len
      Unknown      Down_Inop      Down      S      0      0

```

If the MSS #2 would have had two ATM interfaces, the SCSP protocol function would be operating, and the above display's results would include different information. There would be an entry for the MSS #2 DCS. The DCS ID is the hexadecimal identifier of the distributed ARP server. The HFSM State denotes the status of the Hello Finite State Machine, a task which sends and receives SCSP Hello messages between configured DCSs.

The CAFSM is the task that aligns, or synchronizes the ATMARP cache between DCSs.

8. Check the SCSP cache status of the distributed ATMARP servers.

MSS #1 Test SCSP>**dump**

Network number [0]? **0**

Server Group ID [0]? **1**

Next key to assign = 34C863D6

Key	Origin ID	Seq. No.	Age	Paddr
AC.10.03.03	AC.10.03.03	885548001	0	AC.10.03.03

18.4 APPN Configuration

This example shows how APPN is configured in MSS. This is implemented over ATM connection.

18.4.1 Configuring MSS Server for APPN

The MSS can be configured in one of three ways:

1. Minimum configuration:
 - Allows MSS to accept connection requests from any other node.
 - MSS does not initiate connection requests to other nodes
2. Initiate connection configuration:
 - Allows MSS to accept connection requests from any other node.
 - Enables MSS to initiate connection requests to specified nodes.
3. Controlling connections configuration:
 - Allows MSS to accept connection requests only from specified nodes.
 - Enables MSS to initiate connection requests to specified nodes.

18.4.2 Configuration Steps for APPN

1. Configure LAN Emulation clients
2. Enable MSS Server as an APPN network node and configure:
 - Network ID
 - CP name
 - Route addition resistance
 - XID number for subarea connections
 - Branch Extender support
3. Enable MSS Server as a DLUR node and configure:
 - Fully qualified CP name of primary DLUS
 - Fully qualified CP name of backup DLUS
4. Configure the ports used by APPN/HPR.
5. You may define link stations for the adjacent nodes:
 - To which your MSS Server may initiate a connection.

- That may initiate a connection to your MSS Server.
- 6. Configure information about LEN nodes and their LU resources:
 - Define link station to the LEN node.
 - Define the names of LUs associated with LEN nodes.
- 7. Define attachment to connection networks.
- 8. Define new mode names or mode name to COS mappings.
- 9. Tune the performance of the MSS Server.
- 10. Perform node service trace diagnostics.
- 11. Collect statistics.

18.4.3 APPN Monitoring

You can query the following through talk 5:

- APPN configuration
- APPN port information
- APPN link station information
- CP-CP sessions
- APPN ISR
- Session information (must enable ISR accounting)
- RTP connections

18.5 APPN Configuration Example

APPN Code is a dynamic linking and loading (DLL) function. The code is only loaded when APPN is configured. This can be done with the command line in talk 6 and entering the load add package appn command. Next time the configuration is loaded, the APPN function will be added.

```
MSS_A *t 6

MSS_A CONFIG>load add package appn
appn package configured successfully
MSS_A Config>
```

APPN can also be loaded from the R2.0 Configuration Program by enabling the Enable APPN network node box on the APPN General panel. Then issue a reload command with that configuration to load the APPN function.

APPN configuration is defined in database MSSV2.CDB under appn.

In the MSS 2.0 Configuration program navigating to APPN-General thru NW;Network Device/Protocols/Appn/General brings us to the Network ID and Control Point Name selections. Here we enter USIBMRA to be the SNA/APPN Net-ID, and NNMSS1 to be the MSS NN control point name (CPNAME). We also check mark the box **Enable APPN NN**. If DLUR and Branch Extender supports are needed, the boxes within the the corresponding windows are also check marked. APNN interfaces is also configured in NW;Network Device/Protocols/Appn/Interfaces. **Interface '0'/configure/general** is clicked and

Enable APPN Routing on this port is selected. A port name is given. In the same window under Port definition Local ATM address defined as ESI and selector part of the ATM address.

APPN configuration can be monitored and confirmed using 't 5'.

18.6 Example - RFC 1483 SVC Bridging

18.6.1 Configuration Details

This example uses two MSSs, each creating an ELAN, and each defining an RFC 1483 bridge port to the other using an SVC. Both ELANs support a single subnet, 172.16.1, as shown in Figure 86. The intent of this example is to show all protocols being bridged between the two ELANs. The bridges support source route bridging. The configurations are in MSSV2.CDB and are named RFC 1483 SVC srb MSS #1 and RFC 1483 SVC srb MSS #2.

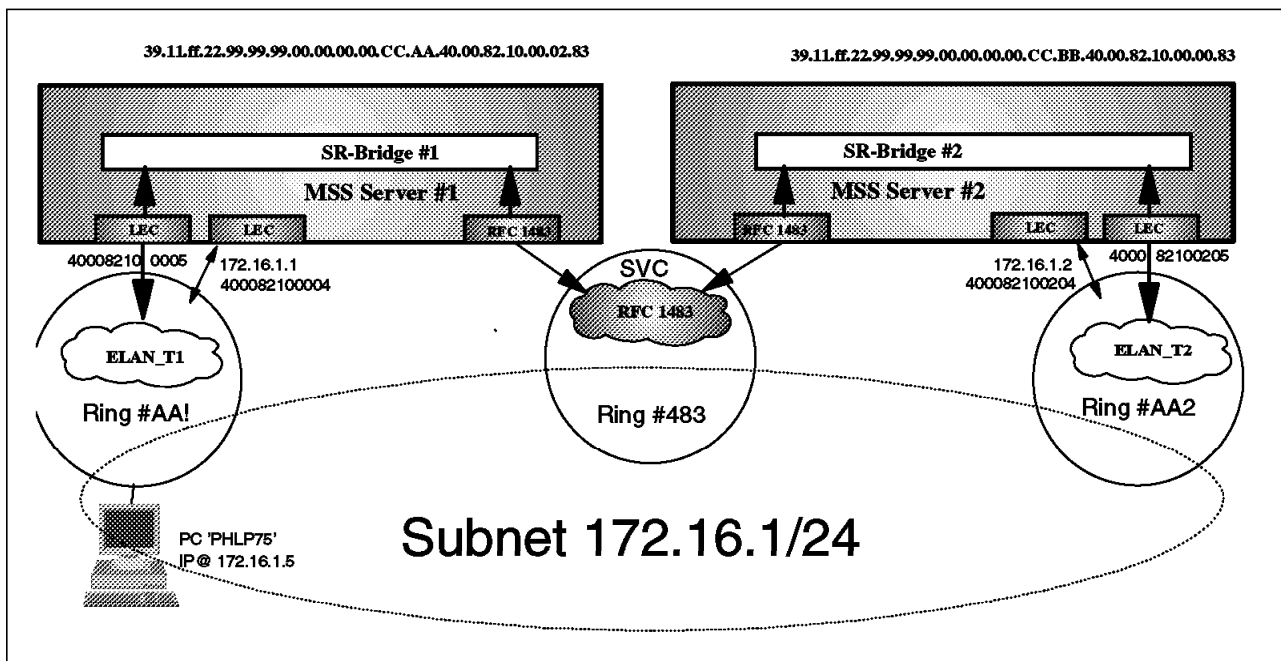


Figure 86. RFC 1483 Bridging Network Example

Figure 86 details the ESI and selectors used for each function. Note that an additional LEC has been added to each MSS. This is needed to accomplish the bridging of all protocols (including IP) between ELANs. Since the LECs using selector 04 have IP addresses defined, they would not be able to bridge IP traffic received from the RFC 1483 bridge port. This is a common function of devices that perform routing. Using the additional LEC (using selector 05) in each MSS, all traffic can be bridged, since this LEC has no IP address defined.

Although this example does not contain an ATMARP client, if the configuration used Classical IP over ATM, the interface 0 would also be routing IP; the RFC 1483 bridge port (since it would by default be defined on interface 0) would not bridge IP. To allow both ATMARP client and RFC 1483 bridge port bridging all protocols, create a virtual interface using the MSS 1.1 new function, and put the ATMARP client on that interface. This actually generates a new virtual interface (similar to creating an LEC instance) and the RFC 1483 port will pass IP and all

the other protocols. The ATMARP client will still support IP properly. Please note that this same issue exists when using IPX. Adding a virtual interface is done via NW;Network Device/Virtual Interface.




Figure 87. RFC 1483 Bridging Example MSS Configuration

The following description applies to MSS #1.

1. Turn on global bridging functions.

The Global Bridging function needs to be enabled. This is done via NW;Network Device/Protocols/Bridging/General, General Tab, Enable Bridging Box. The SRB tab allows configuration of the bridge number, and the internal virtual segment (only used when there are three or more SRB bridge ports).

2. Configure the LEC for bridging (SRB, segment #dd1).

The LEC Bridging parameters for MSS #1 can be viewed in NW;Network Device/Protocols/Bridging, Enable LEC-ELAN_T1 Interface 2, Configure Button.

3. Configure the RFC 1483 bridge port for SVC.

The RFC 1483 port connection is defined as either a PVC (with VPI/VCI defined), or as in this example, an SVC (with the full ATM address of the partner bridge). This is shown in NW;Network Device/Protocols/Bridging/Interfaces ATM 1/1, Configure Button.

4. Enable and configure the RFC 1483 bridge port.

The RFC 1483 bridge virtual port must be enabled and customized. The ESI and selector are also chosen to meet design criteria. In addition, the port type (SRB) and ring number for the RFC 1483 link (segment #483) are defined. The ring number for the RFC 1483 link must match in both MSSs. These can be viewed in NW;Network Device/Protocols/Bridging ATM Virtual Connections, Interface 0 Enable Box, Configure button, Client Addr tab.

5. Configure spanning tree to operate properly With RFC 1483 bridge port.

Since the RFC 1483 PVC port is on interface 0 (the lowest port number, whose MAC address is used for spanning tree BPDUs) and since there is no MAC address on that port, the bridging function for the RFC 1483 port will not become active. A MAC address must be configured for the RFC 1483 bridge port. This is only an issue if spanning tree is enabled (which is the default). Note also that if Talk 6 was used to define MSS, it is possible to have the RFC 1483 bridge port PVC be non-zero (if it was configured last) in which case this step is unnecessary. Adding the MAC address is shown in NW:Network Device/Protocols/Bridging/Spanning Tree Protocol/General, Source Route Bridging window, Bridge address.

18.6.2 Verifying RFC 1483 SVC Bridging

1. Display the bridge configuration.

```
MSS #1 Test *t 5

MSS #1 Test ASRT>ex
MSS #1 Test +p asrt
MSS #1 Test ASRT>li bridge

SRB Bridge ID (prio/add): 32768/40-00-82-10-14-01
Bridge state:           Enabled
UB-Encapsulation:       Disabled
Bridge type:            SRB
Number of ports:        2
STP Participation:      IBM-SRB proprietary
```

Port	Interface	State	MAC Address	Modes	Maximum MSDU	Segment	Flags
1	AT/0:0:156	Up	00-00-00-00-00-00	SR	9234	483	RD
2	TKR/1	Up	02-00-41-08-00-A0	SR	4544	AA1	RD

Flags: RE = IBMRT PC behavior Enabled, RD = IBMRT PC behavior Disabled

```
SR bridge number:      1
SR virtual segment:    DD1 (1:N SRB Not Active)
Adaptive segment:      000
MSS #1 Test ASRT>
```

2. Display the bridge port status and configuration.

```
MSS #1 Test ASRT>li port

Port Id (dec) : 128: 1, (hex): 80-01
Port State : Forwarding
STP Participation: Enabled
Port Supports : Source Route Bridging Only
SRB: Segment Number: 0x483 MTU: 4399 STE Forwarding: Auto
Assoc Interface #/name : 0/ATM/0 Dst:3911FF2299999900000000CCAA40008210000283
+++++
Port Id (dec) : 128: 2, (hex): 80-02
Port State : Forwarding
STP Participation: Enabled
Port Supports : Source Route Bridging Only
SRB: Segment Number: 0xAA1 MTU: 4399 STE Forwarding: Auto
Assoc Interface #/name : 2/TKR/1
Super ELAN bridging: No Super ELAN ID: 0
+++++
MSS #1 Test ASRT>
```

3. Check the two unidirectional SVCs being used for the bridge link.

```
MSS #1 Test ASRT>ex
MSS #1 Test +p arp
MSS #1 Test ARP>di 0
Active Channel List : Net 0
P/S FLAGS LIST VPI/VC I FwdPcr FwdScr MaxSDUsz Control P2P
1) S 00 01 0/155 154999984 154999984 9188 F T
Tgt Addr: 39.11.FF.22.99.99.99.00.00.00.00.CC.AA.40.00.82.10.00.02.83
Client Address (owner): Port no. 1
```



```

2) S 88 01 0/156 155000000 155000000 9188 F T
Tgt Addr: 39.11.FF.22.99.99.99.00.00.00.00.CC.AA.40.00.82.10.00.02.83
Client Address (owner): Port no. 1
New Channel List : Net 0
PVC Channel List : Net 0
MSS #1 Test ARP>

```

4. Display the traffic counters for the SVCs.

```

MSS #1 Test ARP>ex
MSS #1 Test +net 0
ATM Console
MSS #1 Test ATM+int
ATM Interface Console
MSS #1 Test ATM Interface+li cir 0 155

Frames transmitted = 0 Bytes transmitted = 0
Frames received = 6 Bytes received = 516

```

```

MSS #1 Test ATM Interface+li cir 0 156

Frames transmitted = 303 Bytes transmitted = 17242
Frames received = 0 Bytes received = 0

```

```

MSS #1 Test ATM Interface+

```

5. Display the IP address for MSS #1.

```

MSS #1 Test ATM Interface+ex
MSS #1 Test ATM+ex
MSS #1 Test +p ip
MSS #1 Test IP>inter
Interface IP Address(es) Mask(s)
TKR/0 172.16.1.1 255.255.255.0

```

6. Ping From MSS #1 to MSS #2.

```

MSS #1 Test IP>ping 172.16.1.2
PING 172.16.1.1 -> 172.16.1.2: 56 data bytes, ttl=64, every 1 sec.
56 data bytes from 172.16.1.2: icmp_seq=0. ttl=64. time=0. ms
56 data bytes from 172.16.1.2: icmp_seq=1. ttl=64. time=0. ms
56 data bytes from 172.16.1.2: icmp_seq=2. ttl=64. time=0. ms
56 data bytes from 172.16.1.2: icmp_seq=3. ttl=64. time=0. ms
56 data bytes from 172.16.1.2: icmp_seq=4. ttl=64. time=0. ms
56 data bytes from 172.16.1.2: icmp_seq=5. ttl=64. time=0. ms
56 data bytes from 172.16.1.2: icmp_seq=6. ttl=64. time=0. ms

----172.16.1.2 PING Statistics----
7 packets transmitted, 7 packets received, 0% packet loss
round-trip min/avg/max = 0/0/0 ms
MSS #1 Test IP>

```

7. Display the traffic counters for the SVCs again.

```

MSS #1 Test IP>ex
MSS #1 Test +net 0
ATM Console
MSS #1 Test ATM+int
ATM Interface Console
MSS #1 Test ATM Interface+li cir 0 155

Frames transmitted = 0 Bytes transmitted = 0
Frames received = 22 Bytes received = 2475

```

MSS #1 Test ATM Interface+li cir 0 156

Frames transmitted	=	392	Bytes transmitted	=	22787
Frames received	=	0	Bytes received	=	0

MSS #1 Test ATM Interface+

18.6.3 Verifying the Source Routing Information Field (RIF) in the RFC 1483 SVC Bridging Example

This verification was done from an OS/2 PC attached to a token-ring switch in ELAN_T1 as seen in Figure 86 on page 220.

1. Check the IP address of PHLP75.

```
[<phlp75>-C:\]ifconfig lan0
lan0: flags=3063<UP,BROADCAST,NOTRAILERS,RUNNING,BRIDGE,SNAP>
inet 172.16.1.5 netmask ffffffff broadcast 172.16.1.255
```

2. Ping MSS #2.

```
[<phlp75>-C:\]ping 172.16.1.2

PING 172.16.1.2: 56 data bytes
64 bytes from 172.16.1.2: icmp_seq=1. time=0. ms
64 bytes from 172.16.1.2: icmp_seq=2. time=0. ms
64 bytes from 172.16.1.2: icmp_seq=3. time=0. ms
64 bytes from 172.16.1.2: icmp_seq=4. time=0. ms
64 bytes from 172.16.1.2: icmp_seq=5. time=0. ms

----172.16.1.2 PING Statistics----
5 packets transmitted, 5 packets received, 0% packet loss
round-trip (ms) min/avg/max = 0/0/0
```

3. Check the MAC address of the IP routing LEC in MSS #2 (172.16.1.2).

```
[<phlp75>-C:\]arp -a
```

ARP table contents:			
interface	hardware address	IP address	minutes since last use
0	08005a5b12ba	172.16.1.3	0
0	400082100004	172.16.1.1	16
0	400082100204	172.16.1.2	0

4. Use the OS2PING utility to MSS #2's MAC address (IEEE802.2 LLC acknowledgement) for connectivity testing and RIF detail.

```
[<phlp75>-C:\]d:\utils\os2\lan\os2ping -a=400082100204 w=2 r
(C) IBM Corporation 1990
```

OS2PING Version 0.06 - 09/01/92

```
OS2PING: First response received
OS2PING: Addr=400082100204 Route was:
(aa1) 1 -> (483) 2 -> (aa2)
```

OS2PING: 1 response received during 2 second wait

Appendix A. Special Notices

This publication is intended to help people involved in the design, construction, management, and operation of ATM networks using the IBM MSS Server. The information in this publication is not intended as the specification of any hardware or software interface provided by the IBM MSS Server. See the PUBLICATIONS section of the IBM Programming Announcement for the IBM MSS Server 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, 500 Columbus Avenue, Thornwood, NY 10594 USA.

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:

AIX	APPN
AS/400	BookManager
IBM	Nways
OS/2	PowerPC 603
P2P	RS/6000
System/390	

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.

Microsoft, Windows, Windows NT, and the Windows 95 logo are trademarks or registered trademarks of Microsoft Corporation.

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 U.S. and other countries.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Other company, product, and service names may be trademarks or service marks of others.

Appendix B. Related Publications

B.1 International Technical Support Organization Publications

For information on ordering these ITSO publications see "How to Get ITSO Redbooks" on page 229.

- *TCP/IP Tutorial and Technical Overview*, GG24-3376
- *IBM Multisegment LAN Design Guidelines*, GG24-3398
- *ATM Technical Overview*, SG24-4625
- *Using and Understanding the MSS Server*, SG24-4915
- *Campus ATM Design Guidelines*, SG24-5002
- *IBM 8260 As a Campus ATM Switch*, SG24-5003
- *IBM N Ways RouteSwitch Implementation Guide*, SG24-4581

B.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

B.3 Other Publications

These publications are also relevant as further information sources:

- *Quick Guide to the IBM Multiprotocol Switched Services (MSS) Server Release 1.1*
Technical Report TR 29.2260, IBM, April 1997, available from IBM Networking Technical Reports Web page at <http://www.raleigh.ibm.com/tr2/tr2over.html>.
- *8210 Nways MSS Server Setup and Problem Determination Guide*, GA27-4140
- *8210 Nways MSS Server Operations Reference Card*, GX27-4017
- *Nways MSS Server Configuration Guide*, SC30-3821
- *Nways MSS Server Introduction and Planning Guide*, GC30-3820
- *Nways MSS Server Command Line Interface: User's Guide*, SC30-3818
- *Nways MSS Server Command Line Interface: Protocol Configuration Guide*, SC30-3819

- *Nways MSS Server Service Manual*, GY27-0354
- *MSS Server Module Setup Guide*, GA27-4141

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:

- **PUBORDER** — to order hardcopies in United States
- **GOPHER link to the Internet** - type GOPHER.WTSCPOK.ITSO.IBM.COM
- **Tools disks**

To get LIST3820s of redbooks, type one of the following commands:

```
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 one of the following commands:

```
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
```

For a list of product area specialists in the ITSO: type the following command:

```
TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ORGCARD PACKAGE
```

- **Redbooks Web Site on the World Wide Web**
<http://w3.itso.ibm.com/redbooks/>
- **IBM Direct Publications Catalog on the World Wide Web**
<http://www.elink.ibm.link.ibm.com/pb1/pb1>

IBM employees may obtain LIST3820s of redbooks from this page.

- **REDBOOKS category on INEWS**
- **Online** — send orders to: USIB6FPL at IBMMAIL or DKIBMBSH at IBMMAIL
- **Internet Listserver**

With an Internet e-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an e-mail note to announce@webster.ibm.link.ibm.com with the keyword subscribe in the body of the note (leave the subject line blank). A category form and detailed instructions will be sent to you.

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:

	IBMMAIL	Internet
In United States:	usib6fpl at ibmmail	usib6fpl@ibmmail.com
In Canada:	caibmbkz at ibmmail	lmannix@vnet.ibm.com
Outside North America:	dkibmbsh at ibmmail	bookshop@dk.ibm.com

- **Telephone orders**

United States (toll free)	1-800-879-2755
Canada (toll free)	1-800-IBM-4YOU
Outside North America	(long distance charges apply)
(+45) 4810-1320 - Danish	(+45) 4810-1020 - German
(+45) 4810-1420 - Dutch	(+45) 4810-1620 - Italian
(+45) 4810-1540 - English	(+45) 4810-1270 - Norwegian
(+45) 4810-1670 - Finnish	(+45) 4810-1120 - Spanish
(+45) 4810-1220 - French	(+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)	1-800-445-9269
Canada	1-403-267-4455
Outside North America	(+45) 48 14 2207 (long distance charge)

- **1-800-IBM-4FAX (United States) or (+1)001-408-256-5422 (Outside USA)** — ask for:

Index # 4421 Abstracts of new redbooks
Index # 4422 IBM redbooks
Index # 4420 Redbooks for last six months

- **Direct Services** - send note to softwareshop@vnet.ibm.com

- **On the World Wide Web**

Redbooks Web Site	http://www.redbooks.ibm.com/
IBM Direct Publications Catalog	http://www.elink.ibm.link.ibm.com/pbl/pbl

- **Internet Listserver**

With an Internet e-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an e-mail note to announce@webster.ibm.link.ibm.com with the keyword `subscribe` in the body of the note (leave the subject line blank).

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

AAL	ATM Adaptation Layer The layer that adapts user data to/from the ATM network by adding/removing headers and segmenting/reassembling the data into/from cells	CIP	Classical IP An IETF standard for ATM-attached devices to communicate using IP
AAL-5	ATM Adaptation Layer 5 One of several standard AALs, AAL-5 was designed for data communications and is used by LAN Emulation and Classical IP	CIPC	Classical IP Client A Classical IP component that represents users of the Logical IP Subnet (LIS)
A-CPSW	ATM Control Point and SWitch ATM switch module for IBM 8260 hub ARP Address Resolution Protocol, IP ARP translates network addresses into hardware addresses, LE ARP translates LAN Destinations into ATM addresses	DLCI	Data Link Connection Identifier Identifies a specific virtual connection on a frame relay link
ARI	ATM Real Interface MSS interface corresponding to the real ATM attachment	DPF	Dynamic Protocol Filtering A bridge enhancement supporting PVLANS
ATMARP	ARP in Classical IP	ELAN	Emulated Local Area Network A LAN segment implemented with ATM technology
ATM	Asynchronous Transfer Mode A connection-oriented, high-speed networking technology based on cell switching	ESI	End System Identifier A 6-byte component of an ATM address
AVI	ATM Virtual Interface A logical interface to extend MSS Server's routing functions	FDDI	Fiber Distributed Data Interface Standard for 100 Mbps Token-Ring Local Area Networks
ARI	ATM Real Interface The ATM interface corresponding to the real ATM attachment	IBM	International Business Machines Corporation
BBCM	Bridging BroadCast Manager A bridge enhancement designed to limit the effects of broadcast frames	IBUS	Intelligent BUS A LAN Emulation optimization designed to limit the scope of unknown unicast frames sent to the BUS
BCM	BroadCast Manager An IBM extension to LAN Emulation designed to limit the effects of broadcast frames	ICMP	Internet Control Message Protocol A protocol for communicating control information over IP
BUS	Broadcast and Unknown Server A LAN Emulation Service	IEEE	Institute of Electrical and Electronic Engineers An organization involved in establishing Local Area Network standards

IETF	Internet Engineering Task Force		Destinations to ATM Addresses
	An organization that produces Internet specifications	LIS	Logical IP Subnet
ILMI	Interim Local Management Interface	LLC	Logical Link Control
	SNMP-based procedures for managing the User-Network Interface (UNI)		The top sublayer of the Data Link Layer, which is layer 2 of the ISO model
IP	Internet Protocol	LNNI	LANE Network-Network Interface
	A widely-used network layer protocol specified by the IETF		The interface between distributed LE Service entities
ITSO	International Technical Support Organization	LSI	LANE Shortcut Interface
IPX	Internet Packet Exchange		Component of the MSS Server's NHRP implementation that establishes shortcut VCCs to LAN Emulation Clients
ISO	International Standards Organization	LUNI	LANE User-Network Interface
	An organization that specifies international communication standards		The interface between LAN Emulation Clients and the LAN Emulation Service
LANE	LAN Emulation	MAC	Medium Access Control
	An ATM Forum standard supporting communication between legacy LAN applications over ATM networks, the terms LANE and LE are often used interchangeably	MBS	Maximum Burst Size
LE	LAN Emulation	MIB	Maximum length of an uninterrupted data transfer on an ATM Virtual Connection
	An ATM Forum standard supporting communication between legacy LAN applications over ATM networks, the terms LE and LANE are often used interchangeably		Management Information Base
LEC	LAN Emulation Client	MOSPF	A network management database supporting the monitoring and control of network elements
	A LAN Emulation component that represents users of the Emulated LAN		Multicast Open Shortest Path First
LECS	LAN Emulation Configuration Server	MSS	A multicast routing protocol specified by the IETF
	A LAN Emulation Service component that centralizes and disseminates configuration data		Multiprotocol Switched Services
LES	LAN Emulation Server	MTU	A component of IBM's Switched Virtual Networking (SVN) framework
	A LAN Emulation Service component that resolves LAN		Maximum Transmission Unit
			The maximum amount of data that can be transmitted as a single unit (frame) on a communications link (does not include layer 2 headers)

NBMA	Non-Broadcast Multi-Access Attributes characterizing a set of networking technologies that includes ATM	QoS	communications protocol or participation in a particular protocol subnetwork Quality of Service
NHC	Next Hop Client An NHRP entity that initiates establishment of shortcut routes	RIF	Attribute of networks capable of delivering preferred/guaranteed service to specified traffic flows Routing Information Field
NHRP	Next Hop Resolution Protocol An IETF protocol for bypassing routers on NBMA networks	RIP	Sequence of ring and bridge numbers specifying the path between two stations in a Source Route Bridge network Routing Information Protocol
NHS	Next Hop Server An NHRP entity that provides the protocol address-to-NBMA address mappings necessary for establishment of shortcut routes	SAP	A vector-distance routing protocol, versions of RIP are used with IP and IPX Service Advertising Protocol
OSPF	Open Shortest Path First A routing protocol specified by the IETF	SCB	An IPX protocol used to advertise the location of available services Shortcut Bridging
PCMCIA	Personal Computer Memory Card International Association An organization involved in establishing hardware standards that are often associated with miniaturized peripherals	SCR	MSS Server function that enables inter-ELAN data direct VCCs Sustained Cell Rate
PCR	Peak Cell Rate Maximum transmission rate on an ATM Virtual Connection	SDU	Maximum sustained transmission rate on an ATM Virtual Connection Service Data Unit
PPS	Packets Per Second Units commonly used in throughput performance measurements	SLIP	Data as it appears at the interface between a layer and the layer immediately above Serial Line IP
PtP	Point-to-Point A VCC between two parties	SNA	An IETF standard for running IP over serial communication links Systems Network Architecture
PtMP	Point-to-MultiPoint A unidirectional VCC that allows transmissions from the source party to be received by multiple destination parties	SNAP	A networking architecture developed by IBM with a large base of installed systems SubNetwork Attachment Point
PVC	Permanent Virtual Circuit A configured VCC	SNMP	An LLC header extension that identifies the protocol type of a frame Simple Network Management Protocol
PVLAN	Protocol VLAN A specific type of VLAN where membership is based on use of a common		An IETF standard protocol that uses MIBs to control and monitor network elements

SRB	Source-Router Bridging A bridging protocol that is primarily used with Token-Ring Local Area Networks	TLV	Type/Length/Value A generalized information element that may be present in LAN Emulation and NHRP packets
SRT	Source-Route Transparent A bridging protocol for Local Area Networks specified in the IEEE 802.1d standard, SRT bridges support both source-route and transparent bridging on the same port	UNI	User-Network Interface The interface between user equipment and an ATM switch network
SR-TB	Source Route-Transparent Bridge A bridge that connects SR and TB ports	VCC	Virtual Channel Connection A connection between parties communicating via an ATM network
SVC	Switched Virtual Connection A VCC that is dynamically established via signalling protocols	VCI	Virtual Channel Identifier The VPI/VCI pair uniquely identifies a specific ATM connection on a given link
SVN	Switched Virtual Networking The name of IBM's framework for building and managing switch-based networks	VLAN	Virtual Local Area Network A logical grouping of hosts forming a broadcast domain that is independent of the physical topology
TB	Transparent Bridging A bridging protocol for Local Area Networks specified in the IEEE 802.1d standard	VPI	Virtual Path Identifier The VPI/VCI pair uniquely identifies a specific ATM connection on a given link

Index

A

- abbreviations 233
- acronyms 233
- Advanced Peer-to-Peer Networking 39
- AppleTalk 31
- APPN 39
- APPN configuration
 - example 218
- APPN connection networks 41
- ATM connection 139
- ATM real interface 27
- ATM virtual interface 11, 28
- ATMARP Server 114
- ATMARP Server Redundancy 115
- AVI 11

B

- Banyan VINES 35
- BBCM 66
- BCM 63
 - IPX forwarding list 131
- Branch Extender 39
- bridge broadcast management 66
- bridging 28
 - RFC 1483 21
- broadcast control
 - IBM 8274/3 78
 - MSS 78
- broadcast management 62
- BUS
 - adapter mode 125
 - performance 123
 - system mode 125
 - VCC splicing mode 126

C

- cluster control VC 120
- configuration
 - customization 145
 - downloading 146
 - export 146
 - opening 145
 - saving 146
 - XModem 147
- Configuration Examples 201
- configuration examples MSS V1.1 137
- configuration examples V2.0 201
- Configuration Improvements
 - Dynamic Reconfiguration 133
- configure MARS 211

D

- data direct VCCs 57
- Dependent LU Requester 41
- distributed ARP servers in Classical IP over ATM
 - example 213
- distributed ATMARP servers 213
- DPF 66
- Duplicate MAC Address 36
- Dynamic Linking and Loading 134
- dynamic protocol filtering 66
- Dynamic Reconfiguration 133

E

- End Node 40
- ESI selector assignment policy 131
- Ethernet 10
- examples 137
 - APPN configuration 218
 - backward connectivity 171
 - basic connectivity 149
 - distributed ARP servers in Classical IP over ATM 213
 - IP cut-through 164, 202
 - MARS in Classical IP Over ATM 208
 - redundancy 178
 - RFC 1483 Bridging with SVCs 220
 - shortcut bridging 154
- examples 2.0 201

F

- FDDI 7, 11, 27
- frame-relay interworking 23

H

- Hard Drive 10
- hardware 7, 9
- high-performance routing 41
- HPR 41

I

- IBM 2225 23
- IBM 8210 7, 9
- IBM 8260 7, 9, 10
- IBM 8265 9, 10
- IBM 8273 78
- IBM 8274 78
- IBM-compliant LAN emulation 11, 15
- ILMI 140
- in-band 142

- interfaces 11
 - logical 27
 - real 27
 - virtual 27
- IP 30
- IP cut-through 71
 - example 164, 202
- IP gateway
 - redundancy 109
- IP gateway redundancy
 - example 178
 - LAN emulation 103
- IP Multicast over ATM 116
- IP Multicasting 121, 208
- IPX 31
 - forwarding list 131
- IPX Routing Enhancements 35
- ISR 41

L

- LAN emulation 1
- LAN shortcut interface 92
- LE_ARP 130
- LE_ARP Cache 130
- LECS
 - dynamic access 140
 - WKA 140
- low-entry network end nodes 40
- LSI 92

M

- MARS 14, 208
- MARS cluster 117
- MARS in Classical IP Over ATM
 - example 208
- MARS Server and Client 117
- maximum burst size 50
- maximum reserved bandwidth 49
- MBS 50
- MCS 118
- migrating to MSS V2.0 201
- module 7, 9
- MRB 49
- MSS
 - functional overview 1
- multicast 14, 28
- Multicast over ATM 116

N

- network nodes 40
- next hop resolution 87
- NHRP 87
- Nways campus manager 132

O

- one-hop routing 90
- Operational Improvements
 - Dynamic Linking and Loading 134
 - Time Activated Reboot 134

P

- PCMCIA 8, 10
- PCR 48, 49
- peak cell rate 48, 49
- persistent SET 132
- point-to-multipoint VC 118
- point-to-point VC 117
- PVLANS 66

Q

- QoS 43
 - class 50
 - configurable 45
 - LE client 45
 - LECS 45, 51
 - LES 45, 51
 - negotiation 46, 51
 - parameters 45, 46
- quality of service 43

R

- Rapid Transport Protocol 41
- redundancy
 - ARP server 108
 - configuration example 178
 - IP gateway 109
 - LAN emulation 103
- redundancy and multicasting
 - Classical IP 107
- RFC 1483 21
 - bridging 21
 - PVCs sharing 129
- RFC 1483 Bridging with SVCs
 - example 220
- RFC 1483 SVC 26
- RFC 1483 SVC Bridging 220
- RFC 1577+ 113
- RFC 1577+ Client 113
- RIP Version 2 35
- RouteSwitching 95, 202
- RouteSwitching client 96
- routing 29
 - AppleTalk 31
 - IP 30
 - IPX 31
 - zero hop 71
- RSC 96
- RSS 96

S

- SAP address 42
- SCB 56
- SCR 50
- SCSP 113
- security 1
- Server Cache Synchronization Protocol 113
- server control VC 120
- shortcut bridging 56
 - example 154
- Single-slot MSS Server 10
- SLIP 142
- SNMP 143
- source route management 65
- SR-TB Duplicate MAC Address 36
- SRAM 131
- SRM 65
- stand-alone 7, 9
- superVLANs 55
- sustained cell rate 50
- SVC bridging 220
- switched virtual networking 4

T

- T2.1 nodes 39
- Time Activated Reboot 134
- traffic type 49
- translational bridging
 - IP 132

U

- users 146

V

- VCC binding 130
- VLAN 55
- voice 132

X

- XModem 142
 - loading configuration 147

Z

- zero-hop IP routing 202
- zero-hop route 99
- zero-hop route server 95

ITSO Redbook Evaluation

Understanding and Using MSS Release 1.1 and 2.0
SG24-2115-01

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 + 1 914 432 8264
- Send your comments in an Internet note to redbook@us.ibm.com

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!)

