


Troubleshooting IBM LAN/ATM Campus Networks

November 1997



**International Technical Support Organization
Raleigh Center**



International Technical Support Organization

SG24-2105-00

Troubleshooting IBM LAN/ATM Campus Networks

November 1997

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix G, "Special Notices" on page 515.

First Edition (November 1997)

This edition applies to IBM ATM Campus products versions available in July 1997.

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. 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

Figures	ix
Preface	xvii
The Team That Wrote This Redbook	xvii
Comments Welcome	xix
Chapter 1. Introduction	1
1.1 Objective	1
1.2 IBM Campus Network Products	2
1.2.1 Network Interface	2
1.2.2 Functions	3
1.3 Connection Example	5
Chapter 2. Main Concepts of Legacy LANs	7
2.1 The OSI Networking Model	7
2.2 Ethernet LAN Architectures	9
2.2.1 Main Characteristics	9
2.2.2 Ethernet V2 Frame Format (DIX)	11
2.2.3 Ethernet IEEE 802.3 Frame Format	13
2.2.4 Switched Ethernet	15
2.2.5 Fast Ethernet (802.3u)	15
2.2.6 100VG-AnyLAN (802.12)	16
2.2.7 Gigabit Ethernet (802.3z)	16
2.3 Token-Ring IEEE 802.5 LAN Architecture	17
2.3.1 Main Characteristics	17
2.4 FDDI (Fiber Distributed Data Interface) LAN Architecture	22
2.4.1 Main Characteristics	22
2.5 Bridges and Bridging Methodologies	27
2.5.1 Transparent Bridging (TB)	27
2.5.2 Source-Route Bridging (SRB)	31
2.5.3 Source-Route Transparent Bridging (SRT)	32
2.5.4 Source-Route Translational Bridging (SR-TB)	32
2.5.5 Source-Route Switching (SRS)	33
2.6 Protocol Architectures	34
2.6.1 NetBIOS	34
2.6.2 TCP/IP	34
2.6.3 IPX/SPX (Novell's NetWare)	37
2.7 Protocols Used for Network Management and Monitoring	40
2.7.1 Simple Network Management Protocol (SNMP)	40
Chapter 3. Main Concepts of ATM (Asynchronous Transfer Mode)	43
3.1 The Underlying Notions	43
3.1.1 The Structure of ATM Networks	43
3.1.2 ATM Network Characteristics	47
3.1.3 ATM Connections	52
3.1.4 Routing/Switching ATM Cells	54
3.1.5 ATM Cells and Cell Format	56
3.1.6 ATM Signalling	58
3.1.7 ATM Address Format	59
3.2 A Layered View of ATM	61
3.2.1 ATM Layer Model	61

3.2.2	Physical Layer	62
3.2.3	The ATM Layer	64
3.2.4	The ATM Adaptation Layer (AAL)	64
3.3	Networking Models Based on ATM	78
3.3.1	LAN Emulation V1 (LANE V1.0)	78
3.3.2	Classical IP (RFC 1577)	80
3.3.3	Next Hop Routing Protocol (NHRP)	83
3.4	Performance	85
3.4.1	Performance Problems	87
3.4.2	Performance Hints and Tips in ATM Networks	88
 Chapter 4. Problem Determination Guidelines		91
4.1	Problem Source Isolation in a Networking Environment	91
4.2	Before You Begin	91
4.3	Define the Problem	95
4.4	Problem Determination Tools and Procedures	95
4.4.1	LED and Other Visual Indicators	96
4.4.2	Console Commands	96
4.4.3	ATM Forum UNI Cause Codes and LANE Status Codes	103
4.4.4	Internal Traces	103
4.4.5	Protocol Analyzers	103
4.4.6	RMON	104
4.4.7	PING	104
4.5	Connectivity Problems in Intelligent Hubs	104
4.5.1	User Interface	105
4.6	Problem Determination Flowchart	107
 Chapter 5. Starting Problem Isolation in an ATM Network		109
5.1	Introduction	109
5.2	What Is a Healthy ATM Network?	110
5.3	Rules of ATM	111
5.3.1	ATM Forum Physical Interface References	114
5.4	Basic ATM Signalling	115
5.4.2	Requirements to Establish an ATM Call	116
5.4.3	Address Registration Failures	117
5.4.4	ILMI Scenario	123
5.5	Problem Determination Guidelines	133
5.5.1	Problem Phases	133
5.5.2	Hints and Tips with IBM Products	134
5.5.3	Problem Isolation Method	134
5.6	ATM Commons Problems	138
5.7	Tools Used during ATM Troubleshooting Case Studies	145
5.7.1	Fully Using End Devices Information	145
5.7.2	Nways Campus Manager ATM	147
5.7.3	8260 Internal Traces and Dumps	156
5.7.4	An ATM Network Analyzer	156
5.8	ATM Campus Problems Scenarios	157
5.8.1	Case Study 1 - UNI Translation Problem	157
5.8.2	Case Study 2 - Reachable Address Missing after Migration to PNNI	161
5.8.3	Case Study 3 - PVC Setup Error	166
5.8.4	Case Study 4 -Traffic Reachable Address Problem	174
5.8.5	Case Study 5 - Reachable Address Missing for Backup LECS	180
5.8.6	Case Study 6 - Basic PNNI Connection Problem	183
5.8.7	Case Study 7 - Address Not Reachable between Peer Groups	189
5.8.8	Case Study 8 - PNNI Path Selection Tracing	197

5.8.9 Case Study 9 - PNNI Path Selection Troubleshooting	204
5.8.10 About Duplicate Addresses	214
5.8.11 Case Study 10 - Reachability Configuration	225
Chapter 6. LAN Switches in Campus Networks	237
6.1 What Makes a Healthy Switched Campus Environment	237
6.2 Rules for Using LAN Switches	238
6.2.1 Interconnecting Switches	238
6.2.2 Virtual LANs (VLANs)	241
6.2.3 Address Tables and Aging Timers	245
6.2.4 Operating Modes and Performance	246
6.3 Problem Determination Guidelines	247
6.3.1 LAN Switch Configuration Problem Methodology	251
6.3.2 VLAN Assignment Problem Methodology	251
6.3.3 Aging Timers and Address Tables Problem Methodology	252
6.3.4 VLAN Leakage Problem Methodology	252
6.3.5 Diagnosing LAN Switch Problems	253
6.4 Hints and Tips in IBM Switch Environments	254
6.4.1 Hardware Check	254
6.4.2 Login to the User Interface	254
6.4.3 Console Commands	255
6.5 Case Studies for Configuration and VLAN Assignment Problems	257
6.5.1 Network Environment	257
6.5.2 Symptom: Clients Cannot Communicate with Each Other	257
6.5.3 Case Study (1-1): Simple Configuration Problem - Ports in Wrong Groups	258
6.5.4 Case Study (1-2): VLAN Assignment - Devices in Wrong VLANs	265
6.6 Case Studies for Address Aging Time	270
6.6.1 Network Environment	270
6.6.2 Symptom	270
6.6.3 Case Study (2-1): Aging Timer Problem without VLAN Polices	271
6.6.4 Case Study (2-2): Aging Timer Problem with VLAN Polices	273
6.6.5 Case Study (2-3): Auto-Tracker VLAN Aging Timer Problem	277
6.7 Case Studies for VLAN Leakage	283
6.7.1 Network Environment	283
6.7.2 Symptom	284
6.7.3 Case Study (3): VLAN Leakage Problem	284
Chapter 7. ATM Emulated LANs and Logical IP Subnets (LANE 1.0, RFC 1577)	293
7.1 What Makes a Healthy LANE 1.0 or Classical IP Environment	293
7.1.1 Address Resolution	293
7.1.2 Broadcasts	294
7.2 Rules of ATM Forum-Compliant LAN Emulation (LANE 1.0) and Classical IP (RFC 1577) Networks	294
7.2.1 ATM Forum-Compliant LAN Emulation (LANE 1.0) Initialization Rules	295
7.2.2 Differences between IBM-Compliant and ATM Forum-Compliant LAN Emulation	300
7.2.3 IBM MSS Server Performance Extensions	301
7.2.4 LECS and IBM MSS Server LES/BUS Redundancy	302
7.2.5 Classical IP (RFC 1577) Rules	303
7.3 Problem Determination Guidelines	306
7.3.1 Emulated LAN Connectivity Problem Methodology	310
7.3.2 Classic IP Connectivity Problem Methodology	320
7.3.3 Data Transfer Problem Methodology	324

7.3.4 Network Performance Problem Methodology	326
7.3.5 LE Redundancy Failure Problem Methodology	327
7.4 Common Problems Specific to IBM-Compliant LAN Emulation	329
7.5 A Guide for Using Commands with the IBM Nways Multiprotocol Switched Service (MSS) Server	331
7.6 Gathering Information by Using IBM Nways Campus Manager	333
7.7 Hints and Tips with IBM Products	335
7.8 Case Studies Involving Problems with LAN Emulation Connectivity	339
7.8.1 Network Environment	339
7.8.2 Symptom: LAN Emulation Clients Fail to Connect to Their ELANs	341
7.8.3 Troubleshooting Methodology in LAN Emulation Networks	341
7.8.4 Case (1-1): LECS Failure Problem	342
7.8.5 Case (1-2): Wrong ELAN Parameters in LE Client	353
7.8.6 Case (1-3): LES/BUS Failure Problem	360
7.8.7 Case (1-4): Microcode/Driver Problem	370
7.8.8 Case (1-5): LECS Parameters	372
7.8.9 Case (1-6): LES/BUS Parameters	374
7.9 Case Studies Involving Problems with Classical IP Network Connectivity	378
7.9.1 Network Environment	378
7.9.2 Symptom: Failure with a Logical IP Subnet Client-to-Client Connection	380
7.9.3 Troubleshooting Methodology in Classical IP	380
7.9.4 Case (2-1): ATMARP Server Failure Problem	381
7.9.5 Case (2-2): Wrong Subnet Mask in LIS Client	388
7.10 Case Studies Involving Client-to-Client Data Transfer	393
7.10.1 Case (3): LEC-LEC Connection	393
7.11 Case Studies Involving Network Performance Problems	406
7.11.1 Case (4): BUS Performance	406
7.12 Case Studies Involving Redundancy	408
7.12.1 Case (5): LES/BUS Redundancy Problem	408
Chapter 8. ATM Bridging and Routing	417
8.1 What Is a Healthy ATM Network with Bridging and Routing	417
8.2 Rules of ATM Bridging and Routing	418
8.3 PD Guidelines for Bridging and Routing over ATM	418
8.3.1 Categories of ATM Problems with Bridging and Routing	418
8.3.2 PD Methodology for ATM with Bridging and Routing	423
8.4 Hints and Tips with IBM Products	428
8.5 Case Studies for Bridging and Routing over ATM	432
8.5.1 Case Study 1: Bridging and Routing IPX in ATM	432
8.5.2 Case Study 2: Spanning Tree Loop	450
Appendix A. Ports and Cable Pinouts	463
A.1 Pinouts for ATM 25 Mbps versus Common Network Connectors	463
A.2 Other Cabling Considerations	464
A.2.1 Converter Cables	464
A.2.2 Hubs Crossover Wiring	464
Appendix B. Hub Code Level History	467
B.1 8260	467
B.2 8285	468
B.3 IBM Networking Information Home Pages	469
Appendix C. UNI 3.0-3.1 Cause Maintenance Error Codes	471
C.1 ATM Forum UNI Cause Codes	471

C.2 Maintenance Codes Valid on 8260/8285 ATM Hubs	476
C.3 IBM LAN Emulation Server (LES) 8260/8285 Error Reason Codes	477
Appendix D. ATM Forum-Compliant LANE Frame Formats	479
D.1 ATM Forum LAN Emulation Server Parameters	479
D.2 ATM Forum LAN Emulation Client Parameters	480
D.3 Configuration Frame Format	484
D.4 Join Frame Format	485
D.5 Registration Frame Format	486
D.6 Address Resolution Frame Format	487
D.7 Flush Frame Format	490
D.8 Control Frame Status Values	491
Appendix E. Traces and MIBs References	493
E.1 Trace Formatter Object Names	493
E.2 ATM Forum MIB Variable Listing	495
E.3 ILMI Objects	500
Appendix F. ATM/LAN/WAN Analyzers Used	503
F.1 InterWATCH 95000	503
F.1.1 Overview	503
F.1.2 Features	503
F.1.3 Network Interfaces Available	504
F.1.4 Tests and Applications	504
F.2 DA-30C	512
F.2.1 Overview	512
F.2.2 Features	512
F.2.3 DA-30C Analyzer and E3/DS3/OC3 Analysis Applications	513
Appendix G. Special Notices	515
Appendix H. Related Publications	517
H.1 International Technical Support Organization Publications	517
H.2 Redbooks on CD-ROMs	517
H.3 Other Publications	517
H.4 Performance Information	518
How to Get ITSO Redbooks	519
How IBM Employees Can Get ITSO Redbooks	519
How Customers Can Get ITSO Redbooks	520
IBM Redbook Order Form	521
List of Abbreviations	523
Index	527
ITSO Redbook Evaluation	529

Figures

1.	Topology of IBM Campus Network Products	5
2.	The OSI Networking Model	7
3.	An Ethernet LAN	9
4.	Ethernet V2 (DIX) Frame Format	11
5.	Ethernet IEEE 802.3 Frame Format	13
6.	A Token-Ring LAN	17
7.	IEEE 802.5 Token-Ring Frame Format (Part 1 of 2)	19
8.	IEEE 802.5 Token-Ring Frame Format (Part 2 of 2)	20
9.	Elements in an FDDI LAN	22
10.	FDDI Frame Format	25
11.	Transparent Bridging (TB)	27
12.	Bridge Protocol Data Units (BPDU) Used with Spanning Tree	29
13.	Source Route Bridging (SRB)	31
14.	Frame Conversion in Source-Route Translational Bridging (SR-TB)	33
15.	Main TCP/IP Components	35
16.	NetWare's (Novell) LANs	37
17.	NetWare's IPX/SPX Frames (Novell)	39
18.	The Structure of ATM Networks	45
19.	Multicast Tree	49
20.	VPs, VCs, Label Swapping	52
21.	Generic Structure of an ATM Cell	56
22.	Structure of ATM Addresses	59
23.	ATM Layer Model	61
24.	ATM and the Physical Layer	62
25.	ATM Service Classes	65
26.	AAL-3/4 SAR and CPCS PDUs	68
27.	AAL-5 SAR and CPCS PDUs	70
28.	AAL-5 Frames (User Plane)	72
29.	Basic Signalling at the UNI	73
30.	Q.2931 Signalling Frames (Control Plane)	75
31.	The LANE Model (LANE 1.0)	79
32.	The Classical IP Model (RFC 1577)	82
33.	The NHRP Model	83
34.	Example Configuration Sheet for 8260 ATM Switch	92
35.	Example Configuration Sheet for 8260 ATM Switch	93
36.	ATM Address Assignment in Campus Network	94
37.	Show Module Output (8260 Console Command)	96
38.	Show Port Output (8260 Console Command)	98
39.	Show PVC Output (8260 Console Command)	98
40.	Show Device Output (8260 Console Command)	100
41.	Show Reachable_Address Output (8260 Console Command)	101
42.	Show Signalling Command Output (8260 Console Command)	102
43.	8260 Ports Assigned to Different Networks	105
44.	Example of Typical Console Port Settings	106
45.	Flowchart for Problem Isolation	107
46.	Problem Isolation Flowchart	110
47.	ATM Call Establishment	115
48.	Classical IP Network	117
49.	Simple IP Network	118
50.	Show Port Command	119
51.	Normal ILMI (UNI 3.1)	122

52.	ILMI Network Configuration	123
53.	Addresses Reachable before ILMI of Port 13.2	124
54.	Port Status before ILMI	124
55.	ILMI of an 8210 Port Seen from 8260	126
56.	ILMI of an 8210 Port Seen from 8210 ELS	127
57.	Reachable Addresses after ILMI of Port 13.2	129
58.	8260 Port Status after Successful ILMI	129
59.	8210 Physical ATM Interface Status after Successful ILMI	130
60.	8210 ILMI Trace Output Sample	131
61.	Static Address Registration on 8260	132
62.	Problem Isolation (Chart 1 of 3)	135
63.	Problem Isolation (Chart 2 of 3)	136
64.	Problem Isolation (Chart 3 of 3)	137
65.	8285 Show Device Command Result	146
66.	ATM Switch Interface Configuration	148
67.	ATM Call Logging	149
68.	Call Logging Filtered by Interface	150
69.	ATM Call Details	151
70.	ATM SVC List	153
71.	ATM SVC Details Panel	154
72.	ATM SVC Tracking Panel	155
73.	Network Configuration	157
74.	Result of the VSS Command on 8274	158
75.	Event Logging System for LES Subsystem	159
76.	Output of the Event Logging System for SVC Subsystem	159
77.	Add Party Clear Cause	159
78.	Output of the Event Logging System for LES Subsystem	160
79.	Network Configuration	161
80.	Output of the vas Command	162
81.	8260#2 Set LAN Emulation Configuration Server before Migration	163
82.	8260#2 Show Configuration_server Configuration before Migration	163
83.	Links and Routes Configuration for 8260#2 before Migration	164
84.	Foreign Address Registration Example	164
85.	8260#2 Configuration_server Configuration after Migration	164
86.	Route Configuration Needed for 8260#2 after Migration	165
87.	Network Configuration	166
88.	PVC Over IISP Links	167
89.	8210 Bridging Port Status	168
90.	8274 Port Status	169
91.	8210 ATM Interface Status	170
92.	8274 ATM Statistics	170
93.	PVC Status: Failed	171
94.	Addressing Rule for a Soft PVC Configuration	171
95.	The Correct Command for PVC Setup	172
96.	PVC Status: Active	172
97.	8210 ATM Interface Statistics	173
98.	8274 ATM Statistics	173
99.	Network Configuration	174
100.	Cross Connections of Port 13.2 in Hub 8260#3	175
101.	Part of PNNI Path Selection Dump For 8260#3	176
102.	How to Dump PNNI Path Selection	177
103.	Reachability Information for 8260#1	178
104.	Port 2.1 Configuration	178
105.	LAN Emulation Redundancy Network Configuration	180
106.	Reachability Information with a Running LECS	181

107. Reachability Information without LECS	181
108. Setting the Correct Reachability for Backup LECS	182
109. Network Configuration	183
110. Peer Group Members Seen from 8260#3	184
111. Peer Group Members Seen from 8260#1	184
112. Node 0 Configuration for 8260#3	185
113. Node 0 Configuration for 8260#1	185
114. Peer Group Made of Three Nodes	186
115. Peer Group Members	187
116. Neighbors of 8260#1	187
117. Neighbors of 8260#2	188
118. Neighbors of 8260#3	188
119. Network Configuration	189
120. Result of First PING from 8285#1	190
121. Result of PING from 8260#1	190
122. Result of the Second PING from 8285#1	190
123. 8285#1 Configuration	191
124. 8260#1 Configuration	191
125. ARP Server Dump	192
126. 8260(2) Reachable Addresses	195
127. Reachable Address Configuration on a VPC Link	195
128. Network Configuration	197
129. How to Start a Selective Trace on 8260	198
130. NAVTEL InterWATCH 95000 Call Setup Decoding	199
131. NAVTEL InterWATCH 95000 Call Setup Binary Dump	201
132. Output of 8260 Internal Trace	202
133. Network Configuration	204
134. Cross Connections from Port 13.2	205
135. Output of 8260#1 Internal Trace	205
136. PNNI Path Selection Dump for 8260#1	206
137. Changing UBR Path Selection Method	208
138. Administrative Weight Configuration on the Links	209
139. Path Selection Trace when Using Shortest Path	210
140. PNNI Path Selection Dump	211
141. Reachability with Different Network Prefix Length	213
142. Path Selection Trace when Using Different Prefix Lengths	213
143. Multiple 8210 Connected to the Same Hub	214
144. Dynamic Load Balancing of Calls to a Duplicate Address (1 of 2)	217
145. Dynamic Load Balancing of Calls to a Duplicate Address (2 of 2)	218
146. 8260 Configuration For Redundant ARP Servers	219
147. A Network with Multiple LECS	221
148. Creating LAN Emulation Domains	223
149. Network Configuration	225
150. Cross Connections from Port 13.1	226
151. Cross Connections from Port 13.1	227
152. Reachability Configuration	228
153. How to Configure Reachability Information on Parallel IISPs	231
154. Peer Group Connected to a Backbone Peer Group	233
155. Peer Group Connected to a Similar Peer Group	234
156. Peer Group Connected to Two Isolated Peer Groups	235
157. LAN Switch Connections	239
158. Etherpipes and Tokenpipes	240
159. Trunking Used by the RouteSwitch	240
160. PToP Bridging Used by the RouteSwitch	241
161. Port-Based VLANs or VLAN Groups	242

162.	Policy and Port VLANs	243
163.	Common Switch Problems and Our Problem Determination Methodology	247
164.	RouteSwitch Main Console	255
165.	LAN Switch Environment (Physical View)	257
166.	LAN Switch Environment (Logical View Case 1-1)	258
167.	LAN Switch LEDs	259
168.	8274 Command to View the Status of Its Modules	260
169.	8274 Command to View the Assignment of Ports to Groups	261
170.	8274 Command to Show the Defined Groups	262
171.	8274 Command to Change Group Assignment of a Port	263
172.	8274 Command to View the Assignment of Ports to Groups	263
173.	LAN Switch Environment (Logical View Case 1-2)	265
174.	8274 Command to Verify Policy VLANs Defined	266
175.	8274 Command to Verify Multicast VLANs Defined	266
176.	8274 Command to View the VLAN Assignment of Ports	266
177.	8274 Command to Verify the MAC Addresses Learned for a Group	267
178.	8274 Command to Verify the Spanning Tree Parameters	267
179.	8274 Command to Modify the Policies of a VLAN	268
180.	8274 Command to View the VLAN Assignment of Ports	269
181.	LAN Switch Environment for Aging Timer Test	270
182.	8274 Command to Display the MAC address Information For a Slot	271
183.	8274 Command to Display the MAC Address Information for a Slot	272
184.	Network Trace from the Risc System/6000 Port on the 8274 Switch	272
185.	RS/6000 Command to Display Active Processes Running	272
186.	8274 Command to View VLANs Defined	273
187.	8274 Command to View the VLAN Assignment of Ports	274
188.	8274 Command to Display the MAC Address Information for a Port	274
189.	8274 Command to Display the MAC Address Information for a Port	275
190.	Data Flow in the Aging Timeout	276
191.	8274 Command to View VLANs Defined	277
192.	Auto-Tracker VLAN Aging Timer	278
193.	8274 Command to Display the VLAN Membership of Ports	279
194.	8274 Command to Display the VLAN Membership of MAC Addresses	279
195.	8274 Command to Display the VLAN Membership of Ports	280
196.	8274 Command to Display the VLAN Membership of MAC Addresses	280
197.	8274 Command to Display the VLAN Membership of Ports	281
198.	8274 Command to Display the VLAN Membership of MAC Addresses	281
199.	8274 Command to Display the VLAN Membership of Ports	282
200.	8274 Command to Display the VLAN Membership of MAC Addresses	282
201.	LAN Switch Environment for VLAN Leakage Test	283
202.	8274 Command to Modify a Port Setup	284
203.	8274 Command to View VLANs Defined	285
204.	Logical View: Two IP Network-Based VLANs	285
205.	8274 Command to View the VLAN Assignment of Ports	286
206.	8274 Command to Display the VLAN Membership of MAC Addresses	286
207.	8274 Command to Create a VLAN Rule	287
208.	8274 Command to View VLANs Defined	287
209.	Logical View: Two IP Network and IPX-Based VLANs	288
210.	8274 Command to View the VLAN Assignment of Ports	288
211.	8274 Command to Display the VLAN Membership of MAC Addresses	289
212.	Network Trace from Port Slot3/Port1	290
213.	VLAN Leakage	291
214.	ATM LAN Emulation Connection VCCs	296
215.	LE Initialization Process	297

216.	LE Data Transfer Process	299
217.	Classical IP (RFC 1577) Using SVC Connections	304
218.	Problem Determination Methodology	306
219.	MSS Command Menus	332
220.	ATM Interface Submap in IBM Nways Campus Manager	333
221.	Exploded ELAN by IBM Nways Campus Manager	334
222.	LAN Emulation Environment (Physical View)	339
223.	LAN Emulation Environment (Logical View)	340
224.	Troubleshooting LAN Emulation Connectivity Problems	341
225.	MSS Command to Delete the Active LECS	342
226.	MSS Command to Discover Clients Joined to an ELAN	342
227.	MSS Commands to Check Device and Interface Status	343
228.	8274 Command to Check Device and Interface Status	344
229.	8260/8285 Command to Check Module Status	344
230.	8260/8285 Command to Check Port Status	345
231.	8260/8285 Command to Determine ATM Addresses Registered with the Switch	345
232.	8260 Command to Determine the Configured LECS ATM Address	346
233.	8285 Command to Determine the Configured LECS ATM Address	346
234.	MSS Command to View the Status of the Running LECS	346
235.	MSS Command to Check the Saved LECS Configuration	347
236.	MSS Command to Check the LECS ELANs Defined	347
237.	MSS Command to Check the LECS Policy Priorities	347
238.	MSS Command to Determine LECS Policy Values	348
239.	MSS Command to Determine the List of Defined LESs for an ELAN	348
240.	MSS Command to Restart the LECS	348
241.	MSS Command to Check the Active Status of the LECS	349
242.	LECS Configuration in IBM Nways Campus Manager	350
243.	MSS Command to View LECS Statistics	350
244.	MSS "talk 2" Events for the LECS Subsystem	352
245.	MSS Command to View LECs Assigned to an ELAN	352
246.	MSS Command to Determine Which Clients Are Connected to the ELAN	353
247.	8274 Command to Check Its LE Client Configuration	354
248.	MSS Command to List an MSS LE Client's Configuration	356
249.	LEC Configuration in Campus Manager	357
250.	MSS Command to List the LECS Statistics	358
251.	MSS Talk 2 Events for the LECS Subsystem	359
252.	MSS Command to Determine Which LECs Are Attached to an ELAN	359
253.	MSS Command to Determine Which LE Clients Are Attached to an ELAN	360
254.	MSS Command to Verify the Status and Configuration of a LES/BUS	361
255.	LEC Configuration in Campus Manager	362
256.	MSS Command to Check the LECS Statistics	363
257.	MSS talk 2 Command to Examine Events for the LECS Subsystem	363
258.	SVCs Established by DOS Client during LES Failure	365
259.	MSS Command to Restart a LES/BUS	365
260.	SVCs Established by DOS Client after LES Restart	366
261.	MSS Command to List VCCs for a Specific LEC	367
262.	MSS Command to Display the Status and Configuration of an MSS LEC	367
263.	MSS Command to View VCCs for an MSS LEC	368
264.	8274 Command to View a LEC Status and VCCs Established	368
265.	MSS talk 2 Events for the LES Subsystem	369
266.	MSS talk 2 Events for the LES Subsystem	369
267.	8260/8285 Command to View Microcode Levels	370

268.	MSS Command to View Microcode Levels	370
269.	MSS Command to View the LECS Statistics	372
270.	MSS Command to List LECS Status and Configuration	373
271.	MSS Command to Display LES Statistics	375
272.	MSS Command to Display BUS Statistics	376
273.	Classical IP Environment (Physical View)	378
274.	Classical IP Environment (Logical View)	379
275.	Troubleshooting Methodology for Classical IP Connectivity	380
276.	MSS talk 2 Events for the ARP Subsystem	381
277.	8260 Command to Verify ATM Addresses Registered with the Switch	382
278.	ATMARP ATM Connections on the RS/6000	386
279.	MSS talk 2 Events for the ARP Subsystem	387
280.	MSS Command to Determine Which Clients Are Connected to the Emulated LAN	393
281.	MSS Command to Display the LE_ARP Cache for the MSS LEC	394
282.	8274 Command to Verify the LE_ARP Cache for the 8274 LEC	394
283.	MSS Command to Display the VCCs for the MSS LEC	395
284.	ATM SVC List in Campus Manager	396
285.	ATM SVC Show in Campus Manager	397
286.	ATM SVC Tracking in Campus Manager	398
287.	LE_ARP_Request Frame Capture	401
288.	Summary Report Produced from Network Analyzer	402
289.	ICMP Echo Frame Screen 1 of 2	403
290.	ICMP Echo Frame Screen 2 of 2	404
291.	MSS Command to Display the BUS Monitor Statistics	406
292.	MSS Command to Display the Statistics for a Specific LEC	407
293.	MSS Command to Display the Statistics for the BUS	407
294.	LES/BUS Redundancy Test Environment	408
295.	MSS Command to Display LECs Attached to an ELAN	409
296.	MSS Command to Display LECs Attached to an ELAN	409
297.	MSS Command to Verify the Configuration of the LES	410
298.	MSS Command to Verify the Configuration of the LES/BUS	411
299.	MSS Command to Verify the Configuration of the LES/BUS	412
300.	MSS Command to Verify the Configuration of the LES/BUS	413
301.	MSS Command to Display the BUS Statistics	414
302.	MSS talk 2 Display of Events for the LES Subsystem	415
303.	MSS Command to Verify the Configuration of the LES/BUS	415
304.	MSS Command to Display LECs Attached to an ELAN	416
305.	MSS Command to Display LECs Attached to an ELAN	416
306.	Symptoms and Causes for Bridging Problems	421
307.	Symptoms and Causes for Routing Problems	422
308.	IPX Bridging/Routing (Physical View)	432
309.	IPX Bridging/Routing (Logical View)	433
310.	Bridge Status Report (IBM Nways 8281 ATM LAN Bridge)	437
311.	IBM Nways 8281 ATM LAN Bridge: Wrong LAN Emulation Settings	443
312.	IBM Nways 8281 ATM LAN Bridge: Corrected LAN Emulation Settings	444
313.	Spanning Tree Loop (Physical View)	451
314.	Spanning Tree Loop (Logical View)	452
315.	Internal 8281: Bridge Status Report	455
316.	External 8281: Bridge Status Report	456
317.	Wires Crossed between Hubs	465
318.	QoS Test	504
319.	In-Service Test Data Insertion	505
320.	Regenerative Monitoring	506
321.	Out-of-Service Analysis	507

322.	Multiport LAN Monitoring and Analysis	508
323.	ATM Switch Verification	509
324.	Multiport, Multitopology Monitoring	510
325.	Connectivity Testing	511
326.	LAN Emulation Option	511
327.	ATM DA-30C Applications	513

Preface

This redbook will help you secure and troubleshoot your LAN/ATM campus network. It also gives a broad understanding of the LAN/ATM architecture.

In the redbook you will find:

- A listing of IBM product offerings, providing brief information on the interfaces supported and their functions.
- An introduction to the most important legacy LAN concepts, including such important things as frame types, frame conversion, and Spanning Tree. This will be especially useful for a thorough understanding of the most relevant concepts needed by the networking specialist when troubleshooting difficult situations. Needless to say, it will also be of much use to anybody wishing to better understand the real issues in legacy LAN networking, which of course expand also into the ATM-based campus networks.
- An introduction to the most important ATM concepts used in campus networks, with an aim at making them clear cut and easily understandable as they are a much-needed prerequisite to effectively troubleshoot problems in ATM-based networks.
- General rules for how to troubleshoot ATM campus networks. It provides information on the way console commands can be used to collect network status information.
- Provides an approach to isolate and solve problems in ATM IP networks. The same approach can be used when trying to maintain a good working ATM network.
- Guidelines for troubleshooting some of the common problems that can specifically arise when using LAN switches in campus networks, and illustrates these through the use of some problem scenarios.
- Guidelines for problem determination in a LANE 1.0 or Classic IP (RFC 1577) environment. It also uses some typical problem scenarios to illustrate the use of these guidelines in diagnosing problems within IBM networks.
- Methodology for troubleshooting problems commonly found in ATM networks containing bridges and routers. A classification of the symptoms experienced and their possible causes is also worth mentioning. In addition, some case studies are used to illustrate how the methodology can be applied to effectively resolve these kinds of problems.

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.

This project was designed and managed by Georges Tardy of the System Management and Networking ITSO Center, Raleigh. He joined IBM in 1965 and has had several responsibilities in different areas and divisions, such as manufacturing, product engineering field support, and development engineering. He is located in La Gaude and covers the ATM/LAN Campus products for the ITSO.

The authors of this document are:

Thierry Girard is a Senior Network Engineer in France. He has eight years of experience in campus networking and currently works in the Technical Marketing Support department for EMEA. He holds a degree in Networking from Telecom Paris. His areas of expertise include Ethernet, token-ring, ATM and campus network management.

Paul Hamilton is a Senior Network Designer in the UK. He has seven years of experience in campus network design, support and installation and currently works in IBM Global Services NS-CS. He holds a degree in Operational Research and Computer Science from Lancaster University. His areas of expertise include network Outsourcing, Ethernet, token-ring, ATM design, support and installation, network management, LAN Emulation and Classical IP.

Charles M. Magron is an Advisory Network Engineer in Switzerland. He has eleven years of experience in campus networking. He holds a BSEE degree from the Swiss Federal Institute of Technology (ETHZ) in Zurich, Switzerland. He has worked at IBM for eight years, and works currently in the Networking Systems department, providing country level pre- and post-sales support. His areas of expertise include multiprotocol networking with Ethernet, token-ring, ATM, and network management, including such varied activities as network design, education, project management, implementation (four ATM networks implemented by mid 97), protocol analysis, and network troubleshooting in general.

Jun Matsuo is a Network Engineer in Japan. He has three years of experience in the networking field and currently works in the Field Support department. His areas of expertise include multiprotocol networking with hubs, bridges, switches and routers for Ethernet, token-ring and ATM.

Dexter R. Monk is a Staff Engineer in the USA. He has worked at IBM for 23 years. He has nine years of experience in LAN/WAN networking and four years with ATM networking. He currently is working as the Product Engineer for the 8285 and the 8260 ATM hubs.

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

Advanced US Technical ATM/LAN Support

Ray Collins
Volkert Kreuk
Mike Curtis

Systems Management and Networking ITSO Center, Raleigh

Tim Kearby
Martin Murhammer
John Parker

8260 Development IBM France: La Gaude

Minh Tri Do Khac
Elliott Norsa
Mathieu Girard

8260 System Test IBM France: La Gaude

Francois Lemaut
Anne Lise Conjeaud
Thierry Arguillere

8260 Product Engineering IBM France: La Gaude
Olivier Caillau

MSS Development IBM USA
Flo Kandefer

LAN SWITCH Software Development IBM USA
Norm Strole

IBM Dallas Systems Centre
Cindy Young

IBM Switzerland
Yves Haemmerli
Daniel Kipfer

IBM UK
Darius Fariborz

IBM USA
Ed Wagstaff

GNnettest
Mario Pidutti
Larry Scheck

Wandel and Goltermann
Tony Michael

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 529 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@vnet.ibm.com

Chapter 1. Introduction

This chapter explains the objectives of this redbook and lists the IBM product offerings and provides brief information on the interfaces supported and their functions.

1.1 Objective

As diverse mixes of traffic including data, voice, and video begin to move across networks including mixed Ethernet, token-ring, FDDI, and ATM technologies, troubleshooting campus networks becomes more and more difficult. The objective of this redbook is to provide an approach to their debugging, giving rules and advice on how to proceed.

Problem analysis or network performance is treated at a network level and does not replace the dedicated troubleshooting information provided in the user guides or in other publications delivered with the product shipping groups at time of delivery.

For legacy LANs, several complete publications are already available. This redbook only gives an overview of each technology.

Troubleshooting the complex and error-prone world of bridges, routers, gateways and switches is not easy and in some cases the help of an analyzer could be needed. Some examples of their utilization are given.

In this redbook, we provide many scenarios that start with a description or a drawing of the failing network and an explanation of the problem encountered, then a detailed logical narrative that shows how to find the problem.

Scenarios are classified by chapter subjects and include:

1. Basic ATM: Chapter 5
2. LAN Switching on ATM Campus: Chapter 6
3. ATM Emulated LANs and Logical IP Subnets: Chapter 7
4. ATM Bridging and Routing: Chapter 8

1.2 IBM Campus Network Products

This section is intended to list and classify the IBM campus network products. The devices covered are bridges, concentrators, hubs, switches and ATM Service Servers used in campus network environments.

1.2.1 Network Interface

The IBM campus network products are classified by the following network interfaces.

<i>Table 1. Classification of IBM Campus Products by Interfaces</i>					
	Ethernet	token-ring	FDDI	Fast-Ethernet	ATM
1. Bridges					
IBM 8229 Bridge	O	O	-	-	-
8281 Nways ATM LAN Bridge	O	O	-	-	O
2. Concentrators					
8226 Token-Ring RJ45 Connection-001 Multistation Access Unit	-	O	-	-	-
8228 Multistation Access Unit	-	O	-	-	-
8244 FDDI Workgroup Concentrator	-	-	O	-	-
3. Hubs					
8222 Ethernet Workgroup Hub	O	-	-	-	-
8223 Fast Ethernet Workgroup Hub	-	-	-	O	-
8224 Ethernet Stackable Hub	O	-	-	-	-
8225 Fast Ethernet Stackable Hub	O	-	-	O	-
8230 Token-Ring Network Controlled Access Unit	-	O	-	-	-
8237 Ethernet Stackable Hub	O	-	-	-	-
8238 Token-Ring Stackable Hub	-	O	-	-	-
8250 Multiprotocol Intelligent Hub	O	O	O	-	-
8260 Nways Multiprotocol Switching Hub	O	O	O	O	O
8265 Nways ATM Switching Hub	-	-	-	-	O
8276 Nways Ethernet RoutePort	O	-	-	-	-
4. Switches					
8270 Nways Token-Ring Switch	-	O	-	-	O
8271 Nways Ethernet LAN Switch	O	-	-	O	O
8272 Nways Token-Ring LAN Switch	-	O	-	O	O
8273 Ethernet RouteSwitch	O	O	O	O	O
8274 Nways LAN RouteSwitch	O	O	O	O	O
8285 Nways ATM Workgroup Switch	-	-	-	-	O
5. ATM Service Servers					
8210 Nways Multiprotocol Switched Services (MSS) Server	-	-	O	-	O

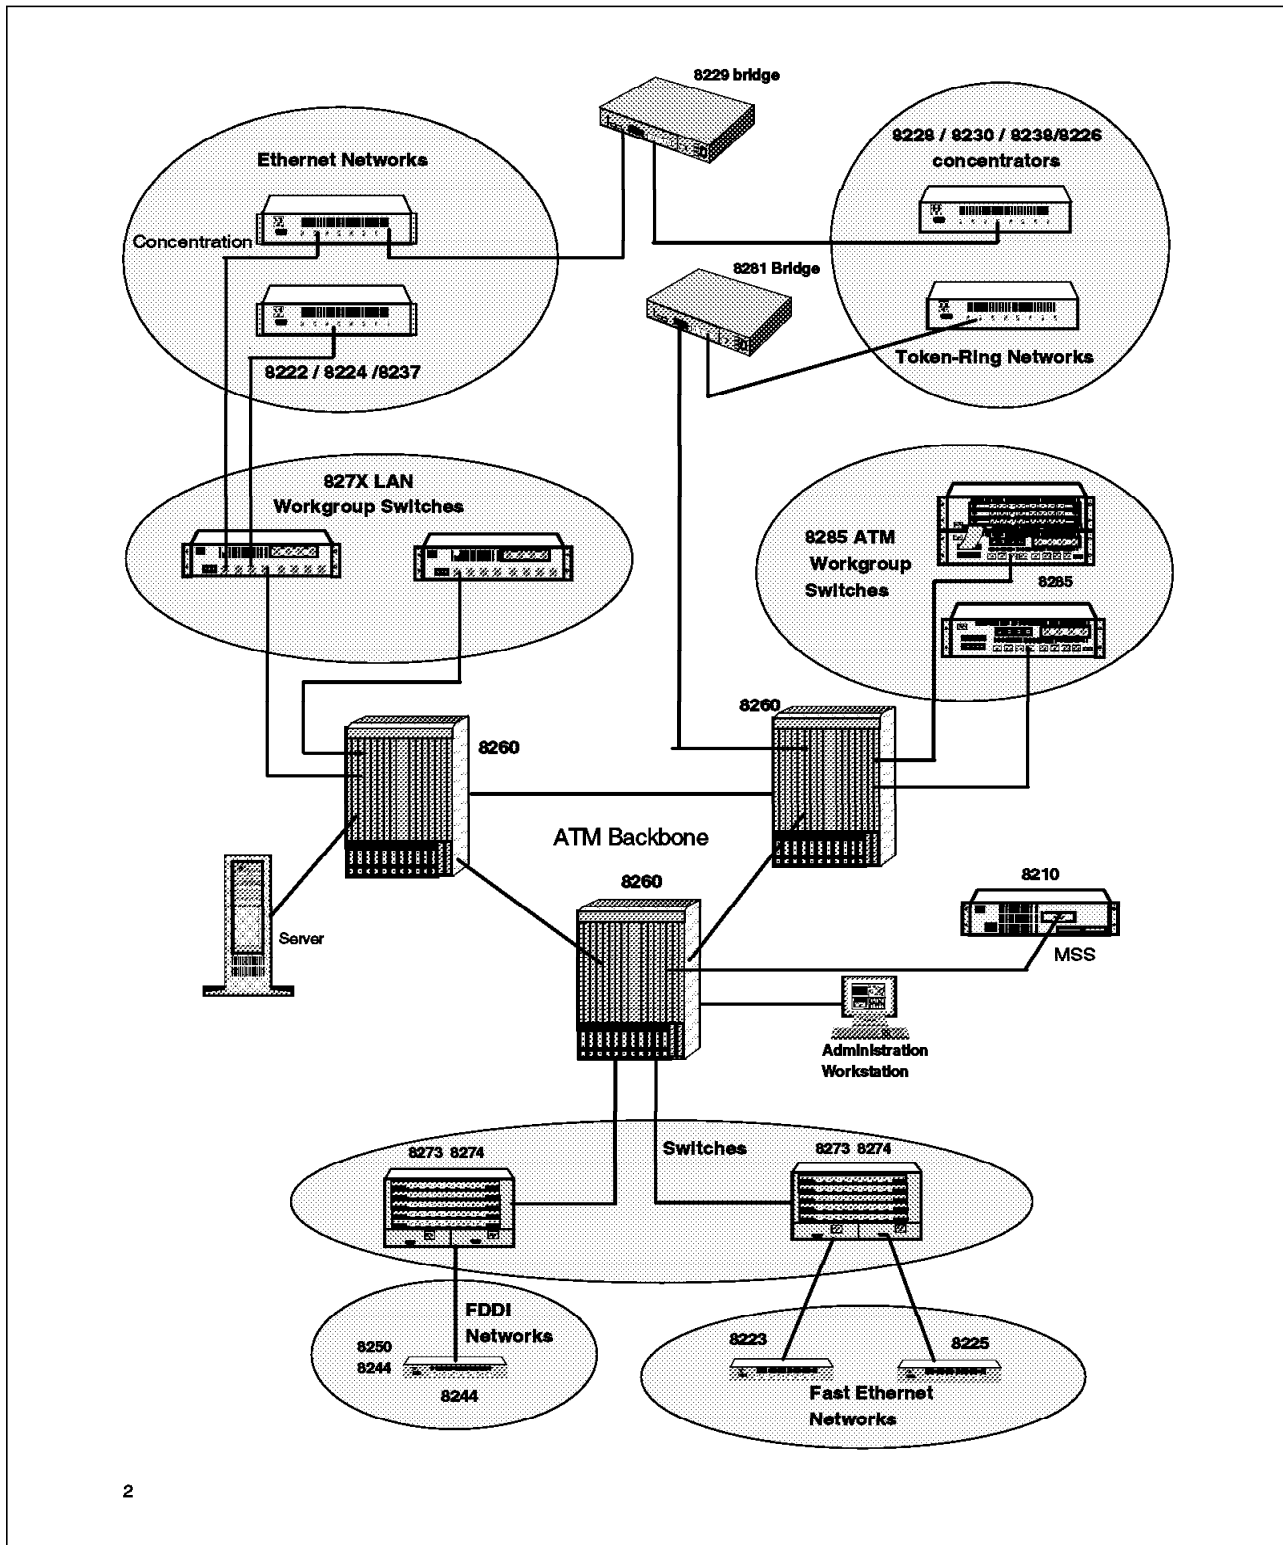
1.2.2 Functions

The IBM campus network products are classified by the following functions:

1. Legacy LAN Bridging: bridging capability between Ethernet, token-ring, FDDI and fast-Ethernet networks
2. WAN attachment
3. Grouping: separate segment/collision-domain in one box per port/slot
4. VLAN: various kinds of policies (port, protocol, MAC address, network address and user defined)
5. ATM Bridging: connection to emulated LANs and legacy LANs to use LAN Emulation, Classical IP and so on
6. ATM Switch: forwarding/connecting between end systems on ATM networks
7. Routing: Not the main function in these products

<i>Table 2. Classification of IBM Campus Products by Functions</i>							
	1	2	3	4	5	6	7
	LAN Bridging	WAN Attachment	Grouping	VLAN	ATM Bridging	ATM Switch	Routing
1. Bridges							
IBM 8229 Bridge	O	O	-	-	-	-	-
8281 Nways ATM LAN Bridge	O	-	-	-	O	-	-
2. Concentrators							
8226 Token-Ring RJ45 Connection-001 Multistation Access Unit	-	-	-	-	-	-	-
8228 Multistation Access Unit	-	-	-	-	-	-	-
8244 FDDI Workgroup Concentrator	-	-	-	-	-	-	-
3. Hubs							
8222 Ethernet Workgroup Hub	-	-	-	-	-	-	-
8223 Fast Ethernet Workgroup Hub	-	-	-	-	-	-	-
8224 Ethernet Stackable Hub	-	-	-	-	-	-	-
8225 Fast Ethernet Stackable Hub	-	-	-	-	-	-	-
8230 Token-Ring Network Controlled Access Unit	-	-	-	-	-	-	-
8237 Ethernet Stackable Hub	-	-	-	-	-	-	-
8238 Token-Ring Stackable Hub	-	-	-	-	-	-	-
8250 Multiprotocol Intelligent Hub	O	-	O	-	-	-	-
8260 Nways Multiprotocol Switching Hub	O	O	O	-	O	O	-
8265 Nways ATM Switching Hub	O	O	-	-	O	O	-
8276 Nways Ethernet RoutePort	-	-	O	-	-	-	-
4. Switches							
8270 Nways Token-Ring Switch	-	-	-	-	O	-	-
8271 Nways Ethernet LAN Switch	-	-	O	-	O	-	-
8272 Nways Token-Ring LAN Switch	-	-	O	-	O	-	-
8273 Ethernet RouteSwitch	O	O	O	O	O	-	O
8274 Nways LAN RouteSwitch	O	O	O	O	O	-	O
8285 Nways ATM Workgroup Switch	-	O	-	-	O	O	-
5. ATM Service Server							
8210 Nways Multiprotocol Switched Services (MSS) Server	-	-	-	-	O	-	O

1.3 Connection Example



2

Figure 1. Topology of IBM Campus Network Products

Chapter 2. Main Concepts of Legacy LANs

This chapter contains an introduction to the most important legacy LAN concepts, including such important things as frame types, frame conversion, and spanning tree. This will be especially useful for a thorough understanding of the most relevant concepts needed by the networking specialist when troubleshooting difficult situations. Needless to say, it will also be of use for anybody wishing to better understand the real issues in legacy LAN networking, which of course expand also into the ATM-based campus networks.

For a more detailed introduction we refer you to the books listed in Appendix H, "Related Publications" on page 517.

LANs are networks used to connect user workstations (for example PCs) in a building or a campus to shared resources such as host applications, printers, and other services. A LAN installation may consist of a single or many segments. Several segments are connected together by using devices called bridges, switches, and routers. When talking about LANs connecting several buildings we are also speaking of campus LANs.

2.1 The OSI Networking Model

The following networking layered model issued by the Open Systems Interconnection (OSI) Committee will be used as a reference when discussing some of the following LAN concepts.

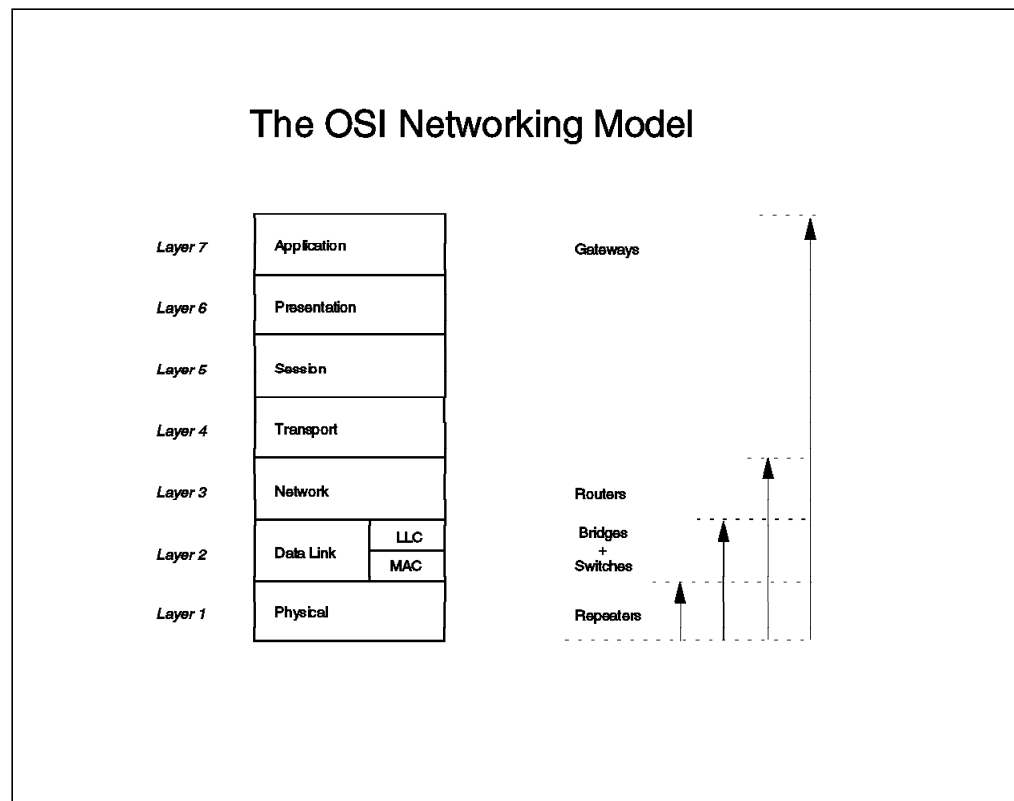


Figure 2. The OSI Networking Model

When discussing LANs, we usually subdivide the OSI layer 2 into two sublayers called:

- Medium access control (MAC) sublayer, which is the one dealing with the details of a certain LAN implementation, for example Ethernet or token-ring
- Logical link control (LLC), which makes operation of the upper layer protocols more independent of the actual MAC layer technology being used

2.2 Ethernet LAN Architectures

The following section is a summary of the key information relative to Ethernet networks.

2.2.1 Main Characteristics

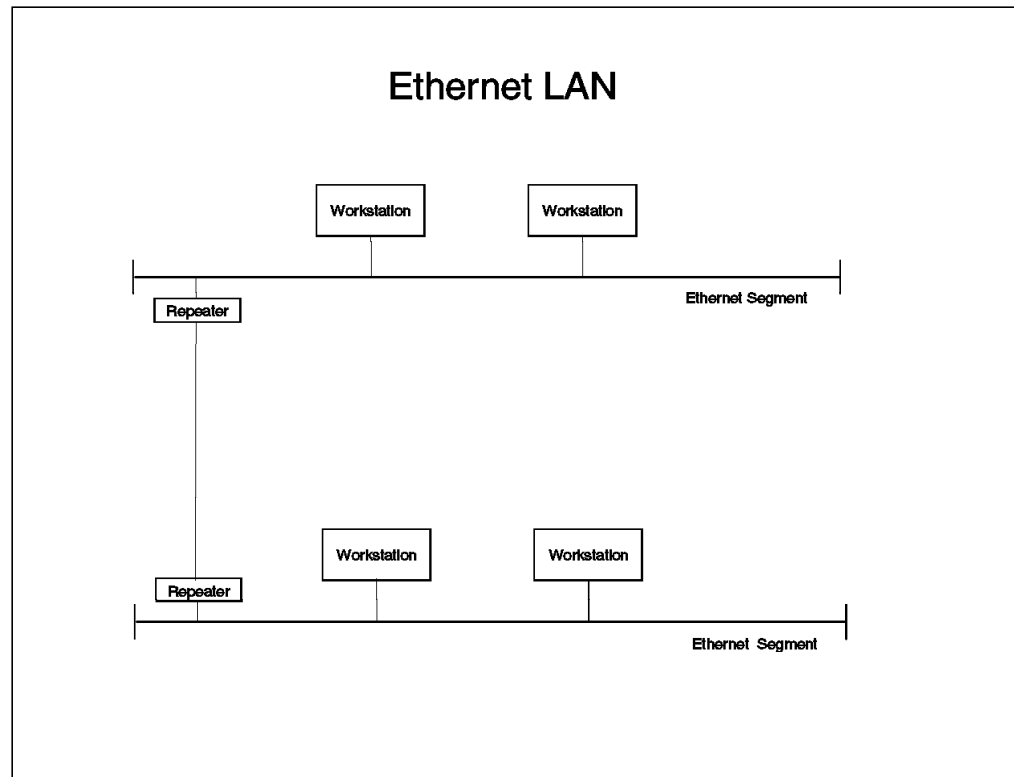


Figure 3. An Ethernet LAN

Workstations attached to an Ethernet LAN share the available segment bandwidth of 10 Mbps among themselves. The workstation access for sending data over the LAN is controlled via a mechanism called carrier sense multiple access with collision detection (CSMA/CD). Since stations will detect attempts for involuntary simultaneous access through actual collisions, we speak here of a *collision domain*. Up to 1024 stations may reside in a collision domain.

Different implementations of Ethernet cabling have become popular, as for example:

- 10Base-2: RG58 coax-based, also known as Thin-Ethernet. Segments may extend over a distance of up to 185 m, with as many as 30 stations per segment (at least 0.5 m apart from each other).
- 10Base-5: Thick coax-based, also known as Thick-Ethernet, sometimes also referred to as *the yellow cable*. Segments may extend over a distance of up to 500 m, and contain as many as 500 stations (attached to the cable every 2.5 m or increments thereof).
- 10Base-T: Uses two twisted pairs of wire (of usually eight wires), and is actually a repeater-like technology. Each station has its own segment, but on the same collision domain. Depending on the technology used in a certain hub, half a repeater-count or none at all will be introduced when connecting

a station to the hub. New installations will very likely use the newest cable for 10Base-T called Category 5 UTP cable, especially because this cable is specified for up to 100 Mbps. This would allow for a later migration to higher transmission technologies than the 10 Mbps used with Ethernet.

- FOIRL: The original repeater standard for connecting Ethernet via optical repeaters. The connection between repeater devices is asynchronous, and runs over multimode fiber.
- 10Base-FL: Newer optical repeater connectivity option than FOIRL. Today possibly the most popular one, since it can be used to connect hubs to hubs (both acting as repeaters), and stations to hubs. The connection between repeater devices is asynchronous, and runs over multimode fiber. FOIRL signalling allows adapters and hubs to easily detect if the partner at the other end is present or not, and is therefore also used to connect station adapters with two FOIRL fibers to two different hub ports (port redundancy), the hub ports usually being located on two different hubs.
- 10Base-FB: The newest optical repeater connectivity option, and the most reliable one. Connection between repeaters is synchronous, and runs over multimode fiber. Some repeater implementations are able to check the implementation of their partner, and go back to 10Base-FL if necessary, but most implementations will not support this.

Segments belonging to the same collision domain can be connected to each other via repeaters (OSI layer 1 devices), which will simply copy frames from one segment to the next. Since all stations in the collision domain are required to recognize collisions (even the ones being farthest from each other) no more than four repeaters in a row may be used to interconnect segments in order to guarantee proper collision recognition (speed of signal propagation is the limiting factor). This is known as the Four-Repeater rule. There is one more restriction, though: There can be a maximum of five possible segments between any two stations, only up to three may be coax cables (10Base-5 or 10Base-2), the other two must be *link segments* (10Base-T, FOIRL, 10Base-FL, or 10Base-FB) without any station connected to them.

For many years the only standard for Ethernet was an industry standard called Ethernet V2 or DIX (for Digital, Intel, Xerox). Later on the Institute of Electrical and Electronic Engineers (IEEE) issued a new specification for Ethernet called IEEE 802.3. The media over which the data traffic runs remained basically the same (apart from some changes to the 10Base-5 physical connections). The largest difference was a redefined frame format known also by the new standard's name.

2.2.2 Ethernet V2 Frame Format (DIX)

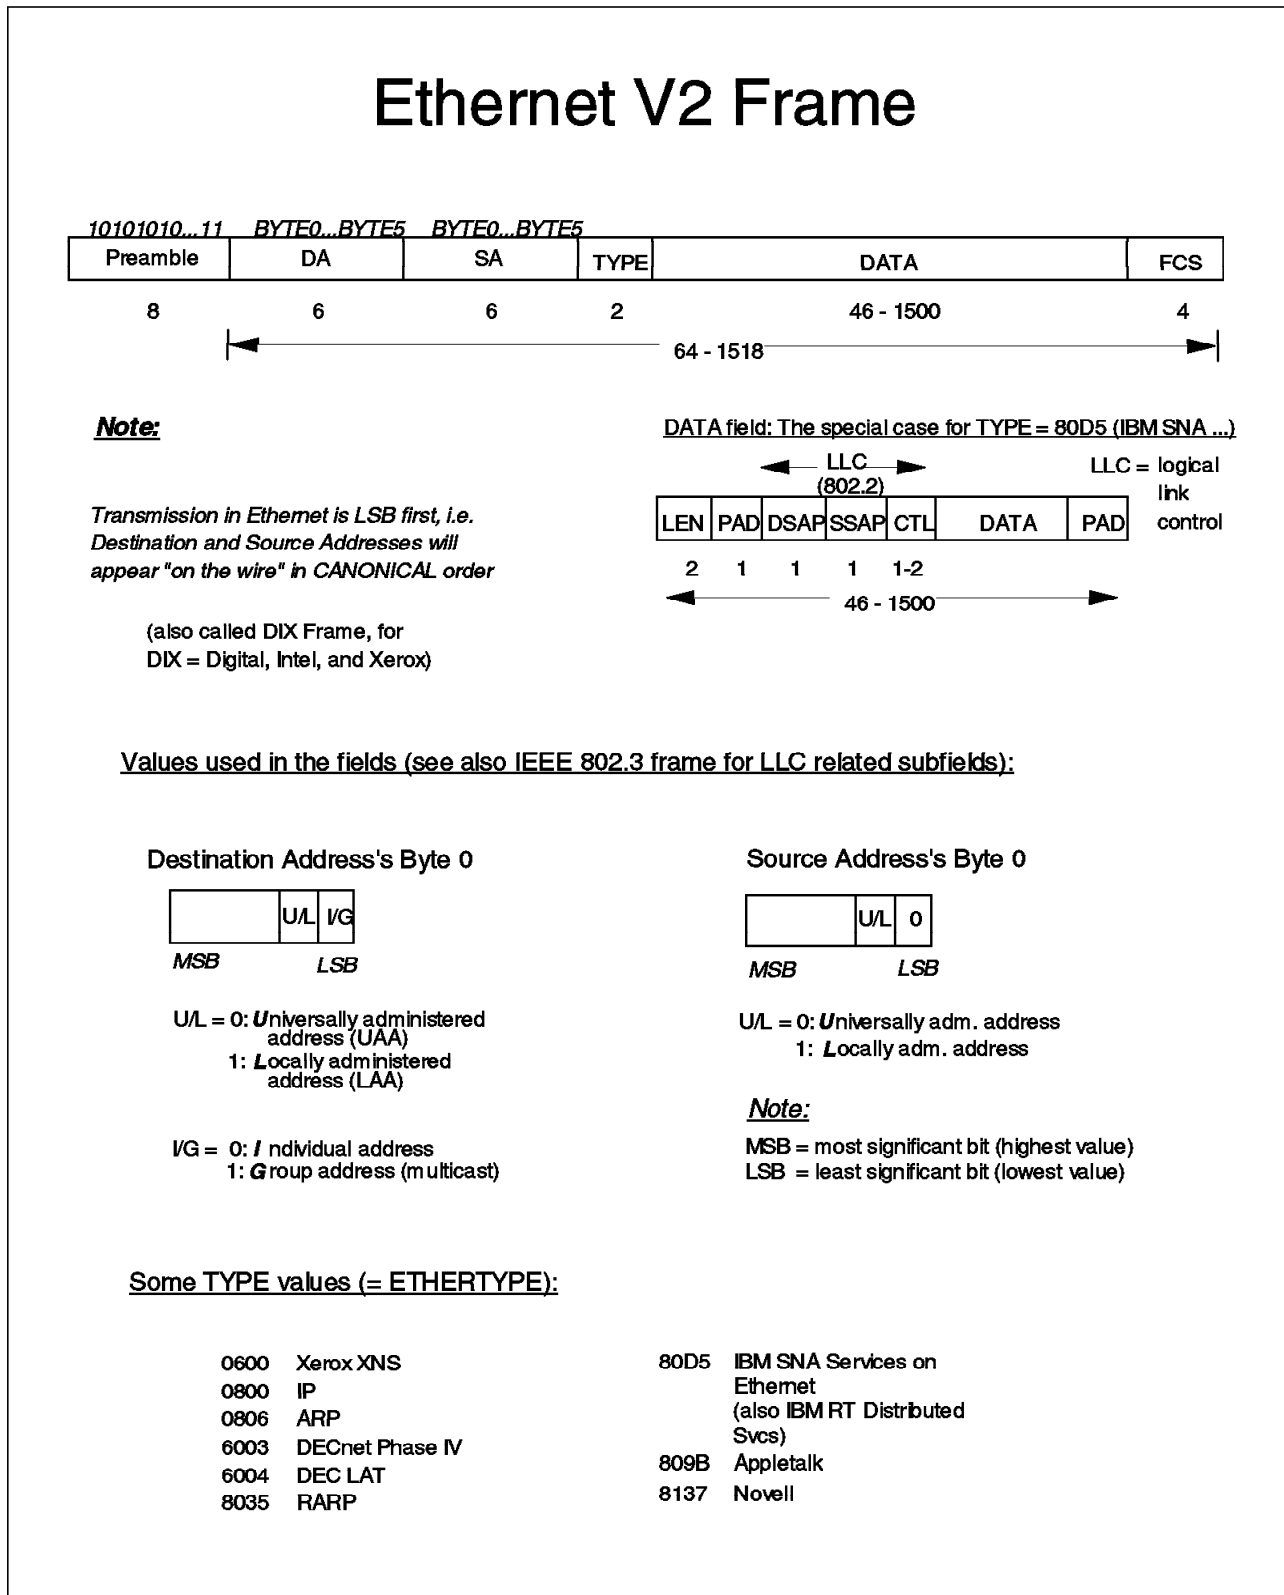


Figure 4. Ethernet V2 (DIX) Frame Format

Figure 4 offers quite a detailed view of the structure of an Ethernet V2 frame. In the bottom of the picture the overall structure is shown:

- With the unique bit combination of the preamble, an Ethernet adapter will be able to recognize the start of a frame.
- The destination address (DA) is the MAC address of the station to which the frame is being sent. It can be either a unicast address (frame sent to a single station), a multicast address (frame sent to a group of stations, for example all Ethernet bridges) or a broadcast address (frame sent to all stations). The broadcast address is a string of 48-bits set all to 1s.
- The source address (SA) is the MAC address, given by the sender. This is the address of the station sending the frame. Like the destination address it is a sequence of 6 bytes or 48 bits.

The bytes of the source and destination addresses are sent over the medium in so-called *canonical order*, that is with the least significant bit (LSB) first. Note that certain bits both in the source and destination addresses have some special meanings. See Figure 4 on page 11 for more details.

- The type field denotes the upper-layer protocol being used for communication between the partner stations for information exchange (such as NetBIOS, TCP/IP, NetWare, etc.). Since the values used in this field are also used in frames of other types of LANs, the field is sometimes referred to as ETHERTYPE.
- 802.2 Logical Link Control (LLC) header: When the mentioned type field contains a hexadecimal value of 80D5, the rest of the frame will contain 802.2 (LLC) encapsulated IBM protocol data, usually SNA or NetBIOS, depending on the DSAP and SSAP values contained in the LLC header. The DSAP and SSAP fields are explained in more detail in the section on Ethernet IEEE 802.3 frames. Please refer to Figure 5 on page 13.

Logical link control (LLC) is used for example (among others) to make sure that a certain sequence of frames has been completely received and is in the proper order; otherwise the partner stations will retry until the desired information has been transmitted correctly. Remember that Ethernet is a best-effort transfer technology with no guarantee for the frames to arrive at their destination, and that SNA was built for bullet-proof transmission with built-in error recovery.

- Frame check sequence (FCS). This field is a 32-bit cyclic redundancy check (CRC) field for the previous contents of the frame (from destination address up to and including all the data). The receiving station will discard any frame if the contents of its FCS field and a new computed CRC value do not match, thus ensuring that only correct frames are forwarded to the higher protocol layers upon reception.

2.2.3 Ethernet IEEE 802.3 Frame Format

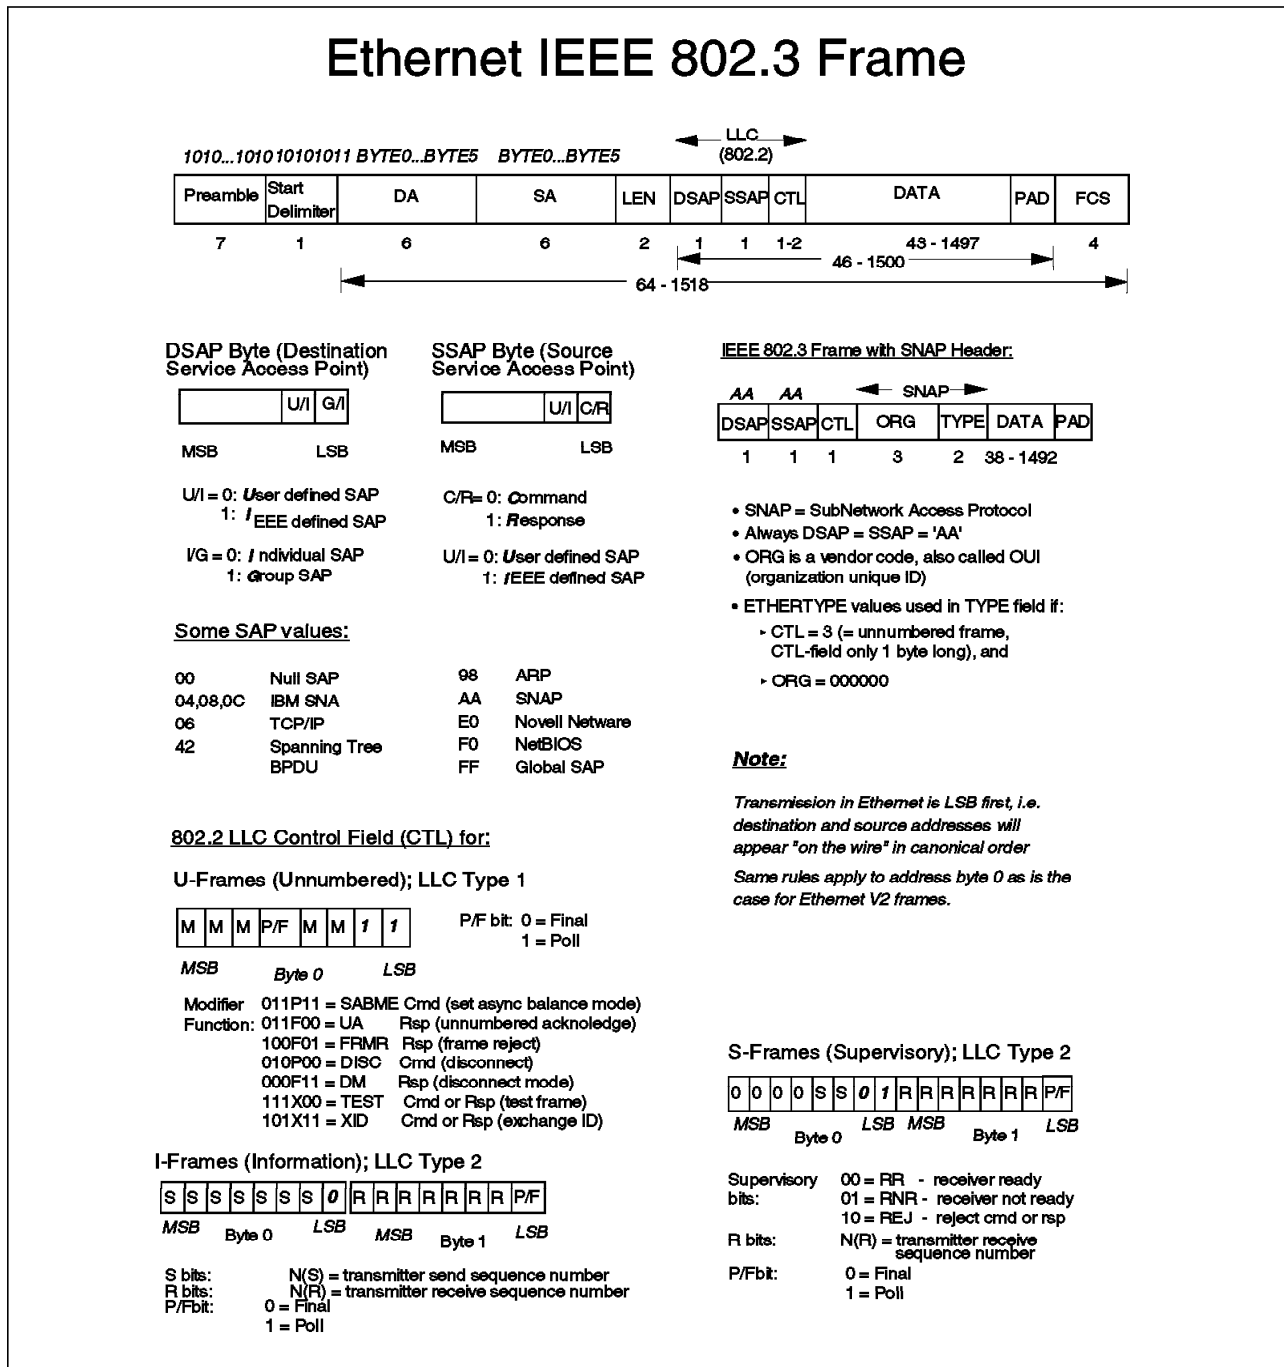


Figure 5. Ethernet IEEE 802.3 Frame Format

As can be seen from Figure 5 the Ethernet IEEE 802.3 is very similar to the Ethernet V2 frame, but with the following changes and or extensions:

- Ethernet IEEE 802.3 preamble is actually equal to the first seven bytes of the preamble found in Ethernet V2 frames (see also page 12).
- Start Delimiter byte (SD) with the same contents as the eighth byte found in Ethernet V2 frames (see also page 12).

- Length field: Indicates the total length of the following fields (DSAP, SSAP, CTL, and DATA, but without PAD and FCS). Note that this field is located at the same place as the type field in Ethernet V2 (see also page 12). If the value of this byte is equal to or less than 1500 (decimal) then the frame is recognized as an IEEE 802.3; otherwise it will be recognized as an Ethernet V2 frame.
- Destination service access point (DSAP): The value contained in this byte is used to recognize the upper-layer protocol entity in the destination station (for example SNA or NetBIOS). This is the first of three fields used for 802.3 logical link control (LLC); the two others are SSAP and CTL (see below). The three mentioned fields together are also referred to as the LLC header. The LLC header is a mandatory field in IEEE 802.3. Although as we shall see later, a proprietary frame type of NetWare known as 802.3 raw frame, which was defined before the IEEE 802.3 standard was in place, does not obey this rule.
- Source service access point (SSAP): This byte is used to denote which protocol entity is sending the frame.
- Control field (CF): This is the last LLC field and it may contain one (LLC type 1, unnumbered frames) or two bytes (LLC type 2, numbered frames). Numbered frames are used for example to guarantee orderly and complete delivery of a sequence of frames, and some higher-level protocols, such as SNA, will require the use of these types of LLC frames.
- Subnetwork access protocol (SNAP): This field is only used if the previous DSAP and SSAP fields contain the hexadecimal value of AA. The need for this field arose due to the size of the SAP fields which are just too short to accommodate many protocols, since each SAP field is actually only 6 bits long (if we overlook the two other special bits in the same byte). Therefore, in IEEE 802.1a a further header was defined, the SNAP header, which is positioned before the actual data. With this enhancement (see type subfield below), a wider selection of protocols can be used with IEEE 802.3 frames than would be possible when using only SAP fields. SNAP frames are unnumbered frames; therefore the previous CTL field will consist only of one byte and will have a value of 3 (for unnumbered frames). As can be seen in Figure 5 on page 13 the SNAP header consists of three subfields:
 - ORG: or organization code. This is the ID allocated by IEEE to a certain vendor (adapter manufacturer, etc.). In IEEE 802.3 SNAP frames the ORG field is set to all 0s. This field is sometimes also referred to as OUI for Organization Unique Identifier.
 - TYPE: This field exactly matches the TYPE (or ETHERTYPE) field of Ethernet V2 frames, and uses the same values as defined for the Ethernet V2 frames (see page 12).
 - PAD: Ethernet IEEE 802.3 frames must be 64 bytes or longer (including DA, SA, FCS) to allow for proper collision detection. Therefore padding characters (this field) are added in order to adjust too short frames (so-called runt frames) to the minimum length of 64 bytes.

2.2.4 Switched Ethernet

Ethernet LAN switches are devices used to pass Ethernet frames from one segment to another very quickly. They differ from Ethernet bridges (see below) in that they may pass frames between several pairs of segments simultaneously. In contrast, bridges pass frames from one segment to another using a store-and-forward strategy; that is, the frame being received through the input port is fully stored in an internal buffer before it gets forwarded through the output port. LAN switches, on the other hand, normally work in cut-through mode; that is, the switch starts forwarding the frame to the output port before the end of the frame has arrived at the input port. An exception to this is when the LAN switch is being used to actually bridge between technologies of different speed such as Ethernet and fast Ethernet (see below for fast Ethernet).

LAN switches work on layer 2 of the OSI model, thus confining unicast traffic to the local segment; therefore they are often used to split overloaded Ethernet segments providing more bandwidth for the local stations. Maximum bandwidth per station can be reached when a single station is directly connected to a port of such an Ethernet LAN switch (which is often the case with file and application servers).

2.2.5 Fast Ethernet (802.3u)

This is a newer high bandwidth technology following the same protocol structure as Ethernet. The big difference is the segment speed which is 100 Mbps, resulting in the name Fast Ethernet. The motivation for the technology was to gradually move the installed base of servers and user workstations to a higher speed network according to communication needs, allowing users to spend their money where the greatest need for upgrading is felt. In any case workstations attaching to Fast Ethernet need new adapters and fast internal workstation buses, which may require a replacement of older workstation equipment to benefit from the higher LAN speed.

Different cabling schemes exist:

- 100Base-TX: For UTP Category 5 cable; it uses two pairs of wires (usually eight wires installed per cable). This is the most popular, and so is the one we will continue to refer to throughout this document.
- 100Base-T4: For UTP Category 3 cable (telephone wire); it requires four pairs of wires, with split bandwidth over the wire pairs.
- 100Base/FX: For optical fiber connections (multimode).

The following different components are also offered in the market to ease migration steps:

- Switches with auto-sensing ports, to allow connection of 100 and 10 Mbps segments.
- Auto-sensing Ethernet cards for workstations to be connected to 10 or 100 Mbps segments according to progress of deployment of the equipment with the higher speed (concentrators, switches) in a given environment.

2.2.6 100VG-AnyLAN (802.12)

This is a technology for 100 Mbps meant as a growth path for both Ethernet and token-ring, though it has not reached the same widespread acceptance as Fast Ethernet. 100VG means 100 Mbps over voice-grade cable, which is a low quality telephony cable known as UTP Category 3. This technology foresees use of all eight wires of the cable, splitting the bandwidth over the wire pairs, and uses a sequential polling mechanism for access control. Due to the higher speed being run over the cable, great care should be taken when migrating from Ethernet or token-ring to 100VG-AnyLAN. Big problems could arise if the installed cable does not really match at least the mentioned UTP Category 3 quality.

2.2.7 Gigabit Ethernet (802.3z)

This is an emerging standard (expected March 1998) meant to increase available bandwidth to 1 Gbps. Due to the increase in bandwidth, maximum segment lengths will decrease to 500 m for multimode fiber and 25 m for UTP Category 5 cabling, making this technology more suitable for data centers than for large distributed user environments. Some pre-standard products are already being offered on the market. The MAC layer continues to use the well-known CSMA/CD access scheme, while the implementation of the PHY layer may vary in the technology used. Some products may start shipping with Fiber Channel Standard (FCS) technology, while others will use Serial HIPPI (serial version of the High Performance Parallel Interface, a bus specification mainly used today as a high-speed computer interconnect). Due to the high speed involved, special attention should be given to the planning and installation of cabling components; troubleshooting problems during operations would probably start by looking at the installed cabling.

2.3 Token-Ring IEEE 802.5 LAN Architecture

This section provides information that must be known to support token-ring networks.

2.3.1 Main Characteristics

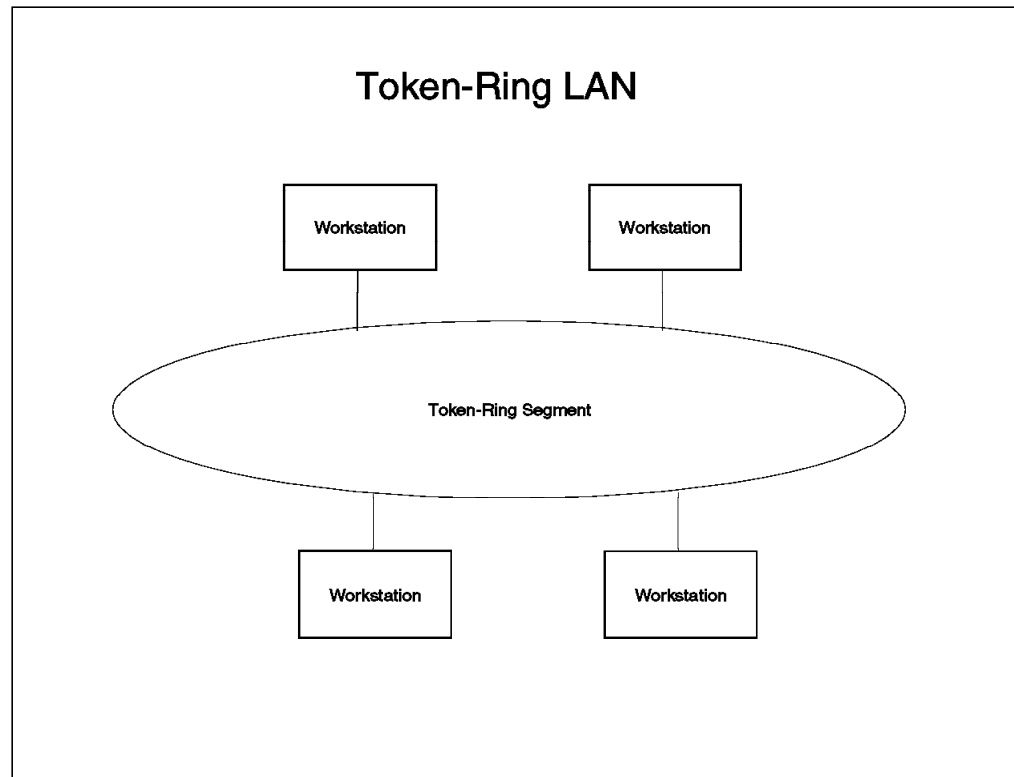


Figure 6. A Token-Ring LAN

Workstations attached to a token-ring LAN share the available segment bandwidth of 4 or 16 Mbps among them. The workstation access for sending data over the LAN is controlled by allocation of a token flowing around the ring. A token is just a special type of frame. When a station receives the token it may then send a frame over the LAN, and release a new token again as soon as it has received back its own frame. Alternatively if the station is not ready to send any data, it will just release the original token for use by another station. Numerous management mechanisms and self-healing strategies have been specified for the token-ring architecture, many of which make use of MAC sublayer frames.

Different implementations of token-ring cables are in use today, for example:

- Shielded twisted pair (STP): Sometimes also referred to as ICS (for IBM cabling system) or Type 1 cable (for the most widely used cable type specification). This cable uses two pairs of wires, and its electrical impedance is equal to 150 Ohms. It is used both for lobe (station) cables, and ring-in/ring-out connections between hubs.
- Unshielded twisted pair (UTP): This cable is available according to different specifications, called Category 3, 4, and 5. Category 3 is basically telephone wire, and Category 5 the one with the best capabilities for data transfer

(longest distances, least cross-talk, etc.). UTP cable uses two pairs of wires (usually eight wires), but with a different pin allocation than Ethernet. The electrical impedance of UTP cable is equal to 100 Ohms. A few countries use 120 Ohms. Depending on the hub implementation, this cable is used both for lobe (station) cables, and also for ring-in/ring-out connections.


- Multimode fiber: Used for ring-in/ring-out connections between hubs to overcome longer distances, or simply to protect the link from electromagnetic emissions in manufacturing environments and the like.

Token-ring cabling is actually a star-wired topology, with the stations being attached to a control access unit (CAU) or a hub using so-called lobe cables. CAUs and hubs will be usually located in a satellite or cabling room. CAUs and hubs are interconnected together with so-called Ring-In/Ring-Out (RI/RO) trunks to form larger token-rings, allowing the connection of a larger number of stations to a single ring. The number of stations attached to a ring depends on many factors such as ring speed, cable types being used, total ring length, etc. For planning purposes, please refer to the books listed in Appendix H, "Related Publications" on page 517.


2.3.1.1 Frame Format

IBM as the inventor of the token-ring has fully adopted the later released IEEE specification for the token-ring frame called IEEE 802.5.


IEEE 802.5 Frame (Token-Ring) (1 of 2)



Fields in MAC Frames:




- MVID: Major Vector ID
- DC/SC: Destination/Source Class
 - Ring Station
 - Ring Error Monitor
 - RPL Server




- C/S: Major Vector ID
- R/O: Required/optional indicator
- P: Code Point bit

MAC Frames:




The architecture describes 28 different MAC frames, each identified by a unique Major Vector Identifier (MVID)

LLC Frames:



Frames with SNAP Header:



- SNAP = SubNetwork Access Protocol
- Always DSAP = SSAP = 'AA'
- ORG is a vendor code, also called OUI (organization unique ID)
- ETHERTYPE values used in TYPE field if:
 - CTL = 3 (= unnumbered frame, CTL-field only 1 byte long), and
 - ORG = 000000

Starting Delimiter (SD):




MSB LSB
 J,K: Differential Manchester code violations

Access Control (AC):



MSB LSB
 P: Priority bits
 T: Token bit
 M: Monitor bit
 R: reserved bits

Frame Control (FC):




MSB LSB
 F: Frame type
 00 = MAC frame
 01 = LLC frame
 10 = undefined
 11 = undefined
 Z: Control bits (used with MAC frames)
 0000 = normally buffered frame
 0001 to 1111 = express-buffered frame

Destination Address's special bits:



- I/G = 0: Individual address
- 1: Group address (multicast)
- U/L = 0: Universally administered address (UAA)
- 1: Locally administered address (LAA)
- FAI : Functional Address Indicator (meaningful only with group addresses, i.e. if I/G-bit = 1)
- = 0: functional address
- = 1: group address

Source Address's special bits:




- R/I = 0: no RIF (Routing Information Field) follows
- 1: Routing Information Indicator (means that RIF follows)
- U/L = 0: Universally administered address (UAA)
- 1: Locally administered address (LAA)

Note:


Transmission in Token-Ring is MSB first, i.e. Destination and Source Addresses will appear "on the wire" in NON-CANONICAL order.

End Delimiter (ED):



MSB LSB
 J,K: Differential Manchester code violations
 I: Intermediate frame bit
 0 = single or last frame
 1 = first or intermediate frame
 E: Error detected bit
 0 = no error detected
 1 = error detected

Frame Status (FS):



MSB LSB
 A: Address recognized bits
 0 = address not recognized
 1 = address recognized
 C: Frame copied bits
 0 = frame not copied
 1 = frame copied

Frame Check Sequence (FCS):

A 32-bit CRC checksum of frame contents

Figure 7. IEEE 802.5 Token-Ring Frame Format (Part 1 of 2)

IEEE 802.5 Frame (Token-Ring) (2 of 2)

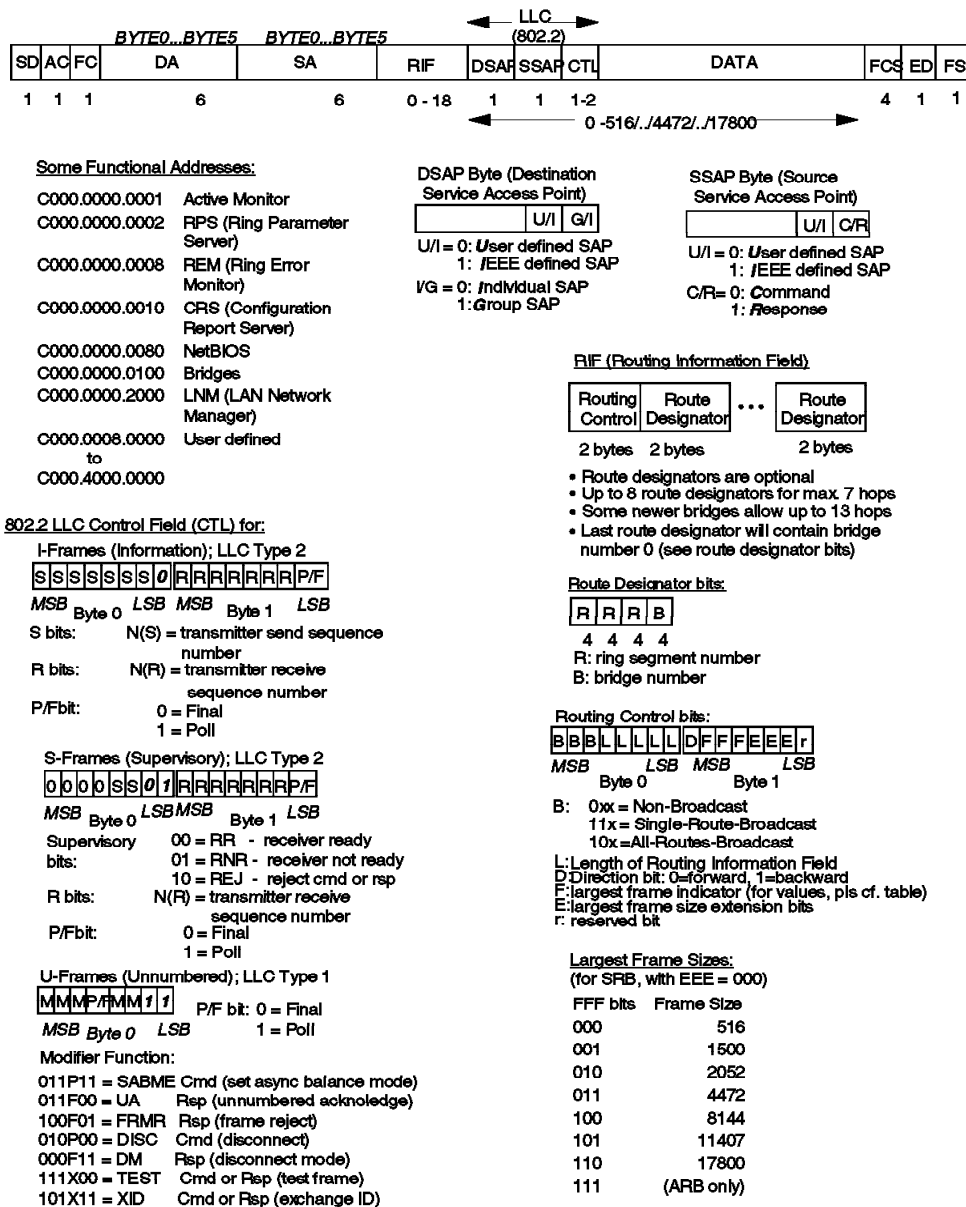


Figure 8. IEEE 802.5 Token-Ring Frame Format (Part 2 of 2)

A quite complete overview of the contents of a token-ring frame can be seen in Figure 7 on page 19 and Figure 8. Although the LAN architectures are very different, some fields in the token-ring frame look similar to those used in Ethernet IEEE 802.3 frames. Therefore we take a closer look only at the more important aspects differentiating both architectures, and do not repeat those items already covered in the Ethernet sections.

- Starting Delimiter (SD): This is used by the network adapter to recognize the beginning of a new token-ring frame.

- Access Control (AC): This field indicates the overall type of the frame (token, monitor, or normal MAC or LLC frame). A token will consist solely of starting delimiter (SD), access control (AC), and end delimiter (ED) fields.
- Frame Control (FC): This field indicates whether the frame is a MAC (mainly for ring management) or an LLC frame.
- Destination and source addresses (DA and SA): The bytes of destination and source MAC addresses are sent over the medium in so-called *non-canonical order*, that is with the most significant bit (MSB) first. Note that certain bits both in the source and destination addresses have some special meanings (see Figure 7 on page 19 for more details).
- Routing Information Field (RIF): This is appended and updated by source-route bridges along the route of the frame over several segments and bridges. A more detailed explanation of how this works can be found in 2.5.2, "Source-Route Bridging (SRB)" on page 31.
- MAC frame fields: These are MAC vectors and subvectors used in frames such as monitor present frames to preserve and/or rebuild a healthy token-ring. For further details refer to the books listed in Appendix H, "Related Publications" on page 517.
- LLC frame fields: These are similar to those used in Ethernet IEEE 802.3 frames, although there is a very significant difference, namely the size of token-ring frames may be much larger than those used with Ethernet (DIX or IEEE 802.3). The maximum possible size over a chain of rings and bridges will be learned by means of the bridges adapting the contents of the RI field at the time of path discovery to the destination station (more details found in 2.5.2, "Source-Route Bridging (SRB)" on page 31). When bridging Ethernet and token-ring LANs, the frame size should be defined in such a way, that the frames can really be forwarded by the bridge, which means that the length of the DSAP, SSAP, CTL, and DATA fields together should not exceed 1500 bytes. Therefore the DATA part may be a maximum of 1492 bytes long when using SNAP frames.

2.3.1.2 Switched Token-Ring

Token-ring LAN switches are similar in concept to the Ethernet LAN switches regarding their use of multiport and cut-through architectures, but they may differ in the OSI layer 2 methods used for forwarding frames. Depending on the implementation, token-ring LAN switches may use either transparent bridging, source-route switching or source-route bridging techniques to accomplish this. Please refer to 2.5, "Bridges and Bridging Methodologies" on page 27 for an explanation of the terminology.

2.4 FDDI (Fiber Distributed Data Interface) LAN Architecture

This section contains information that must be known to support FDDI networks.

2.4.1 Main Characteristics

FDDI has been standardized by the X3T12 American National Standards Institute (ANSI) committee and later ratified by the International Organization for Standardization (ISO).

An FDDI ring may be built from concentrators and stations, which all share the bandwidth of 100 Mbps. The FDDI dual ring may be up to 200 km (125 miles) in length, which includes the total length of primary and secondary rings in case of a wrap around of ports due to primary ring disruption (see also Figure 9). This means that when designing an FDDI ring, the total cable length should not exceed 100 km (62.2 miles). An FDDI ring may contain up to 500 devices (stations and concentrators) attached to both primary and secondary rings.

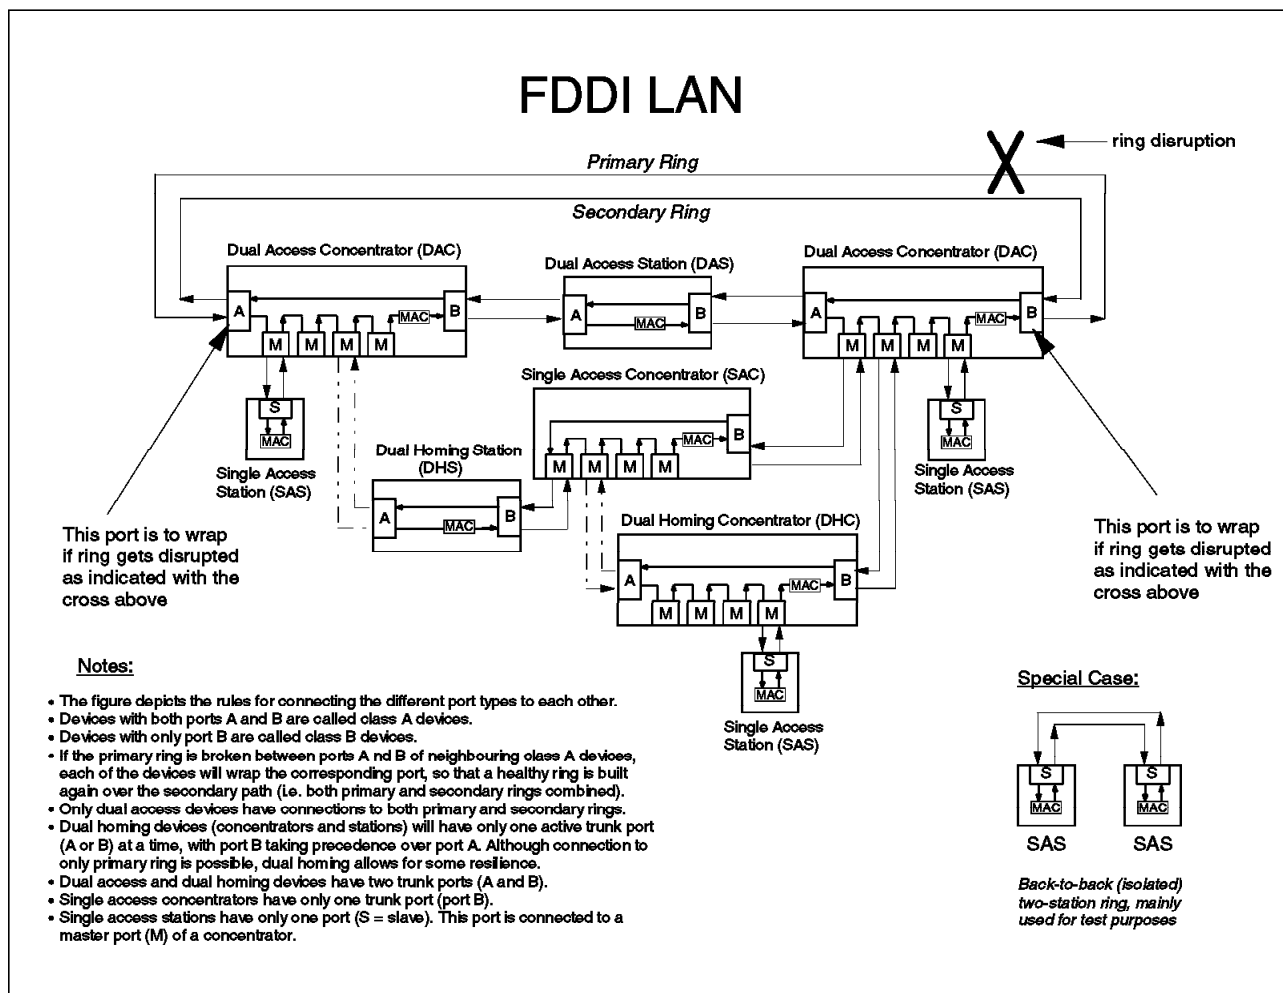


Figure 9. Elements in an FDDI LAN

As can be seen from Figure 9, the dual ring structure is built using so-called Class A devices, each with both an A and a B port:

- Dual Access Concentrators (DAC), which also have master ports (M) for attachment of single access devices to only the primary ring (see also

below), the master ports usually being used to connect user workstations to the FDDI LAN.

- Dual Access Stations (DAS), which in most cases would probably be LAN and applications servers with high availability requirements. In some LANs these could also be for example routers and bridges.

Note that port A has a primary ring in and a secondary ring out connection, and port B has a primary ring out and a secondary ring in connection.

There are also the following class B devices which connect only to the primary ring:

- Single Access Concentrators (SAC) with a B port to connect to the primary ring, and with many master ports (M) to connect stations to it. They are used to expand the number of available master ports (M).
- Single Access Stations (SAS) with a slave port (S), being normal LAN participants such as user workstations, and possibly routers and bridges. The S port is connected to the M port of a concentrator (either DAC or SAC).

Under normal conditions, only the primary ring is used for traffic flow. The secondary ring is used as the backup path in case the primary ring is disrupted at some point, in a similar manner as with token-ring. In Figure 9 on page 22 we can see the primary and the secondary ring at the very top of the picture. Unless the ring gets broken between ports connecting *dual access* devices, the data traffic will continue to flow only over the primary ring, and the secondary ring will remain unused.

Let us assume that the ring gets broken at the place marked with the big cross (indicated as ring disruption). Then port B of the right-most DAC (dual access concentrator) and port A of the left-most DAC will both wrap around. This way data flow will continue over the secondary ring from the right DAC over the DAS (dual access station) to the left DAC. Remember that only dual access devices (DAC and DAS) are connected to both the primary and the secondary ring; all other devices are only connected to the primary ring.

In order to avoid ring disruption by a failing dual access station or one just being switched off, optical bypasses can be used to connect the DAS to the dual ring. The optical bypasses will act as a signal front end to the station, and will keep the dual ring intact in case the station is switched off for any reason. Optical bypasses introduce a much larger signal loss than the one occurring with a normal station. Therefore use of only a certain number of optical bypasses will be possible in a given environment. The maximum number depends on the signal loss characteristics of the optical bypasses, which may vary from vendor to vendor.

Although there are many more differences than just the speed of 100 Mbps being used in FDDI, the FDDI architecture resembles in many aspects the architecture of token-ring:

- FDDI uses a token-passing mechanism very similar to the one used in the IEEE 802.5 Token-Ring architecture. Even the wrapping of a broken ring is part of the architecture.
- Some self-healing mechanisms used in FDDI are very similar to those used in token-ring (for example beaconing).

- As in token-ring, FDDI frames may be longer than those used with Ethernet: FDDI frames may be up to 4500 bytes long. Special attention should be paid when bridging between FDDI and other LAN types, since bridging of some protocols may not be supported by certain device implementations. This is especially true for protocols for which there is no standard in place for frame fragmentation. There is a standard way of fragmenting frames in TCP/IP. Therefore most installations will choose to connect FDDI and other LANs via routers instead of bridges.

The following cabling schemes exist for FDDI:

- CDDI: This is copper based FDDI, using UTP Category 5 (unshielded) cable, being the ANSI standard for copper medium.
- SDDI: This is copper based FDDI, using STP (shielded) cable, and as such a pre-standard implementation, but being used extensively due to existing cabling installations.
- MMF: For multimode optical fiber connections.
- SMF: For singlemode optical fiber connections for much larger distances.

FDDI has its own management standard called station management (SMT). Nevertheless, an SNMP MIB has been defined, so that SNMP management stations will be able to talk to a suitable SNMP agent, which in turn would be watching the FDDI ring using the SMT procedures (probably participating in the ring as a dual access station).

2.4.1.1 Frame Format

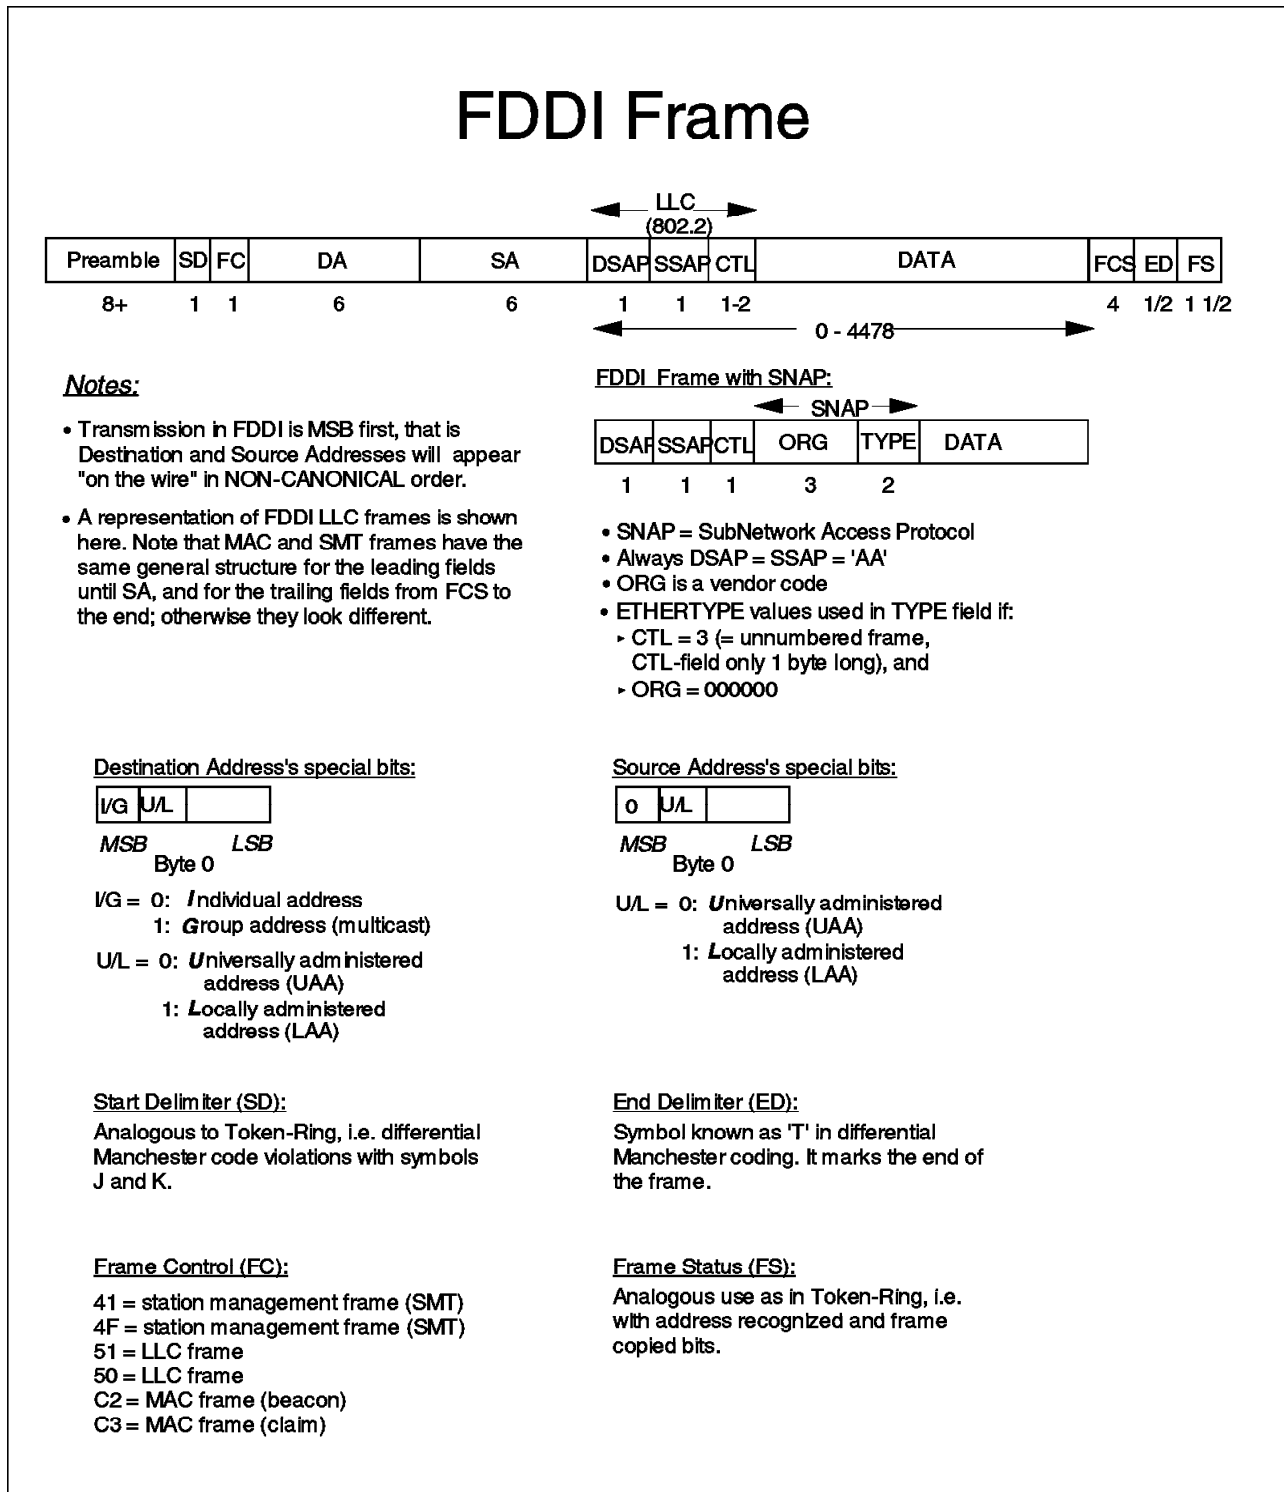


Figure 10. FDDI Frame Format

As seen in Figure 10 FDDI frames resemble token-ring frames. The bytes are transmitted over the medium with the most significant bit (MSB) first, so that MAC addresses will appear in non-canonical order (as in token-ring).

One big difference is that FDDI has MAC and LLC frames as does token-ring, but FDDI has in addition SMT frames (SMT stands for station management). The frame type is denoted by the value found in the Frame Control (FC) field. Another difference is that FDDI frames have no routing information field (RIF). In Figure 10 on page 25 only the LLC frame has been drawn. For a detailed outline of MAC and SMT frames please refer to the books mentioned in the reference section.

2.5 Bridges and Bridging Methodologies

Bridges are devices used to pass Data Link Layer information (OSI Layer 2 frames) from one LAN segment to another. Frames will be passed or not depending on destination MAC address, RI field (Ring Information field, only in token-ring), and protocol filters used. Depending on the technology used, bridges may connect like segments (for example two Ethernets) or segments of different kinds (for example Ethernet and token-ring). In passing frames between segments of different kinds, the frames may need to be converted (translated).

2.5.1 Transparent Bridging (TB)

Transparent bridges are used to interconnect Ethernet segments. Some implementations also exist to interconnect Ethernet LANs to FDDI backbones. In this case, we speak of hybrid or interconnect networks. In the latter case, care should be taken in defining the proper frame conversion (DIX and/or IEEE 802.3 to FDDI and vice versa) should the bridge device not take care of this itself.

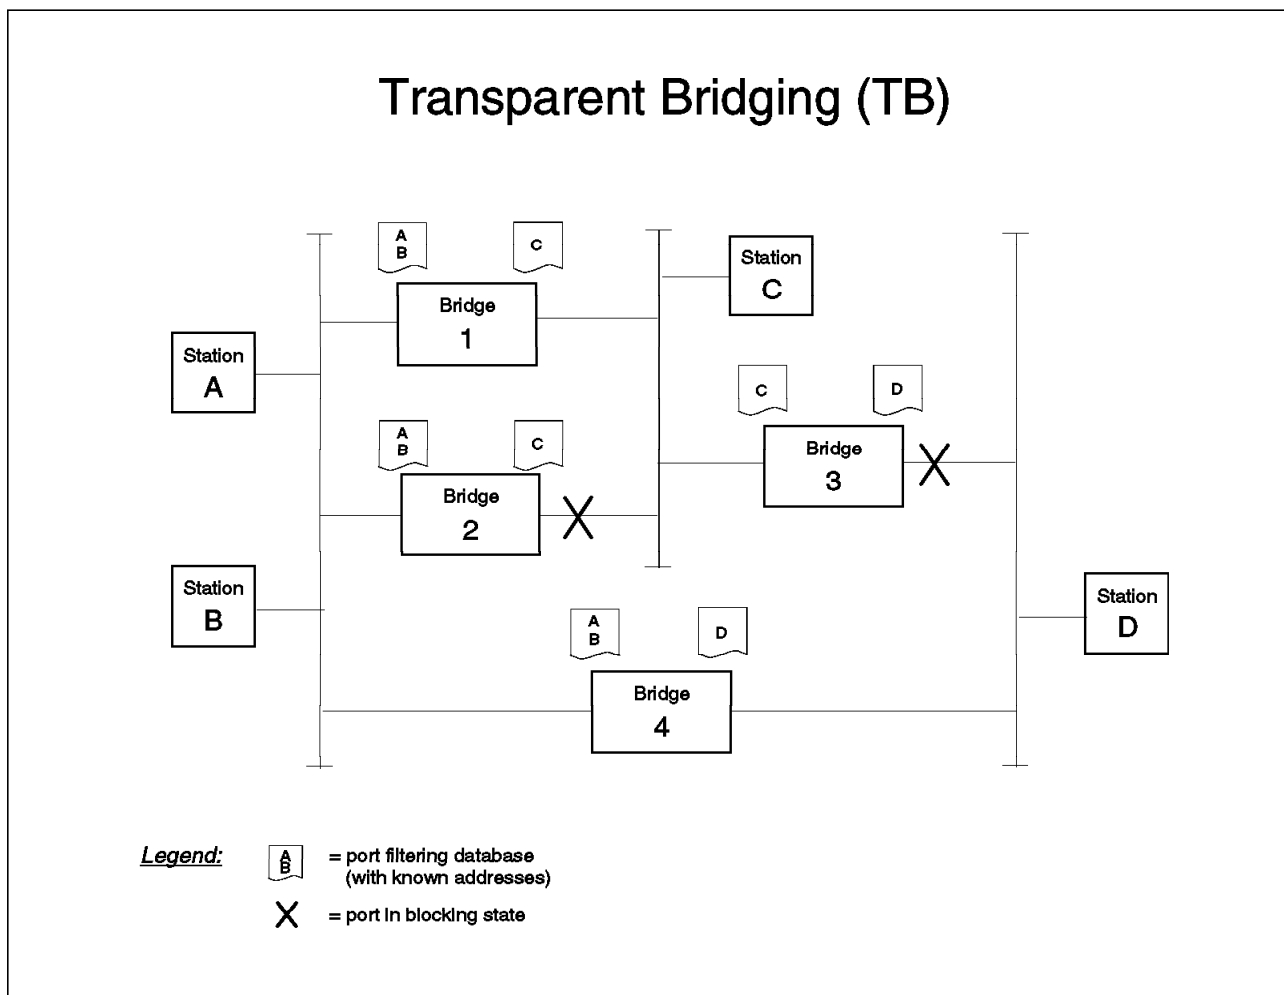


Figure 11. Transparent Bridging (TB)

When starting operations a transparent bridge will go through several states:

- Disabled: Initial state with no activity at the bridge ports.
- Blocking: Bridge participates in building the spanning tree (see 2.5.1.1, "Spanning Tree Protocol (STP)" on page 28), so it will send and receive

special frames called bridge protocol data units (BPDUs), and will not do anything else.

- Listening: In transition from blocking to learning, awaiting stable spanning tree.
- Learning: Bridge learns MAC addresses of stations attached to segments on the specific ports by reading the source addresses in the frames it receives. A filtering database with the learned addresses is built per port; no frame forwarding during this phase.
- Forwarding: Bridge forwards frames, learns new MAC addresses, ages out those MAC addresses not being active for a long time, and continues to participate in the spanning tree process.

In Figure 11 on page 27, bridge 1 has added MAC addresses of stations A and B to the filtering database of its left port, and MAC address of station C to the filtering database of its right port. The other bridges will behave similarly. In the case where station B sends a frame to station A, both bridges 1 and 4 will not forward it to their ports at the right, since the bridges know from their respective filtering databases that station A is on the same segment as station B.

On the other hand, in case station B sent a frame to station C or D, both bridges 1 and 4 would forward the frame through their respective right ports, because they know that neither station C nor station D is on the left segment.

2.5.1.1 Spanning Tree Protocol (STP)

The spanning tree algorithm and the spanning tree protocol (STP) are defined by the IEEE 802.1d standard. The spanning tree algorithm enables transparent bridges to establish a loop free topology (spanning tree), by having some bridges keep their ports in the blocking state. This way, it is ensured that no station will receive replicated frames circulating endlessly, and no station MAC address will be added to the filtering databases of both ports of any bridge in the network.

Every bridge participating in the spanning tree algorithm will initially assume it is the *root bridge*. Therefore it will send, at every Hello timer interval, Hello BPDUs (with its own address as the root bridge) through all of its ports, addressing all other 802.1d bridges (group address X'0180C2000000' in canonical format, or X'800143000000' in non-canonical format). A bridge will continue doing so until it gets a Hello BPDU with a higher priority, after which it will stop sending its own Hello BPDUs, and will only forward Hello BPDUs with the higher priority than its own. A higher priority BPDU is based on the lowest root ID (concatenated bridge priority and bridge MAC address), lowest path cost to the root bridge, lowest transmitting bridge ID (concatenation of bridge priority and bridge MAC address), and lowest port ID (concatenation of port priority and port number).

The bridge will send the higher priority BPDUs through all its ports except the one where the bridge itself received the higher priority BPDU. After a while, only (highest priority) root bridge BPDUs will be being forwarded by the bridges in the network. At this time the spanning tree has *converged*. Also each LAN will then have a designated bridge, being the only one that forwards the root bridge BPDUs into the segment. The port of the designated bridge from where it received the root bridge BPDU is called the root port, and the ports where the BPDUs are forwarded are called the designated ports; all other ports will just stay in the blocking state.

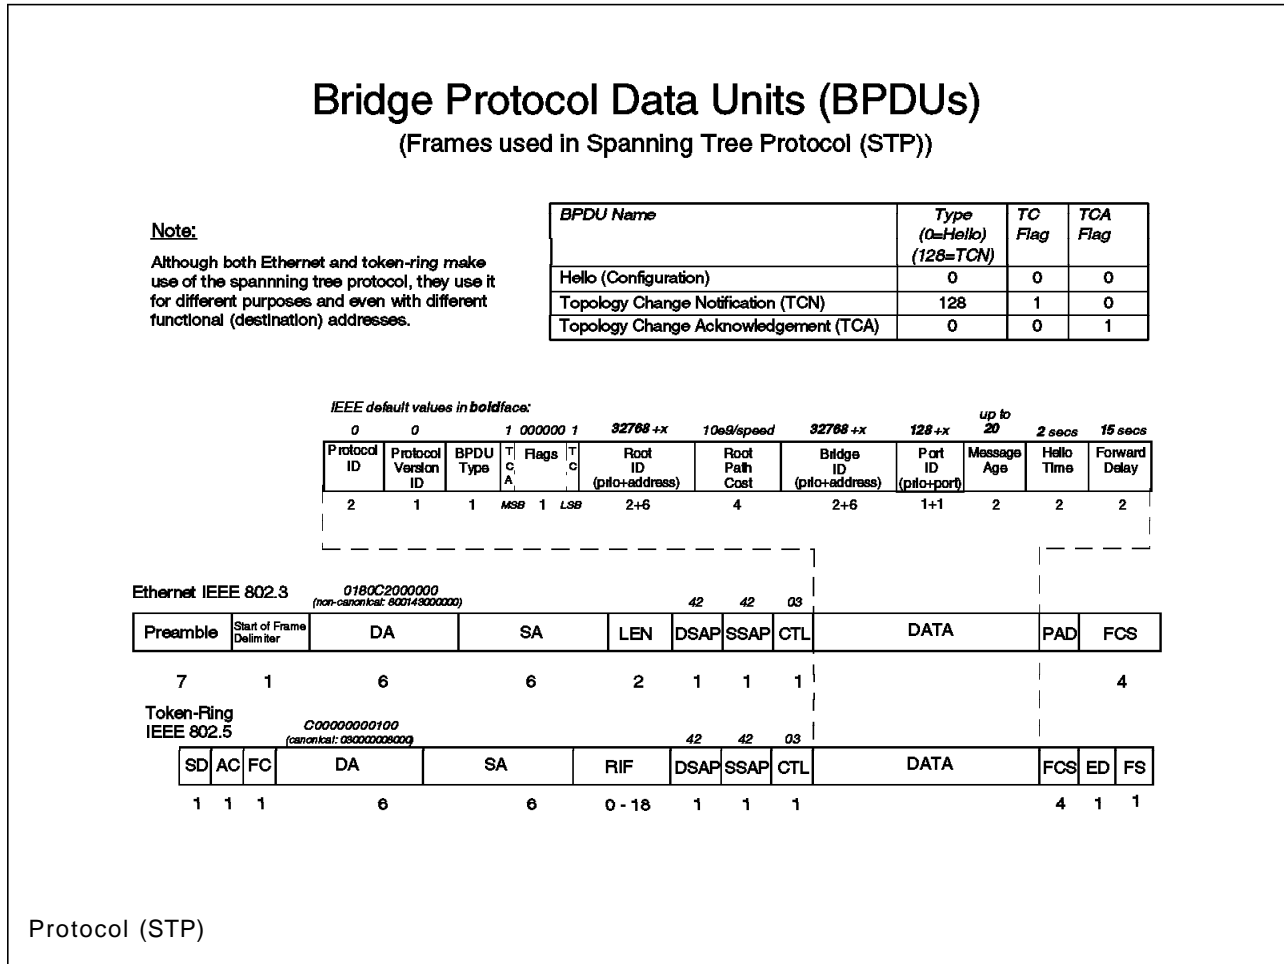


Figure 12. Bridge Protocol Data Units (BPDU) Used with Spanning Tree

2.5.1.2 Topology Change Notification

At some times a bridge or just a port may become inoperable, or a port may enter the forwarding state. Due to Hello BPDUs being sent regularly by the root bridge (and expected by the others), some bridges will notice the change first. It immediately sends a topology change notification (TCN) BPDU through its root port, addressing all 802.1d bridges (group address). The designated bridge on that segment will acknowledge this frame by sending back a Hello BPDU known as topology change acknowledgment (TCA) with the TCA bit set to 1. The designated bridge then issues its own TCN BPDU through its root port. The process continues until a TCN BPDU reaches the root bridge. The root bridge then starts sending Hello BPDUs with the TCN bit set to 1 for a period of time equal to the sum of *forward delay time* and *maximum age time*. The bridges receiving the Hello BPDU with TCN set to 1 then start aging out their filtering databases with a shorter timer (*forward delay*) than usual. They do so until they receive a Hello BPDU with TCN set to 0. By then the spanning tree will have reconfigured itself.

2.5.1.3 Most Important Spanning Tree Parameters

The most important parameters in a bridge affecting its behavior in the spanning tree are:

- Bridge priority: The lower the value, the higher the priority.
- Bridge MAC address: The lowest MAC address of all bridge ports. This might be a Universally Administered Address (UAA) set by the bridge manufacturer, or a Locally Administered Address (LAA) set by the network manager responsible for the bridge in the network. The UAA is also called burnt-in address (BIA).

Note that the bridge ID in a BPDU consists of the bridge priority and the bridge MAC address concatenated.

- Port priority: used by the bridge to choose which port to block if there is a loop situation.
- Path cost: Before forwarding a BPDU in a segment, a bridge adds the path cost to the "path cost to the root bridge" field within the BPDU being forwarded.

The path cost to the root bridge is used in electing the designated bridge of a segment, which is the one with the lowest path cost to the root bridge. Note that a higher port cost should be used for slow links and for segments you want to have at the periphery of the spanning tree.

- Hello time: The root bridge sends Hello BPDUs every Hello time interval to the other bridges in order to keep the spanning tree alive.
- Max. age: All bridges store the root bridge's Hello BPDU (received, for example, every 2 seconds) before forwarding it to other segments in order to maintain the health of the spanning tree. If the Hello BPDU is not received during a certain time (some Hello BPDUs might get lost), then the bridges send the stored Hello BPDU instead. The age of the stored BPDU is incremented continually so the BPDU is kept not longer than the Max. age parameter indicates. In case the stored BPDU is purged after Max. age, a new spanning tree will be built.
- Forward delay: Before changing a port state from blocking to forwarding, the bridge port is kept in the listening and learning states, each for the length of time specified by forward delay. In addition, this time interval is used to age out MAC addresses in the filtering database in the event of a topology change.

Note that the bridge MAC addresses and port priorities of token-ring bridges are defined in a different way from those of Ethernet bridges, but their utilization when calculating the spanning tree is exactly the same for both LAN types (see also 2.5.2.1, "Spanning Tree Protocol (in Token-Ring)" on page 32).

2.5.2 Source-Route Bridging (SRB)

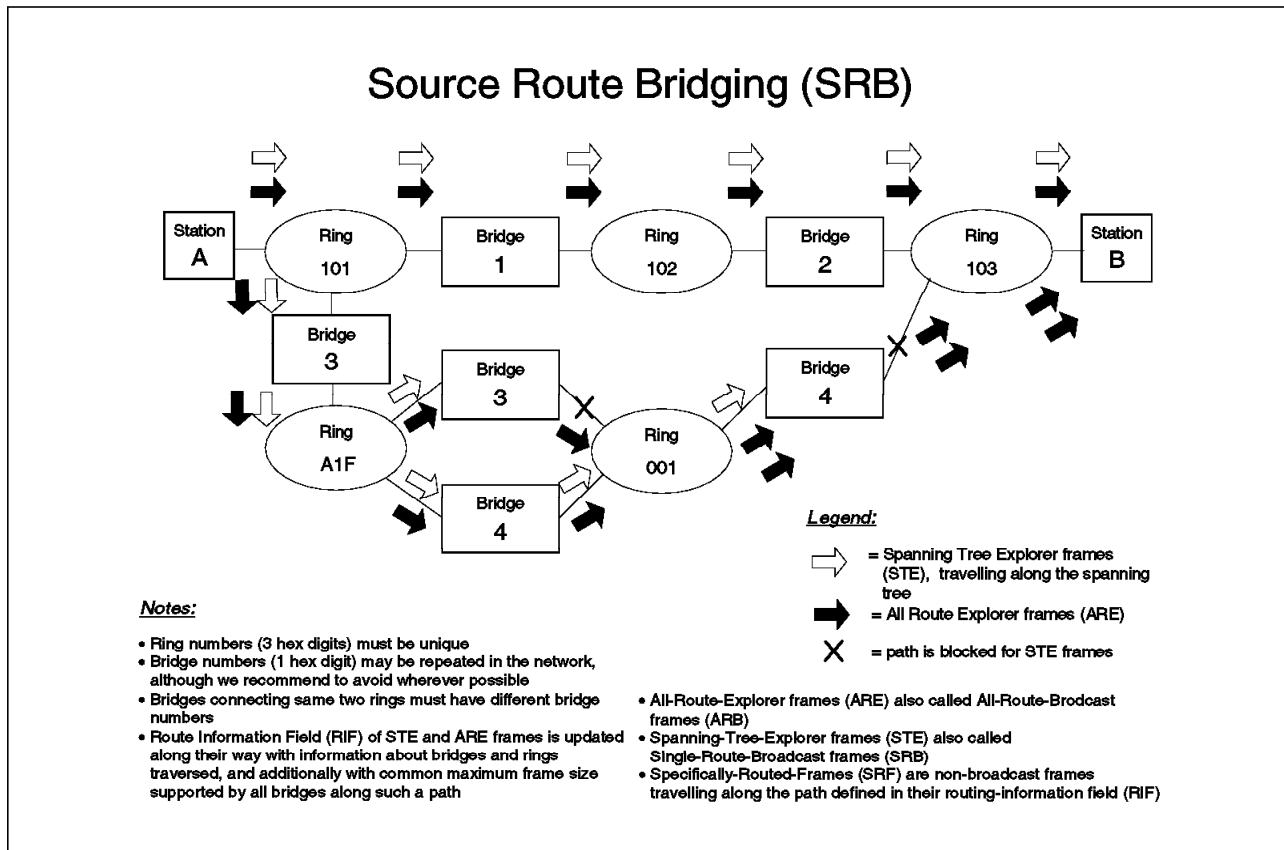


Figure 13. Source Route Bridging (SRB)

This is the bridging scheme used with token-ring. All bridge ports are always in forwarding state; that is no port is blocking (unlike transparent bridges). As shown in Figure 13 station A will send either a spanning tree explorer (STE) frame or an all route explorer (ARE) frame in order to find the best way to its partner station B. An STE frame is one that travels along the spanning tree until it reaches its destination (also known as Single Route Broadcast frame, SRB). An ARE frame is one sent over all possible paths (also known as All Route Broadcast frame or ARB). Whenever a bridge forwards an explorer frame, it will add a route designator to the routing information field (RIF) in the frame being forwarded. See Figure 8 on page 20 for a detailed structure of the RIF. The route designator consists of the ring number from which the frame comes and the bridge number of the bridge itself. This way, when the frame arrives at station B, it will be carrying in its route information field (RIF) a complete route of the path it has taken. Note that the size of the RIF is limited (known as maximum hop count), and therefore explorer frames will be discarded if this maximum distance is reached. Depending on the bridge model the hop count may be set up to 7, or with some newer bridges up to 13 (for some protocols).

Station B will respond to an STE frame with an all route broadcast frame, which will result in station A receiving many answers (and using probably the first one). On the other hand, if station B receives ARE frames, it will respond with one specifically routed frame (SRF) for each ARE frame. An SRF is a frame travelling back the same route its ARE came from; it will be able to do so using the information which was recorded in the RIF field during the exploring process.

Here station A will probably use the first response again. The path learned from the chosen response will then be used for all the following data exchange between stations A and B.

2.5.2.1 Spanning Tree Protocol (in Token-Ring)

In token-ring a spanning tree is also built, but only for proper forwarding of STE frames. The scheme used is the same as the one used in Ethernet (see also 2.5.1.1, "Spanning Tree Protocol (STP)" on page 28). The key difference is that token-ring uses a different group address than Ethernet (802.1d) bridges. In token-ring the functional address being used is X'C0000000100' in non-canonical or X'030000008000' in canonical format (see also Figure 12 on page 29). Another difference is the way the bridge MAC address and port priority are defined; this is explained next.

In token-ring bridges the bridge MAC address is the MAC address of the port with the lowest port ID. The port ID is the concatenation of ring and bridge number in the form X'RRRB', where 'RRR' is the ring number and 'B' is the bridge number (both in hexadecimal).

Note that in calculating the spanning tree, *canonical* MAC address values are used by the bridges, no matter how LAAs were entered by the network manager.

2.5.3 Source-Route Transparent Bridging (SRT)

This scheme is used especially in some mixed environments (Ethernet and token-ring, possibly with FDDI backbone). Bridges using this scheme will forward frames based on the presence or absence of a routing information field (RIF); see also page 21). Frames without a RIF will be forwarded using transparent bridging, while frames with a RIF will be forwarded using source route bridging. Special care should be taken when designing and or troubleshooting networks with both SRB and SRT bridges; otherwise, spanning tree problems could arise. Another point to look at is the fact that SRB and SRT bridge implementations could use different frame size limitations (for details refer to the redbooks mentioned in Appendix H, "Related Publications" on page 517).

2.5.4 Source-Route Translational Bridging (SR-TB)

This bridging scheme is used to connect Ethernet and token-ring LAN environments.

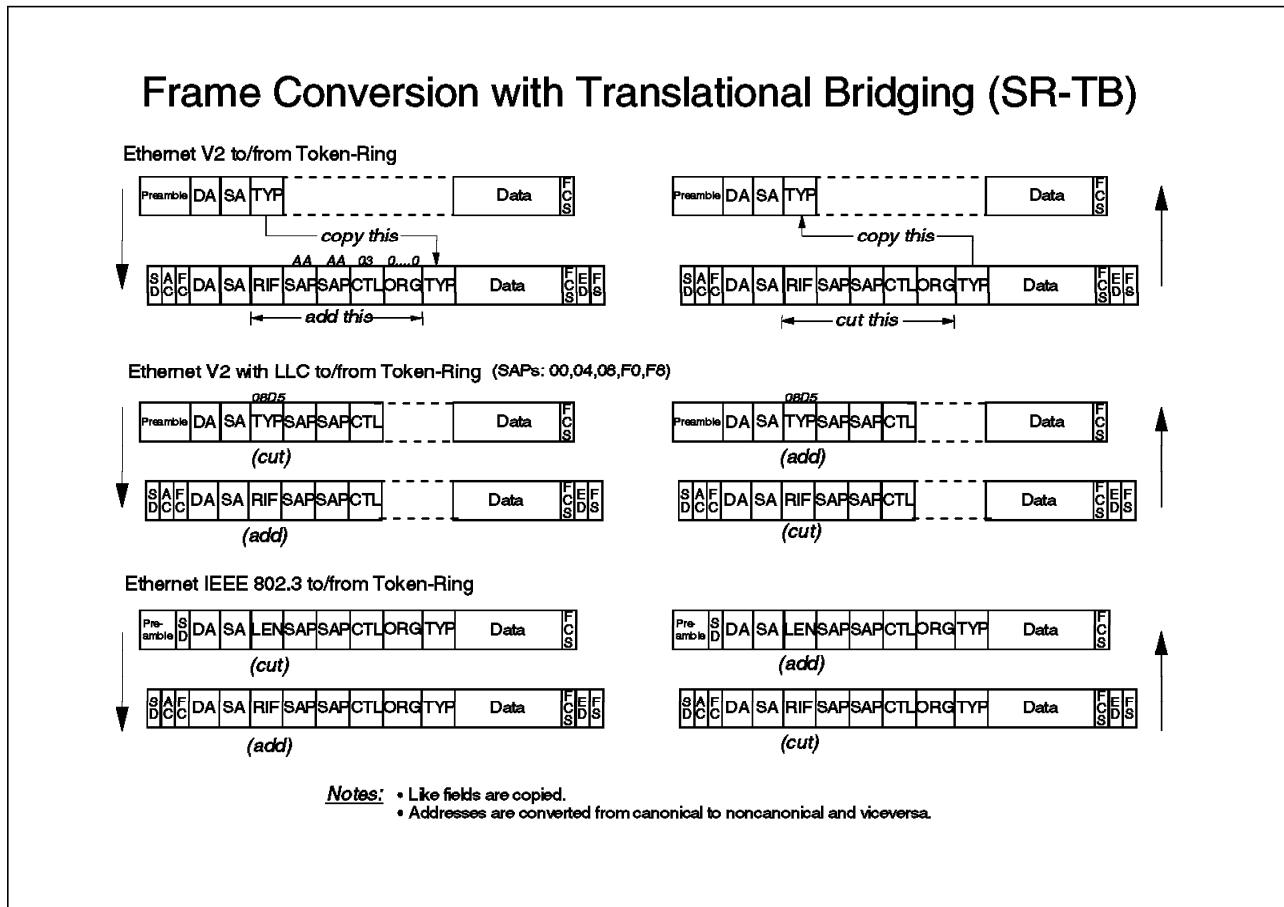


Figure 14. Frame Conversion in Source-Route Translational Bridging (SR-TB)

Bridging of protocols where MAC addresses are used only by layer 2 protocols is straightforward to configure and to use. On the other hand, bridging of protocols which include the MAC address within the higher layer information requires a careful setup of the translational bridge. This is the case for example with Novell's NetWare, which uses the MAC address of the network adapter also as part of the IPX network address (layer 3). Then, the source-route translational bridge must also translate the parts in the data field of IPX/SPX frames where the IPX network address is stored; this becomes necessary since the MAC address is either in canonical (Ethernet) or noncanonical (token-ring) format.

2.5.5 Source-Route Switching (SRS)

A source-route switch learns and forwards frames based on source-route descriptors rather than MAC addresses, for stations one or more source-route bridge hops away. When compared to using transparent switching (analogous to transparent bridging), a switch working with SRS eliminates the need for big tables to cache MAC addresses (as needed with transparent bridging). Since the SRS switch works only on source-route descriptors, table space will fill up less quickly than with transparent switching.

2.6 Protocol Architectures

Brief explanations are given below for the main protocols used today.

2.6.1 NetBIOS

This popular LAN protocol is used mainly between file and application servers and their clients. Note that this protocol is not routable (or not a routed protocol).

Stations communicating with each other will know their partner by using symbolic names of up to 32 characters in length. The symbolic names may be unique or group names. Unique names represent a certain user workstation or a server. Group names are used to offer common services on different servers to a group of users in a LAN.

A machine wishing to be known by a certain name makes this name known to other LAN participants by means of a broadcast message during its initialization phase. In cases where there is already another machine using the same name, this mechanism keeps the new machine from completing its initialization phase and from staying in the same LAN with a duplicate name.

2.6.2 TCP/IP

TCP stands for *Transmission Control Protocol* and IP for *Internet Protocol*. This protocol family has enjoyed enormous acceptance especially since the advent of the Internet. The standard body is the Internet Engineering Task Force (IETF), and the protocols are described in request for comments (RFCs). In TCP/IP we speak of a host when we mean any kind of device addressable via an IP address.

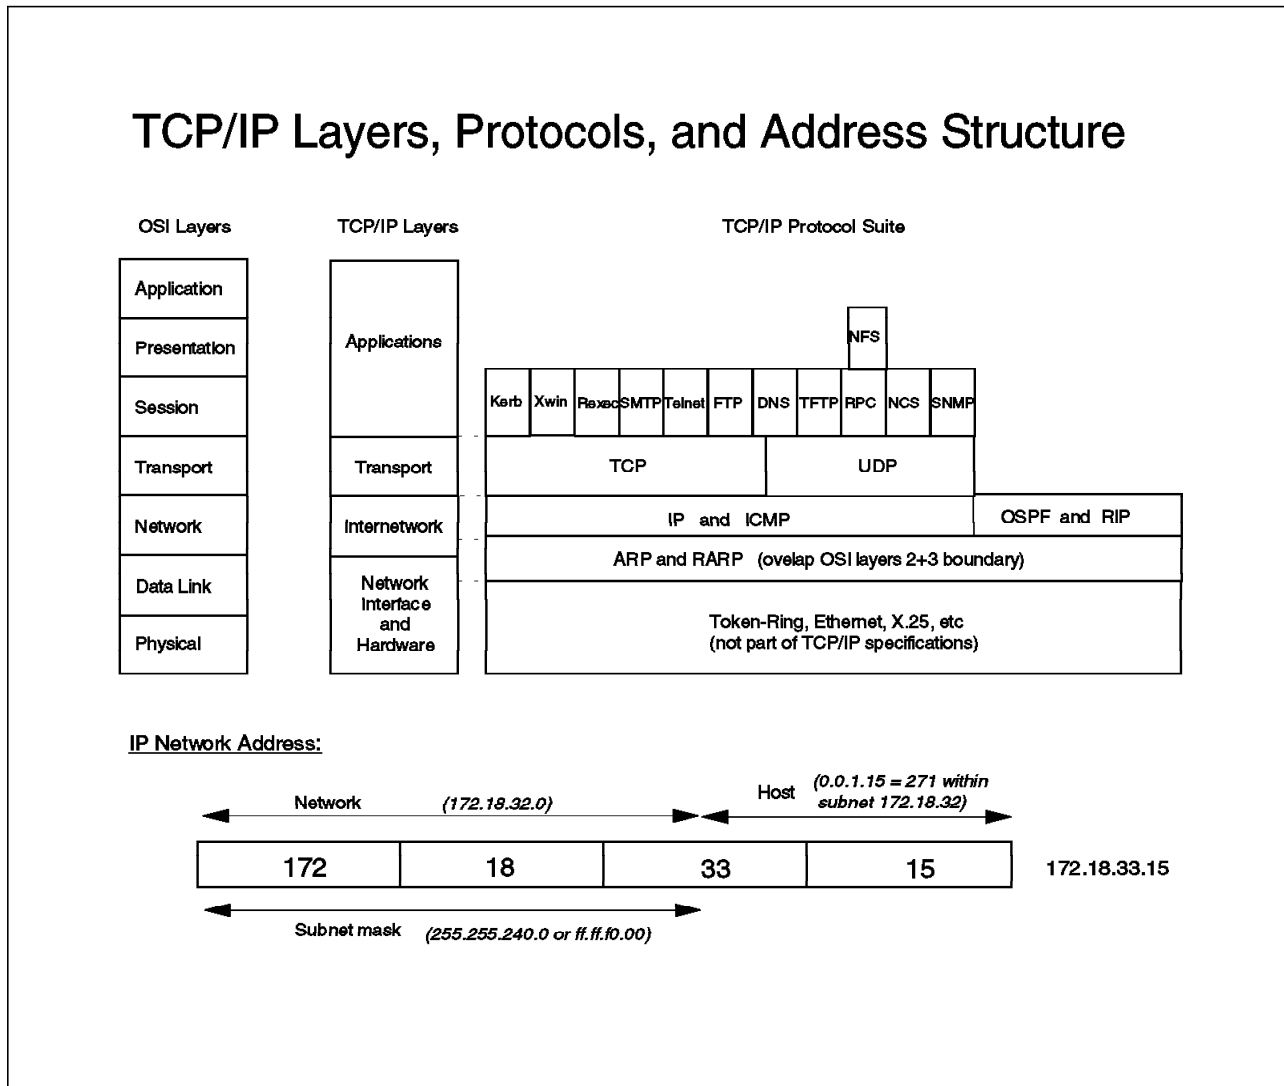


Figure 15. Main TCP/IP Components

Figure 15 offers a comparison between TCP/IP and OSI layers. Additionally, some important protocols are listed, showing in which layer they appear:

- Address resolution protocol (ARP) and RARP (reverse ARP) are used as *the glue* between the technology dependant layers (token-ring, Ethernet, etc.) and the upper layers, since they help with translation from MAC to IP address and vice versa. In every host participating in a network, a so-called ARP table (or ARP cache) is maintained per network adapter using TCP/IP. The main function of the table is to cache MAC and IP address pairs, which are subsequently used for the address translation operations. In non-broadcast environments such as Classical IP over ATM, the RARP function is accomplished by a similar protocol called InARP (Inverse ARP). See Chapter 7, "ATM Emulated LANs and Logical IP Subnets (LANE 1.0, RFC 1577)" on page 293 for a deeper look at Classical IP.
- IP is the protocol used to forward data packets between the higher layer protocols of different hosts. Its current implementation is sometimes referred to as IPv4 (for IP version 4). Some work is now being done to define a new release (see below for IPng, or IP next generation).

- Internet control message protocol (ICMP) is more of a maintenance protocol, but a very important one. The well-known PING operation (sending a message to another host and awaiting its echo) is nothing more than an ARP request (translating the IP address of the target host into a MAC address) and a subsequent ICMP echo request. The target host will answer with an ICMP echo reply, thus completing the PING operation.
- RIP and OSPF are interior routing protocols, which means they are mainly used to connect somewhat related sites (a campus, a company) that may or may not be using WAN links. RIP, as the older protocol, has been very widely implemented, reworked with RIP V2 to support subnetting (see explanation of IP addresses below for subnetting), and is a so-called distance vector routing protocol (routing table contains calculated distances to other networks and hosts); the latter includes support of subnets (check below for an explanation of subnets).
- OSPF is a link state routing protocol (routing tables contain a full view of the network link topology). This is newer than RIP and is thus the routing protocol of choice when building new or migrating old TCP/IP networks.
- BGP-4 is an exterior routing protocol, since it is mainly used to connect otherwise unrelated sites, so that minimum routing information is exchanged between the two. Use of BGP-4 may also be required if connecting over low-speed lines, so that minimum administrative (routing) data makes use of the small available bandwidth.
- TCP is a connection-oriented protocol and as such is used by higher layer protocols when setting up sessions for reliable data transfer.
- UDP is the datagram (connectionless) protocol and as such is used by higher protocols when communicating with partner hosts without needing reliable data transfer at this level. Reliability checks are then done by higher level protocols.
- There are numerous other application-oriented protocols; for these, refer to the TCP/IP introductory documentation.

Figure 15 on page 35 shows the structure of an IP address which is 32 bits or 4 bytes long. The address is usually written using decimal numbers separated by a dot (period) between the corresponding bytes. The leading bits of the address form the network address, and the trailing bits the host address. All hosts having the same network address can communicate with each other without the need of a router in-between. Class A hosts have a one-byte long network address, while class B hosts have two, and class C hosts three bytes in their respective network addresses.

A subnet mask is used to subdivide class A, B or C networks into smaller networks to accommodate the actual population of hosts in a LAN, or just to save IP address space. The subnet mask is a 32-bit number with all leading bits set to 1, and all trailing bits set to 0. The subnet mask is written either in decimal notation (such as with IP addresses) or in hexadecimal notation, but both notations use a dot (period) between the bytes. When the subnet mask is ANDed with the IP address, the remaining bits (leading part) form the actual network address. *Natural subnets* are equal to class A (subnet mask 255.0.0.0 or ff.00.00.00), class B (subnet mask 255.255.0.0 or ff.ff.00.00), and class C (subnet mask 255.255.255.0 or ff.ff.ff.00) networks.

Example: In Figure 15 on page 35 the IP address of a certain host is 172.18.33.15 and the subnet mask is 255.255.240.0 (or ff.ff.f0.00). When ANDing the IP address

with the subnet mask we obtain the number 172.18.32.0 which is then the network address. The trailing bits after the actual network address look like 0.0.1.15, which means that this host has the address 271 (=1*256+15) within the given network (that is the subnet).

Due to the enormous popularity of TCP/IP the number of still free or available IP address ranges is rapidly becoming less and less. Therefore a new addressing scheme is being worked on, referred to sometimes as IPng (for IP next generation) or IPv6 (IP Version 6), and based on 16-byte addresses (128 bits). The new standard is not restricted to just enhancing the address space, but is looking into many more areas where new needs are expected.

2.6.3 IPX/SPX (Novell's NetWare)

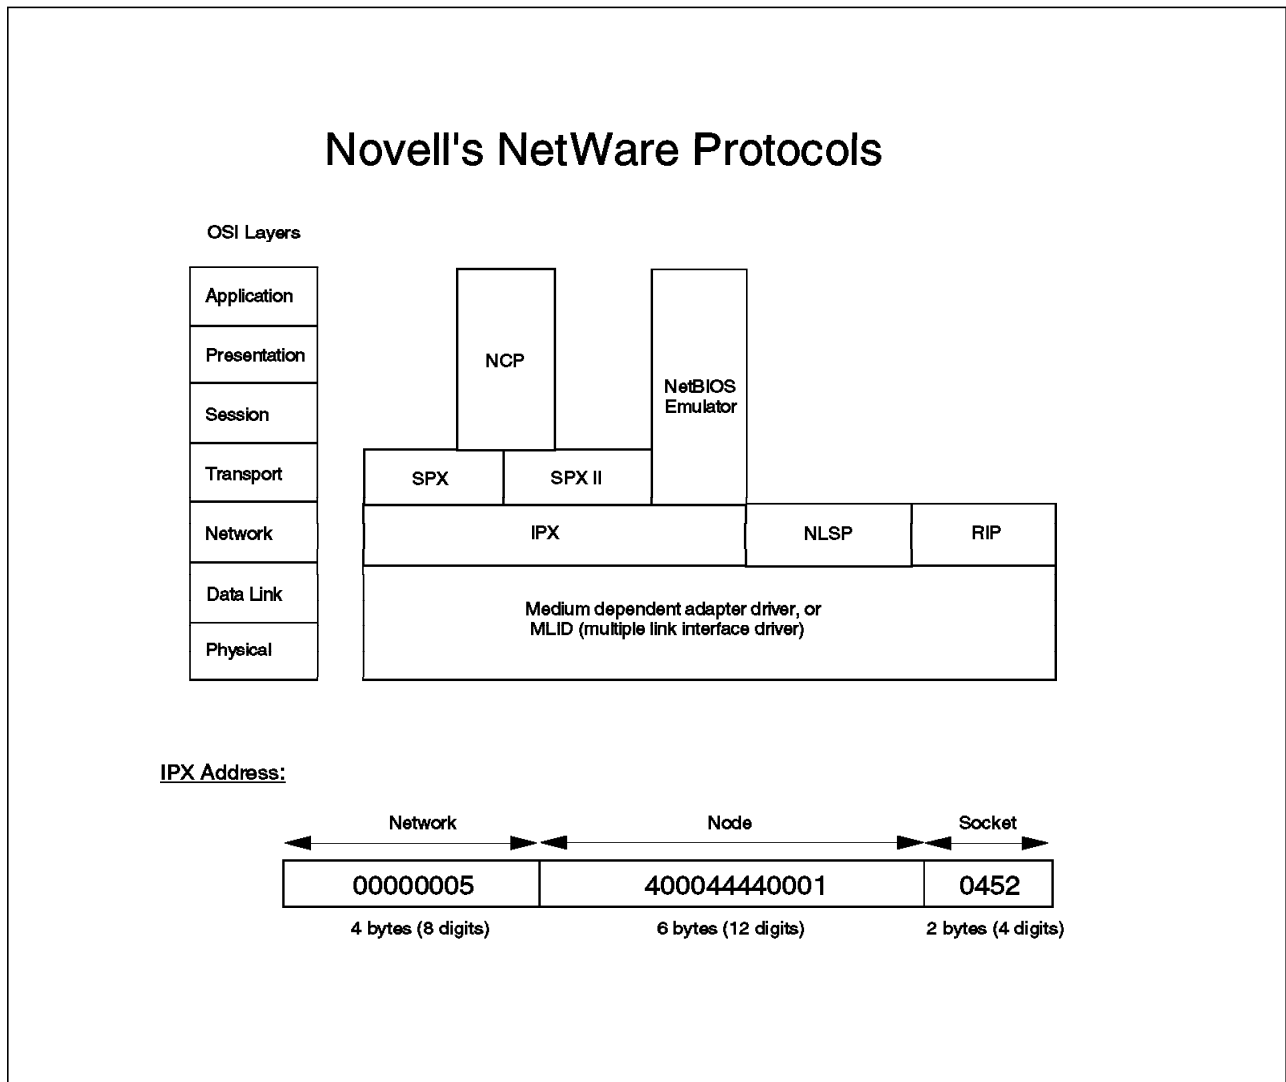


Figure 16. NetWare's (Novell) LANs

Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX) is a popular protocol suite used in LAN environments where LAN clients utilize services offered by NetWare (Novell) servers. IPX is a connectionless (that is datagram) protocol acting as an envelope for higher layer NetWare protocols such as SPX, SAP, and NCP, and is used to carry their information over the

network. IPX is a routed (or routable) protocol, and works as such in the network layer.

Novell servers announce their services every 60 seconds by means of Service Advertising Protocol (SAP) broadcast frames. 60 seconds is the default value for a configurable parameter which must be equal for all servers. All servers and routers in the network are supposed to listen to these broadcasts and to collect that information (into their services list). Thus, when a client sends a SAP request looking for its nearest server, it will potentially receive SAP responses from many servers and routers, and will do a default login to the first one to respond; afterwards the client asks this default server to provide it with the address of its preferred server (if the name of the preferred server was configured in the client configuration file).

NetWare routers exchange topology information every 60 seconds by means of RIP packets. Here again the 60 seconds are a default value, which might have been changed for all NetWare routers in a certain environment. Thus, they will be able to respond to client RIP requests asking for the fastest route to a certain server (for example, the preferred server, as configured in the client config file). NetWare routers also exchange RIP packets whenever a network change occurs. Note that any packet travelling along a NetWare routed network will be discarded when reaching the 16th router, since by then the architecture assumes there is a loop in the network.


SPX is a connection-oriented protocol meant to build more sophisticated applications which need to rely on a properly functioning underlying connection-oriented layer. Those applications do not worry about acknowledging messages, requesting a resend of those messages that are lost, or other communication functions, but they do want to focus on the conversation-like application protocols needed to accomplish their task. NCP (NetWare Core Protocol) builds upon SPX, and is used between servers and clients to exchange information (for example, file or printing data).

NetWare addresses are 12 bytes in length. Figure 16 on page 37 gives an overview of the address structure. Note that a client obtains the network part of its address from a server (or a router) in the same (layer 2) network. Then the client appends its MAC address to it as its node address. The socket number in a target address identifies the service addressed by the request.

Since the MAC address is used as part of the network address (it is equal to the node address), some special configuration of bridging devices may be needed when bridging between unlike topologies (for example between token-ring and Ethernet). For this reason, and because different frame formats might be in use in a certain network, depending on NetWare versions and other protocols used, an overview of the NetWare frames has been included in Figure 17 on page 39 to make initial installation and later troubleshooting of the network a little bit easier.

NetWare's IPX/SPX Frames (Novell)

Ethernet



Token-Ring




Figure 17. NetWare's IPX/SPX Frames (Novell)

2.7 Protocols Used for Network Management and Monitoring

The following protocols are used by network management to track and report events.

2.7.1 Simple Network Management Protocol (SNMP)

This protocol operates over UDP, which is itself a protocol of the TCP/IP protocol suite. Some (proprietary) implementations may use a different transport protocol than TCP/IP.

The SNMP protocol is used between a network management station and the so-called SNMP agents within the devices being managed. Those devices might be bridges, routers, hubs, or just any kind of workstation or host in the network.

The SNMP agent keeps a record of the status of the device in a structure called the Management Information Base (MIB), sometimes also referred to as the MIB database. Both the management station and the SNMP agent need to agree upon the structure of the MIB, or more simply said, they must both be using the same version of the MIB.

The management station uses the SNMP protocol to interrogate the SNMP agent about the status of its device. It does so by asking for the value of a certain parameter in the MIB of the device being monitored, and by asking for the value of a MIB item. This operation is called an SNMP request. Then the SNMP agent sends back the requested information to its management station, an operation called an SNMP response.

Network devices might also send unrequested SNMP messages to their management stations, which is commonly referred to as issuing traps. From the perspective of the management station, these are said to be external events. For example, a device will issue a trap when it starts its operation (so-called cold-start trap), or when somebody tries to modify or intrude into the device and modify its configuration characteristics (whether done intentionally or not). Depending on the configurable characteristics and traps available on a certain device, there may be many other types of traps that can be sent to the management station.

As its name implies, the primary focus of SNMP was to quickly implement an easy and flexible structure and protocol for the management of network components, which could be implemented into the managed devices quickly, with little effort and at a reasonable price, both from the perspective of the resources needed in the device (CPU, memory, etc.), as well as from the development effort for both managed and management components.

In the early days of SNMP no special emphasis was put into specifying characteristics such as robust security (anybody can read the messages in a LAN), nor more complex retrieval of device statuses (for example whole collections of device data, perhaps stored as vectors or matrices). But more recently many users, vendors, and other organizations started to think of this and to request it. Therefore some development is today in place towards a specification called SNMP-II, which should fill the gaps perceived with the former (or current) SNMP, sometimes also called SNMP-I or Version 1.

2.7.1.1 Remote MONitoring (RMON)

RMON is a standardized network management discipline. There are different RMON groups defined, both for Ethernet and token-ring. The RMON probe (a device within the LAN segment) collects the relevant data (number of frames, number of octets, bandwidth utilization in percent, addresses and ranking of top users in a LAN segment, etc.) according to the mentioned RMON groups, and according to those groups specifically configured to collect the needed data. The network management station periodically requests the data being collected by the RMON probe via SNMP and uses it for validation of network health, display and analysis of statistical network data, etc.

2.7.1.2 Station Management (SMT)

SMT is the native network management provided by the FDDI architecture, now in Version 7.3. Since nowadays there are also SNMP management implementations available on FDDI gear, we refer the reader interested in more information about SMT to the following WWW URL:

<http://www.iol.unh.edu/consortiums/fddi/index.html>

Chapter 3. Main Concepts of ATM (Asynchronous Transfer Mode)

This chapter contains an introduction to the most important ATM concepts used in campus networks, with an aim at making them clear cut and easily understandable as they are a much needed prerequisite to effectively troubleshoot problems in ATM-based networks.

Other sources of information may be found in Appendix H, "Related Publications" on page 517.

3.1 The Underlying Notions

The primary building blocks of ATM that must be understood for effective troubleshooting of ATM networks include:

- The Structure of ATM Networks
- ATM Network Characteristics
- ATM Connections
- Routing/Switching ATM Cells
- ATM Cells and Cell Format
- ATM Signalling
- ATM Address Format

3.1.1 The Structure of ATM Networks

There are two main types of ATM networks:

- Private ATM networks, managed by their respective owners (banks, hospitals, universities, manufacturing companies, etc.), or possibly managed by a service company in lieu of the network owner (for example, over an outsourcing contract)
- Public ATM networks, managed by the respective carriers who own those networks

The ITU-T (former CCITT) is responsible for all ATM standards in the public domain. The ITU-T also works on and issues standards for the private ATM networks. Where there is no standard yet for the private domain, the ATM Forum (a consortium of vendors and user representatives) publishes its own specifications which might be later refined and possibly replaced by ITU-T standards. The aim here is to get the industry moving ahead on a common path, and to later converge on official standards when they become available.

In this book we will be dealing only with private ATM networks, but since they might be interconnected over public ATM carriers, it is better to understand the terminology of the whole picture. Let us therefore take a look at Figure 18 on page 45.

On the left hand side of the picture we see at the bottom the ATM network maintained by XYZ Corp. in its offices near San Francisco (this is a fictive company we use as an example). The ATM switches are connected to each other over so-called *private network-to-network interface (NNI)* links. In the following sections we will refer to them just as NNI links, keeping in mind that these are the only architectural type of links used to connect switches to each other (although, as we shall shortly see there are various flavors of NNI). An NNI link can be either a pair of multimode or single-mode fibers or a copper cable

(UTP category 5), depending on the connectors available at the switches and on the distances between them.

Since ATM is capable of carrying data, voice, and image, there are many types of end user devices which can attach to the ATM switches. These could be for example PCs, workstations, hosts, video equipment for teleconferences, and even PBXs. The end user devices use the so-called *private user-to-network interface (UNI)* to connect to the ATM switch next to them. Here also we refer commonly to this interface as the UNI (that is we omit the *private* prefix). Apart from the single-mode fibers, we can use the same types of physical links for UNI connections as the ones we use for NNI, that is multimode fiber or copper cables (UTP Category 5). The reason for this is that the switch ports can be used for either type of connection, depending only on the definitions set at the switch.

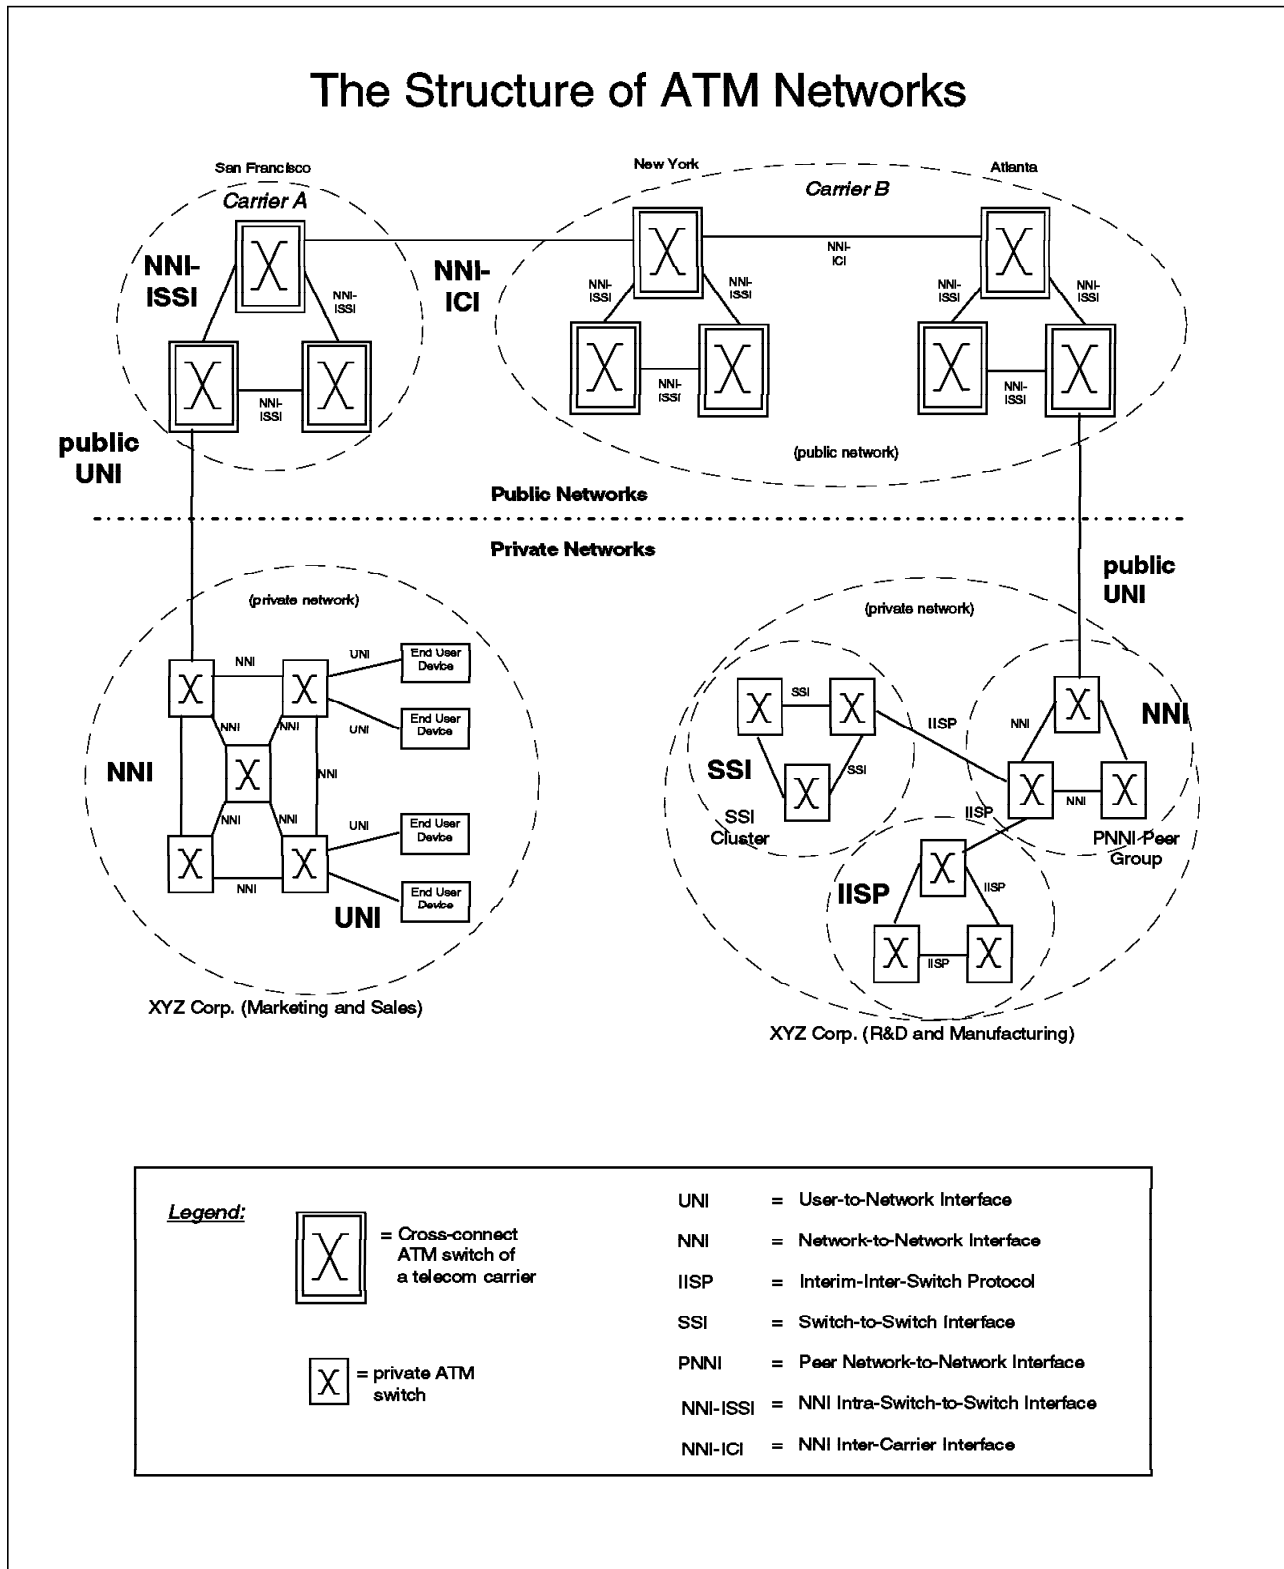


Figure 18. The Structure of ATM Networks

Let us now look at the other ATM networks used by XYZ Corp. in Atlanta, as shown in Figure 18 (right-hand side, below). There we see that the switches are grouped, and that each group uses apparently different types of links to interconnect those switches (these are the NNI flavors we mentioned earlier). Nevertheless let us emphasize and keep in mind the fact that all those links are

NNIs, and that the differences exist mainly for historical reasons due to the rapid evolution of ATM. Having said this let us explain the different types of switch groups and their connections:

- **IISP:** Switches connected via interim inter-switch protocol links. The IISP was the first ATM Forum specification used to connect switches, and is so to say the lowest common denominator to interconnect switches of different makes (that is any ATM switch should be able to use IISP to interconnect to another switch). IISP is sometimes also referred to as PNNI-0 (PNNI Phase 0; see below for PNNI). The disadvantage of this early specification is that static routes need to be defined for every connection to other switches, no matter if they are neighbors or far apart. Additionally, since we work here with static routes, there is no way to define redundant paths to dynamically overcome failure of links.
- **SSI Clusters:** SSI stands for Switch-to-Switch Interface, and is an IBM proprietary architecture based on an early draft of the PNNI-1 specification. The aim here was to provide automatic route discovery (no need for static route definitions), at a time when there was no standard available to provide for this. The SSI implementation allowed for easy and fast deployment of ATM-based networks almost from the earliest stage of the ATM evolution. The switches within a cluster use a link state protocol (similar to OSPF used by routers) to maintain a common view of the network. All switches in a cluster share a common portion of the ATM address (leading 12 bytes of the overall 20-byte ATM address), while the 13th byte is used to identify the switch within the cluster. Therefore any different address combination within the first 12 bytes would be recognized to belong to an external cluster (formed by a single switch or a group of switches). SSI clusters may be connected to each other via IISP links making possible the design and implementation of much larger networks.
- **PNNI Peer Groups:** These are groups of switches interconnected using the latest specification of the ATM Forum for NNI links called Private Network-to-Network Interface, Phase 1 (PNNI-1). As with SSI, there is no need to define static routes since the switches interchange topology information automatically. All switches within a certain PNNI Peer Group have the same Peer Group Identifier, which might be a common portion of the ATM address (some of the 104 leading bits, referred to as the *level id*, or some other user-specified hexadecimal string). Although PNNI-1 defines also the possibility for building routing hierarchy levels, most vendors implement only a flat (one level) hierarchy in the first release of their switch code. Therefore different peer groups need to be interconnected via (static) IISP links. IBM's PNNI-1 implementation allows interconnecting Peer Groups over redundant IISP links with recognition of link failures and dynamic rerouting, making PNNI-based networks more robust from the very beginning.

Note that end user devices may connect to any of the above types of switch groups via the **user-to-network interface (UNI)**. Due to the evolution of the ATM specifications there are also different flavors of the UNI, differing mainly in the type of signalling they use: UNI 3.0 and UNI 3.1. The IBM ATM switches automatically translate this signalling (3.0 to 3.1 or vice-versa) so end user devices using different versions can nevertheless communicate with each other. A further refinement of portions of the UNI 3.1 has been made by the ATM Forum lately. Those are specifications such as ILMI 4.0 and some others. Note that a full replacement of UNI 3.1 has not taken place at the time of this writing, but only some pieces have evolved so far. In addition, due to some common signalling needs of the links, some of the signalling specifications prepared

originally for the UNI have also been used (slightly modified) in the different flavors of NNI.

Now, for the sake of completeness, let us have a look at the components in the public ATM networks. For this we refer again to Figure 18 on page 45. If XYZ Corp. wishes to connect both locations in San Francisco and Atlanta, it will connect each of the sites to the respective local ATM carrier using the *public UNI* interface; that is, the carriers will have the impression that end user devices are being connected to their networks. Therefore the carriers will set up a permanent and transparent path between the two customer locations to let those devices (actually the private ATM networks of XYZ Corp.) communicate together as if they were connected over a local NNI link. Since the private networks are attached to the public ATM networks over the public UNI, the private switches will not exchange topology information with the carrier switches. As we have seen on page 46 topology exchange is done among switches only over NNI links (SSI or PNNI). This way the carriers are able to offer the same kind of service to many customers, knowing that the customer (private) networks will not interfere with each other nor with the carrier's own network.

On the other hand, the remote switches of XYZ Corp. will interoperate over the transparent path offered by the carriers as if they were connected over a local NNI connection, allowing XYZ corporation to set up a big common network (for example, one SSI cluster or one PNNI peer group) or to interconnect different switch domains (SSI clusters, PNNI groups) depending on XYZ's needs and the design and size of the switch domains. Anyway, the NNI signalling exchanged between the two remote switches of XYZ Corp will be carried transparently over the path the ATM carriers have set up.

In Figure 18 on page 45 we also see that the switches of the carriers use other interfaces than the ones used in the private ATM networks. Nevertheless, the philosophy remains the same, that is all of them are of a certain NNI type, as specified by ITU-T. The switches in the same hierarchy structure (in US this is Local Access Transport Area or LATA) will use an interface called NNI intra-switch-to-switch interface (NNI-ISSI), while switches in the above hierarchies (long distance and international connections) will use the NNI inter-carrier interface (NNI-ICI). Since the scope of this book is private ATM networks, we refer you to specialized literature if you want to learn more about public networks.

3.1.2 ATM Network Characteristics

Now that we know the general structure of ATM networks, let us look at their characteristics which are both important to understand when designing and/or troubleshooting ATM networks. They are the following:

- **Connection-Oriented:** Communication between end user devices in an ATM network is always connection-oriented. This means that there is no way of sending data over the network before either a *switched* connection or a *permanent* connection has been set up. Switched connections are built call-by-call by means of signalling procedures. Permanent connections are pre-established connections, which have been set up by the network operator using certain administrative procedures (permanent configuration of connections on the switches along the path).

In order to utilize current LAN schemes over ATM, some networking models have been developed to implement the broadcast mechanisms inherent to those LAN schemes. All of these (connectionless) networking models

operate at layers above the (connection-oriented) ATM transport layers. Among them are LAN Emulation and Classical IP, which will be covered later in this chapter. Using these networking models, current applications can already today take advantage of the much higher speeds of ATM, without a need for any modifications in those applications.

- **Guaranteed In-Sequence Delivery:** As we will see later, any kind of information (data, voice, image) is sent over the ATM network packed in fixed length cells (see 3.1.5, "ATM Cells and Cell Format" on page 56). Since these cells are delivered over a virtual connection (switched or permanent), each cell travels along the same route, and all cells will be transferred to the destination end station in the same order as they were presented to the network.
- **Multipoint Connections:** The natural connection between two end user devices is a connection-oriented point-to-point virtual connection. Nevertheless, ATM also permits you to define point-to-multipoint connections. These types of connections might then be used by implementations of (connectionless) network models, operating above the ATM layers, to broadcast or multicast messages from the root to the leafs of such a point-to-multipoint connection. This has led to the widespread use of the name "multicast trees" as a synonym for the point-to-multipoint connections. Such a point-to-multipoint tree structure is illustrated on Figure 19 on page 49.

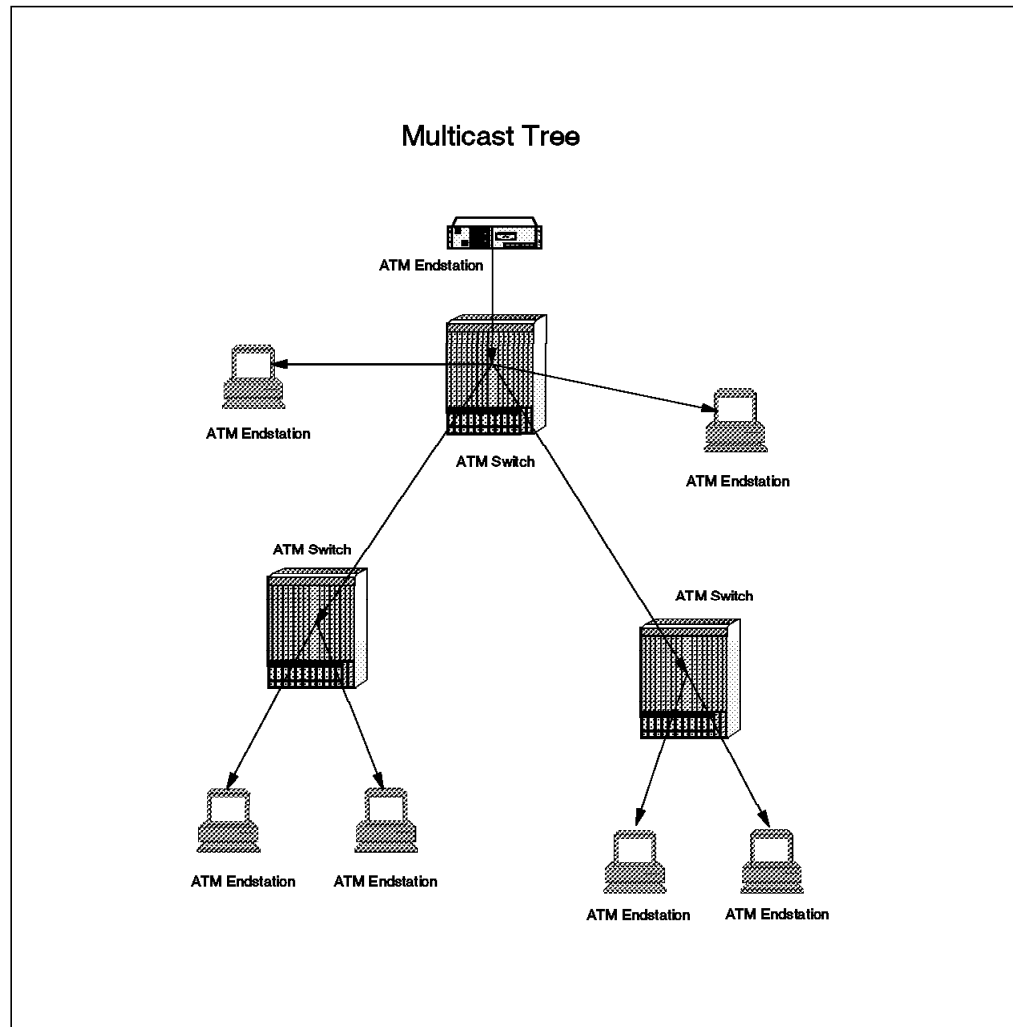


Figure 19. Multicast Tree

Point-to-multipoint connections are first established as a single point-to-point connection between the root end system and one leaf. Once this connection is established, a second leaf is connected to it using the optimal route from the established connection. This algorithm is used until the last leaf is connected to the tree.

The following are some characteristics of a multicast connection:

- Communication is available from the root to the leaf.
- Data may be sent from leaf to the root, but does not allow leaf-to-leaf communication over this connection.
- The multicast tree may be set up by signalling or by the network administrator as a permanent connection.
- **Quality of Service (QoS):** Each ATM virtual connection has quality of service characteristics associated with it. During congestion, when a network cannot recover from an overload, it discards only the cells marked as low priority. The network can select which cells to discard depending on the QoS characteristics of the virtual connection.

The QoS parameters defined by the ITU-T are as follows:

- Cell transfer delay (network latency)

- Cell delay variation (jitter)
- Cell transfer capacity (speed - average and peak allowed rates)
- Cell error ratio
- Cell loss ratio
- Cell misinsertion rate

As we will see later, information is transported in ATM along virtual paths. Within these virtual paths, there are virtual connections along which the actual cells are transported. Each virtual path also has a QoS (Quality of Service) associated with it. The QoS of a virtual connection within a virtual path may be lower than that of the virtual path, but cannot be higher.

- **Cell Loss and Cell Discard:** Cells may be lost or discarded by an ATM network. The network does not detect the loss of cells, and does not signal the user when it has discarded cells from a particular connection.

Remember that in fast cell-based networks, congestion is handled by discarding cells, and recovery is accomplished by retransmission of the full block rather than individual cells.

Some variable bit rate applications for voice and video can produce two types of cells. Standard cells contain basic information, and optional cells contain less important information used to enhance the quality of the information contained in the standard cells (better pictures, higher voice quality). If the voice or video user equipment being used can mark the optional cells accordingly (with the Cell Loss Priority indicator, or CLP), it can avoid the loss of essential information during network congestion, since the optional cells will be discarded first.

- **Congestion Control:** ATM networks do not have flow control of the kind found in traditional packet switching networks. This is because traditional windowed link protocols are no longer effective at high-link speeds.

In ATM the parameters for a connection (for example, requested bandwidth, QoS) are examined before the connection is established, and the connection is only allowed if the network can support the desired parameters (checked by the network with so-called Connection Acceptance Control, or CAC, procedures). The network allocates resources on a statistical basis. It allows for the possibility that demand may exceed available network resources, in which case the network will discard cells.

- **Input Rate Policing:** At the entry point of the network (for example, at the entry from a private network into a public carrier network), the entry ATM switch monitors the rate of data arrival for a virtual connection or virtual path according to the negotiated QoS parameters for the connection. It will take action to prevent an ATM endpoint from exceeding its allowed limits using a technique called back-pressure.

In the case of overload, depending upon the network configuration, the network may either:

- Discard cells received over the allowed maximum rate.
- Mark the overloaded cells with the cell loss priority indicator (CLP), which means those cells will be discarded first under congestion conditions.

- **No End-to-End Data Integrity:** ATM does not provide end-to-end data integrity. This function is the responsibility of end-user equipment or of a higher layer protocol.
- **No Priorities:** There are no priorities defined within an ATM network. The Cell Loss Priority indicator (CLP) is not a real priority mechanism. Despite its name, it is an indicator simply used to mark those cells which may be discarded first in the event of congestion.

3.1.3 ATM Connections

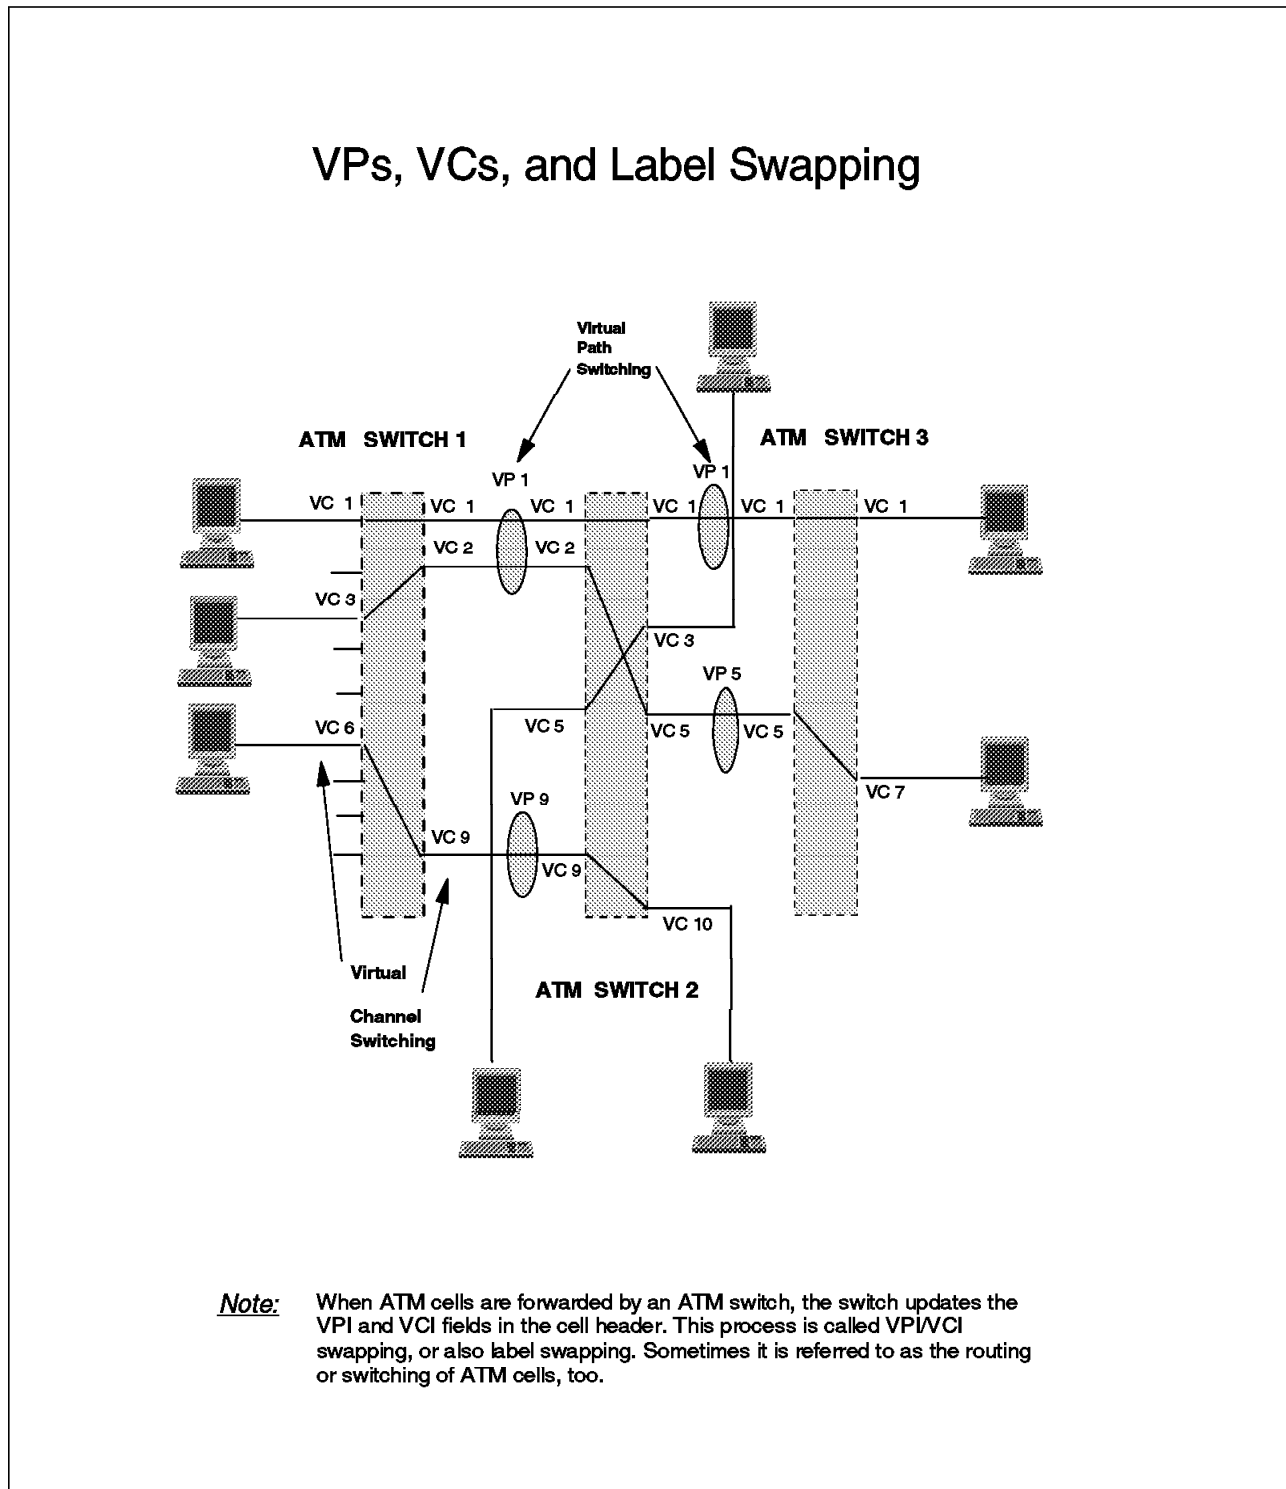


Figure 20. VPs, VCs, Label Swapping

ATM differs from existing LAN networks, because it uses connection-oriented technology. The connection, in ATM terminology, is a point-to-point or point-to-multipoint link from one end system to another across a series of ATM switches in a network.

This connection-oriented technology simplifies the routing of cells across the ATM network. Station destination and source addresses do not need to be carried in each ATM cell, only the connection identifier is required by each ATM switch to route the cell correctly. Because the information between end systems is sent over the single route described by the connection identifier, the information is received in the same order as it was sent. This in-sequence delivery is required especially for voice and video traffic and also simplifies the processing of data traffic.

Before data can be transferred, a virtual connection needs to be established between the end systems either by using a pre-determined fixed path or by means of the signalling protocol. The connection may, therefore, be permanently established by the network operator (*Permanent Virtual Connection, or PVC*), or temporarily established on request from an ATM end system by signalling (*Switched Virtual Connection, or SVC*).

To handle each connection, ATM uses the concept of virtual paths and virtual channels. These are described in the following sections.

Virtual Path (VP) and Virtual Path Indicator (VPI): As illustrated in Figure 20 on page 52 a virtual path is an aggregate route through a network representing a group of virtual channels (VCs). An example in the mentioned figure is VP1 between ATM switches 1 and 2, where VP1 contains VC1 and VC2.

VPs may exist:

- Between ATM end systems
- Between ATM end systems and ATM switches
- Between ATM switches

The Virtual Path Indicator (VPI) denotes which virtual path is to be used by a cell and is contained within each cell travelling along an ATM network.

Virtual Path Link (VPL): A virtual path link exists between the points where a VPI value is assigned or where it is translated or determined. Typically these points would be switches in the ATM network.

Virtual Path Connection (VPC): A virtual path connection is the concatenation (sequence) of VPLs that extends between virtual path terminations.

Virtual Path Connection Identifier (VPCI): This identifier of a VP connection is returned by the ATM network when call setup is performed by a user device. It is 16 bits long, and used by the signalling protocol instead of the VPI, which is unique only within a single ATM link.

The VPCI concept is especially important when connecting private ATM networks over public ATM carriers (VP tunneling). The VPCI remains the same at both ends of the public connection, while the corresponding VPIs might be different (and changed internally within the public network) due to VPI swapping.

Virtual Channel (VC) and Virtual Channel Indicator (VCI): A virtual channel is defined in ATM as a unidirectional connection between user devices. The virtual channel indicator (VCI) is the indication of the virtual channel to be used by a cell and is contained within each cell in the network.

Virtual Channel Connection (VCC): A virtual channel connection is the end-to-end connection along which a user device sends data. Because,

strictly speaking, virtual channels are unidirectional, a VCC would normally consist of two virtual channels to provide complete full-duplex data transfer.

Virtual Channel Link (VCL): A virtual channel link exists between the points where a VCI value is assigned or where it is translated or determined. Typically these points would be switches in the ATM network.

A virtual channel link is a separately defined data flow within a link or virtual path. A virtual channel connection (VCC) through a network is a sequence of interconnected (concatenated) VCLs.

There are limits on the number of VPs, VCs, etc. within an ATM network as follows:

- The maximum number of VPs on links is determined by the number of bits allocated to address the VPs in the cell header (VPI). This is either 8 (at UNI) or 12 bits (at NNI). See also Figure 21 on page 56.
- The maximum number of VCs within a VP is determined by the number of bits allocated to address the VCs in the cell header (VCI). This is 16 bits.
- There may be additional limits imposed by the capacity of specific ATM equipment.

3.1.4 Routing/Switching ATM Cells

An ATM cell is transmitted along a virtual channel connection according to the routing information contained in its header. This routing information is swapped at every switch along the path of the connection, to enable the routing of the cell to the next switch along the connection. This process is referred to as label swapping (see also Figure 20 on page 52).

The routing information in the cell consists of the VPI and the VCI fields in the cell header (see also Figure 21 on page 56). The definition of the VPI/VCI mapping is established when the connection is established. The mapping for the connection is held at every intermediate switch and consists of VPI/VCI input fields mapped to VPI/VCI output and port output fields.

This means that each VPI/VCI pair is associated with a particular port on a switch and each VPI/VCI associates a cell with an input link and a corresponding output link. Based on the VPI/VCI in the header of the cell an ATM switch can identify the output link across which the cell is to be routed and give the new link identifiers to the cell. Figure 20 on page 52 shows some examples of mapping VPI/VCI. Let us take a closer look at one of them following the path of cells travelling from the lower station attached to ATM switch 1 to the lower station attached to ATM switch 2:

- Cells forwarded by the station to ATM switch 1 are carried over VCI 6.
- The ATM switch 1 receives those cells and updates both VPI and VCI before sending them to ATM switch 2 (so-called VPI/VCI or label swapping). The original VPI (most probably VP 0, although not represented in the figure) is changed to VPI 9, and the VCI 6 is changed to VCI 9.
- When receiving the cells via VP/VC with VPI 9 and VCI 9 the ATM switch 2 forwards them to the target end station over a VP/VC with VP 0 and VCI 10, doing again label swapping.

Switching of all virtual channels within a virtual path can take place inside a switch. In this case the switch unpacks the virtual path into unique virtual

channels, and using its routing table groups them together again, forming new outgoing virtual paths as (virtual channel switching).

The upper portion of Figure 20 on page 52 shows an example where the VPI is switched without unpacking it (virtual path switching). In this example, all VCs keep their VCIs throughout the entire path. The only thing that may change is the VPI, when the switches do VP switching. In our figure even the VPIs remain the same (just by chance).

3.1.5 ATM Cells and Cell Format

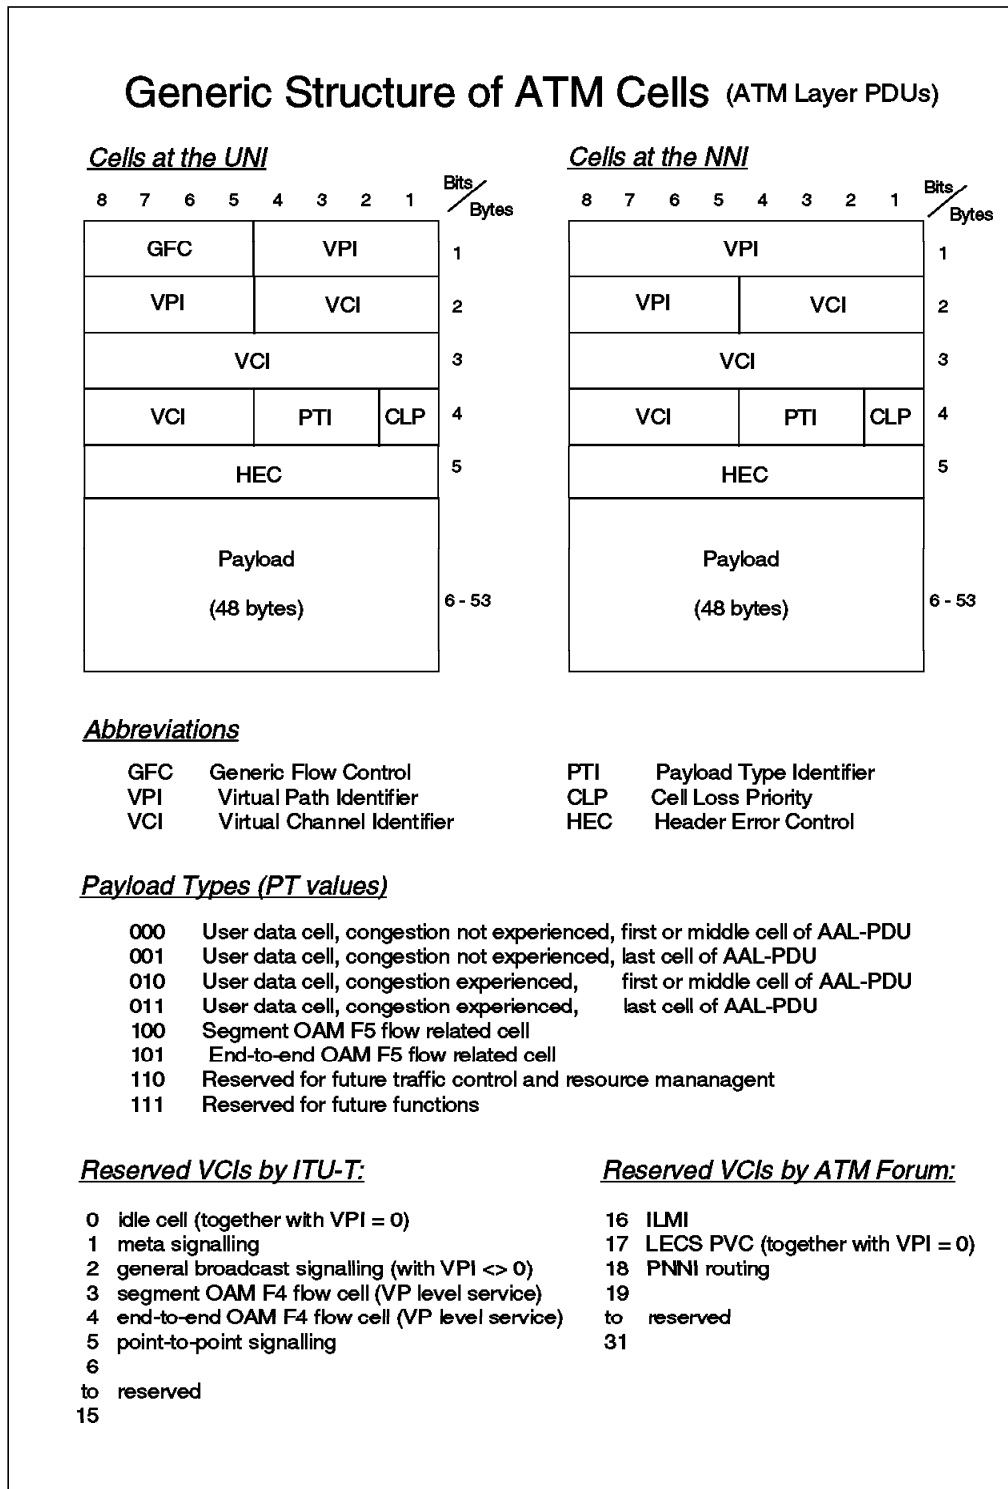


Figure 21. Generic Structure of an ATM Cell

In ATM networks, information is transmitted in cells. Cells are fixed-length packets. Each packet is 53 bytes long: 48 bytes are the payload with a 5-byte header. The header contains control information, including the VPI/VCI route identifier that defines cell route information. The header is error checked to avoid errors being propagated over the network. The payload is user

information which is not protected by error checking at the ATM network level. The size of 48 bytes for the payload was chosen by the ITU-T as a compromise between data and voice requirements. Longer cells for data transfer reduce the overall transmission overhead, while smaller cells for voice transfer minimize the possibility of unnatural *silence gaps* and other undesired phenomena.

The cell formats for the user-to-network interface (UNI) and the network-to-network interface (NNI) are illustrated in Figure 21 on page 56. The fields within the cell are described below.

Generic flow control (GFC): Since this field is present only in the UNI cell header, it has local significance only. The UNI 3.1 specification says only that "it can be used to provide standardized local functions (for example, flow control) on the customer site.

VPI and VCI: These two fields, Virtual Path Identifier (VPI) and Virtual Connection Identifier (VCI) are arguably the most important fields in the cell header. Together they identify the logical connection (the virtual connection) over which the cell is travelling. Some VPI/VCI values are reserved for signaling (for example connection establishment), and for maintenance and resource management (see also the list of reserved VCI values in Figure 21 on page 56).

If the VPI and VCI values are set to zero, it means that the cell is empty or idle. Empty cells may be required in a network to maintain physical link protocols.

Payload type indicator (PTI): This is a 3-bit field (bits number 2, 3, and 4 within the byte as shown in the mentioned figure; bit number 4 being the MSB within the field).

- Bit 4 determines if the cell is user data (bit 4 = 0) or operations, administration and maintenance (bit 4 = 1).
- Bit 3 is used when the cell carries user data. Bit 3 = 0 means no congestion, while bit 3 = 1 means that congestion was experienced somewhere along the route passed by the cell.
- Bit 2 is used by higher layer processing. Bit 2 = 0 means that the cell is the first or the middle cell of the user data frame. Bit 2 = 1 means that the cell is the last cell of the frame.

Cell loss priority (CLP): When set to 1, this bit indicates that the cell is a low priority cell. When the network needs to discard cells in a congested situation, these cells should be discarded first.

Header error check (HEC): This field allows the correction of all single-bit errors or the detection of multi-bit errors in the header part of the cell. It is used as CRC value (cyclic redundancy check) for the preceding fields of the cell header. Note that there is no similar error check for the cell payload (higher protocol layers are expected to use their own error check mechanisms).

In Figure 21 on page 56 operation and management (OAM) cells are mentioned. These are cells related to the management plane. For further reference please check the UNI specification of the ATM Forum.

3.1.6 ATM Signalling

ATM signaling is the process used for dynamic setup and clearing of ATM switched virtual connections (VPCs and VCCs) at the UNI interface. Permanent virtual connections are established by a network operator. If a connection is permanent, and if the network fails and is restarted, the circuit is reestablished. In case of recovery from failure the network will not re-establish lost switched virtual connections; this is the responsibility of the end stations.

The key elements of ATM signalling at the UNI are the following (applies with some modifications to the NNI, too):

- Signalling takes place on separate VCCs from those used by user data. The same principle is used as in narrow-band ISDN where the D channel is used to initiate connections.
- There is a method to set up additional signalling channels besides the predefined channels (although this is not used in current ATM implementations). It is known as meta-signalling (VCI = 1).
- Point-to-point signaling is the default signalling method (uses VCI = 5), but broadcast signalling may be implemented in the future (using VCI = 2). Point-to-multipoint connections are established using point-to-point signalling.
- Class B (AAL-2) services have not yet been defined, so they are not supported by the signalling protocol.
- The Class D protocol is also not supported, because clients must have connections to connectionless servers, so the calling procedure is not required.

Some more information regarding signalling at the UNI can be found in 3.2.4.6, "Signalling AAL (SAAL), Control Plane" on page 73.

The routing framework developed for the routing of ATM connections as used in PNNI-1 networks is based on an extension to the OSPF (Open Shortest Path First). While OSPF networks exchange links statuses for maintaining a view of the network topology, PNNI networks exchange information regarding the state of the nodes (ATM switches) in addition to the link statuses. This is done via PNNI signalling to the immediate neighbors (while each ATM switch will in turn propagate perceived topology changes to the whole switch network).

3.1.7 ATM Address Format

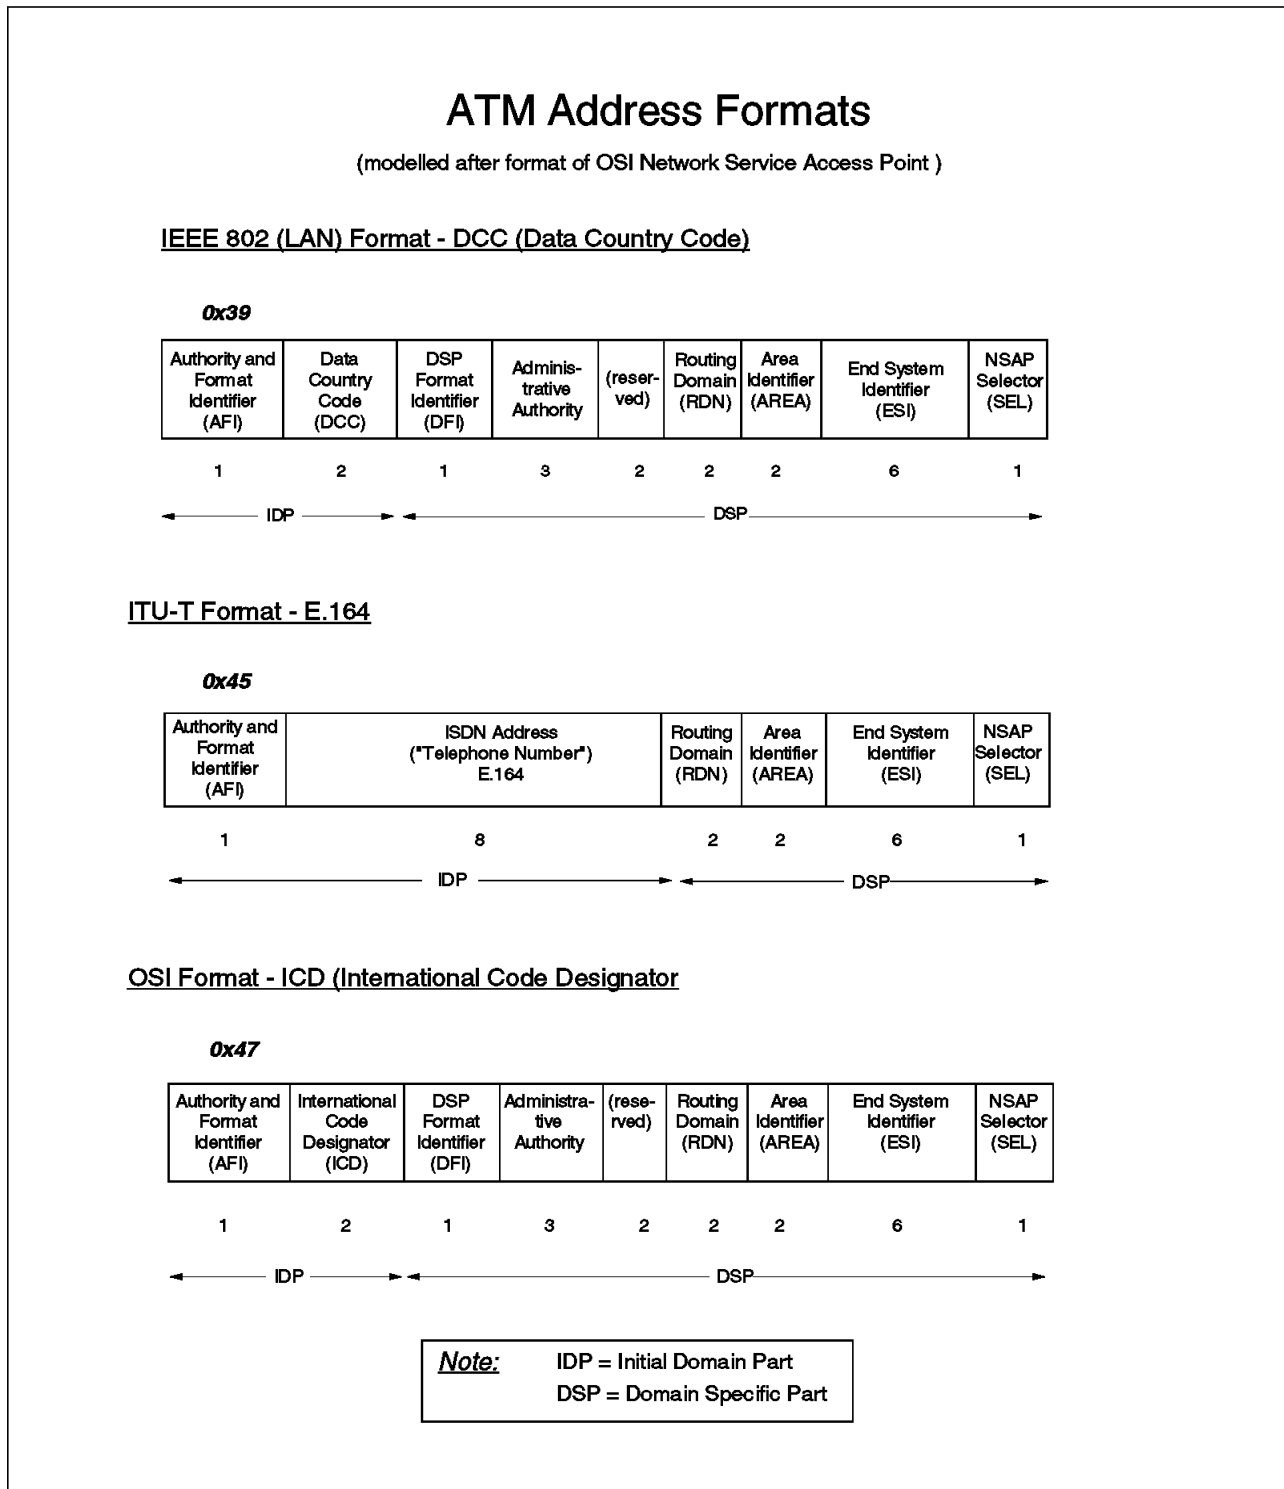


Figure 22. Structure of ATM Addresses

The address formats used in ATM are shown in Figure 22. The ATM address is built from the following main parts:

- The leading 13 bytes contain the network part of the address. These leading 13 bytes are looked at differently, depending on whether SSI or PNNI is used as the routing protocol between a group of switches:
 - Within an SSI cluster these leading bytes include fields for the standard network prefix (bytes 1 to 9), routing domain (bytes 10 and 11), cluster number (byte 11), as well as hub or switch number (byte 12) as shown in the mentioned figure. All switches within an SSI cluster share the identical 11 leading bytes.
 - With a PNNI peer group all switches share the same prefix (some number of leading address bits from 0 to 104), and the remaining bits of the leading 13 bytes are used to uniquely identify the ATM switch within the peer group. Additionally, PNNI uses the notion of hierarchical peer groups. At the lowest hierarchy level a Peer Group consists of single switches each maintaining a view of the topology of the whole group. At higher hierarchy levels a peer group consists of a group of peer groups, where one switch of each immediately lower peer group maintains a view of the topology of the higher-hierarchy peer group only. The advantage of this approach will be to reduce the amount of routing information being exchanged among the switches.
- The end system address identifies the end system uniquely among all end systems connected to the same ATM switch. The field is 7 bytes long and consists of an ATM end system address (similar to a LAN MAC address) and a single-byte end system identifier (ESI) that identifies a sub-component of an end system. Using narrow-band ISDN as an analogy, the first six bytes of this field would identify the house, and different ESIs (last byte) would be used to identify several phone sets, fax devices, and other devices used within the same house.

The following are the three different formats for ATM addresses, each controlled by a different authority:

- **ITU-T (E.164) Format:** This format is essentially the same as telephone style addressing. It is specified by the ITU-T and will be used by public (carrier provided) ATM networks.
- **Data country code (DCC) format:** This format carries LAN addresses as specified in IEEE 802 recommendations.
- **IDC format:** This format is specified by the ISO for OSI.

The ATM Forum specifies that equipment in a private network must support all three formats.

As of today, the International Organization for Standardization (ISO) is the authority granting official ATM addresses in ICD format (first byte = X'47') to US organizations and companies. In the UK this is done by the British Standards Institute (BSI). In other countries the corresponding standard bodies will most probably be the entitled authorities for ATM address registration within their respective geographies.

3.2 A Layered View of ATM

3.2.1 ATM Layer Model

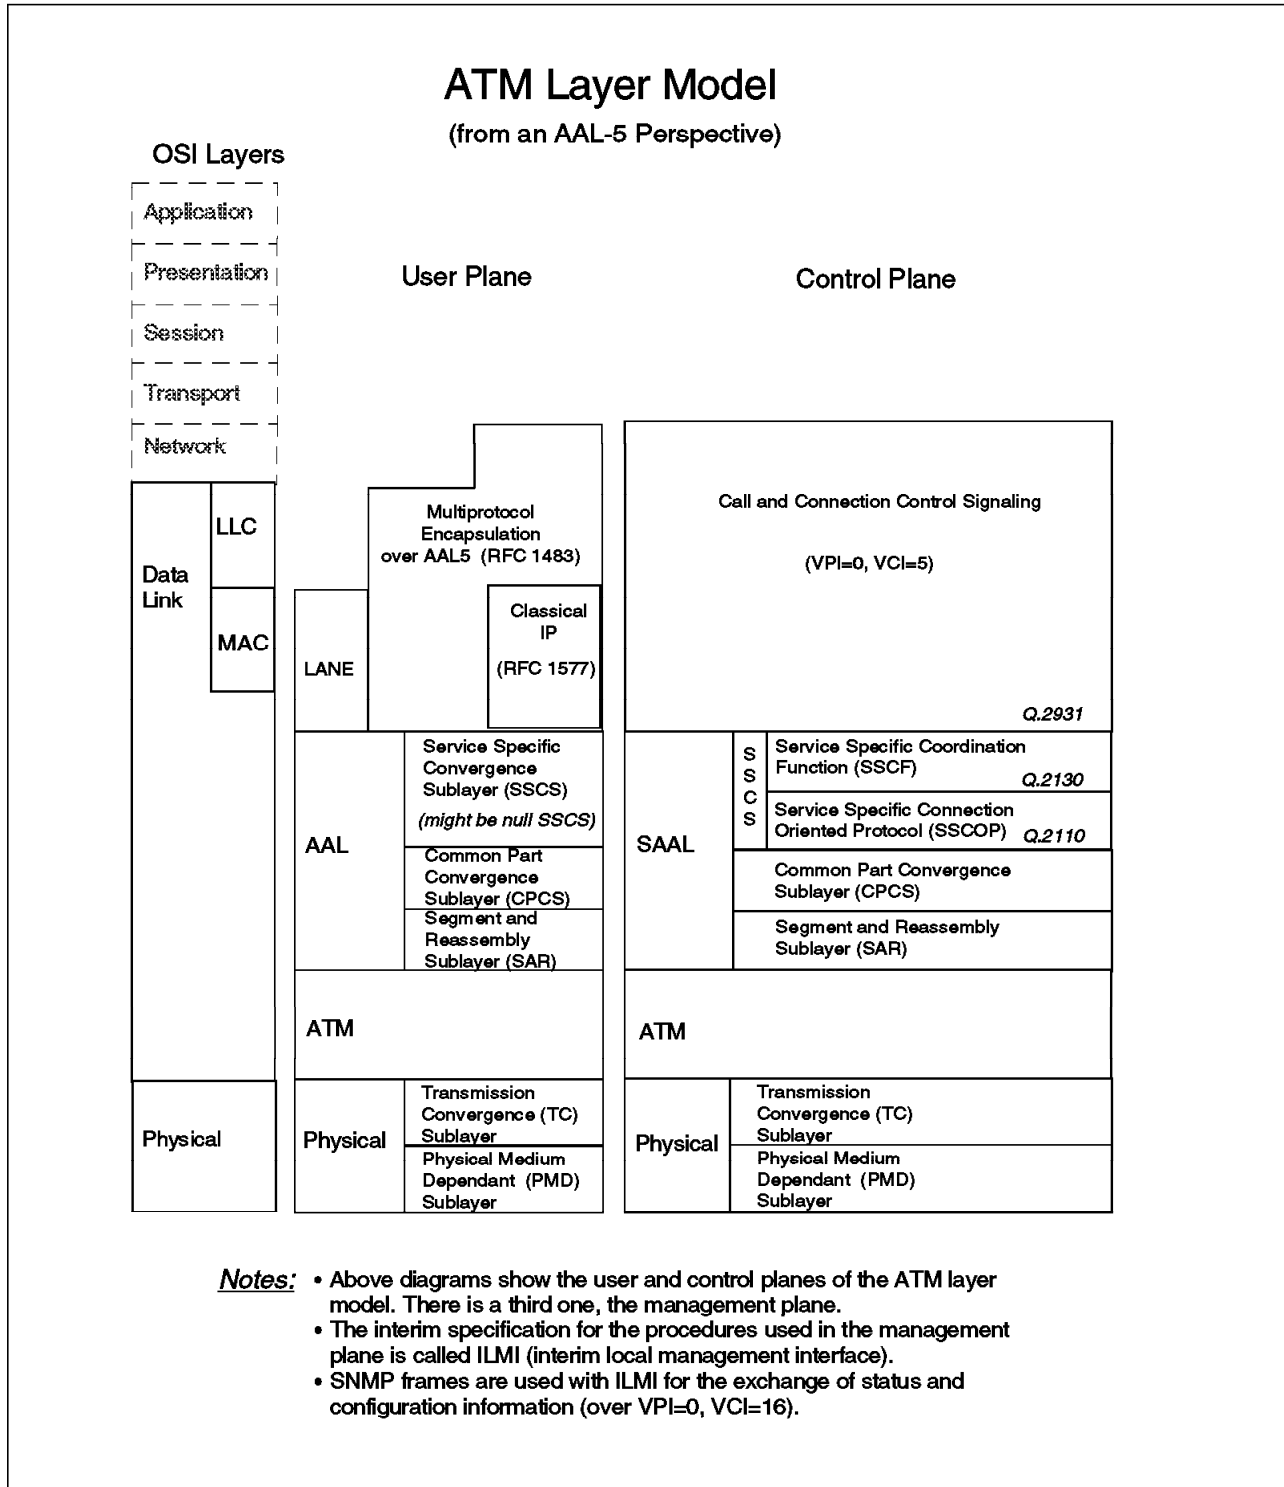


Figure 23. ATM Layer Model

3.2.2 Physical Layer




Figure 24. ATM and the Physical Layer

3.2.2.1 ATM Physical Interfaces

SONET and SDH: Synchronous Optical Network (SONET) is a US standard for the internal operation of PTT optical fiber networks. It relates closely to a system called Synchronous Digital Hierarchy (SDH) which has been adopted by the ITU-T for the internal operation of carrier (PTT) optical fiber networks worldwide.

Traditionally, PTT networks have been built by using a cascade of bandwidth multiplexors at each end of the high-speed connection. This resulted in more and more stages of multiplexing to provide faster links, the internals of which were generally proprietary. For example, the US used a different structure from Europe, and both the US and Europe were different from Japan.

Both SONET and SDH, which was developed from it, eliminate the problems illustrated above by providing a standardized method of internal operation and management and worldwide compatibility, while enabling existing speeds to be accommodated. They permit many levels of multiplexing and demultiplexing to be achieved in a single step, and allow many different speed channels to be carried through the system. With this new scheme, it is possible to access low bandwidth channels without having to demultiplex the whole bandwidth stream.

The basic structure in SONET is a frame of 810 bytes, which is sent every 125 micro seconds. This allows a single byte within a frame to be a byte in a 64 kbps digital voice channel. Since the minimum frame size is 810 bytes, this means that the minimum speed at which SONET will operate is 51.84 Mbps:

- $810 \text{ bytes} \times 8000 \text{ frames/second} \times 8 \text{ bits} = 51.84 \text{ Mbps}$

The basic SONET frame is called the Synchronous Transport Signal Level 1 (STS-1), as shown in Figure 24 on page 62.

It is conceptualized as containing 9 rows of 90 columns each with the following attributes:

- The first three columns of every row are used for administration and control of the multiplexing system. They are called overhead in the standard but are very necessary for successful systems operation.
- The frame is transmitted row by row, from the top left of the frame to the bottom right. One frame is transmitted every 125 micro seconds.
- It is important to remember that the representation of the structure as a two dimensional frame is purely a method of showing a repeating structure. In reality it is just a string of bits with a defined repeating pattern.

Multiple STS-1 frames can be byte multiplexed together to form higher speed signals. When this is done, they are called STS-2, STS-3, etc., where the numeral suffix indicates the number of STS-1 frames that are present. For example, STS-3 is three times an STS-1 or 155.52 Mbps.

In Figure 24 on page 62 the structure of other frames used at different speeds with ATM are also shown. The most popular speed is probably 155 Mbps, as used commonly to connect switches within private campus networks. Other speeds are also becoming increasingly popular such as 25 Mbps for ATM to the desktop, and 622 Mbps for switch interconnection within the campus.

SONET LITE: This is not an officially accepted term, but is a convenient description of the underlying technology upon which it is based. It is the standard proposed by the ATM Forum for use with multimode fiber in the local LAN environment, and is based on the SONET STS-3c standard. With SONET LITE some of the overhead bytes are just not used or interpreted (but still present, and sent as empty or meaningless bytes).

The term "LITE" comes from the fact that most of the management information flows of STS-3c have been replaced. This is possible because of

the relatively short distances involved in LAN environments which make them inherently more reliable than WAN connections. The approach of minimizing some of the frame overheads, significantly reduces the cost of implementation making it an attractive implementation option.

The ATM cells from the ATM layer (see next section) are packed in sequence into the above frame types for transmission over the physical links, and extracted from the frames upon reception at the destination side.

3.2.3 The ATM Layer

The ATM layer multiplexes the cells it receives from the higher AALs (ATM Adaptation Layers) and sends them to the physical layer. If there are not enough cells from the AALs to be sent, the ATM layer adds so-called idle cells to fill the gaps (with VPI = 0 and VCI = 0). On the receiving side the ATM layer discards idle cells and forwards cells with actual payload to the appropriate AALs.

3.2.4 The ATM Adaptation Layer (AAL)

The network characteristics required by various types of traffic over an ATM network are provided by an ATM adaptation layer which is found in each end system and, in special forms, in switches also.

The ITU-T has defined different generic service classes of network traffic, each of which must be treated differently by an ATM network. These classes are designated Class A to Class D and four different types of ATM adaptation layers (AAL) have been defined to realize the necessary network characteristics to handle them. Class X with AAL-0 is a non-defined class, meaning that it is defined and understood by the user application implementing it. An overview of these service classes is shown in Figure 25 on page 65.

Service Classes					
<i>Service Class</i>	Class X	Class A	Class B	Class C	Class D
<i>AAL type</i>	AAL-0	AAL-1	AAL-2	AAL-5	AAL-3/4
<i>Bit rate</i>	(unspecified)	constant	variable	variable	variable
<i>Connection mode</i>	(unspecified)	connection oriented	connection oriented	connection oriented	connectionless
<i>Timing (end-to-end)</i>	(unspecified)	required	required	not required	not required
<i>Application example</i>	user defined	isochronous voice/video (circuit emulation)	compressed voice/video	data (LANE, Classical IP, etc.)	data (SMDS)

Figure 25. ATM Service Classes

The following sections provide a short description of these service classes:

- Class A (circuit emulation)

This service emulates a leased line and is used for traffic that has a constant bit rate, for example voice and video.

The characteristics of Class A are the following:

- A constant bit rate at source and destination
- A timing relationship between source and destination
- A connection between end users of the service

To realize these functions, the adaptation layer must perform the following services:

- Segmentation and reassembly of data frames into cells
- Handling (by buffering) of cell delay variations
- Detection and handling of lost, discarded, misrouted or duplicated cells
- Recovery of the source clock frequency
- Detection of bit errors in the user information field

- Class B (Variable Bit Rate Services)

This service is intended for isochronous voice and video traffic which may be coded as variable rate information, and requires a timing relationship between the ends of the connection. The service is strictly connection oriented.

The services provided by the Class B adaptation layer are the following:

- Transfer of variable rate information between endpoints.
- Transfer of timing between source and destination.

- No indication is provided of lost or corrupted information.

Some video applications (such as videoconferencing) require synchronization between voice and video, while others can be transmitted in multicast mode (just like a film) and are not sensitive to network delay.

To design and control a network for Class B traffic is very challenging because of the unpredictable way high bandwidth is required, and because the data rates are often near to the peak rate capability of the network.

- Class C (connection-oriented data)

Class C traffic is traditional data traffic, such as SNA and TCP/IP. The service offered for it is connection oriented and it supports variable rate information flow.

The services provided by the Class C adaptation layer are the following:

- Segmentation and reassembly of frames into cells
- Detection and signalling of errors in user data frames
- Possible multiplexing and demultiplexing of multiple end-user connections into a single ATM connection

This is the service class most widely implemented and used nowadays.

- Class D (connectionless data)

The class D service is connectionless and also supports variable rate information flow. It is intended to support connectionless protocols such as TCP/IP.

The services provided by the class D adaptation layer are the following:

- Segmentation and reassembly of frames into cells
- Detection and signalling of errors in user data frames
- Multiplexing and demultiplexing of multiple end-user connections into a single ATM connection
- Network layer addressing and routing

- Class X (user defined)

This is a connection-oriented ATM transport service where the network characteristics are user defined. Only the required bandwidth and the QoS parameters are used by the network.

3.2.4.1 AAL-0

This is the null AAL and corresponds to a process that connects the AAL service interface directly to the ATM networking service.

3.2.4.2 AAL-1

AAL-1 is used for Class A (constant bit-rate) traffic. In practice this may take the form of data in an SDH or PDH frame where the frame rate is constant but data exists in a specific part of the frame, so that it arrives at the network in short periodic bursts.

When these cells are transported through the network, depending upon the current load on the network, they may be delayed. This delay is not consistent, and can introduce jitter. The receiver has to buffer the received data to avoid problems of overrun or underrun of data frames, with consequent additional delay in the data stream.

At the beginning of the AAL-1 cell payload there are two additional control fields, each one 4 bits long.

- **Sequence Number (SN):** This field contains a 3-bit sequence number field that operates in cyclic fashion, with a 1-bit CS indicator (CSI bit), which is used depending upon the type of the traffic.
- **Sequence Number Protection (SNP):** This field is a CRC error check protection for the SN field.

In order to minimize the delay caused by assembly and playout, it is possible to send cells that are not full. This fact must be specified between the end systems at the time of circuit establishment.

3.2.4.3 AAL-2

AAL-2 is used for Class B traffic. AAL-2 processes data streams similar to AAL-1. The only difference is the variable bit rate due to compression. One of the key problems of AAL-2 is to handle the skew between voice and video information. AAL-2 is currently absent from the draft standards. The detailed description being delayed because of the above mentioned problems.

3.2.4.4 AAL-3/4

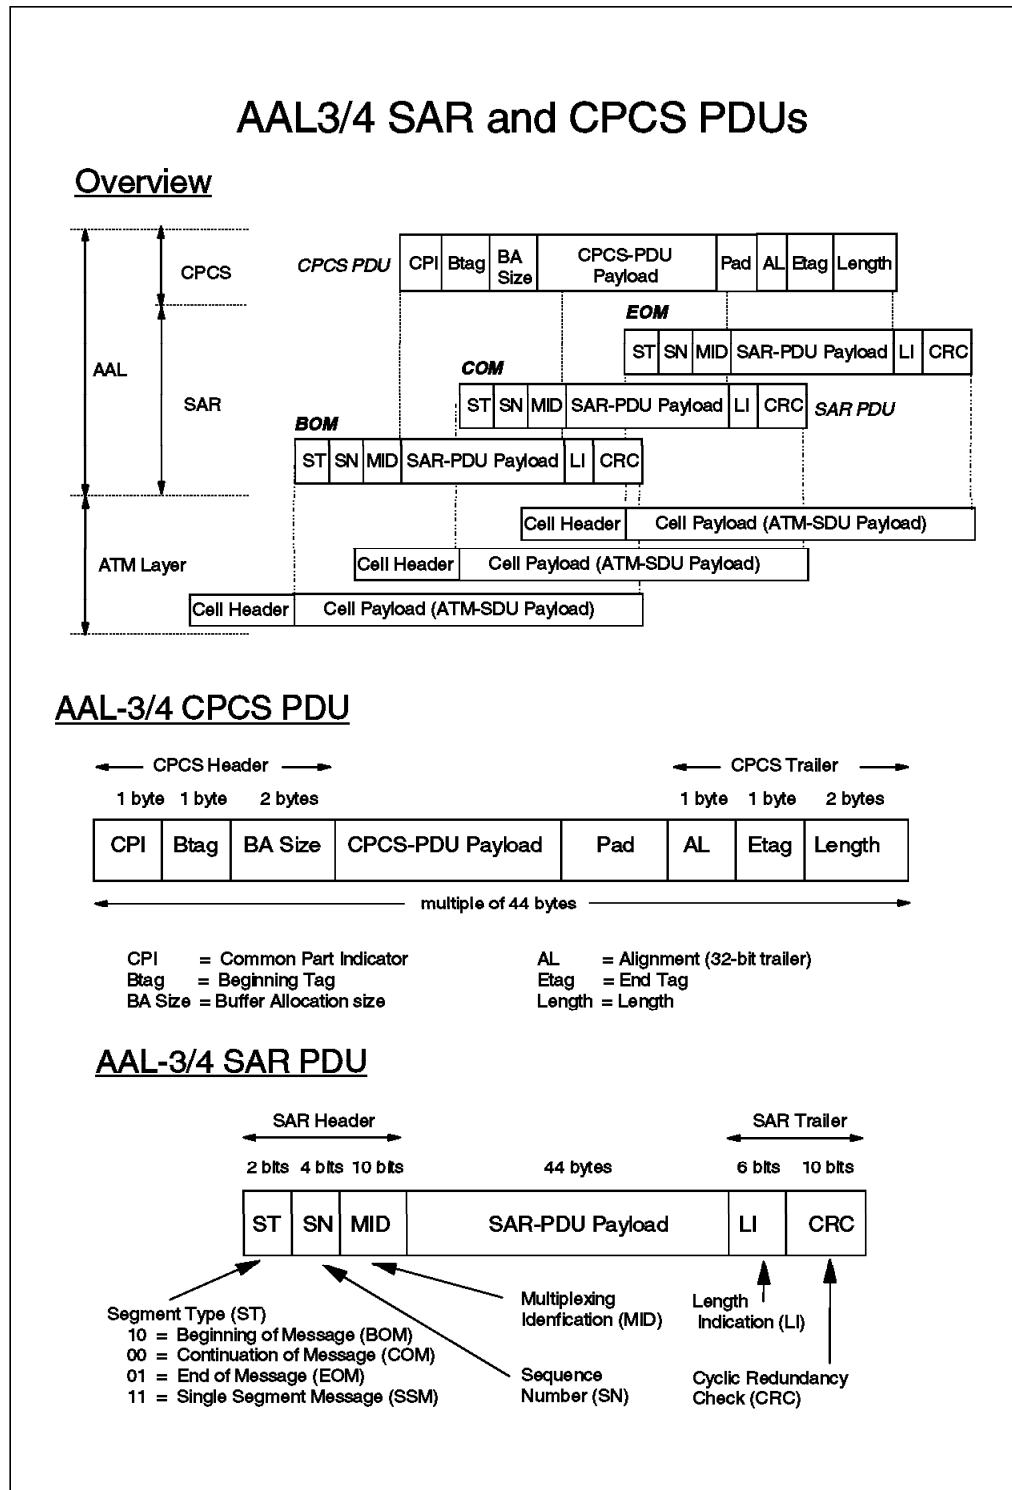


Figure 26. AAL-3/4 SAR and CPCS PDUs

AAL-3/4 is used for Class C and Class D traffic.

It is a relatively complex, high-function AAL that will offer ensured data delivery. When the AAL detects corruption of a data frame, it will be able to re-transmit the data frame, although the details of this operation have not been defined.

- In blocking mode short data frames are blocked into a longer data unit for transmission through the network, and the multiplexing of several AAL-to-AAL connections onto a single VCC connection is possible.
- Point-to-multipoint connections are also possible using AAL-3/4. In this mode several additional control information bits are placed in the payload field.

The format of the SAR and CPCS PDUs for AAL-3/4 is shown in Figure 26 on page 68. The upper part of the picture depicts the process of segmentation and reassembly (SAR), where the payload of ATM cells is used to build the SAR protocol data units (one to one relationship for AAL-5). Then the payload of several SAR SDUs (SAR SDU = SAR PDU payload) is concatenated to build a CPCS PDU. The latter reaches further into the SSCS sublayer or (in case of null SSCS) directly into the corresponding protocol above the AAL layer. The middle and bottom part of the picture show the structure of the AAL-3/4 CPCS PDU and SAR PDU respectively.

The fields have the following meaning:

- **Segment Type (ST):** These two bits indicate where the content of this cell is located in the data frame.
- **Sequence Number (SN):** This 4-bit field is a module four counter to verify the correct sequence of the cells.
- **Multiplexing Identification Field (MID):** In the case of multiplex mode this field indicates the connection to which the cell belongs.
- **Length Indicator (LI):** This field specifies how much data is in the variable part of the cell. All except the last cell of the data frame should be full.
- **Cyclic Redundancy Check (CRC):** This is a polynomial CRC to protect the entire cell payload except the CRC itself.

3.2.4.5 AAL-5

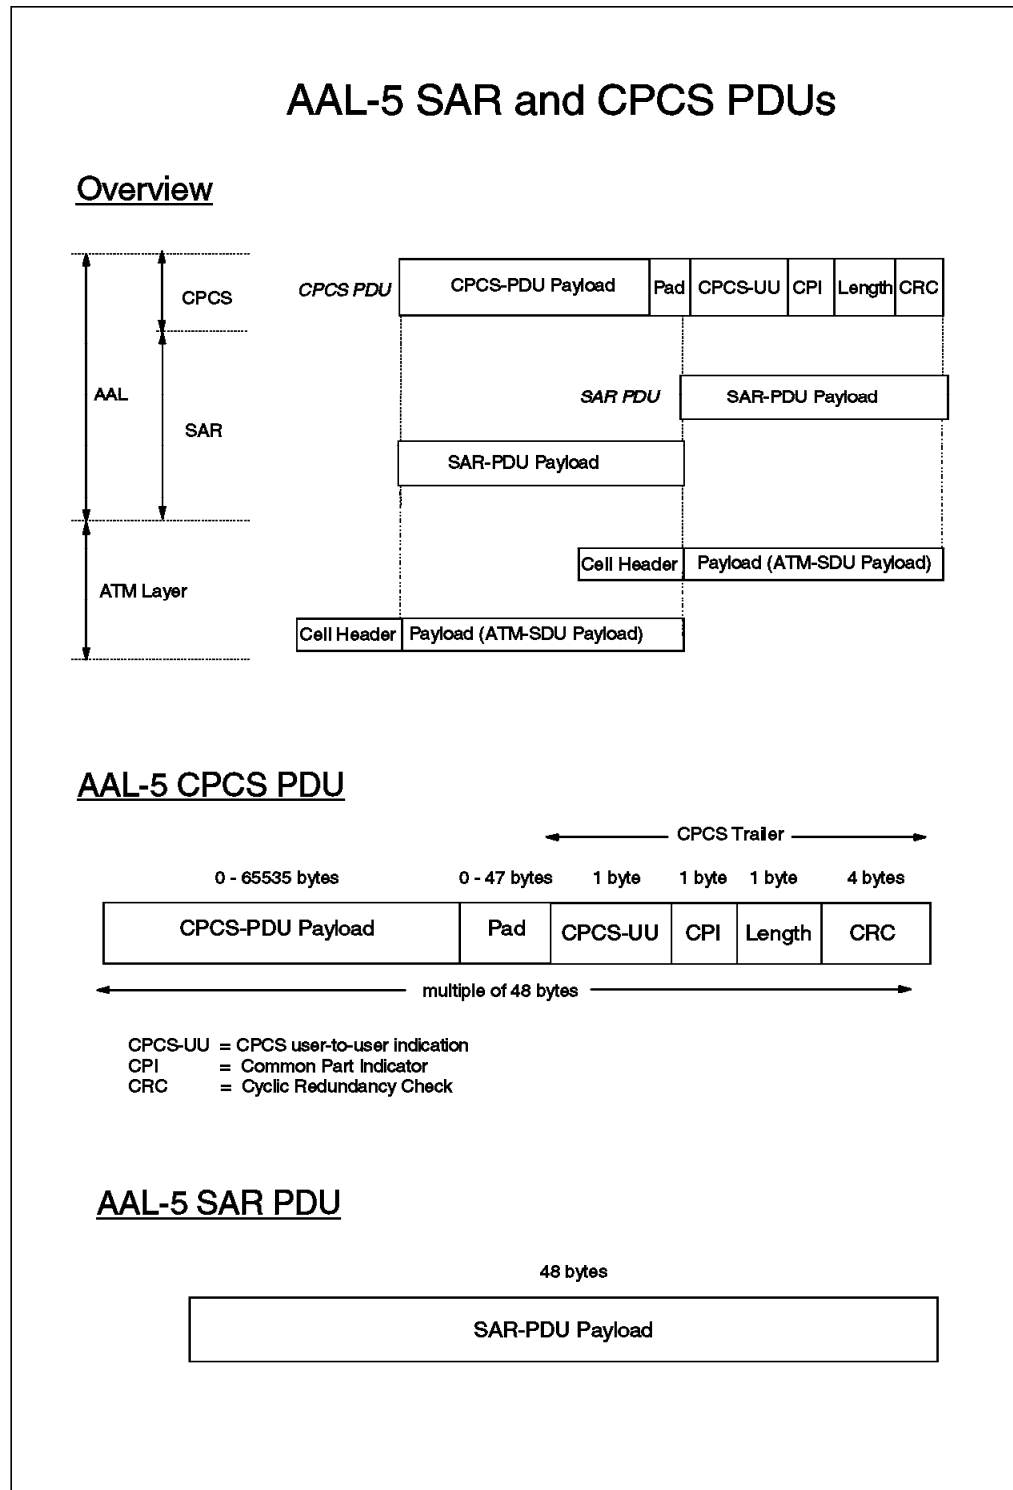


Figure 27. AAL-5 SAR and CPCS PDUs

This AAL is often called SEAL (Simple and Efficient Adaptation Layer). It is designed to operate significantly more efficiently than AAL-3/4 but, with the exception of connection multiplexing, has the same functions as AAL-3/4.

The format of the SAR and CPCS PDUs for AAL-5 is shown in Figure 27. The upper part of the picture depicts the process of segmentation and reassembly (SAR), where the payload of ATM cells is used to build the SAR protocol data units (one to one relationship for AAL-5). Then the payload of several SAR SDUs (SAR SDU = SAR PDU payload) is concatenated to build a CPCS PDU. The latter reaches further into the SSCS sublayer or (in case of null SSCS) directly into the corresponding protocol above the AAL layer (for example LANE or Classical IP). The middle and bottom part of the picture show the structure of the AAL-5 CPCS PDU and SAR PDU respectively.

In the case of AAL-5 the whole user frame is protected by CRC, and there is no length field in every cell as in the case of AAL-3/4. In this way the whole 48 bytes of the payload can be used for data transmission. AAL-5, therefore, can recognize corrupted or missing data only when the whole data frame is received at the destination end system.

In Figure 28 on page 72 we can see how the more familiar frames are packed for transport over ATM. A more detailed overview of the structure of LANE frames can be found in Appendix D, "ATM Forum-Compliant LANE Frame Formats" on page 479.

AAL-5 Frames (User Plane)

- Notes:**
- Neither LANE nor RFC 1483/1577 make use of SSCS
 - LE Header is 2 bytes long: either LEC-ID or X'0000'
 - Type of SVC connection (LANE, 1483, 1577, etc) negotiated between parties at call setup time (control plane), or agreed by default (PVCs)

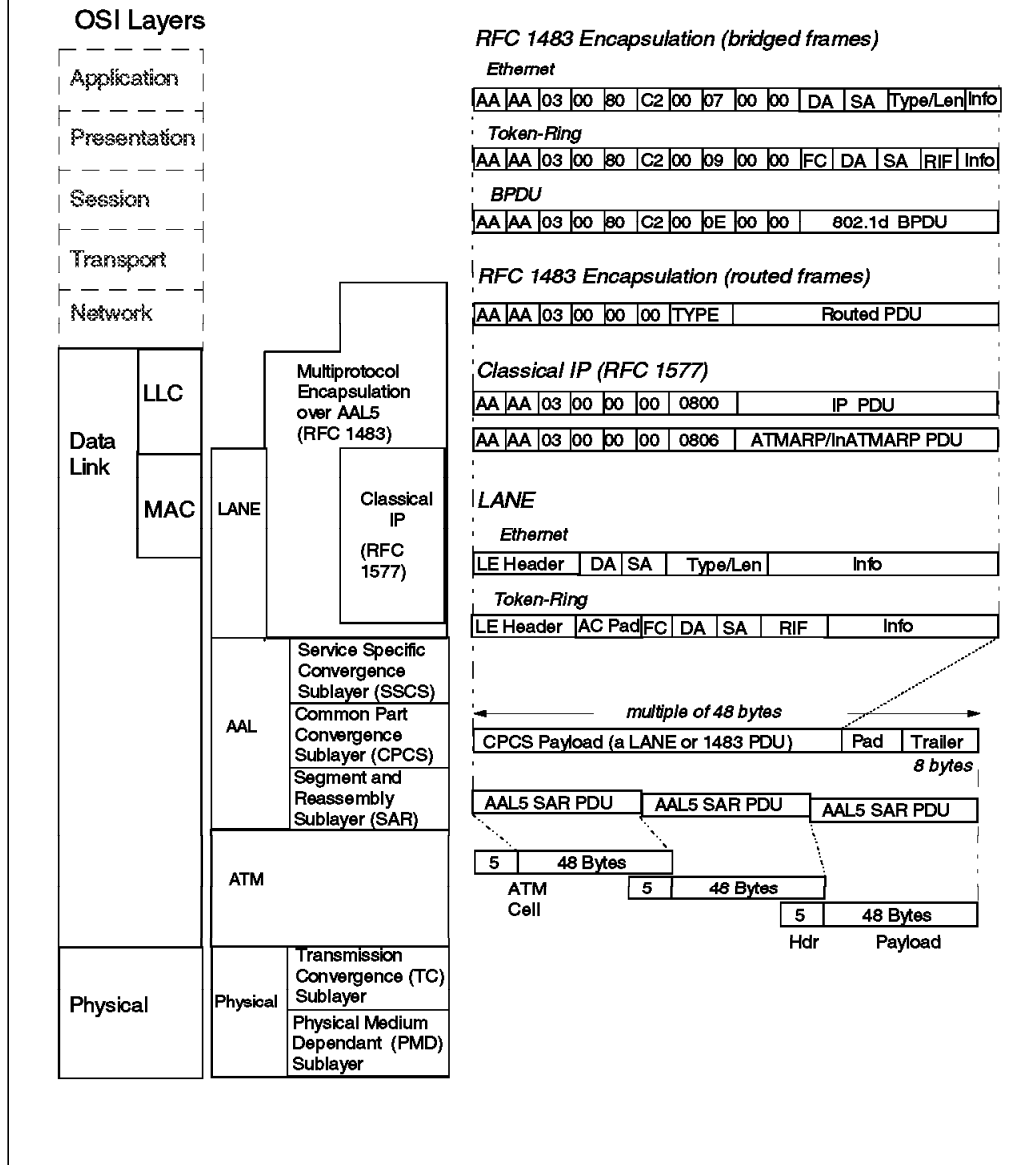


Figure 28. AAL-5 Frames (User Plane)

3.2.4.6 Signalling AAL (SAAL), Control Plane

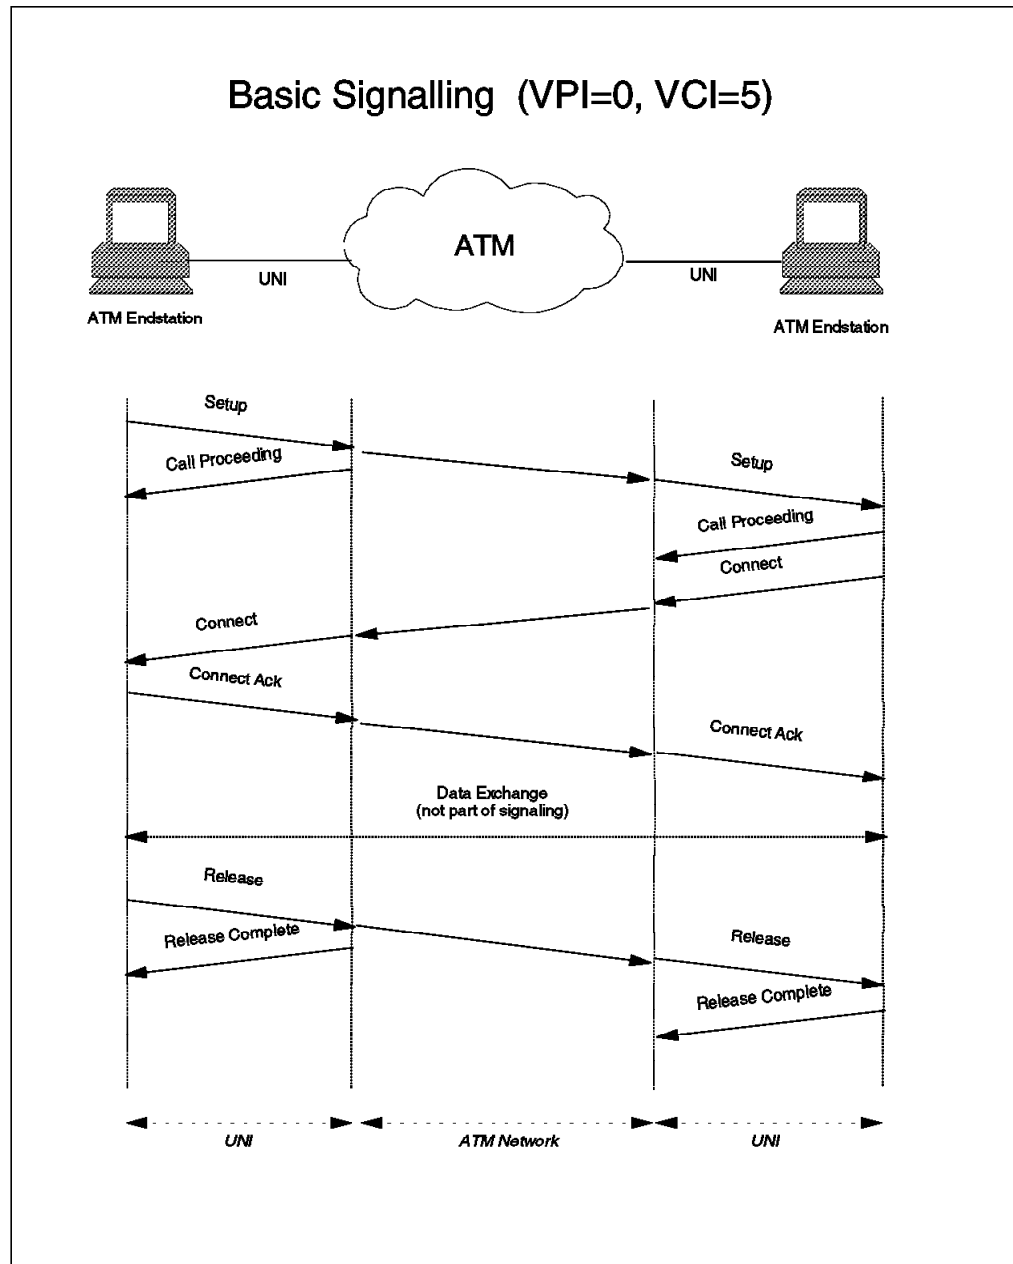


Figure 29. Basic Signalling at the UNI

Connection setup, control and tear-down is done via signalling over the VPI=0, VCI=5 channel across the UNI. The following functions and sub-functions are defined for ATM signalling:

1. Call Establishment
 - Setup
 - Call Processing
 - Connect
 - Connect acknowledge
2. Call clearing
 - Disconnect

- Release
 - Release complete
3. Status
- Status enquiry
 - Status
4. Point-to-Multipoint Messages
- Add party
 - Add party acknowledge
 - Add party reject
 - Drop party
 - Drop party acknowledge

Typical connect and disconnect procedures are shown in Figure 29 on page 73.

Signalling Frames (Control Plane)

Q.2931 Signalling Frames (VPI=0, VCI=5) ("messages" in UNI terminology)

Protocol discriminator (X'09')				1
0	0	0	0	Length of call ref. val. 2
Flag	Call reference value			3
Call reference value (continued)				4
Call reference value (continued)				5
Message type				6
Message type (continued)				7
Message length				8
Message length (continued)				9
Variable length information element(s) (IEs)				etc.

Information Elements (IEs) (within Q.2931 frames)

8	7	6	5	4	3	2	1	
Information element identifier								1
1 ext	Coding Standard	IE Instruction Field						2
		Flag	Res.	IE Action Indicator				
Length of information element								3
Length of information element (continued)								4
Contents of information element								etc.

Note: The Q.2931 frame shown above is the frame used to setup, control and tear down connections. The call connection and control function delivers it to the SAAL layer (see also ATM layer model for reference of the layers involved).

Figure 30. Q.2931 Signalling Frames (Control Plane)

In Figure 30 the structure of the frames exchanged between the end system and the network when establishing and disconnecting connections is shown. These frames are called *messages* in UNI terminology. All these frames have the same structure: A basic part identifying the type of message, plus none or some additional information packed in so-called information elements (IE). The optional inclusion of IEs depends on the type of the message being sent. For example the IEs used at connection setup may contain information such as

- AAL parameters, stating the AAL type, CPCS-SDU forward and backward sizes, etc.
- ATM traffic descriptor, indicating desired/supported Peak Cell Rate, Sustainable Cell Rate, Maximum Burst Size, Best Effort Indicator, etc.

- Broadband low layer information, indicating the type of the desired connection such as LE config direct VCC, IEEE 802.5 LE data direct VCC, etc.
- Called party number (ATM address of target end system)
- Calling party number (ATM address of source end system, the originator of the call)
- QoS parameters

and some others (check the UNI specification for further reference).

In case of network problems, and just after the end system has issued a setup message, the switch might reject the connection while returning a corresponding message to the end system. The returned message will contain a cause code for the rejection of the connection. This cause code will be packed in a so-called Cause IE. A list of the cause codes supported by IBM switches can be found in Appendix C, "UNI 3.0-3.1 Cause Maintenance Error Codes" on page 471.

Note that although the UNI interface is local between the end system and the switch it is connected to, some of the parameters contained in the signalling frames must be sent to the target end system. This is the case for example with the broadband low layer information contained in the corresponding IE at connection setup.

Generic Flow Control (GFC): There is another type of local signalling used for traffic management at the UNI. The GFC is a field in the header of UNI cells (see also Figure 21 on page 56). Using this field one-way flow control is defined from the ATM end system to the ATM switch. There is no control in the opposite direction.

According to the UNI standard three queues may be defined at the end system, one for uncontrolled, and two for controlled traffic. Usually, though, only one queue is used for controlled traffic.

- **Controlled traffic:** This is the traffic for which the GFC mechanism is defined. It is usually all the non-reserved bandwidth (NRB) traffic on the interface.

Controlled traffic is distinguished in the cell header by the presence of a non-zero GFC field.

- **Uncontrolled traffic:** This traffic is not subject to GFC control and is treated as having a higher priority than the controlled traffic. It would normally be traffic for which there is a reserved bandwidth (RB).

The flow control mechanism uses windowing. Each queue has a window that represents the number of cells that it is allowed to send before the network must respond to give permission to continue. The window is maintained as a counter in the end system. Each time a cell is sent the counter is decremented and the end system is allowed to send cells as long as the value of the counter is not zero. If the counter reaches zero the end system must stop transmitting until the counter has been reset to an initial value. During normal operation the switch sends reset signals fairly often so that the counter never reaches zero.

Use of the GFC field for controlled traffic:

GFC Field towards the network (outbound):

- Bit 0 is unused (MSB of the GFC field).
- Bit 1 indicates that the cell is flow controlled by Q1.

- Bit 2 indicates that the cell is flow controlled by Q2 (in the case where only one queue exists this bit is always zero).
- Bit 3 (LSB of the GFC field) indicates if the equipment is controlled (1) or not controlled (0).

GFC field away from the network (inbound):

- Bit 0 means HALT (1) or NOHALT (0). If HALT the network is unable to receive input from the end system, even uncontrolled reserved bandwidth traffic.
- Bit 1 when set means reset the counter of Q1.
- Bit 2 when set means reset the counter of Q2. If Q2 does not exist it must be 0.
- Bit 3 is reserved for future use.

3.2.4.7 The Interim Local Management Interface (ILMI)

Standards for management of ATM networks are still in the formative stage. In the interim period until these are completed, the ATM Forum has defined an interim specification for the management plane based on the use of existing SNMP technology and an ATM UNI management MIB. It is called the Interim Local Management Interface (ILMI) specification, and was included in the documentation of the UNI specification.

ATM switches communicate their network address prefixes to end systems using the ILMI, and end systems communicate the ESI portion of their network address to the switches at initialization time (so-called end system address registration).

ILMI uses SNMP flows between the end system and the switch. The flows are encapsulated in AAL-5 and a standard channel, VPI=0, VCI=16 is used, although the specification indicates that this should be a programmable option. Message formats are as defined in RFC 1157, that is according to SNMP Version 1, not SNMP Version 2.

ILMI provides access to a range of information that is provided in the MIB defined by the ATM Forum. Examples of the categories of information held within the MIB are as follows:

- Physical layer
- ATM layer
- ATM layer statistics
- Virtual path (VP) connections
- Virtual channel (VC) connections
- Address registration information

This information is used in a variety of ways including end system address registration, monitoring of interface states, determination of VP and VC availability and many others.

3.3 Networking Models Based on ATM

There are many networking models based on ATM technology, all with the common goal to provide for as direct as possible connections between communication partners, eliminating the overhead and delay of intermediate equipment. Among them we find *LAN Emulation V1 (LANE 1.0)*, *Classical IP (RFC 1577)*, *Next Hop Routing Protocol (NHRP)*, *LANE V2*, *Multiprotocol over ATM (MPOA)*, etc. In the following sections we shall take a closer look at those which are already in widespread use, that is LANE 1.0, Classical IP, and NHRP.

3.3.1 LAN Emulation V1 (LANE V1.0)

LAN emulation enables the implementation of emulated LANs over an ATM network. An emulated LAN provides communication between end systems similar to that over a real LAN but using the facilities of the ATM network. LAN emulation software in each end system is provided to allow existing network system software to continue to operate in the new ATM environment; thus existing application software can continue to be used without any changes. This very fact makes LAN emulation powerful and so an important concept for ATM today, especially prior to widespread availability of applications using native ATM APIs and heavily exploiting ATM capabilities such as quality of service and others.

As mentioned earlier ITU-T (former CCITT), being a sub-organization of the Organization of the United Nations (ONU), is the overall standardization body for ATM. The ATM Forum, as a consortium of vendors and user representatives fills the gaps where there are no ITU-T standards yet in place. LAN Emulation version 1 (LANE 1.0) is such a specification from the ATM Forum. As of today the ATM Forum is already working on the next release of this specification due to be called LANE 2.0.

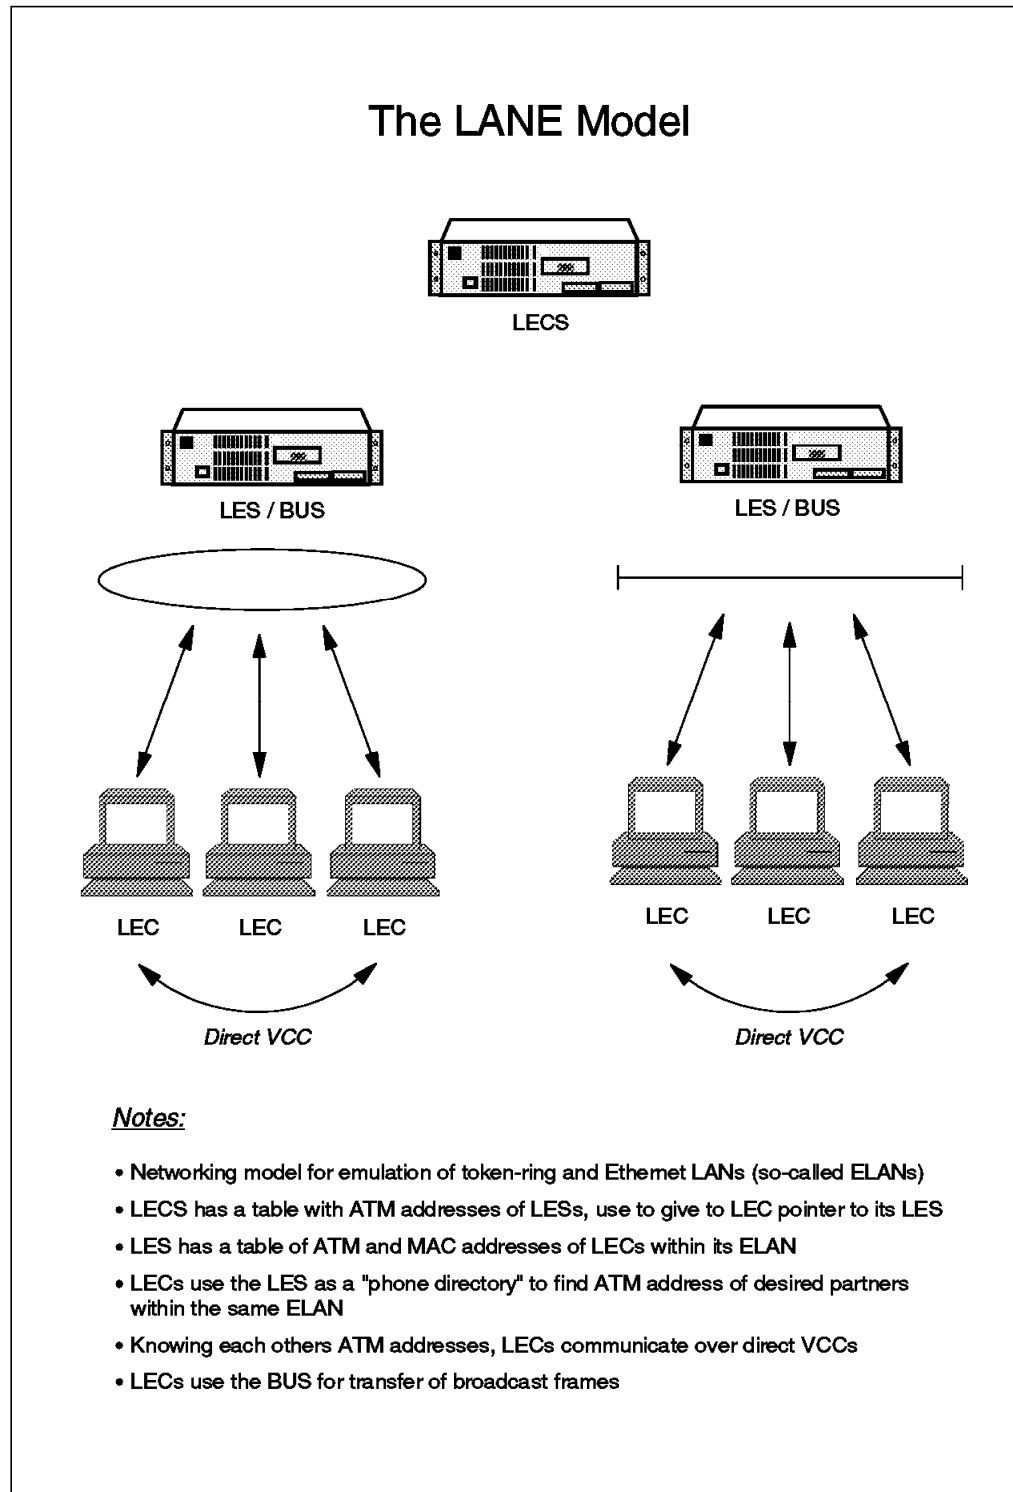


Figure 31. The LANE Model (LANE 1.0)

Figure 31 gives an overview of the components of LAN Emulation over ATM (LANE 1.0). There we recognize the following entities:

- LAN emulation configuration server (LECS): This is an administrative entity of central importance for LANE. It contains:
 - A list of pointers (ATM addresses) to the LAN emulation servers (LESs) in an ATM network, and

- A list of so-called policies per LAN emulation server. The policies are conditions a LAN Emulation Client (LEC) is expected to fulfill before it is allowed to join the emulated LAN (*ELAN*) of a certain LES.
- LAN emulation Server (LES): The LES fulfills two main functions:
 - It helps a LAN emulation client (LEC) to join the ELAN.
 - It holds a table needed for address translation between MAC and ATM addresses. When a LEC (source) wants to contact another LEC (target), and it knows only the MAC address of the target LEC, the source LEC will submit to the LES the MAC address of the target LEC asking for the ATM address it translates to. Upon reception of this ATM address, the source LEC will contact the target LEC directly. The mechanism used for address translation resembles very much the mechanism used in TCP/IP for address resolution (ARP and RARP protocols).
- Broadcast and unknown server (BUS): This component takes care of the broadcasts. Since in an ATM environment we work with dedicated bandwidth connections, there is no way to reach all other participants of an ELAN directly. Therefore a LEC wishing to broadcast a frame (sending it to all others) will send the broadcast frame only to the BUS, and the BUS will then forward it to all others. For this purpose, the BUS maintains and uses a special type of connection called a multicast tree. The multicast tree is a point-to-multipoint connection where the root is able to send messages to the leaves. During its initialization phase, a LEC will request the BUS to add it to the bus's multicast tree, so it becomes available for reception of later broadcast frames.
- LAN emulation client (LEC): This is a participant in an emulated LAN (ELAN). There are two types of LAN Emulation clients:
 - Non-proxy LEC: This is a normal workstation or a host.
 - Proxy LEC: This is a LAN/ATM bridge (or a router with bridging functions), representing the non-ATM workstations attached to legacy LANs (token-ring, Ethernet, etc.).

At initialization time, a LEC does not yet know the ATM address of the LES representing the ELAN the LEC is supposed to join. Therefore the LEC will ask the LECS for the ATM address of its corresponding LES. In doing so, the LEC will provide to the LECS some (optional) data such as an ELAN name, type of emulated LAN it wishes to join (token-ring or Ethernet), etc. The LECS then checks through its list of policies to find a match. If the data given by the LEC fulfills the criteria of a certain LES, the LECS will provide the ATM address of this LES to the LEC. Then the LEC will request the LES to let it join the desired ELAN.

For more detailed information about LAN Emulation, please refer to Chapter 7, "ATM Emulated LANs and Logical IP Subnets (LANE 1.0, RFC 1577)" on page 293.

3.3.2 Classical IP (RFC 1577)

Classical IP is the term used for operation of conventional IP and ARP over an ATM network. The standards for doing this have been defined in RFC 1577.

Classical IP defines the use of IP and ARP in terms that are almost identical to their use on conventional broadcast LANs. ATMARP is used to resolve an IP address to an ATM address, allowing an end system to establish a direct

point-to-point connection with the destination end systems. This process is identical to that used by conventional ARP to resolve IP addresses to LAN MAC addresses.


Classical IP operates at the logical IP subnetwork (LIS) level (logical because there may be many such subnetworks on an ATM network). Each LIS must have an ATM ARP server to deliver the ARP service to end systems in the network, and all clients must be configured with the ATM address of the ARP server. The ARP server can be a dedicated server or it can be integrated with other functions.

Client end systems in the LIS must be registered with the ARP server at initialization time. The ARP server, therefore, is able to build up a table of ATM addresses and corresponding IP addresses. When queried by a client system it can therefore respond with the destination ATM address for the target IP address. If it does not have the required information for the response, it forwards the ATMARP request to all other clients with which it has a connection. In this way, there is the maximum chance of the IP address being resolved with the LIS.

Once the client end system has the ATM address of the destination end system, it establishes a direct connection over which data frames can be passed. Unlike other architectures, for example LAN emulation, there is no capability for the ARP server to forward data traffic.

The components of a classical IP network and their interaction are shown in Figure 32 on page 82.

The Classical IP Model (RFC 1577)



Notes:

- Networking model for ONE subnet only, a so-called LIS (logical IP subnet)
- ARP server has a table with IP and ATM addresses of ARP clients
- ARP clients use the ARP server as a "phone directory" to find ATM address of desired partners
- Knowing each others ATM addresses, ARP clients communicate over direct VCCs (little overhead within frames)

Figure 32. The Classical IP Model (RFC 1577)

For more more information about Classical IP, please refer to Chapter 7, "ATM Emulated LANs and Logical IP Subnets (LANE 1.0, RFC 1577)" on page 293.

3.3.3 Next Hop Routing Protocol (NHRP)

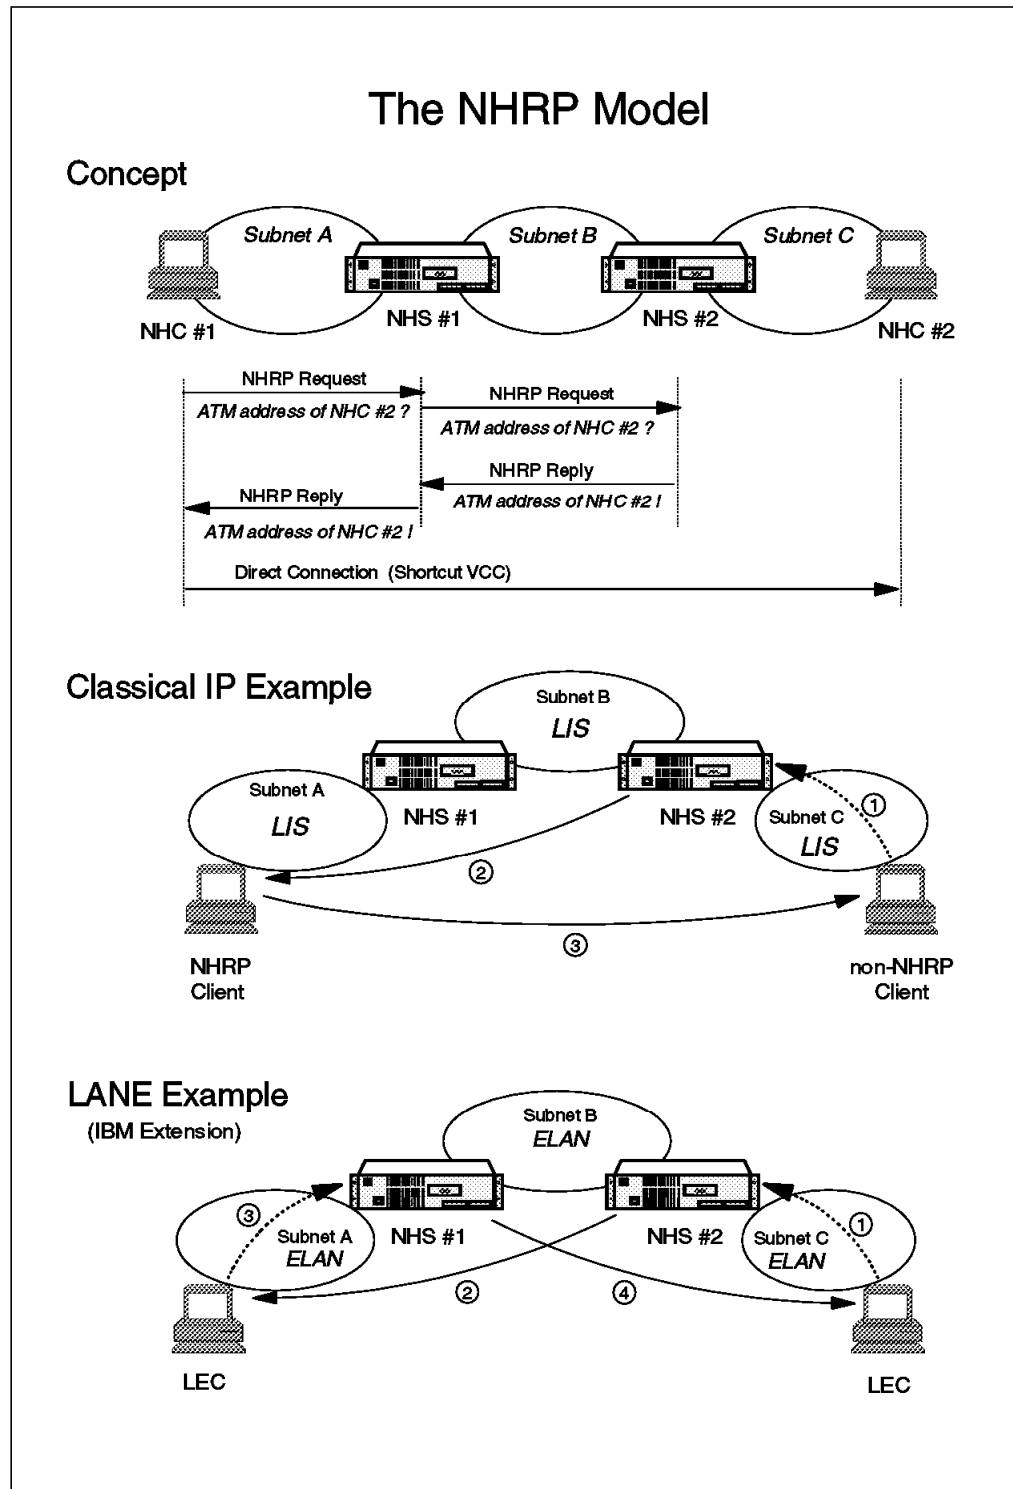


Figure 33. The NHRP Model

Classical IP communication is restricted to participants within the same subnet. This is made possible using an Address Resolution Protocol (ARP) mechanism. NHRP is similar in concept, but it allows communication over the subnet boundary, while keeping the advantage of using direct VCCs, that is without

incurring the need of intermediate equipment to forward the frames exchanged between the communication partners.

Figure 33 on page 83 shows the mechanism used between two Classical IP clients having also NHRP client (NHC) functionality, when they want to exchange data (upper part of the picture). This is also known as zero-hop routing. The mechanism is pretty similar in concept to the one used with the ARP protocol. An NHRP client (NHC) wanting to communicate with another client gets the ATM address of its remote partner from its immediate NHRP server (NHS) so it can build a direct VCC to the destination end station.

In the middle of the picture we see what happens when only one of the communication partners is not NHRP enabled: it still needs the help of a router (same device is acting as NHRP server, or NHS) to forward its frames to the destination. This is known as one-hop routing. The NHRP client, on the other hand, still can make use of a direct VCC to forward its frames to the destination (non-NHRP client) with the least possible overhead.

IBM has also developed extensions to the NHRP specification in order to take advantage of NHRP benefits within a LANE environment. This is shown in the lower part of the mentioned picture. The LECs are non-NHRP clients and will make use of the NHRP servers (NHS) as their routers to forward frames in the most direct way as possible; that is, each NHS will establish a direct VCC to the destination in lieu of its LEC. Here again we have one-hop routing.

3.4 Performance

The end-to-end performance observed by a user at the application level can be affected by many factors. The network is only one of these. The application throughput observed by a user will typically depend on the lowest performing component used in the application path. This component is referred to as the bottleneck.

Today's legacy LAN networks were often the bottleneck affecting data transmission. However, with the advent of new switching and high-speed network technologies like ATM, other components can often become the limiting factor. To achieve optimum performance all components must be tuned for your specific network environment.

Some of the most important factors affecting end-to-end application performance are:

- Processor speed

Less powerful machines are generally incapable of delivering frames at ATM media speeds. Maximum throughput is achieved for only the most powerful machines. Handling the transmission or reception of a frame involves overhead, so typically the larger the Maximum Transmission Unit (MTU), the less processing that is required.

- Device driver, protocol, application and operating system efficiency

The adapter device driver must compete for processor cycles with the operating system, protocol and user applications. The efficiency with which the adapter driver and protocol stack use the system processor may determine the final end-to-end performance.

- Workstation internal BUS

The industry standard architecture (ISA) bus has a maximum theoretical capacity of about 64 Mbps. However, its typical practical throughput given arbitration/transmission is only about 30 Mbps. Though not a problem in traditional legacy networks in ATM networks with typical capabilities of 100 or 155 Mbps this becomes a severe limitation. When the shared bus is either Micro Channel (maximum throughput: 40-160 Mbps) or Peripheral Component Interconnect (PCI maximum throughput: 132 Mbps) the impact is reduced but still a significant problem.

- Protocols (overhead/window size/buffering/MTU)

Tuning is extremely important for the effective use of different protocols over different types of networks. The defaults for most protocols are configured for their efficient use over lower speed legacy LAN environments. In a high speed ATM environment this can place significant overheads on the use of different protocols making the tuning of these protocols essential.

Flow control and lost data recovery are the two mechanisms that have the greatest impact on measured system performance. The window size of a higher layer protocol defines the maximum amount of data that will be transmitted before an acknowledgment of receipt is required. This places, and therefore potentially limits, the amount of data transmitted. Typically the larger the window size the greater the throughput. Larger window sizes also require additional buffering of unacknowledged frames. Peak performance requires the transmitter to have buffering that is equivalent to the maximum number of packets that can be generated while the system is waiting for

acknowledgment of receipt at the other end. Buffering must also occur at the receiver to accommodate any differences between the network delivery rate and the rate at which the adapter can receive the data and then process it.

For TCP the window size is defined by the `tcp_sendspace` and the `tcp_recvspace` parameters and for UDP by the `udp_sendspace` and the `udp_recvspace` parameters. IPX usually uses a window size equal to the MTU, which can have dramatic effects on performance. The use of burst mode defines a window size of several frames and can drastically increase performance.

All protocols introduce network overheads. Each protocol has a header in its protocol data unit (PDU). The smaller the data frame the larger the amount of wasted data due to protocol overhead. ATM uses 53-byte cells for transmission of data on the network and uses a 5-byte header. Frames are sent as a number of ATM cells. The last cell is padded. Additional overhead costs are also inherent in LANE 1.0 and Classical IP.

Protocol overheads are not restricted to ATM networks. Headers also introduce overheads in all the traditional Legacy LAN networks (token-ring, Ethernet and FDDI) MAC frames and in the higher layer protocols.

The overhead for IPX is 30 bytes per PDU. Other higher layer protocol headers (for example, those associated with NetWare Core Protocol, NCP) can further increase this overhead.

The maximum transmission unit can make a significant difference to minimizing the effect of protocol overheads. The larger the data frame the lower the proportion of network bandwidth used for header traffic. Different networks support different possible values for Maximum Transmission Unit.

- Operating system overheads

The operating system used can have a significant impact on performance. Two examples of this are DOS and OS/2. DOS was originally designed for the Intel 8086, 20-bit architecture. This architecture could therefore not address more than 1 MB of real memory. Later processors were capable of addressing more memory, leading to segmentation of memory into REAL (less than 1 MB) and PROTECT mode regions. Throughput can be dramatically reduced by the device driver being forced to switch between modes.

In OS/2 and other multitasking environments, tasks must share their use of the processor. The efficiency of allocating priority to different applications can therefore significantly improve performance.

- Network capacity (available network bandwidth and throughput capabilities)

The network bandwidth available on the connection path between the communicating workstations is of considerable importance. On traditional legacy LANs (token-ring, Ethernet and FDDI) bandwidth is shared between all the devices on each network segment. Stations must compete for transmission time on the network. This means collisions, network errors or time spent waiting for a free token all cause delays and re-transmissions. Splitting the network into separate segments can have a dramatic, positive effect on performance.

In switched networks, bandwidth is dedicated to each switch port. Devices connected to these ports do not share their bandwidth with other devices. The capacity of such networks is therefore considerably higher, however

bottlenecks can still occur. The available bandwidth of each switch port is shared between all of the various applications on a device which are competing for transmitting or receiving data. The available bandwidth of any inter-switch, bridging or routing connection in the connection path between the communicating devices is shared between other workstations communicating across this link and competition for these resources might also be a bottleneck.

For more information on the factors affecting performance refer to the *Factors Influencing ATM Adapter Throughput* document mentioned in Appendix H, "Related Publications" on page 517

3.4.1 Performance Problems

To truly detect network performance problems it is important to baseline your networking environment. If you do not know how your network performs under normal working conditions you cannot evaluate whether it is performing worse than normal.

Baselining your network involves measuring the data transfer rates of typical user transactions, under the normal working load for your network. User transactions can be varied and difficult to measure; hence this will normally involve simulating the data traffic characteristics of the user transactions using simple utilities or more advanced testing programs.

Performance problems occur when a user's perceived end-to-end application performance is reduced. This is usually due to a reduction in the performance of the potential bottlenecks mentioned previously. Most commonly this is the network capacity available to the client.

In most networks, resources are shared between clients using the network. In a traditional LAN environment (token-ring, Ethernet or FDDI) the available bandwidth is shared between all clients on a network segment. In a switched environment the situation is improved but there are still shared resources.

It is difficult to predict therefore how much capacity of the shared resource one client will manage to use at any moment in time. The end-to-end performance of an application will therefore fluctuate if the bottleneck is the network. It is quite possible for one or more clients using a specific network resource to considerably affect the performance experienced by another client.

Effective network design, using small segments, high-speed switches and high bandwidth networks, like ATM, can make dramatic improvements. ATM, with its Quality of Service parameters, can also significantly improve this situation.

The key to effective network design is load balancing your network. This is the practice of distributing your network requirements evenly across your network. Make sure all your power clients are not attached to the same legacy network segment or accessing their server via the same inter-switch, bridge or router link. Ensure your ATM servers are attached to different media modules in your hub to provide the maximum bandwidth available to each. Servers typically transmit traffic for most of the time and clients typically receive traffic for most of the time. Load balancing your network can significantly increase the available bandwidth to users and reduce the chance of network bottlenecks.

3.4.2 Performance Hints and Tips in ATM Networks

Choose a network server workstation that is capable of delivering enough power to manage the network traffic. A x86-based machine is unlikely to ever saturate a 155 Mbps Adapter. An IBM RISC System/6000 has the extra processing power and BUS speed required to deliver data at media speed.

Tune your higher layer protocols for the most efficient use over ATM networks. Use larger window sizes, for example, by using the packet burst function with IPX and increasing the read and write window sizes, or the `tcp_sendspace`, `tcp_recvspace`, `udp_sendspace`, `udp_recvspace` parameters with TCP and UDP. Ensure your client workstations have sufficient send and receive buffers to handle the extra traffic queues resulting from increased window size.

RFC 1323 defines some extensions for TCP for high speed networking. If possible turn these extensions on by using the `rfc1323` parameter. This is usually beneficial when using either large datagrams or for networks with larger round-trip delays.

Possibly increase the maximum transmission unit (MTU) used to reduce processing power and increase data transmission. Standard MTU sizes are defined for LANE 1.0, Classical IP, token-ring and Ethernet. Using larger MTU sizes can increase performance but may cause interoperability problems, which means the entire MTU will be retransmitted even if only one cell in the MTU is lost. The MTU used for clients in a LANE 1.0 or Classical IP environment must match the value defined in the client's LES or ATMARP server. Increasing the MTU size does little good if your applications never generate a datagram of that size. TCP/IP tries to compensate by concatenating small datagrams to create transfers closer to the full MTU size. Use the `tcp_nodelay` parameter to tune this effect.

In AIX Classical IP performance tests the optimal values for the TCP and UDP parameters were found to be:

- `tcp_sendspace` = `tcp_recvspace` = 65,536 for an MTU size of 9180 bytes
- `tcp_sendspace` = `tcp_recvspace` = 655,360 for an MTU size of 59K bytes
- `udp_sendspace` = `udp_recvspace` = 16,384

Increases in UDP values were not found to significantly increase throughput.

- `sb_max` = 600 KB for an MTU size of 9180 bytes
- `sb_max` = 6 MB for an MTU size of 59 KB or higher
- Recommended MTU = 9180 bytes (default)

For more information refer to *The IBM Turboways 155 PCI ATM Adapter: Classical IP and LAN Emulation Performance for AIX* document mentioned in Appendix H, "Related Publications" on page 517.

Table 3 on page 89 shows the relationship between the frame or transmission size and the data and media speeds for IBM's RISC System/6000 running Classical IP and LANE 1.0. Since Ethernet frames are limited to 1500, neither emulated Ethernet nor 100BaseT (Fast Ethernet) can achieve as high throughput as emulated token-ring or Classical IP. If users must run LANE 1.0, a token-ring emulated LAN is a better performance choice than Ethernet.

<i>Table 3. Relationship between Transmission Size and Throughput</i>			
Protocol	MTU	User Data Rate	Media Rate
Classical IP	9180	133 Mbps	152 Mbps
	1500	125 Mbps	143 Mbps
Token-Ring ELAN	9200	133 Mbps	152 Mbps
	4500	129 Mbps	147 Mbps
Ethernet ELAN	1500	125 Mbps	143 Mbps
Fast Ethernet	1500	< 90 Mbps	100 Mbps

For more information on network performance, refer to the reference material mentioned in Appendix H, "Related Publications" on page 517. Many useful hints and tips for RISC System/6000s using ATM LANE 1.0 and Classical IP networks can be found in *Banking on ATM Networking - Real LAN Emulation Interoperability Scenarios*.

Chapter 4. Problem Determination Guidelines

This chapter details how to proceed with troubleshooting ATM campus networks. It provides information on the way console commands can be used to collect network status information.

4.1 Problem Source Isolation in a Networking Environment

When faced with multiple problems, generally you would want to correct the one with the greatest impact on the network first. Exceptions are when a less critical problem can be quickly and easily resolved.

4.2 Before You Begin

Effective network problem determination requires a good knowledge of the network topology involved. In that regard, the more information the troubleshooter can get, the better. Having the following information available will save much time and effort for the network troubleshooter:

- Create network layout diagrams

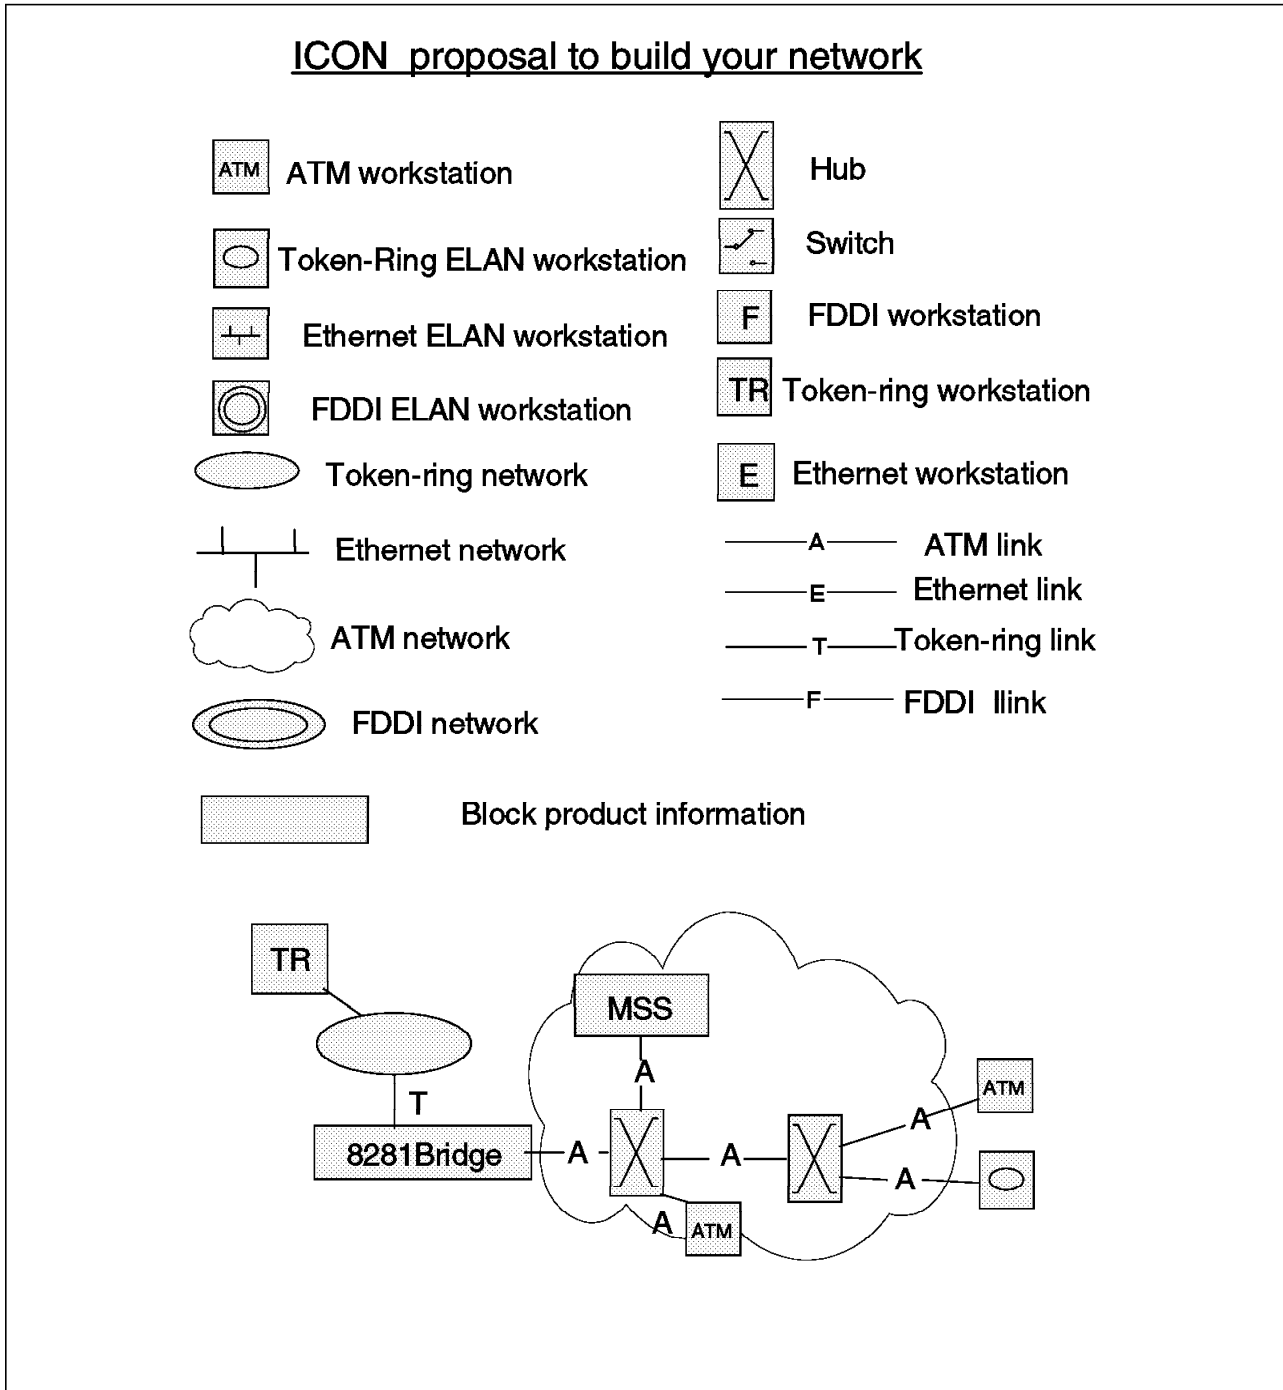


Figure 34. Example Configuration Sheet for 8260 ATM Switch

These drawings should document the logical and physical connections throughout the network in question. The physical diagrams should include detailed connectivity information for media runs from devices (workstations, servers, routers, etc.) to wall plates, patch panels and hub ports. These will be especially useful for connectivity problems since many connection problems are still the result of bad or loose media cable connections, or as is often the case with hubs, the result of disabled or misconfigured port settings.

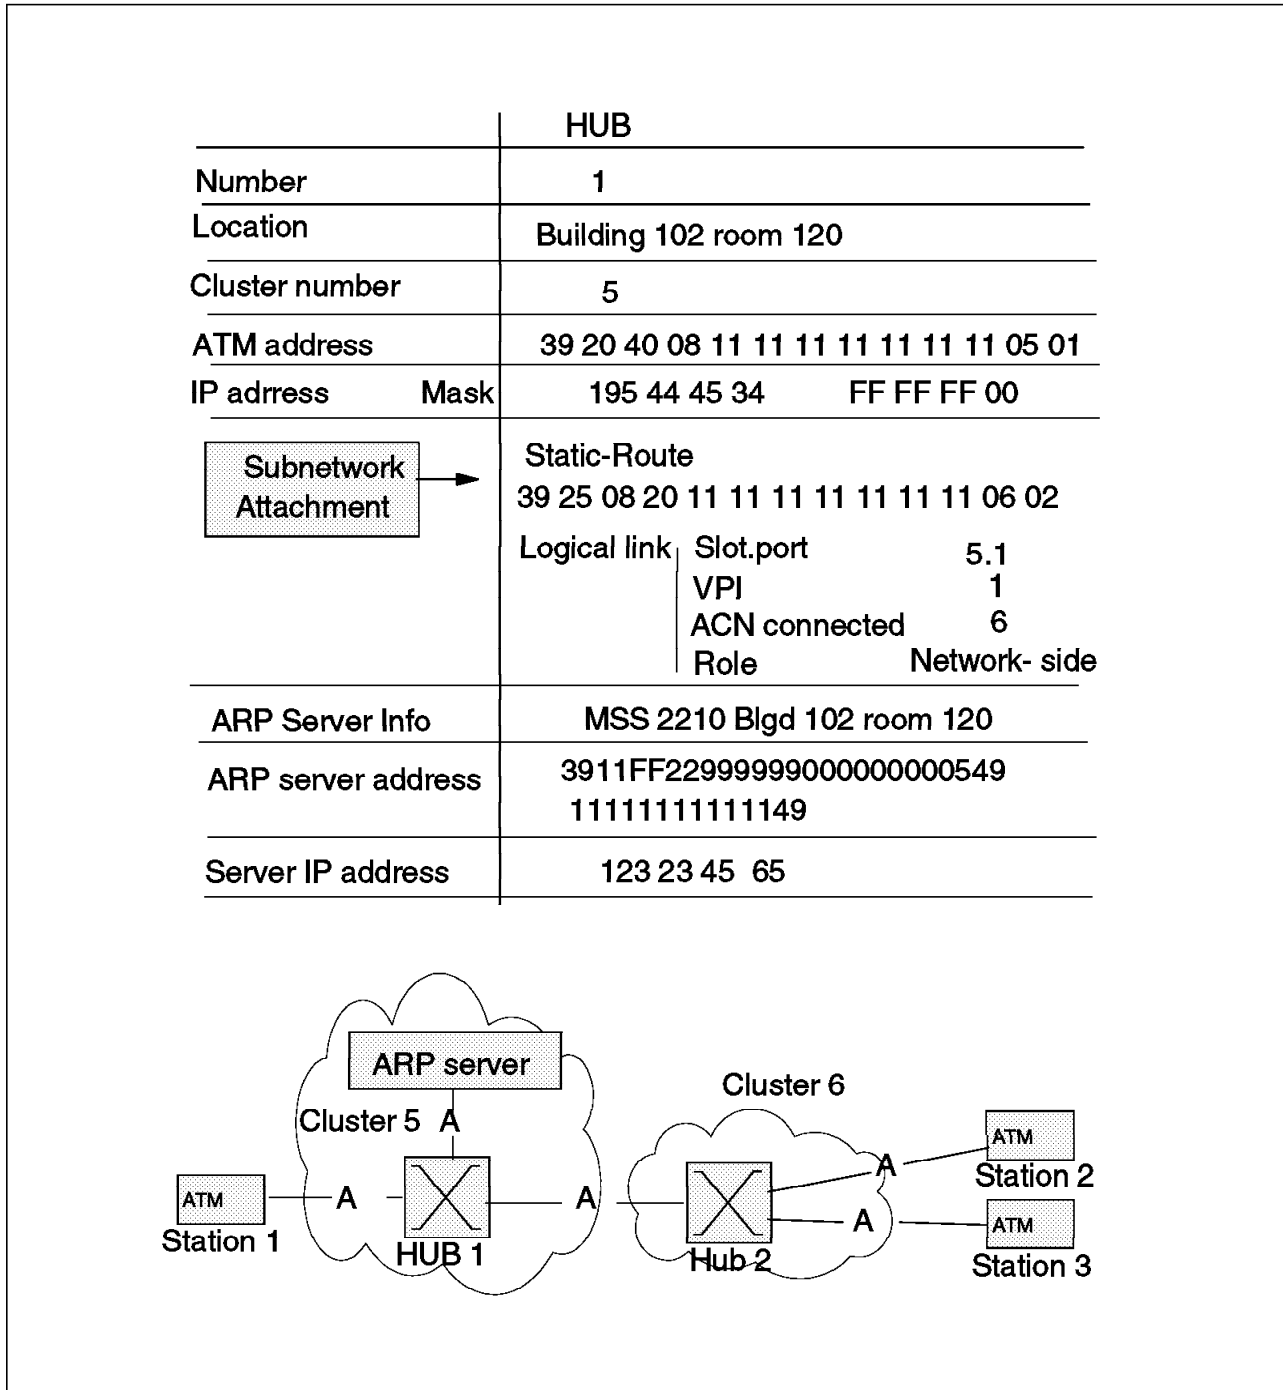


Figure 35. Example Configuration Sheet for 8260 ATM Switch

Figure 35 shows an example of a configuration sheet that can be used to record information about 8260 ATM switches, while Figure 34 on page 92 illustrates some sample icons that can be used to create network diagrams.

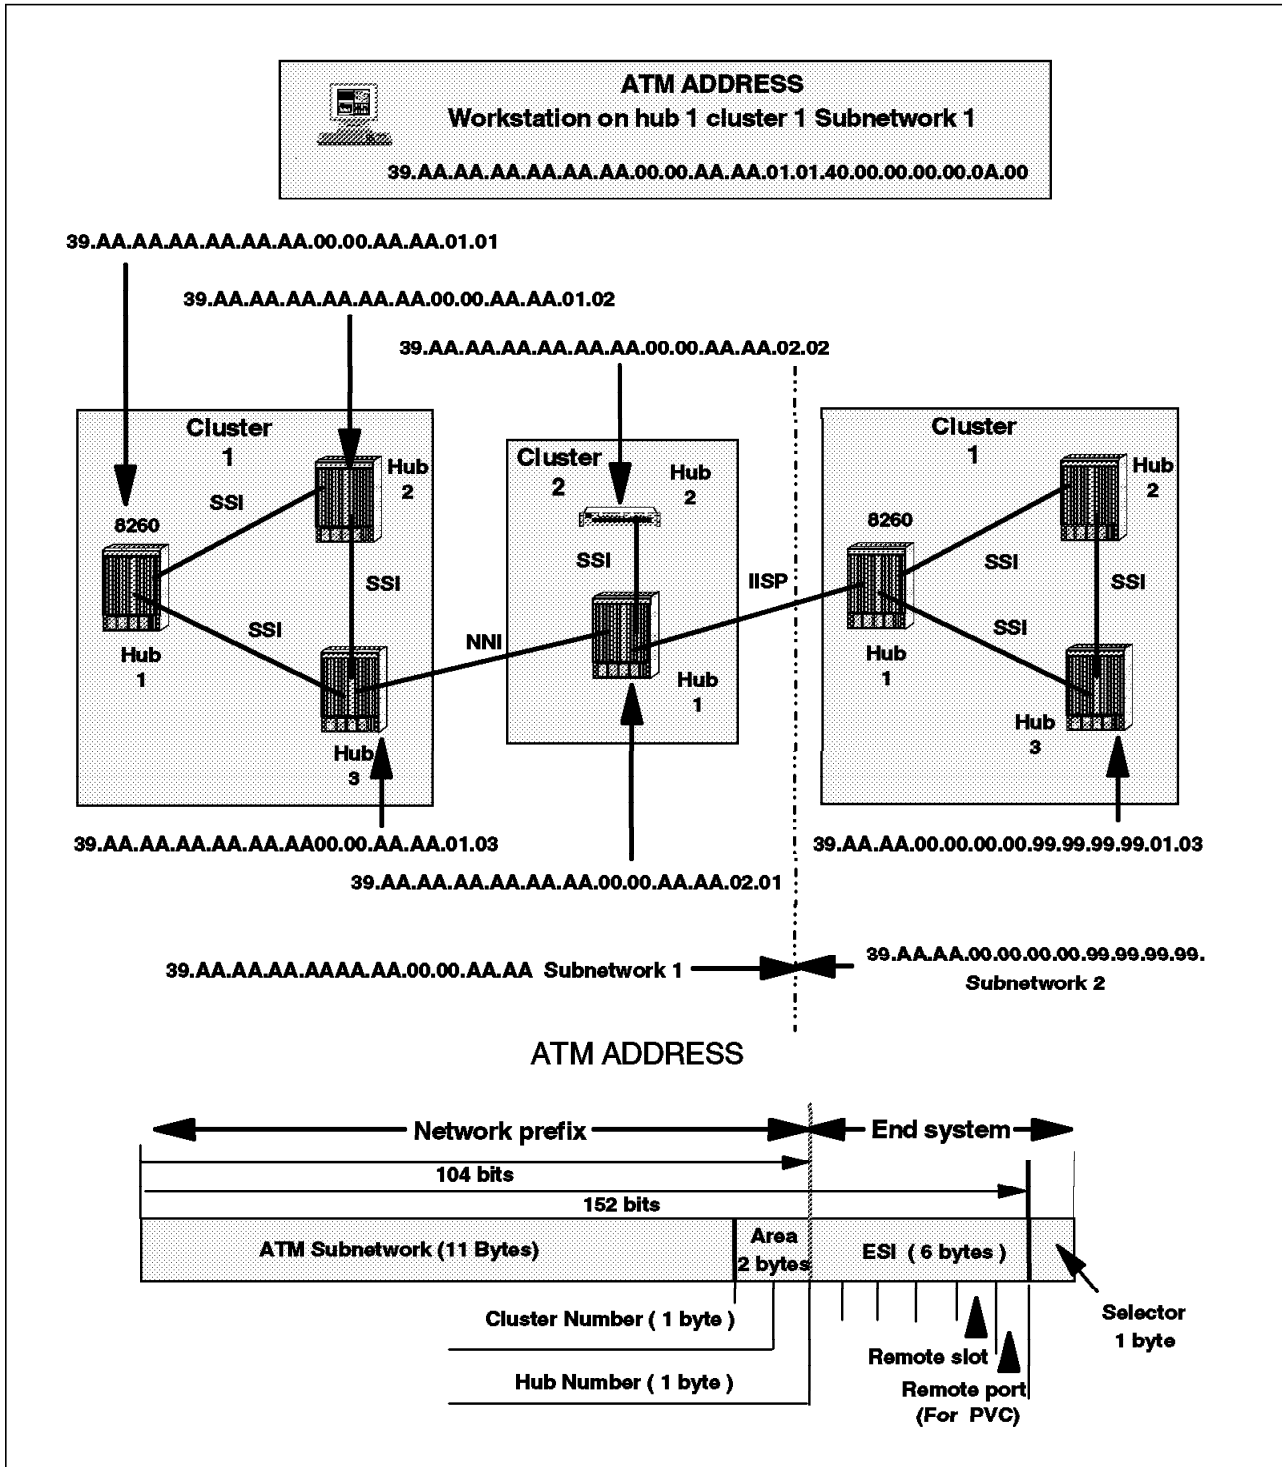


Figure 36. ATM Address Assignment in Campus Network

Figure 36 shows how ATM addresses can be set in clusters or peer groups to help in the isolation of hubs or switches.

- Collect user manuals for each component (installation, customization, and problem determination manuals).
- The network administrator should also have a knowledge of the baseline performance of the network to use as a reference for performance-related problems.

Plan your troubleshooting strategy. If the network is not totally usable, you should plan how to attack the problem when maintenance becomes available so as to make the most productive use of your time.

4.3 Define the Problem

Quick troubleshooting requires that the problem be defined as accurately as possible. Certain basic questions need to be asked and answered to help narrow the problem search area:

- What changed?

The biggest cause of network problems is changing something.

- Has it ever worked correctly?
- What is wrong?

Define what you cannot do. Put it in writing if you can; then you can stare at it and contemplate solutions.

- What is the scope of the problem?

For example, you discovered that a device cannot connect to a server. Is it confined to one device not being able to connect to a server or is it that no devices can connect to the server? Can the device connect to any server? Is the problem confined to one segment or does it span multiple segments?

- Is the problem related to a specific protocol or application?
- Is the problem intermittent?
- Is it more prevalent at a particular time of day or under load?
- Can the problem be duplicated?

4.4 Problem Determination Tools and Procedures

Tools to gather data are varied depending on the problem, but generally will include the following:

1. Interrogation of affected users
2. Visual aids (LEDs, power indications, etc.)
3. Console commands from the affected devices

Initial problem isolation activity will likely start with the ATM switch. The 8260 and 8285 have several commands that allow a network administrator to quickly and easily target potential problem areas.

4. Protocol commands such as IP PING, traceroute or IPX PING
5. ATM Forum UNI Cause Codes
6. ATM Forum LANE Status Codes
7. Internal traces
8. Protocol analyzers
9. RMON

4.4.1 LED and Other Visual Indicators

LED indicators on some products provide indications that will help quickly pinpoint a hardware problem and possibly a congested link (an activity LED always on solid). The network administrator should be familiar with the normal indications of the LEDs on his or her products and refer to the product user documentation when LED indications are abnormal.

The observant troubleshooter should also look for loose connectors on adapters, hubs, and patch panels, poorly seated adapters and power connections.

4.4.2 Console Commands

Console commands will be the most useful tool to the network troubleshooter. Switches such as the control point and switch function in the IBM 8260 gives the user a wide range of commands that will help her/him to easily determine the status of ports and links between devices and nodes. The user can then use that information to quickly formulate an action plan to correct the problem.

4.4.2.1 Checking for Hardware Problems

```

8260ATM#1> show module 13 verbose
Slot Install Connect Operation General Information
-----
 13      Y      Y      Y      8260 ATM 2 Ports LAN 155 Mbps Module

status: connected / hardware okay
        enable / Normal
P/N:51H3635 EC level:E28056 Manufacture:VIME
Operational FPGA version : B50
Backup FPGA version : B40
Type Mode      Status
-----
13.01:PNNI enabled UP
13.02:UNI enabled UP

```

Figure 37. Show Module Output (8260 Console Command)

The Show Module command (Figure 37) can be used to check the status of the module. If the module is not okay, then the port won't be either, so this command should be used to check for hardware connection problems related to the module and hub backplane. The command not only verifies whether or not there is a hardware problem on the module but also gives some status of the ports. This is also where the network administrator can learn if someone has neglected to connect the module to the ATM subsystem as mentioned on page 133. For the 8260, the good response to the Show Module command for a particular module should show

- Installed=Yes
- Connect=Yes
- Enabled=Yes
- Status=Hardware okay

Anything else could represent a potential problem. Refer to Table 10 on page 138 or the product user guide for additional detailed assistance.

4.4.2.2 Checking for End Device to Switch Connection Problems

Many things can be learned about a connection by looking at the status of the two endpoints of the physical connection. In hubs this means looking at the status of the port on which the link is connected. Some examples of the information to be found there are the actual port type (UNI, IISP, PNNI, etc.), if there is a hardware problem on the port, if the fiber or copper link is good, if address registration was done on the port, and what UNI signalling version is being used the ports. To get this information, on most switches the user can issue a console command. On the IBM 8260 and IBM 8285 that command is Show Port. Figure 38 on page 98 is the output of the Show Port command issued for port 13.2 and port 17.2.

Take a minute to look at the type of information that can be learned about a port from a console command such as this. Note that port 17.2 is a UNI port; it's status is "UP" which means that the port detects that a fiber link is installed and that the remote device connected to it is signalling okay. However, there was no address registration completed by the switch and the remote device. This could be cause for alarm if the remote device that is connected to the link was expected to do automatic address registration via ILMI. For PVC connections, address registration is not necessary. However, as we see in Figure 39 on page 98, port 17.2 is configured for a PVC connection.

Commands such as Show Port and Show PVC should be the first to be used to troubleshoot connection problems from a hub.

```

8260ATM#1> show port 13.2 verbose
      Type Mode      Status
-----
13.02:UNI enabled UP
Signalling Version : Auto
> Oper Sig. Version : 3.1
ILMI status        : UP
ILMI vci           : 0.16
RB Bandwidth       : unlimited
Signalling vci     : 0.5
Administrative weight: 5040
VPI.VCI range      : 15.1023 (4.10 bits)
Connector          : SC DUPLEX
Media              : multimode fiber
Port speed         : 155000 kbps
Remote device is active
Frame format       : SONET STS-3c
Scrambling mode    : frame and cell
Clock mode         : internal

8260ATM#1> show port 17.2 verbose
      Type Mode      Status
-----
17.02:UNI enabled UP:No Address Registration
Signalling Version : Auto
> Oper Sig. Version : 3.1
ILMI status        : UP:No Address Registration
ILMI vci           : 0.16
RB Bandwidth       : unlimited
Signalling vci     : 0.5
Administrative weight: 5040
VPI.VCI range      : 15.1023 (4.10 bits)
Connector          : SC DUPLEX
Media              : multimode fiber
Port speed         : 155000 kbps
Remote device is active
Frame format       : SONET STS-3c
Scrambling mode    : frame and cell
Clock mode         : internal

```

Figure 38. Show Port Output (8260 Console Command)

```

8260ATM#1> show pvc all
-----
PVC:Port 13.02 (id=69,Primary,BE) PTP-PVC VP/VC=0/82
-> Party:(id=0) VP/VC=0/74 STATUS:Active

   39.99.99.99.99.99.00.00.99.99.01.02.42.00.00.00.11.02.00(port 17.02)
-----

PVC:Port 17.02 (id=1001,Secondary,BE) PTP-PVC VP/VC=0/74
-> Party:(id=0) VP/VC=0/82 STATUS:Active

   39.99.99.99.99.99.00.00.99.99.01.02.42.00.00.00.0D.02.00(port 13.02)

```

Figure 39. Show PVC Output (8260 Console Command)

4.4.2.3 Checking for Address Problems

Three of the initially important things to know about are:

1. Is the address of the end device registered with its switch?
2. What is the address of the switch?
3. What is the address of the ATM ARP server (RFC1577) or LES or LECS (LANE)?

To check the ATM addresses of the switch, use a command such as the 8260 Show Device command whose output is shown in Figure 40 on page 100. This command gives the address of the ATM ARP server the switch uses as well as information about how well the clients did connecting to RFC1577 and LANE VLANs.

All ATM switches also have a way to determine what end system addresses are registered. The Show Reachable_Address command (Figure 41 on page 101) shows representative output that can be used to verify end system registration with the local switch. Note that the output gives the port number **1** on the left and then the address(es) that registered on that port **2**, and whether those addresses are active **3**. For proper call setup via SVCs, the addresses of end devices need to be registered with the switch. This can be done either dynamically by ILMI, or it can be entered by the network administrator.

```
show device
8260 ATM Control Point and Switch Module
Name : 8260ATM
Location :
  12 Jun 97
Manufacture id: VIME
Part Number: 51H3659 EC Level: E28028
Boot EEPROM version: v.3.0.0
Flash EEPROM version: v.3.0.0
Flash EEPROM backup version: v.2.5.1
Last Restart : 19:06:05 Fri 30 May 97 (Restart Count: 9)
A-CPSW
-----
> Subnet atm: Up
IP address: 192.2.1.8. Subnet mask: FF.FF.FF.00
> Subnet lan emulation ethernet/802.3
  Not Started
  Name : ""
  MAC Address: 000000000000
  IP address : 0.0.0.0. Subnet mask: 00.00.00.00
  ATM address :39.99.99.99.99.99.99.00.00.99.99.01.02.40.00.82.60.00.07.00
  Config LES addr:none
  Actual LES addr:00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00
  BUS ATM address:00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00
  Config LECS add:none
  Actual LECS add:00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00
  LEC Identifier: 0. Maximum Transmission Unit: 0
> Subnet lan emulation token ring
  Not Started
  Name : ""
  MAC Address: 000000000000
  IP address : 0.0.0.0. Subnet mask: 00.00.00.00
  ATM address :39.99.99.99.99.99.99.00.00.99.99.01.02.40.00.82.60.00.07.01
  Config LES addr:none
  Actual LES addr:00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00
  BUS ATM address:00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00
  Config LECS add:none
  Actual LECS add:00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00
  LEC Identifier: 0. Maximum Transmission Unit: 0
Default Gateway : OK
  IP address 192.2.1.69
ARP Server:
-----
  ATM address: 39.99.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.20.69.20

Device configured for PNNI port capability. No LES can start.
Dynamic RAM size is 16 MB. Migration: off. Diagnostics: enabled.
Device defined as primary.
Duplicate ATM addresses are allowed.
```

Figure 40. Show Device Output (8260 Console Command)

```
8260ATM#1> show reachable_address
Enter module: all
Port Len Address Active Idx VPI
1 2 3 -----
13.02 152 39.99.99.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.82.10 Y 3 -
14.04 96 39.99.99.99.99.99.99.99.00.00.99.99.70. . . . . N 2 -
17.01 96 39.99.99.99.99.99.99.99.00.00.99.99.50. . . . . Y 1 -
3.01 152 39.99.99.99.99.99.99.99.00.00.99.99.01.02.02.00.82.71.00.01 Y Dyn 0
3.02 152 39.99.99.99.99.99.99.99.00.00.99.99.01.02.08.00.5A.99.86.DC Y Dyn 0
4.01 152 39.99.99.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.00.69 Y Dyn 0
4.01 152 39.99.99.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.20.69 Y Dyn 0
4.01 152 47.00.79.00.00.00.00.00.00.00.00.00.00.00.00.00.A0.3E.00.00.01 Y Dyn 0
6.02 152 39.99.99.99.99.99.99.99.00.00.99.99.01.02.40.00.82.81.00.66 Y Dyn 0
6.03 152 39.99.99.99.99.99.99.99.00.00.99.99.01.02.40.00.05.70.00.51 Y Dyn 0
13.02 152 39.99.99.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.00.72 Y Dyn 0
13.02 152 39.99.99.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.01.72 Y Dyn 0
13.02 152 39.99.99.99.99.99.99.99.00.00.99.99.01.02.40.82.10.14.83.00 Y Dyn 0
13.02 152 47.00.79.00.00.00.00.00.00.00.00.00.00.00.00.00.A0.3E.00.00.01 Y Dyn 0
```

Figure 41. Show Reachable_Address Output (8260 Console Command)

```

8260ATM#1> show signalling cross_connections port 13.2

In: slot.port vpi.vci  type  Out: slot.port vpi.vci  type      Conn  Cat
-----
      13.2      0.75   SVC      13.1      0.76   SVC      P2P   UBR
      13.2      0.76   SVC      13.1      0.75   SVC      P2P   UBR
      13.2      0.77   SVC      4.1       0.835  SVC      P2P   UBR
      13.2      0.78   SVC      13.1      0.79   SVC      P2P   UBR
      13.2      0.79   SVC      13.1      0.78   SVC      P2P   UBR
      13.2      0.80   SVC      13.1      0.81   SVC      P2P   UBR
      13.2      0.81   SVC      13.1      0.80   SVC      P2P   UBR
      13.2      0.82   PVC      17.2      0.74   PVC      P2P   UBR
      13.2      0.83   SVC      3.1       0.92   SVC      P2M   UBR
      13.2      0.83   SVC      4.1       0.851  SVC      P2M   UBR
      13.2      0.83   SVC      13.1      0.84   SVC      P2M   UBR
      13.2      0.84   SVC      13.1      0.83   SVC      P2M   UBR
      13.2      0.85   SVC      4.1       0.845  SVC      P2M   UBR
      13.2      0.117  SVC      3.1       0.95   SVC      P2P   UBR

Total number of cross connections = 14

8260ATM#1> show signalling cross_connections port 17.2

In: slot.port vpi.vci  type  Out: slot.port vpi.vci  type      Conn  Cat
-----
      17.2      0.74   PVC      13.2      0.82   PVC      P2P   UBR

Total number of cross connections = 1

8260ATM#1> show signalling atm_interface port 17.2

Interface Type: privateUNI
Sig Version   : UNI 3.1

Sig Side      : Network
tg Protocol   : N/A
Sscop State   : Idle
Max Vpi Bits  : 4
Max Vci Bits  : 10
Active Vps    : 0
Active Vcs    : 1
Nb Connexions : 0

```

Figure 42. Show Signalling Command Output (8260 Console Command)

Figure 39 on page 98 and Figure 42 are examples of ways to get more specialized information about a port. The Show Signalling command gives more information about the various connections that are going through a particular port (Figure 42). As the output shows, the user can learn the VP/VC associated with each connection going through a particular port as well as the number and type of connections established. Similar information can be gathered about PVC links by using the Show PVC command (Figure 39 on page 98).

4.4.3 ATM Forum UNI Cause Codes and LANE Status Codes

When a call setup attempt fails many times there is either a UNI cause code or LANE status code that is returned that will indicate the cause of the problem. Depending on the device, the code can be retrieved by console commands. The 8260 Show Device command is an example of a command that can be used to show the ATM Forum UNI cause code or the LANE status code if one of the clients is not able communicate with its respective VLAN. Some of the problems related to cause codes can be found in Table 11 on page 143.

Refer to Table 33 on page 471 for a complete listing of the UNI cause codes and to Table 43 on page 491 for a listing of the LANE status codes. For more information and assistance on LANE or Classic IP problems, refer to Chapter 7, "ATM Emulated LANs and Logical IP Subnets (LANE 1.0, RFC 1577)" on page 293.

4.4.4 Internal Traces

Most devices have internal traces and event logs that can be activated to help isolate a problem. Operation and analysis of these tools usually require more specialized knowledge than the console commands do, but for the problems that require their use, these tools can be invaluable. These traces are particularly useful for address registration and ILMI failures. Internal traces are also very good for isolating LANE problems regarding setting up the proper VCCs to allow clients to join their ELAN and work properly. Refer to the particular product documentation to learn more about the use of its traces and logs.

4.4.5 Protocol Analyzers

ATM protocol analyzers are very expensive so they are not widely available. Fortunately, most problems can be resolved without them. Protocol analyzers will likely be the tool of choice when trying to solve performance related problems. A problem such as intermittent cell loss over a public link would be very difficult to prove without an external trace tool. Use of a protocol analyzer and analysis of its output requires specialized skills but if he or she can afford it, it would be very advantageous for the network administrator to learn how to use one. Two things that should be kept in mind when planning to use a protocol analyzer:

1. Since ATM is point to point, most analyzers need to be serially connected into the link they need to monitor, which of course, requires the troubleshooter to schedule time to take the link down to install the analyzer. Fiber optical splitter cables are available that help eliminate this restriction. The splitter cables are designed so they can be inserted at certain trace points in their network and left there so that the analyzer can later be "inserted" into the network non-disruptively. These splitter cables take a small percentage of light from the main fiber path to the analyzer so this signal loss needs to be included in any budget calculations.
2. Typically, the LES and BUS are implemented in the same hardware. Therefore, protocol analyzers will not be able to trace VCC activity between the LES and BUS.

4.4.6 RMON

RMON is not available for ATM but it can be very useful to the network administrator for isolating problems on token-ring, Ethernet and FDDI. RMON in conjunction with an SNMP manager station can assist with baselining network performance as well as detecting potential problems before they become critical.

4.4.7 PING

The IP PING command is just one of the protocol commands that can be issued from one station on the network to another to determine connectivity. Others are IPX PING, AIX traceroute, or SNMPget. For some other protocols verifying connectivity could be as simple as using a workstation to try to connect to its server.

These commands are easy to use and can be used to help define the scope of the connection problem and to subsequently verify a fix.

4.5 Connectivity Problems in Intelligent Hubs

When diagnosing connection problems, we should first verify the connection between any failing device and its entry point into the network. Since hubs usually have status LEDs and sometimes activity LEDs for each port, a quick visual check often can be useful to determine if there is media connection between the connecting device and the hub port. If so, it is then a simple matter of troubleshooting the physical connection. The troubleshooter can then refer to the hub and device user documents for more information to assist in this problem determination process.

Another thing to consider about connectivity issues involving intelligent hubs is that many of them have the capability to have multiple segments on the same backplane. These segments can be of the same protocol or of different protocols and they give the user the capability and responsibility to assign media modules or ports to these backplane segments as they see fit to develop the network topology desired for their organization. For proper connectivity it is required then that ports and modules be configured properly (for example, token-ring speed) and assigned to the appropriate backplane segments. In Figure 43 on page 105 port 16.01 is assigned to network Ethernet_4 and 16.02 and is isolated (not assigned to a network) while ports 16.03 and 16.04 are assigned to network Ethernet_2. The assignments are fine as long as port 16.01 was not mistakenly assigned to Ethernet_4. However, as they are currently configured, a device attached to port 16.01 will not be able to communicate to the device on port 16.03 without the benefit of bridging or routing. A device on port 16.02 would not be able to communicate with devices on ports 16.01 and 16.03 because port 16.02 is not assigned to a network.

```
8260> show port 16.all

Port Display for Module 1 E24PS-6/8 :

Port  Mode      Status           Network
-----
16.01  ENABLED  OKAY             ETHERNET_4
16.02  ENABLED  OKAY             ISOLATED_1
16.03  ENABLED  OKAY             ETHERNET_2
16.04  ENABLED  LINK FAILURE    ETHERNET_2
16.05  ENABLED  OKAY             ETHERNET_2
16.06  ENABLED  LINK FAILURE    ETHERNET_2
16.07  ENABLED  LINK FAILURE    ETHERNET_2
16.08  ENABLED  LINK FAILURE    ETHERNET_2
16.09  ENABLED  LINK FAILURE    ETHERNET_2
16.10  ENABLED  LINK FAILURE    ETHERNET_2
16.11  ENABLED  LINK FAILURE    ETHERNET_2
16.12  ENABLED  LINK FAILURE    ETHERNET_2
16.13  ENABLED  LINK FAILURE    ETHERNET_2
16.14  ENABLED  LINK FAILURE    ETHERNET_2
16.15  ENABLED  LINK FAILURE    ETHERNET_2
16.16  ENABLED  LINK FAILURE    ETHERNET_2
16.17  ENABLED  LINK FAILURE    ETHERNET_2
16.18  ENABLED  LINK FAILURE    ETHERNET_2
16.19  ENABLED  LINK FAILURE    ETHERNET_2
16.20  ENABLED  LINK FAILURE    ETHERNET_2
16.21  ENABLED  LINK FAILURE    ETHERNET_2
16.22  ENABLED  LINK FAILURE    ETHERNET_2
16.23  ENABLED  LINK FAILURE    ETHERNET_2
16.24  ENABLED  LINK FAILURE    ETHERNET_2
```

Figure 43. 8260 Ports Assigned to Different Networks

4.5.1 User Interface

The user interface to the hub can be used to verify that these settings are correct. Access to the hub user interface can usually be gained by direct attachment through a serial console management port, remotely by a dial-up modem connected to the serial port, or by an inband TELNET session. Figure 44 on page 106 shows typical settings for the serial management console of the IBM 8260.

```
8260> show terminal

Terminal Session Parameters:
  Prompt:      8260>
  Timeout time: 0

Console Port Parameters:
  Baud:        19200
  Data bits:   8
  Parity:      NONE
  Stop bits:   2
  Hangup:     DISABLED
  Mode:        COMMAND LINE
  Terminal:    VT100

Auxiliary Port Parameters:
  Baud:        9600
  Data bits:   8
  Parity:      NONE
  Stop bits:   2
  Hangup:     DISABLED
  Mode:        COMMAND LINE
  Terminal:    VT100
```

Figure 44. Example of Typical Console Port Settings

4.6 Problem Determination Flowchart

After gathering as much information about the problem, many times the solution becomes obvious. For those problems where that is not the case, the reader should go to Chapter 5, "Starting Problem Isolation in an ATM Network" on page 109.

Figure 45 is a short flowchart that should assist the reader trying to use this book.

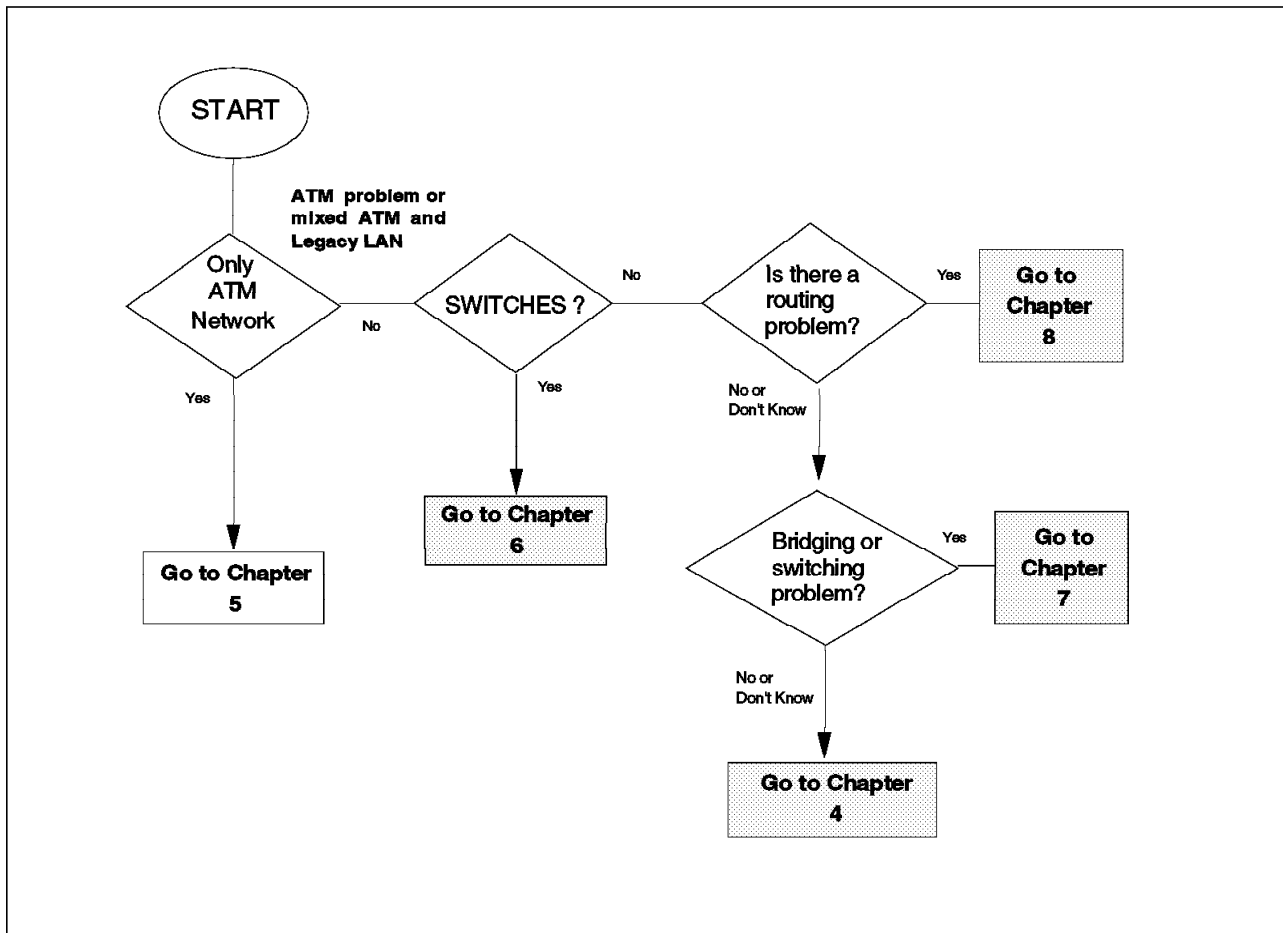


Figure 45. Flowchart for Problem Isolation

Chapter 5. Starting Problem Isolation in an ATM Network

This chapter provides an approach to isolate and solve problems that may be encountered when trying to maintain a good working ATM network. Refer to Chapter 3, "Main Concepts of ATM (Asynchronous Transfer Mode)" on page 43 for an introduction to ATM concepts.

5.1 Introduction

The approach to problem determination this redbook will use is the one outlined in Chapter 4, "Problem Determination Guidelines" on page 91. Hopefully, the initial investigation was enough to determine that the problem may be associated with your ATM network and that you were referred to this chapter because of that.

As Figure 46 on page 110 shows, this chapter primarily tries to help with isolating problems that involve the ATM components of the network connections. That is, those parts of the network that involve establishing UNI or PNNI links or completing address registration and that do not specifically apply to LANE, Classic IP, bridging, switching or routing in the ATM network. After the basic connectivity part of the network has been verified to be okay, the reader is referred to other chapters for additional assistance, depending on what the problem requires.

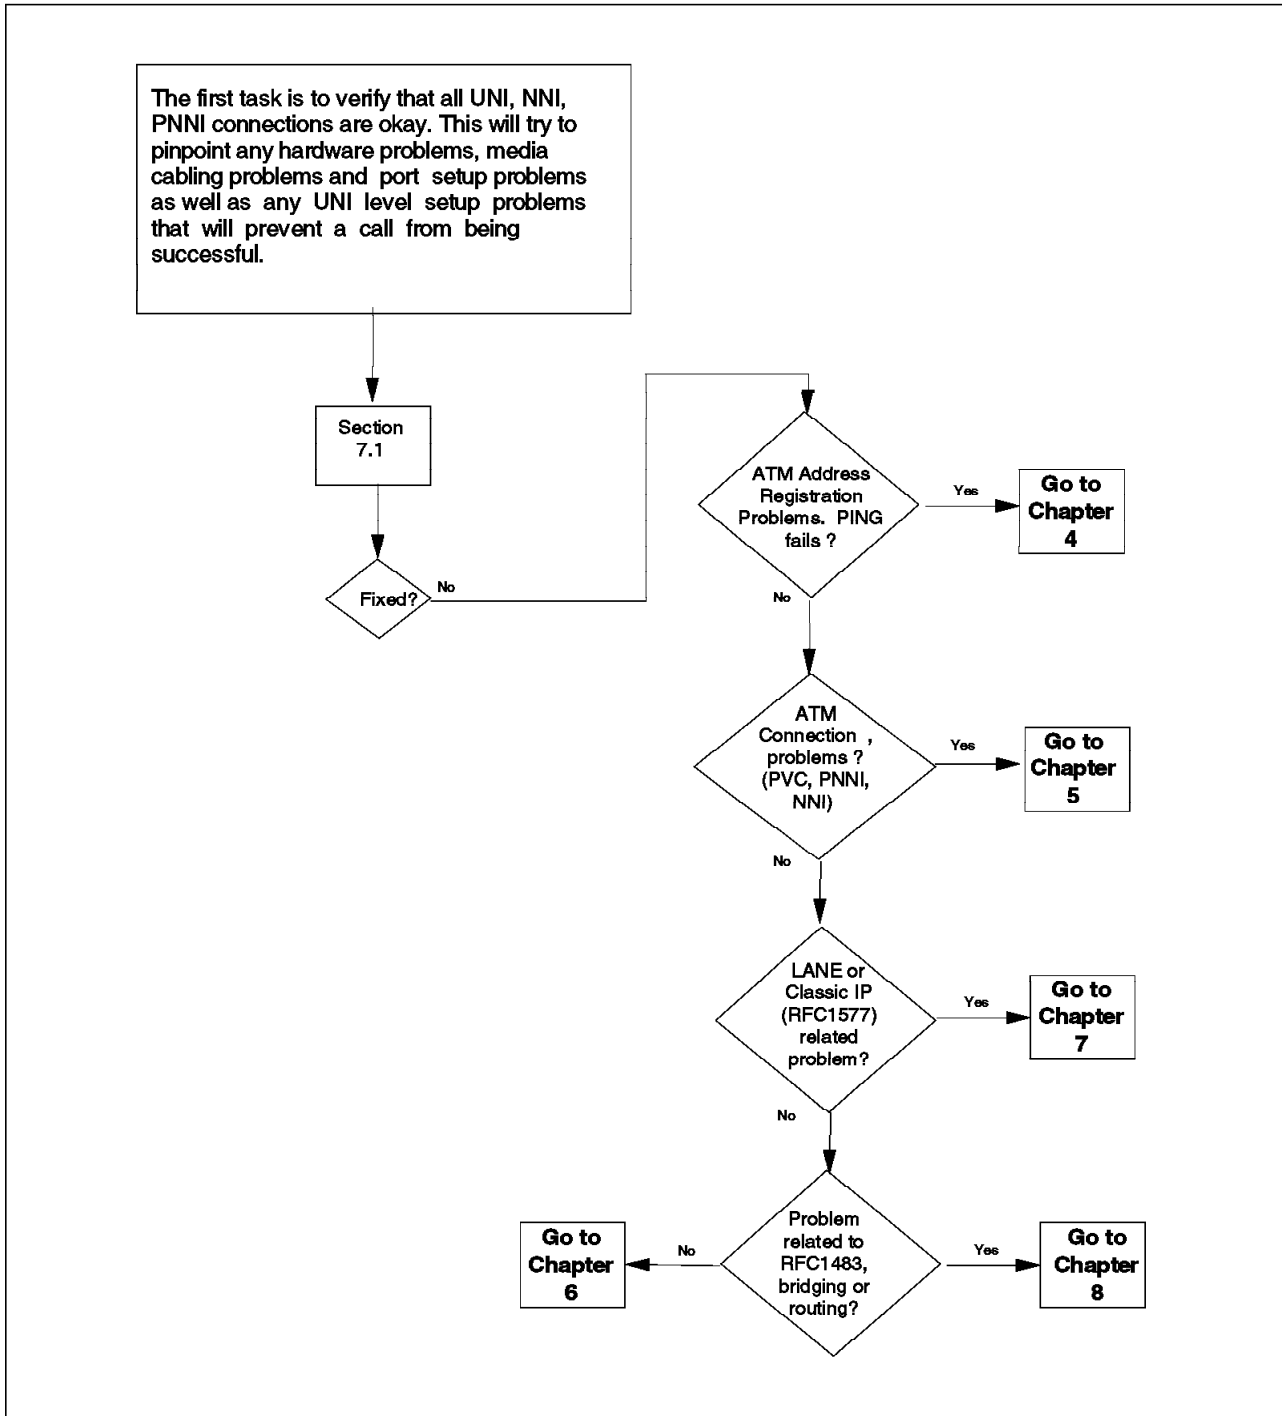


Figure 46. Problem Isolation Flowchart

5.2 What Is a Healthy ATM Network?

A healthy ATM network is one that can provide dynamic and transparent connectivity between wide area and local area domains as well integrate private and public networks. The network should be able to support needs for high transmissions and large capacities with the flexibility of bandwidth on demand, end-to-end quality of service and effective management capabilities.

5.3 Rules of ATM

Table 9 on page 114 lists some of the physical layer interfaces defined by the ATM Forum each of which has its own distance limitations. Most other limitations are specific to the particular manufacturer. For example, the maximum number of connections possible per node or the maximum number of registered addresses per switch node are examples of functionality that differ from node to node depending on the manufacturer of the product and even depending on what level of that manufacturer's microcode may be installed. Functionality built into ATM products is changing rapidly, and as vendor product offerings mature and improve, these choices will be broadened.

Refer to Table 4, Table 5, Table 6 on page 112, Table 7 on page 112, and Table 8 on page 112. These show a representative list of some the network rules given for the IBM 8260 as of this writing. The reader is cautioned to refer to product specifications and release notes for more current information.

<i>Table 4. 8260 Network Rules</i>	
Maximum Number of Connections	
The following maximum number of connections apply to Switched Virtual Circuits (SVC) and Permanent Virtual Circuits (PVC).	
Maximum number of virtual connections per switch	6000
Maximum number of PtP plus PtM connections per port	4064
Maximum number of PtP plus PtM connections per module	4064
Maximum number of parties over point-to-multipoint connections	4000
Maximum number of point-to-multipoint	127
Maximum number of PVC, PVP and PARTY that can be defined per switch PVC definitions are stored in non-volatile RAM for automatic restart.	100
Maximum number of ILMl registered addresses per switch	512
Connection identifiers (VPI:VCI) for SVCs	VP=0 is selected

<i>Table 5. 8260 Switch-to-Switch Interface Rules</i>	
The Switch-to-Switch Interface (SSI) is an IBM proprietary interface for connecting switches in a cluster. A cluster is defined as a group of switches connected by SSI links and whose first 12 bytes of the ATM address are the same.	
Bandwidth limitation	The bandwidth you specify, or which is taken by default, must be identical at both ends of the SSI link.
The maximum bandwidth budget of the SSI ports defined on an ATM media module	212 Mbps Parallel SSI links can be connected between hubs to give more bandwidth if necessary.
ATM address	The ATM addresses for all switches connected by SSI links must have the same first 12 bytes of their ATM address.

<i>Table 6. IP Over ATM (RFC 1577) Client</i>	
The A-CPSW supports an IP client implementation to be managed over ATM (SNMP, Telnet, TFTP, PING).	
MTU size support for IP client	944 bytes
Maximum connections for IP client	64

<i>Table 7. 8260 Network-to-Network Rules</i>	
The Network-to-Network interface (NNI) defines the interface between two nodes to different ATM networks. NNI links are supported both over physical links and virtual path connections (VP tunneling). Parallel NNI links can be enabled between two clusters.	
Logical links per NNI port	64
Maximum logical links for the 12-port 25 Mbps	16
Maximum bandwidth reserved for NNI per module	180 MBps Maximum reserved bandwidth is limited to 85% of the full bandwidth of the media module per port.
The maximum number of static route to ATM Cluster Number (ACN) associations that can be defined	64
The maximum number of logical links that can be defined	64
Connecting hubs in different clusters (same first 11 bytes of ATM address)	Set LOGICAL LINK is required. One side of the link is defined as network side and the other as user side.
Connecting hubs to different subnetworks (different first 11 bytes of ATM address)	Set LOGICAL LINK is required and Set STATIC_ROUTE is required. One side of the link is defined as network side and the other as user side.
Connecting NNI link to PNNI peer group.	The link at the PNNI switch is defined as IISP. One side is of the link is defined as network side and the other as user side.

<i>Table 8 (Page 1 of 2). 8260 PNNI Network Rules</i>	
An ATM Peer Group Intraconnection (PNNI) An ATM Peer Group is a group of ATM hubs interconnected by Private Network-to-Network Interfaces (PNNI). The PNNI protocol supports networking functions such as routing, node failure recovery, backup and topology management.	
Bandwidth	<ol style="list-style-type: none"> 1. The bandwidth you specify or taken by default must be identical at both ends of the PNNI link. 2. The bandwidth budget of the PNNI ports defined on a ATM media module must not exceed 212 Mbps.
Switches per peer group	Depending on network topology and complexity, the peer group can have up to 100 nodes.
Maximum number of PNNI physical links and VPCs per hub	32
An ATM Peer Group Intraconnection (IISP) Interim Interswitch Signalling (IISP) defines the interface between two ATM hubs belonging to different ATM clusters in the same subnetwork or in different networks. IISP links are supported both over physical links and virtual path connections (VP tunneling).	

<i>Table 8 (Page 2 of 2). 8260 PNNI Network Rules</i>	
Bandwidth	The total bandwidth reserved for IISP links is limited to 85% of the full bandwidth of the media port. The limitation is 180 Mbps per module.
Maximum number of reachable addresses that can be defined per 8260 hub.	64
Defining PVCs	PVCs are defined over IISP links by defining a PVC on each switch involved in the connection.

5.3.1 ATM Forum Physical Interface References

These documents can be obtained from the ATM Forum FTP site:
<ftp.atmforum.com/pub/approved-specs>

<i>Table 9. ATM Forum Physical Interfaces</i>	
Interface	ATM Forum Document File
44.736 DS3 Mbps Physical Layer	af-uni-0010.002
100 Mbps Multimode Fiber	
Interface Physical Layer 155.52 Mbps SONET STS-3c Physical Layer	
155.52 Mbps Physical Layer	
ATM Physical Medium Dependent Interface Specification for 155 Mbps over Twisted Pair Cable	af-phy-0015.000
DS1 Physical Layer Specification	af-phy-0016.000
Utopia Level 1 v2.01	af-phy-0017.000
Mid-range Physical Layer Specification for Category 3 UTP	af-phy-0018.000
6,312 Kbps UNI Specification	af-phy-0029.000
E3 UNI	af-phy-0034.000
Utopia Level 2 v1.0	af-phy-0039.000
Physical Interface Specification for 25.6 Mbps over Twisted Pair	af-phy-0040.000
622.08 Mbps Physical Layer	af-phy-0046.000
155.52 Mbps Physical Layer Specification for Category 3 UTP (See also UNI 3.1, af-uni-0010.002)	af-phy-0047.000
120 Ohm Addendum to ATM PMD Interface Spec for 155 Mbps over TP	af-phy-0053.000
DS3 Physical Layer Interface Spec	af-phy-0054.000
155 Mbps over MMF Short Wave Length Lasers, Addendum to UNI 3.1	af-phy-0062.000
WIRE (PMD to TC layers)	af-phy-0063.000
E-1	af-phy-0064.000

5.4 Basic ATM Signalling

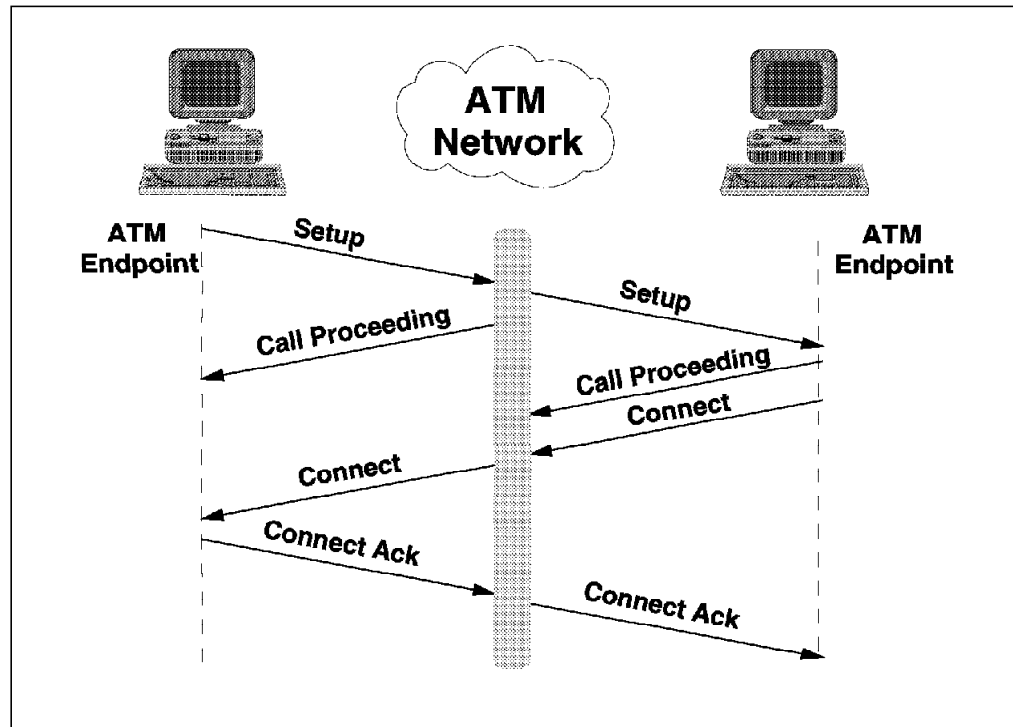


Figure 47. ATM Call Establishment

For an endstation to communicate in a switched environment such as ATM, it must register with the network, request a connection when necessary, and clear the connection when through. For native ATM endstations, this is done by the following:

5.4.1.1 Initial Registration

When an endstation wishes to enter the network, it must first register its full ATM address with its associated switch. This signalling process is described in ATM UNI Specification 3.0 (based on ITU-T Q.93B recommendations), or more recently, in ATM UNI Specification 3.1 (based on ITU-T Q.2931 recommendations) and is performed when the endstation is activated. During this process, the workstation receives its 13-byte network prefix from the switch, appends its own local address (ESI plus selector), and registers its complete ATM address with the switch.

5.4.1.2 Connection Setup

When an endstation wishes to communicate with another endstation, it must first establish a connection to it. It does this by issuing a SETUP request to the ATM network.

If the requested address is local, the switch acknowledges the request by issuing a CALL PROCEEDING response to the requesting endstation and forwarding the SETUP request to the requested endstation, which acknowledges receipt with a CALL PROCEEDING response.

If the requested endstation is not local, the switch will forward the request to the correct switch based on routing information compiled and maintained by the 8285 ATM Control Point's Topology and Routing Services (TRS) subsystem. The path will be selected based on the *widest* path (not the *shortest*) available between the end points. This path information is appended to the setup request and is used by intermediate switches to determine the next hop through the network. There can be no more than 15 hops in any given path.

If the requested workstation is able to accept the incoming connection, it issues a CONNECT response to the network, which forwards it back to the requesting workstation, where it is acknowledged by issuing a CONNECT ACK response to the network which forwards it to the destination endstation to complete the call setup process.

5.4.1.3 Connection Tear-Down

When an endstation wishes to end a connection, it issues a DISCONNECT request to the network. The network acknowledges the request by returning a RELEASE response (instructing the requesting endstation to drop all resources associated with the call), and by forwarding the DISCONNECT on to the destination workstation, which acknowledges the request by returning a RELEASE command to the network. The process is completed when the requesting endstation returns a RELEASE COMPLETE to the network, which forwards it to the destination endstation, indicating that the call has been dropped and the associated resources freed.

5.4.2 Requirements to Establish an ATM Call

For an ATM connection to be established it must satisfy the following general criteria determined by the network and end systems:

- Basic ATM service support
- VC availability
- Physical and virtual network resource availability to provide the quality of service requested
- End system resource availability to provide the quality of service requested
- End-to-end compatibility

In order to establish an ATM connection at the UNI, both the user and the network must know the ATM address(es) which are in effect at the UNI. The calling party must then be able to get the ATM address of the destination device and proceed to set up the call. The components involved and the processes they require to resolve addresses of remote endpoints and subsequently establish the connections differ depending on whether LANE or Classical IP is used, but common requirements remain. All required components, that is, the ARP server for IP over ATM or LES/BUS, and the optional LECS for LANE, must be operational and the physical and logical links between them must be good. This means all UNI connections and interconnecting NNI or PNNI links must be verified to be up and configured correctly. Referring to Figure 48 on page 117, we see three devices connected to hub ports depicting a simple Classical IP network. These devices are communicating together via Classical IP using the ARP server (shown connected to port 6.1). Physical connection from each device to the hub port must be good. Each device (including the ARP server) needs to have its address properly registered with the switch if SVCs are going to be used.




Figure 48. Classical IP Network

5.4.3 Address Registration Failures

When two or more stations cannot communicate, it is generally due to an address registration problem. Following is an example of a registration failure that shows an approach isolating and solving such a problem.

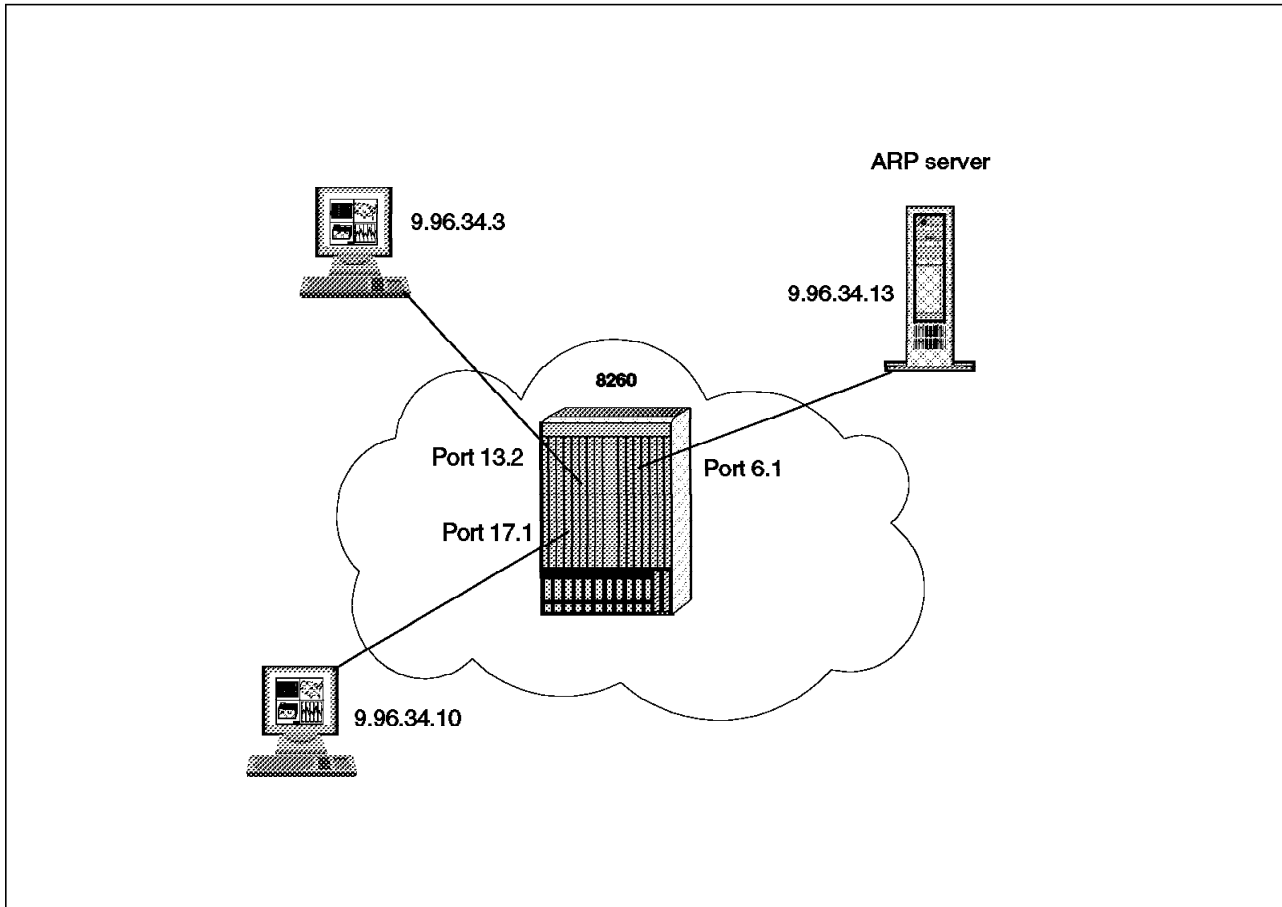


Figure 49. Simple IP Network

5.4.3.1 Symptom and Status Summary

The device on port 13.2 (refer Figure 49) is booted up but cannot communicate on the ATM network.

Complementary information:

- A PING or similar test command from the device to any other device through the ATM interface fails.
- Investigating for the status of port 13.2 at the switch shows NO ACTIVITY or NOT IN SERVICE or NO ADDRESS REGISTRATION.
- We will also discover that the ESI is not registered at the switch. This can be checked on the IBM 8260 with the command: SHOW ATM_ESI ALL.

5.4.3.2 Methodology

We first need to check LEDs for status indicating problems with the hardware of the switch. Also check the workstation adapter and media cable connections.

Use console commands as on the switch to get the status of ports connecting to the station. For example, entering the command show port 13.2 verbose at the 8260 console will result in the display shown in Figure 50 on page 119.

```
8260ATM#1> show port 13.2 verbose
  Type Mode      Status
-----
13.02:UNI enabled NOT IN SERVICE
Signalling Version : Auto
> Oper Sig. Version : 3.1
ILMI status        : UP
ILMI vci           : 0.16
RB Bandwidth       : unlimited
Signalling vci     : 0.5
Administrative weight: 5040
VPI.VCI range      : 15.1023 (4.10 bits)
Connector          : SC DUPLEX
Media              : multimode fiber
Port speed         : 155000 kbps
Remote device is active
Frame format       : SONET STS-3c
Scrambling mode    : frame and cell
Clock mode         : internal
```

Figure 50. Show Port Command

Information that can generally be gathered from the switch regarding the ports include the present configuration of switch parameters regarding ILMI signalling version supported on the port and the status of the port after the last registration attempt by device. The console commands will also indicate if there is some internal problem with the switch or port that may have prevented completion of the ILMI. In our case, we see that the port is okay except that address registration did not complete as is indicated by the status of NOT IN SERVICE. If we wanted to check further to verify if the address is registered, we could use another switch console command. For the 8260, one such command is Show Reachable_Address (Figure 41 on page 101).

Internal traces or protocol analyzers are great for checking address registration. Figure 51 on page 122 is the output of a formatted internal trace from the 8260 that shows a normal address registration via ILMI.

5.4.3.3 Causes for Failure of Address Registration

Causes are:

- End device hardware
Check for errors at the end device.
- End device configuration
Adapter drivers and procedures to configure and install them are many and varied. It is important to read and understand the installation instructions and follow them carefully. Many adapters also include README files and Release Notes and other information that may have been discovered too late to be included in the published documentation. Specifically check whether the endstation supports ILMI and whether the UNI version it supports matches that of the switch port it is attached to.
- Bad cable/fiber or connection between switch and device
Check the cable/fiber and any connections between the end device and switch.
- Switch port hardware

Try another port.

- Switch port configuration

Make sure the port is configured as a UNI port.

Make sure the switch port is configured to match the capabilities of the endstation in terms of UNI signalling version supported and whether ILMI is supported or not.

- Endstation address not in access control table

When access control servers are used by the switch, make sure that the endstation device address has been entered in the access control table and has been entered in the proper format.

- Adapter does not support LECS well-known address

Some adapters have trouble completing ILMI if a LECS address is defined in the switch because they do not support the LECS well-known address. These adapters will likely need to be upgraded either by swapping the hardware or installing new driver code. However, you can sometimes work around the problem by removing the LECS addressing from the switch and hardcoding the LES address.

- End device or switch does not comply with ATM Forum UNI specifications

Interoperability has long been an important issue with networking products and so it is with ATM. ATM networking is relatively new and standards are still evolving. Vendors sometimes do not implement the complete standard in their products (especially early release products) and this leads to the possibility that one vendor's adapter may not interoperate properly with another vendor's switch.

These types of problems can usually best be identified by using a protocol analyzer to get a trace of the problem and then contacting the vendors involved. They will likely be able to correct the problem with a new version of microcode for the switch or end device.

5.4.3.4 Address Registration Using ILMI

Automatic address registration of end devices to the switch is the function of ILMI.

The Interim Local Management Interface (ILMI) was defined by the ATM Forum to define information to be collected by agents called UNI Management Entities (UME) and to define the protocol to be used in communicating this information with a peer UME.

ILMI uses the Simple Network Management Protocol (SNMP) for data transfer across the UNI and for the information formats used.

The ILMI MIB includes information pertaining to the physical layer, the ATM layer configuration (UNI interface port type, number of VPs and VCs, etc.), ATM layer statistics, information about any VPCs and VCCs, and information about address registration information.

Address registration is important here because an end device cannot establish a call using SVCs without first having its address successfully registered with the switch.

Figure 51 on page 122 shows typical flows on a UNI 3.1 ILMI sequence captured by the internal traces of an 8260.

One can see the communication between the switch and the end device as the switch gathers information about the device's characteristics.

The process culminates by the switch passing its NetPrefix (first 13 bytes of its address) to the end device, after which the end device then validates the NetPrefix, appends its ESI and selector byte to it, and returns the resulting 20-byte address to the switch to be registered in the switch's address table.

Once the switch validates the address, the registration on the end device is complete.

The ATM Forum UNI specifications do allow a device to register other addresses after initial registration is complete, so anyone analyzing a trace of the ILMI process will likely see this additional address registered.

```

09:50:20 TX TRAP      TYPE= COLDSTART
09:50:20 TX GETNEXT   VAR(1)= NetPrefixTable.0.0.0.0.0.0.0.0.0.0.0
09:50:21 RX TRAP      TYPE= COLDSTART
09:50:21 RX GETNEXT   VAR(1)= AddressTable.0.0
09:50:21 TX RESPONSE VAR(1):
    atmfSrcvRegATMAddress.0.10.1.3.6.1.4.1.353.1.5.1.1 =
    39.99.99.99.99.99.01.01.40.00.82.10.22.22.00.39.99.99.99
09:50:21 RX GET      VAR(1)= UniVersion.0
09:50:21 TX RESPONSE VAR(1): UniVersion.0 = 3
09:50:21 RX GET      VAR(1)= MyIpNmAddress.0
09:50:21 TX RESPONSE VAR(1): MyIpNmAddress.0 = 192.2.2.7
09:50:21 RX GET      VAR(1)= atmfPortMyIfName.0
09:50:21 TX RESPONSE VAR(1): atmfPortMyIfName.0 = 'at203.2.7'
09:50:25 TX GETNEXT   VAR(1)= NetPrefixTable.0.0.0.0.0.0.0.0.0.0.0 0
09:50:25 RX RESPONSE VAR(1): NetPrefixTable.0.0.0.0.0.0.0.0.0.0.0
09:50:25 TX GET      VAR(1)= UniVersion.0
09:50:25 RX RESPONSE VAR(1): UniVersion.0 = 3
09:50:25 TX GET      VAR(1)= MaxVpiBits.0 VAR(2)= MaxVciBits.0
09:50:25 RX RESPONSE VAR(1): MaxVpiBits.0 = 0 VAR(2): MaxVciBits.0 = 9
09:50:25 TX GET      VAR(1)= sysDescr.0 VAR(2)= sysObjectID.0
09:50:25 RX RESPONSE VAR(1): sysDescr.0 VAR(2): sysObjectID.0
09:50:25 TX GET      VAR(1)= sysName.0
09:50:25 RX RESPONSE VAR(1): sysName.0
09:50:25 TX GET      VAR(1)= sysLocation.0
09:50:26 RX RESPONSE VAR(1): sysLocation.0
09:50:26 TX GET      VAR(1)= MyIpNmAddress.0
09:50:26 RX RESPONSE VAR(1): MyIpNmAddress.0
09:50:26 TX GET      VAR(1)= atmfPortMyIfName.0
09:50:26 RX RESPONSE VAR(1): atmfPortMyIfName.0
09:50:26 TX GET      VAR(1)= MyOsiNmNsapAddress.0
09:50:26 RX RESPONSE VAR(1): MyOsiNmNsapAddress.0
-----
09:50:26 TX SET
VAR(1): NetPrefixTable.39.99.99.99.99.99.0.0.99.99.1.1 = VALID
09:50:26 RX RESPONSE
VAR(1): NetPrefixTable.39.99.99.99.99.99.0.0.99.99.1.1 = VALID
-----
09:50:26 RX SET      VAR(1):
AddressTable.39.99.99.99.99.99.0.0.99.99.1.1.40.0.3.50.1.47.0 = VALID
09:50:26 TX RESPONSE VAR(1):
AddressTable.39.99.99.99.99.99.0.0.99.99.1.1.40.0.3.50.1.47.0 = VALID

```

Figure 51. Normal ILMI (UNI 3.1)

5.4.4 ILMI Scenario

This scenario is aimed at highlighting the role of ILMI in communication setup, and how to investigate a potential failure at this level. ILMI should start when the presence of a remote device has been detected by the physical layer. That is, the switch is powered on, the device is powered on, the physical connecting media is good and the port is enabled and configured correctly as a UNI port.

5.4.4.1 Network Investigations

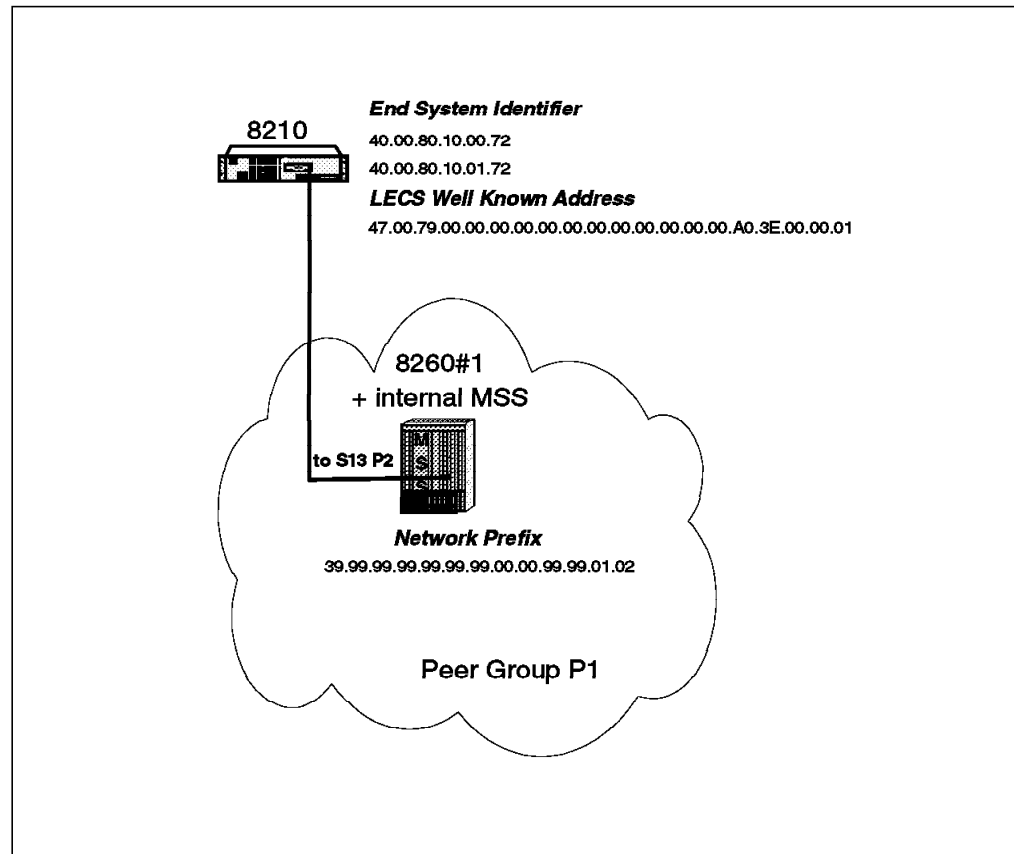


Figure 52. ILMI Network Configuration

Initial Switch Configuration (8260): The remote device connected to port 13.2 is an 8210 with two ESI and some internal LECs that use auto-configuration for their ELAN assignment.

An integrated MSS module is installed in slot 4 and as such will use port 4.1 in the 8260.

Figure 53 on page 124 depicts the reachability information found in the switch before the connection of port 13.2. The figure also shows the LECS addresses on port 4.1 that will be sent by the switch to the device when requested through ILMI.

In this scenario the switch is PNNI-compatible code levels, thus the use of the SHOW REACHABLE_ADDRESS command.

On a switch running a version lower than 3.0, an equivalent command would be show atm_esi all.

```

8260ATM#1> show reachable_address all
Port Len Address Active Idx VPI
-----
14.04 96 39.99.99.99.99.99.99.00.00.99.99.70. . . . . N 2 -
17.01 96 39.99.99.99.99.99.99.00.00.99.99.50. . . . . Y 1 -
2.01 152 39.99.99.99.99.99.99.00.00.99.99.01.02.02.00.82.71.00.01 Y Dyn 0
2.02 152 39.99.99.99.99.99.99.00.00.99.99.01.02.08.00.5A.99.86.DC Y Dyn 0
4.01 152 39.99.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.00.69 Y Dyn 0
4.01 152 39.99.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.20.69 Y Dyn 0
4.01 152 47.00.79.00.00.00.00.00.00.00.00.00.00.00.00.00.A0.3E.00.00.01 Y Dyn 0
6.02 152 39.99.99.99.99.99.99.00.00.99.99.01.02.40.00.82.81.00.66 Y Dyn 0
6.03 152 39.99.99.99.99.99.99.00.00.99.99.01.02.40.00.05.70.00.51 Y Dyn 0
17.02 152 39.99.99.99.99.99.99.00.00.99.99.01.02.00.20.DA.72.09.D0 Y Dyn 0

```

Figure 53. Addresses Reachable before ILMI of Port 13.2

The 8260 will respond with two ATM addresses of LECS through ILMI.

```

8260ATM#1> show lan_emul configuration_server
Index ATM address
-----
1 39.99.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.00.69.00
2 39.99.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.00.72.00

```

These addresses correspond to the two LECS connected in this hub in both ports 4.1 and 13.02.

5.4.4.2 Port Status Command Information

Figure 54 depicts the 8260 port status before ILMI. This status may also be seen when the remote device does not start ILMI.

```

8260ATM#1> show port 13.2 verbose
Type Mode Status
-----
13.02: UNI enabled DOWN: Not in service

Signalling Version : Auto
ILMI status : DOWN: Not in service
ILMI vci : 0.16
--More--
Remote device is active

```

Figure 54. Port Status before ILMI

5.4.4.3 ILMI Problem Investigation

To investigate an ILMI failure we must first gather information at each end regarding status of the devices and ports or connect a network analyzer on the link to capture ongoing activity.

Data Capture: To illustrate the troubleshooting that can be performed on ILMI, Figure 55 on page 126 gives the result of an internal trace captured in the 8260 and Figure 56 on page 127 shows the results of the 8210 Event Logging System started on the ILMI subsystem.

The following keywords are used in Figure 55 on page 126:

- The five leftmost columns are the time stamp.
- TRAP is an SNMP TRAP.
- GET stands for SNMP GET.
- GETN stands for SNMP GETNEXT.
- SET is an SNMP SET.
- RESP stands for an SNMP RESPONSE. This is the response to one of the following three commands:
 1. SET
 2. GET
 3. GETNEXT
- NOSUCHNAME means that no data was provided in the SNMP response because the variable does not exist in the MIB.
- RX means that a frame has been received by the switch.
- TX means that a frame has been sent by the switch.

As an example, TX GETN means that the switch has SENT an SNMP GET NEXT to the remote device.

Figure 55 on page 126 and Figure 56 on page 127 contain labels from **1** to **9**. Each of these correspond to a particular phase of the scenario which will be explained in the next paragraph.

```

49:40 RX GETN AddressTable.0.0                                1
49:40 TX RESP atmfSrvRegATMAddress.0.10.1.3.6.1.4.1.353.1.5.1.1
= 39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.00.69.00
49:42 TX TRAP TYPE= COLDSTART
49:42 TX GETN NetPrefixTable.0.0.0.0.0.0.0.0.0.0.0.0.0
49:42 RX RESP NOSUCHNAME NetPrefixTable.0.0.0.0.0.0.0.0.0.0.0.0.0
49:42 TX GET UniVersion.0                                     2
49:42 RX RESP UniVersion.0 = 2
49:42 TX GET MaxVpiBits.0 MaxVciBits.0                       3
49:42 RX RESP MaxVpiBits.0 = 8 MaxVciBits.0 = 16
49:42 TX GET sysDescr.0 sysObjectID.0                         4
49:42 RX RESP NOSUCHNAME sysDescr.0 sysObjectID.0
49:42 TX GET sysName.0
49:42 RX RESP NOSUCHNAME sysName.0
49:42 TX GET sysLocation.0
49:42 RX RESP NOSUCHNAME sysLocation.0
49:42 TX GET MyIpNmAddress.0
49:42 RX RESP MyIpNmAddress.0 = 192.2.1.72
49:42 TX GET atmfPortMyIfName.0
49:42 RX RESP atmfPortMyIfName.0 = 'at1.2.1.72'
49:42 TX GET MyOsiNmNsapAddress.0
49:42 RX RESP NOSUCHNAME MyOsiNmNsapAddress.0
49:42 TX SET NetPrefixTable.39.99.99.99.99.99.0.0.99.99.1.2    5
49:42 RX RESP NetPrefixTable.39.99.99.99.99.99.0.0.99.99.1.2
49:42 RX SET AddressTable.39.99.99.99.99.99.99.0.0.99.99.1.2.40.0.82.10.0.72.0 6
49:42 TX RESP AddressTable.39.99.99.99.99.99.99.0.0.99.99.1.2.40.0.82.10.0.72.0
49:42 RX SET AddressTable.39.99.99.99.99.99.99.0.0.99.99.1.2.40.0.82.10.1.72.0
49:42 TX RESP AddressTable.39.99.99.99.99.99.99.0.0.99.99.1.2.40.0.82.10.1.72.0
49:42 RX SET AddressTable.47.0.79.0.0.0.0.0.0.0.0.0.0.0.0.a0.3e.0.0.1.0
49:42 TX RESP AddressTable.47.0.79.0.0.0.0.0.0.0.0.0.0.0.0.a0.3e.0.0.1.0
49:46 RX GETN AddressTable.0.0                                7
49:46 TX RESP AddressTable.39.99.99.99.99.99.99.0.0.99.99.1.2.40.0.82.10.0.72.0
49:47 RX GETN atmfSrvRegATMAddress.0.10.1.3.6.1.4.1.353.1.5.1.1 8
49:47 TX RESP atmfSrvRegATMAddress.0.10.1.3.6.1.4.1.353.1.5.1.1
= = 39.99.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.00.69.00
49:47 RX GETN atmfSrvRegATMAddress.0.10.1.3.6.1.4.1.353.1.5.1.1
49:47 TX RESP atmfSrvRegATMAddress.0.10.1.3.6.1.4.1.353.1.5.1.2
= 39.99.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.00.72.00
49:47 RX GETN atmfSrvRegATMAddress.0.10.1.3.6.1.4.1.353.1.5.1.2
49:47 TX RESP atmfSrvRegATMAddress.0.10.1.3.6.1.4.1.353.1.5.1.2
49:47 TX GET MaxVpiBits.0 MaxVciBits.0                       9
49:47 RX RESP MaxVpiBits.0 = 8 MaxVciBits.0 = 16
49:51 RX GETN AddressTable.0.0
49:51 TX RESP AddressTable.39.99.99.99.99.99.99.0.0.99.99.1.2.40.0.82.10.0.72.0
49:52 TX GET MaxVpiBits.0 MaxVciBits.0
49:52 RX RESP MaxVpiBits.0 = 8 MaxVciBits.0 = 16
49:56 RX GETN AddressTable.0.0
49:56 TX RESP AddressTable.39.99.99.99.99.99.99.0.0.99.99.1.2.40.0.82.10.0.72.0

```

Figure 55. ILMI of an 8210 Port Seen from 8260

```

ILMI.020: nt0 snt Trap 1
ILMI.020: nt0 snt GetNext
ILMI.008: nt0 rcv Trap
ILMI.008: nt0 rcv GetNext
ILMI.009: nt0 Snd GetNextRsp NoSuchName, state=WAITING_FOR_PREFIX
ILMI.020: nt0 snt Get Response
ILMI.008: nt0 rcv Get 2
ILMI.011: nt0 Snd Rsp UNI Vers, state=WAITING_FOR_PREFIX
ILMI.020: nt0 snt Get Response
ILMI.008: nt0 rcv Get 3
ILMI.011: nt0 Snd GetRsp Vpi+Vci, state=WAITING_FOR_PREFIX
ILMI.020: nt0 snt Get Response
ILMI.008: nt0 rcv Get 4
ILMI.009: nt0 Snd GetRsp NoSuchName, 2 vars, state=WAITING_FOR_PREFIX
ILMI.020: nt0 snt Get Response
ILMI.008: nt0 rcv Get
ILMI.009: nt0 Snd GetRsp NoSuchName, state=WAITING_FOR_PREFIX
ILMI.020: nt0 snt Get Response
ILMI.008: nt0 rcv Get
ILMI.009: nt0 Snd GetRsp NoSuchName, state=WAITING_FOR_PREFIX
ILMI.020: nt0 snt Get Response
ILMI.008: nt0 rcv Get
ILMI.011: nt0 Snd GetRsp MyIpNmAddress, state=WAITING_FOR_PREFIX
ILMI.020: nt0 snt Get Response
ILMI.008: nt0 rcv Get
ILMI.011: nt0 Snd GetRsp MyIfName, state=WAITING_FOR_PREFIX
ILMI.020: nt0 snt Get Response
ILMI.008: nt0 rcv Get
ILMI.009: nt0 Snd GetRsp NoSuchName, state=WAITING_FOR_PREFIX
ILMI.020: nt0 snt Get Response
ILMI.008: nt0 rcv Set 5
ILMI.020: nt0 snt Get Response
ILMI.011: nt0 setting Prefix valid, state=WAITING_FOR_PREFIX
ILMI.011: nt0 Set state, state=WAITING_FOR_PREFIX
ILMI.021: nt0 net pref=39999999 99999900 00999901 02
ILMI.020: nt0 snt Set 6
ILMI.020: nt0 snt Set
ILMI.008: nt0 rcv Get Response
ILMI.020: nt0 snt Set
ILMI.008: nt0 rcv Get Response
ILMI.008: nt0 rcv Get Response
ILMI.020: nt0 snt GetNext 7
ILMI.008: nt0 rcv Get Response
ILMI.023: nt0 ntrd func alloc_atm_addr, addr=40008210 0072,sel=04
ILMI.004: nt0 ntrd func SendGetNextForLeCsAdrs, index=-1 8
ILMI.020: nt0 snt GetNext
ILMI.017: nt0 get_LECS_list: in progress
ILMI.008: nt0 rcv Get Response
ILMI.016: nt0 Snd GetNext for more LECS, cnt=1
ILMI.008: nt0 rcv Get Response
ILMI.004: nt0 ntrd func SendGetNextForLeCsAdrs, index=2
ILMI.020: nt0 snt GetNext
ILMI.016: nt0 Snd GetNext for more LECS, cnt=2
ILMI.008: nt0 rcv Get Response
ILMI.008: nt0 rcv Get 9
ILMI.011: nt0 Snd GetRsp Vpi+Vci, state=ILMI_PREFIX_READY
ILMI.020: nt0 snt Get Response
ILMI.020: nt0 snt GetNext
ILMI.008: nt0 rcv Get Response
ILMI.008: nt0 rcv Get
ILMI.011: nt0 Snd GetRsp Vpi+Vci, state=ILMI_PREFIX_READY

```

Figure 56. ILMI of an 8210 Port Seen from 8210 ELS

During this particular scenario, there are nine major operations performed by the ILMI protocol. What follows is a brief explanation of traces captured from the 8260 and the 8210.

1 This is the initialization step. Once a device is ready it sends a COLD START TRAP to inform remote device of its presence. On reception of a COLD START TRAP, the user must clear its Network Prefix table, while the network switch must clear its address table. Subsequent GETNEXTs are sent on these tables to verify that they are really empty. As illustrated in step **1**, the switch is sending a COLD START TRAP, followed by an SNMP GETNEXT on the Network Prefix table. The 8210 responds with NOSUCHNAME, indicating that there is no instance of the variable in the table. The table is empty. The ILMI process then proceeds. If the device does not respond to the GETNEXT, the switch will issue another GETNEXT. If still there is no response the switch can issue another COLDSTART TRAP and start the process again or declare the UNI down (NOT IN SERVICE).

2 The switch performs UNI level detection. It asks for the UNI level. In this scenario, the response is 2 (MIB value) which indicates a port configured as UNI 3.0. This information will then be used by the signalling component.

3 The switch gets the number of bits for VPI and VCI supported by the remote device. This information is then used by the signalling component. When receiving a call setup, it is now able to allocate VPI and VCI values suitable for both end of the link.

4 During this phase, the switch tries to gather neighboring information about the device connected. This is particularly important for network management applications, since this will allow the management to build the topology of the network, track a connection from one device to another, and so on.

5 This is the first part of address registration. The switch sends its network prefix to the attached device. The network prefix is the first 13 bytes of the ATM address.

6 This is the second part of the address registration. The end device appends the ESI (6 bytes) and a selector (the last byte) and builds a full ATM address that is sent to the switch for registration. In this scenario, it registers two regular addresses and the LECS well-known address.

7 ILMI has now completed. The remote device starts the regular polling on AddressTable to check that the switch is still alive.

8 The remote device contains auto-configure LECs. These LECs try to get the address of the LECS using ILMI, and the switch returns the two addresses that were configured as configuration servers.

9 This is the steady state flow. The switch issues a periodic GET on MaxVpiBits and MaxVciBits to check remote device presence, while the end device periodically issues a GETNEXT on the AddressTable to validate registration.

5.4.4.4 Checking Successful Completion

Following are some commands that allow confirmation of a successful ILMI:

1. Reachable addresses on 8260.

```
8260ATM#1> show reachable_address all
Port Len Address Active Idx VPI
-----
14.04 96 39.99.99.99.99.99.00.00.99.99.70. . . . . N 2 -
17.01 96 39.99.99.99.99.99.00.00.99.99.50. . . . . Y 1 -
2.01 152 39.99.99.99.99.99.00.00.99.99.01.02.02.00.82.71.00.01 Y Dyn 0
2.02 152 39.99.99.99.99.99.00.00.99.99.01.02.08.00.5A.99.86.DC Y Dyn 0
4.01 152 39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.00.69 Y Dyn 0
4.01 152 39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.20.69 Y Dyn 0
4.01 152 47.00.79.00.00.00.00.00.00.00.00.00.00.00.00.A0.3E.00.00.01 Y Dyn 0
6.02 152 39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.81.00.66 Y Dyn 0
6.03 152 39.99.99.99.99.99.00.00.99.99.01.02.40.00.05.70.00.51 Y Dyn 0
13.02 152 39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.00.72 Y Dyn 0
13.02 152 39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.01.72 Y Dyn 0
13.02 152 47.00.79.00.00.00.00.00.00.00.00.00.00.00.A0.3E.00.00.01 Y Dyn 0
17.02 152 39.99.99.99.99.99.00.00.99.99.01.02.00.20.DA.72.09.D0 Y Dyn 0
```

Figure 57. Reachable Addresses after ILMI of Port 13.2

Note that three more addresses are now reachable in this switch and registered at port 13.2.

2. 8260 port status

```
8260ATM#1> show port 13.2 verbose
Type Mode Status
-----
13.02: UNI enabled UP
Signalling Version : Auto
> Oper Sig. Version : 3.0
ILMI status : UP
ILMI vci : 0.16
--More--
Remote device is active
8260ATM#1>
```

Figure 58. 8260 Port Status after Successful ILMI

The port status and ILMI status are both UP.

3. 8210 ATM interface status

```

ATM Interface+
ATM Interface+list all
***** USERS *****
--More--

***** ADDRESSES *****

          ATM Address
      Network Prefix          ESI          SEL
-----
39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.00.72.02
39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.00.72.04
39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.01.72.02
39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.00.72.00

***** VCCS *****

Conn  Conn  Sap  Frames  Frames
Handle Type  VPI  VCI  FrameSap  Type  Transmitted  Received
-----
--More--
   1   SAAL  0    5    F464A0  Buff      227          219
   2   ILMI  0   16    F464A0  Buff       49           46

ATM Interface+

```

Figure 59. 8210 Physical ATM Interface Status after Successful ILMI

The 8210 has ATM addresses, and the ILMI indicates frames exchanged with the switch.

5.4.4.5 Further Investigation

If this level of tracing has indicated an abnormal ILMI, more investigation can be performed on both ends. Many trace tools will allow the troubleshooter to capture more detailed data in either hexadecimal format or in a formatted trace output. Figure 60 on page 131 is an example of an 8210 ILMI frame trace. This particular frame was the first SNMP SET sent by the 8210 in step 6.

The procedure for capturing this trace on the MSS is by starting traces on the ILMI subsystem, and selecting VPI 0, VCI 16 for ATM interface VCC tracing.


```
#65 Dir
:OUTGOING Time:0.30.29.10 Trap:4851
Comp:ILMI Type:ATM_CHARM Port:0 Circuit:0x000010 Size:75
SEQ (30) (49)
INTEGER (02) ( 1) 0 = 0(d)
OSTRING (04) ( 4) 49 4C 4D 49 = 'ILMI'
SET REQ (A3) (3E)
INTEGER (02) ( 2) 1 0 = 256(d)
INTEGER (02) ( 1) 0 = 0(d)
INTEGER (02) ( 1) 0 = 0(d)
SEQ (30) (32)
SEQ (30) (30)
OID (06) (2B) 2B 6 1 4 1
            82 61 2 6 1
            1 3 0 14 39
            81 19 81 19 81
            19 81 19 81 19
            81 19 0 0 81
            19 81 19 1 2
            40 0 81 2 10
            0 72 0 =
1.3.6.1.4.1.353.2.6.1.1.3.0.20.
57.153.153.153.153.153.0.0.153.153.1.2.64.0.130.16.0.114.0
INTEGER (02) ( 1) 1 = 1(d)
```

Figure 60. 8210 ILMI Trace Output Sample

5.4.4.6 Conclusion

This example gives a high level view of the ILMI protocol. It allows the troubleshooter to verify that all functions expected from ILMI were performed and that the result is the expected one. For instance, with this scenario, we know the method used by the remote device to reach the LECS. In this example, 8210 uses ILMI, and the LECS that will be called first will be the one at the first address returned **8**. This address is reachable from port 4.1 in the same switch. This means that the internal LECs of the 8210 connected to slot 13.2 uses the configuration server located in the MSS connected to port 4.1.

This type of data capture can also highlight an invalid or incomplete ILMI scenario such as:

- No response from remote device
- Incomplete address registration
- No use of ILMI for LECS address retrieval

Bypassing ILMI: 8210/MSS services will not start without a successful ILMI.

There are two main reasons to bypass ILMI:

1. The remote device does not support it.
2. The ILMI scenario cannot be completed due to interoperability issues.

In either case, ILMI can be disabled and bypassed. However, if the port is to be used for SVCs, address registration must be manually done. Figure 61 on page 132 shows the command used to perform static address registration on an 8260 running PNNI. Note that the prefix length should be 152 bits (19 bytes) since

this command will actually give reachability to a single address. On an 8260 switch running a code version prior to V3.0, the command is set atm_esi.

```

8260ATM#1> set reachable_address
Enter slot: 14.2
Enter prefix length: 152
Enter reachable address :
39.99.99.99.99.99.00.00.99.99.01.02.40.00.32.23.32.23
Entry set.
8260ATM#1> show reachable_address all
Port Len Address Active Idx VPI
-----
14.02 152 39.99.99.99.99.99.99.99.00.00.99.99.01.02.40.00.32.23.32.23 N 3 -
14.04 96 39.99.99.99.99.99.99.99.00.00.99.99.70. . . . . N 2 -
17.01 96 39.99.99.99.99.99.99.99.00.00.99.99.50. . . . . Y 1 -
4.01 152 39.99.99.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.20.69 Y Dyn 0
4.01 152 47.00.79.00.00.00.00.00.00.00.00.00.00.00.00.00.A0.3E.00.00.01 Y Dyn 0
6.03 152 39.99.99.99.99.99.99.99.00.00.99.99.01.02.40.00.05.70.00.51 Y Dyn 0
17.02 152 39.99.99.99.99.99.99.99.00.00.99.99.01.02.00.20.DA.72.09.D0 Y Dyn 0

```

Figure 61. Static Address Registration on 8260

If ILMI has been disabled and the device is running LAN Emulation, then the only way to reach LECS will be by the well-known address or by configuring the LECS real address in the attached device. Refer to Chapter 7, “ATM Emulated LANs and Logical IP Subnets (LANE 1.0, RFC 1577)” on page 293 for more information about LAN Emulation.

5.5 Problem Determination Guidelines

This section gives guidance and a method for how to start to troubleshoot a network.

5.5.1 Problem Phases

We can break problem determination into four phases. These can be defined as follows:

- **Phase 1:** From unpacking the product until application of power.

In this phase media cables have not yet been attached.

Problems occurring in phase 1 are generally addressed by diagnostic chapters of the user documents of the specific product. The failures usually result in an LED indication that can be used for problem isolation.

- **Phase 2:** From active LEDs up to the end of ATM configuration.

In the case of hubs, in this phase the network administrator would assign the proper addresses of modules as required, enable ports, and connect (attach) modules to the ATM subsystem network. For example, ATM modules inserted into the IBM 8260 are not automatically attached to the ATM network to prevent possible unintentional disruption of the network. This requires that the network administrator use a console command to attach the module to the ATM subsystem before it will be usable on the network. Figure 37 on page 96 shows output of the SHOW MODULE command which shows the status of a module after it has been properly connected to the ATM subsystem network. In this phase media cables have not yet been attached. At the completion of phase 2 the normal port status will be NO ACTIVITY.

Like phase 1 problems, the problems in phase 2 can also be found using the specific product user documentation. Common causes of problems at this stage are:

- Configuration console problems. The user cannot get a proper connection to the configuration console or it does not respond to its commands.
 - Modules do not respond to configuration commands.
 - Ports do not show proper status after being configured.
- **Phase 3:** From the end of ATM configuration *but* still no ATM traffic.

At this phase, devices are connected to ATM ports via the appropriate cabling but at this point there is still no traffic. At the completion of phase 3, the normal status of ports and links will be OKAY. For UNI ports, address registration is completed.

Common problems at this phase are endstation problems such as code level incompatibility between devices and the switch, defective media cabling or connectors, address registration problems and other configuration problems.

- **Phase 4:** There is ATM traffic.

These are problems that occur after all ATM devices are successfully attached to ATM media ports and ATM traffic is started in the network. The problems in this phase concern interruptions to the normal operation of the network. The most common cause of problems in this phase is that something in the network is changed by someone. Problems may occur

when the maximum number of virtual connections (VCs) allowed on a switch or an individual media module is exceeded.

5.5.2 Hints and Tips with IBM Products

Table 10 on page 138 and Table 11 on page 143 contain some representative problems that may be seen as well as some recommendations on what to do to fix them. These examples were taken from troubleshooting chapters of the user guides of the 8260 and 8285 and can be used as examples for what to do when the particular fault is seen in any network. The tables are not a complete listing, so refer to the user guide of the particular product involved for a complete list of symptoms and fixes.

5.5.3 Problem Isolation Method

For simplicity, most problems can be considered connectivity problems (device A cannot communicate with device B) or performance problems (that is, slow response). Connectivity problems can be further divided into hardware problems (that is, red LEDs, no power, bad ports or adapters), defective media links, or configuration problems. Using the network topology diagrams and a knowledge of what type of problem is involved from data gathered per Chapter 4, "Problem Determination Guidelines" on page 91, the reader will then be able to start from one end of the connection to determine how much of the network is good and in the process determine the scope of the problem.

The task for the network administrator is to first determine that the expected UNI connections and switch node to switch connections are set up correctly. An approach that works well is to start at the switch that is closest to the failing device and then test connectivity in the direction of the device.

5.5.3.1 Problem Isolation Flows

A series of flowcharts for problem isolation are presented on the following pages.

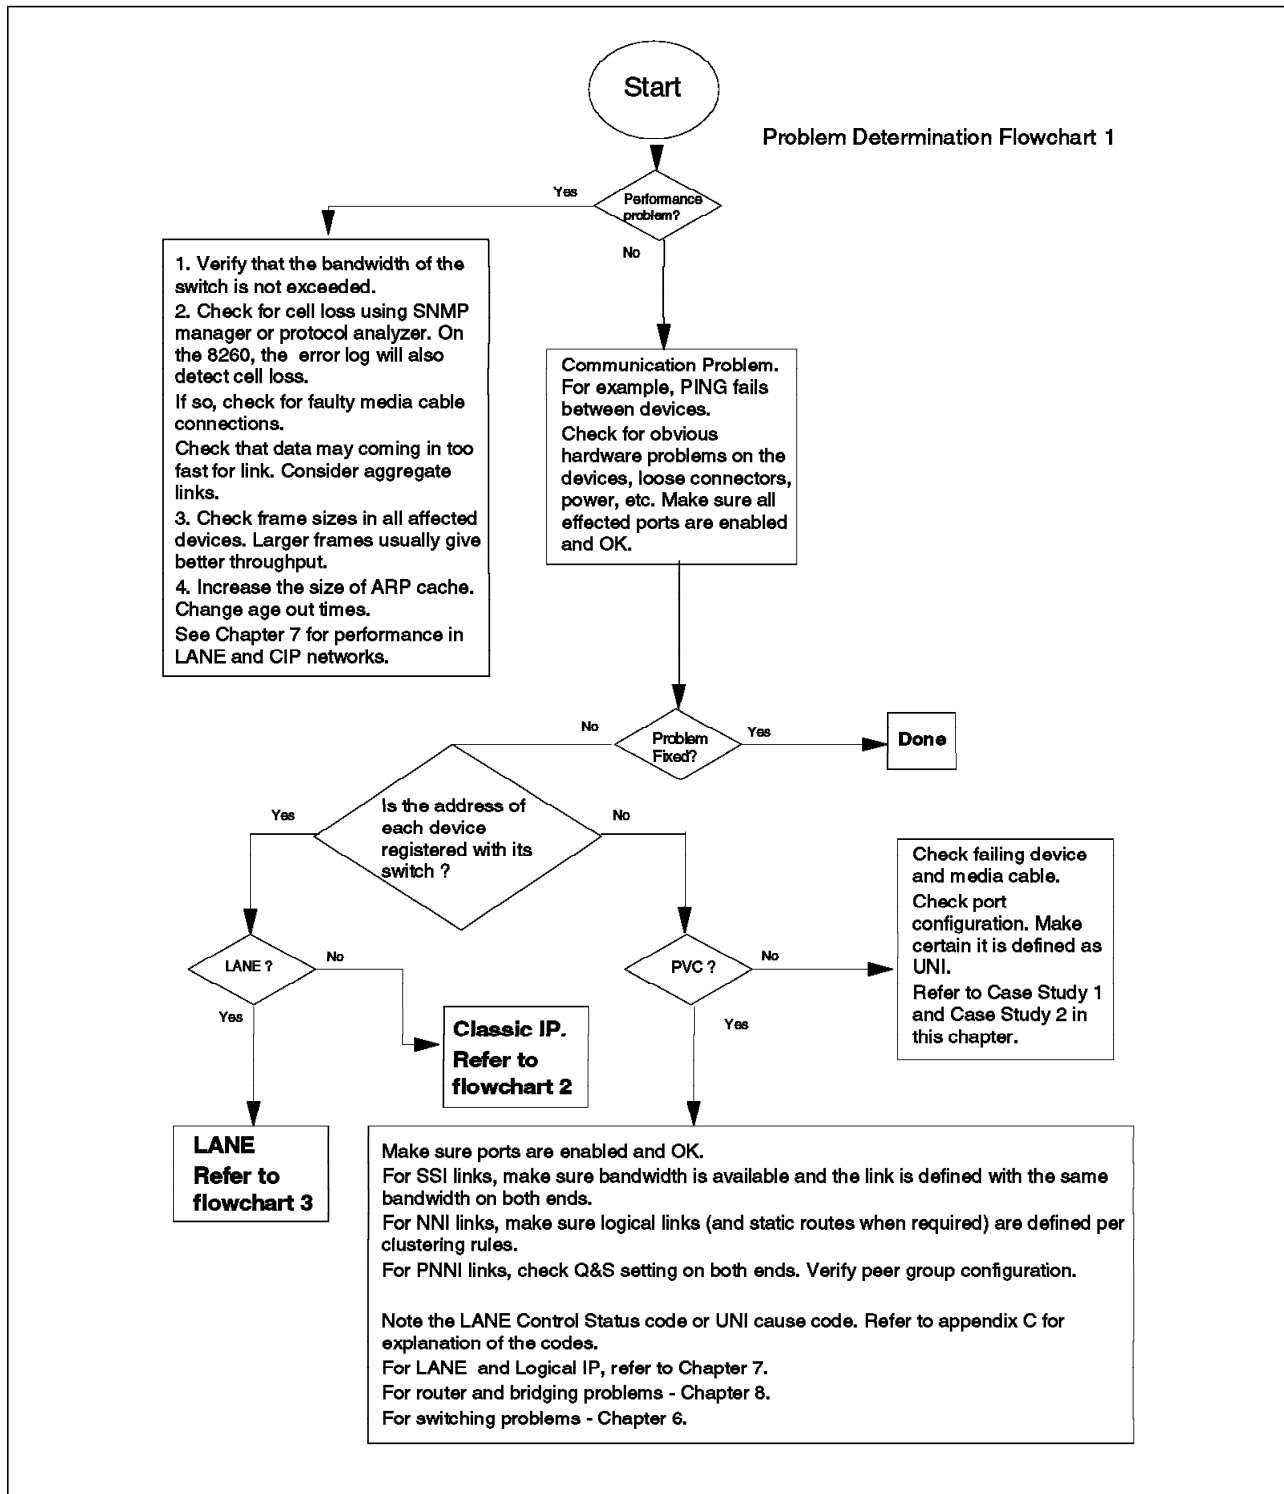


Figure 62. Problem Isolation (Chart 1 of 3)

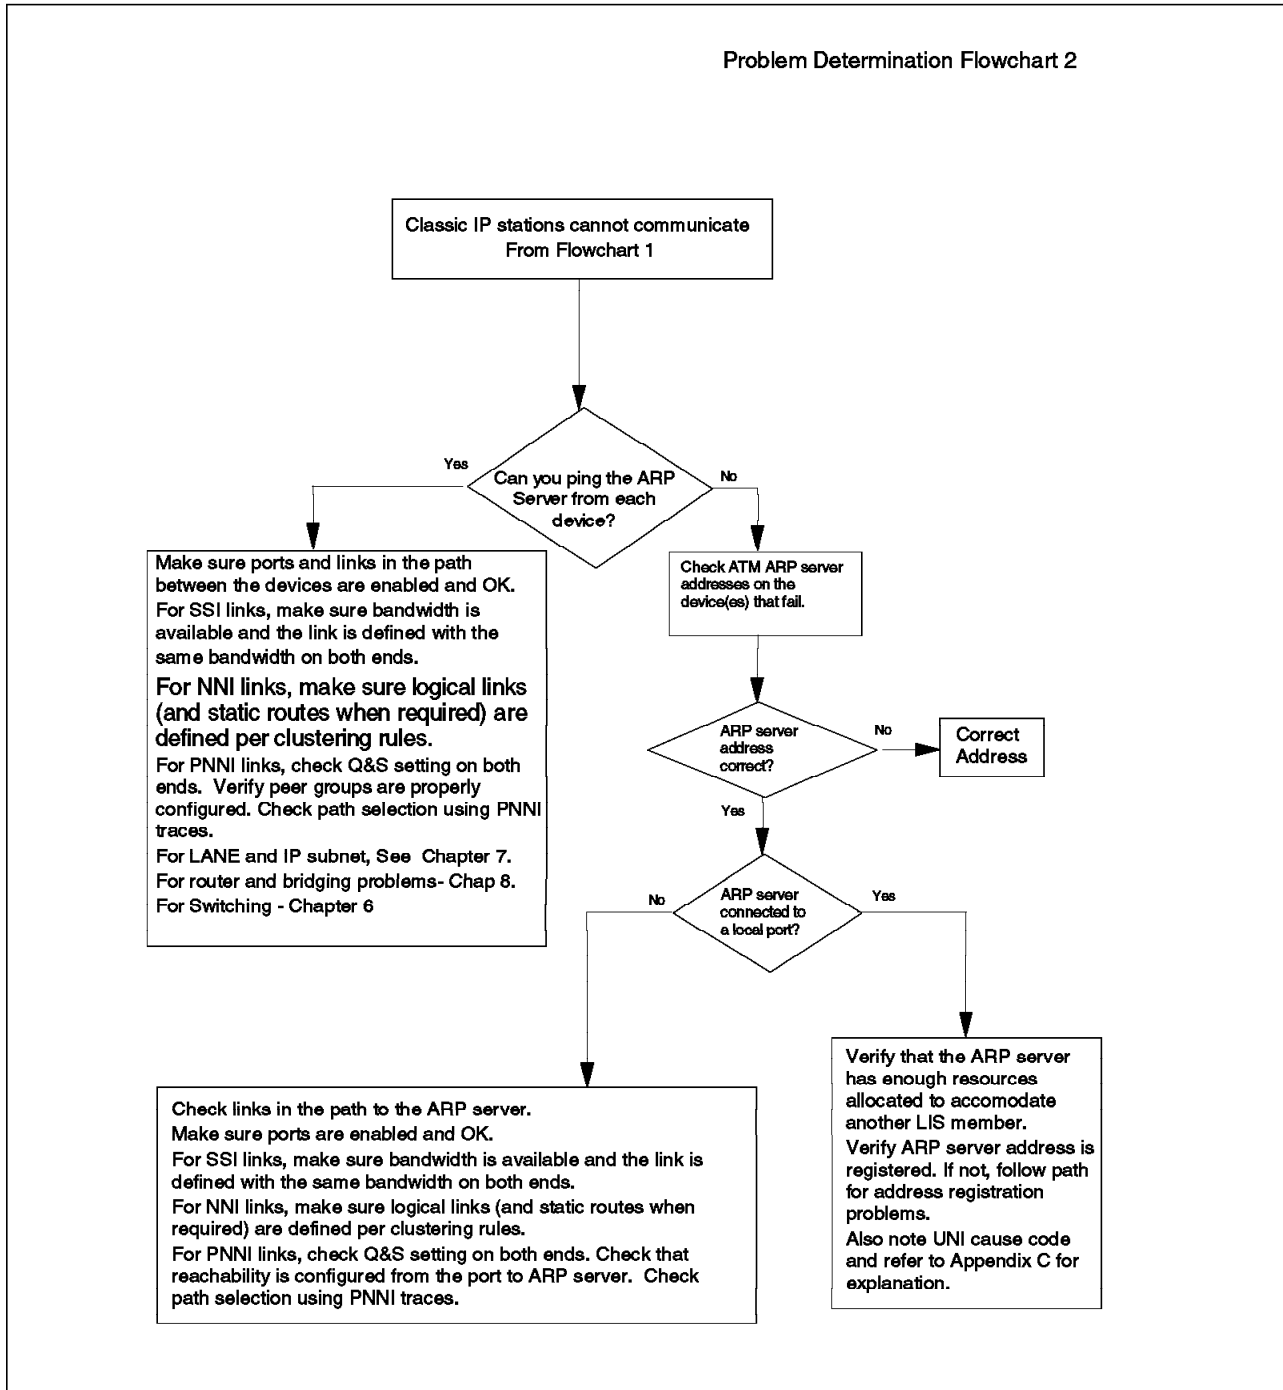


Figure 63. Problem Isolation (Chart 2 of 3)

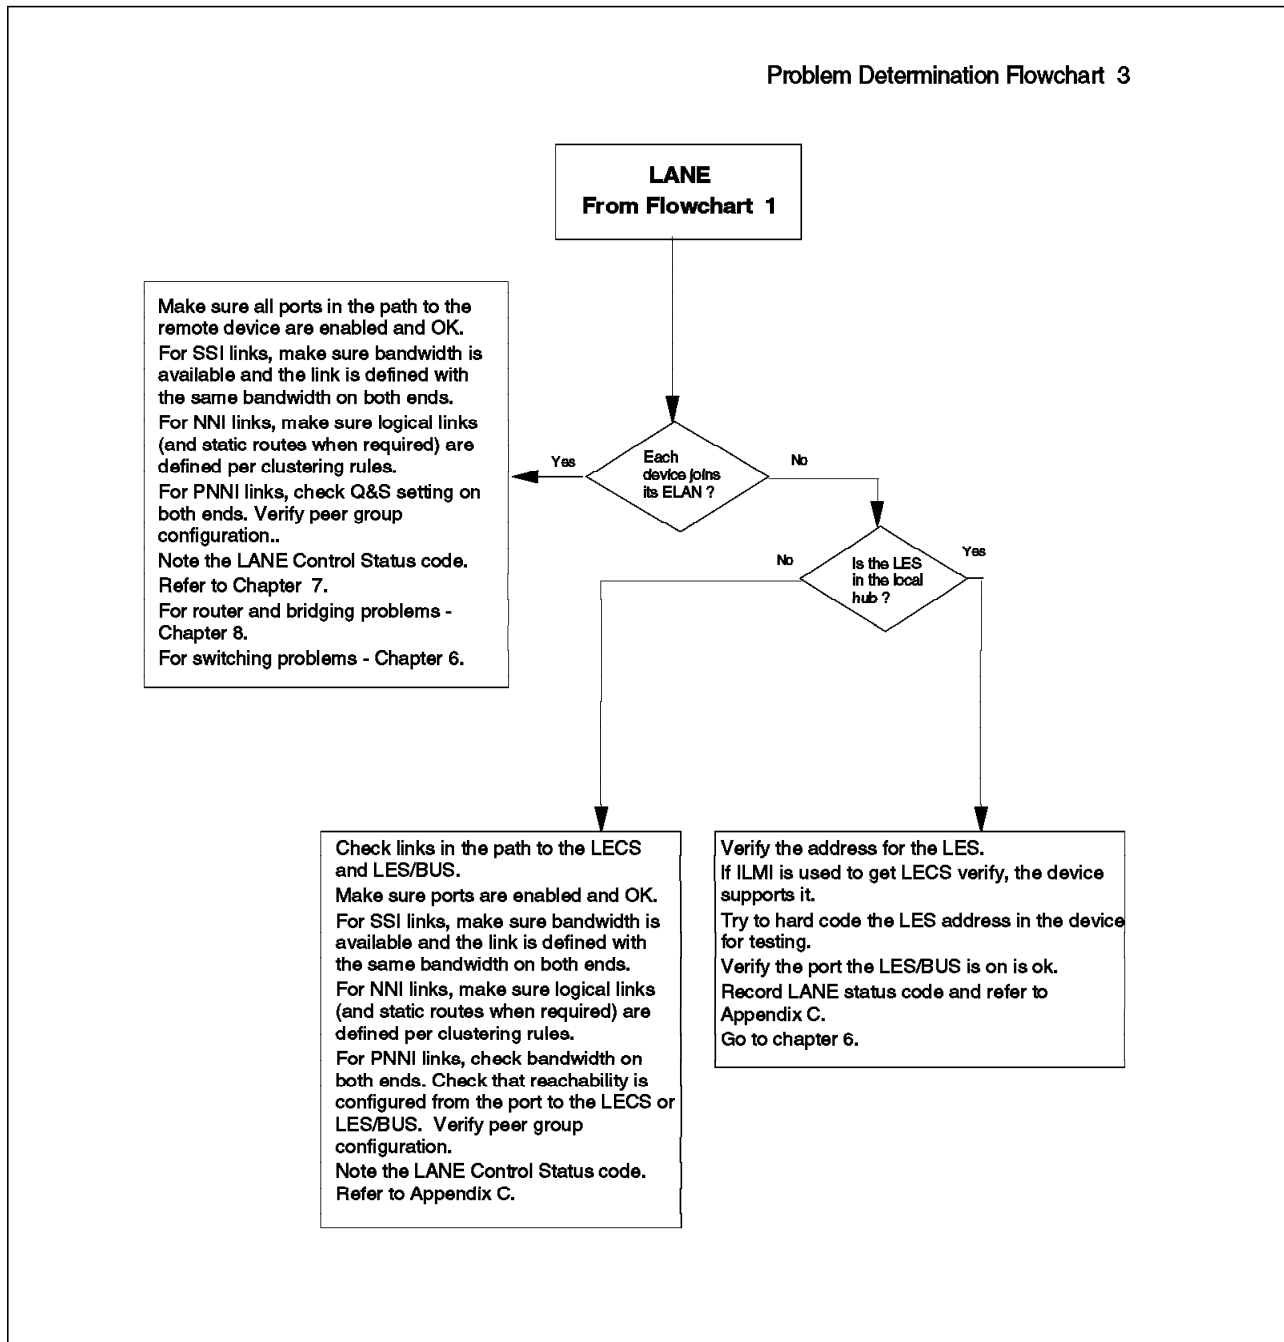


Figure 64. Problem Isolation (Chart 3 of 3)

Starting with Figure 62 on page 135 the reader will be directed to the problem area for most common networking problems.

5.6 ATM Commons Problems

Table 10 shows some typical problems seen at the first three phases. The IBM 8260 and 8285 are used as a reference but much of the information and problem isolation process will apply to other ATM network problems as well.

Table 10 (Page 1 of 5). Typical Problems Seen in Phases 1-3

Symptom	Possible Causes	Recommended Actions
RED LED indications that appear abnormal	Possible hardware error	Go to the product's user documentation
No Power	Powered off or bad power	Check power source
	Hardware problem	Call service provider
Wrong slot LED lit on the ATM blade in the IBM 8260	Blade is not plugged in the right slot. For for IBM 8260 the Control Point/Switch module (CPSW) which is a 2-slot module should be plugged into slot 9-10 or 11-12. Other ATM media modules can only be plugged into slots 1 through 8 and slots 12 through 17. Note that slot 12 is available for a media module only if it is not used by a CPSW.	Plug the ATM media modules in any slot except slots 9, 10, 11. Plug the CPSW in the slots 9-10 . Note: The slots 11-12 (on the 8260 model A17) are used by the backup CPSW; to use them, make sure the slots 9-10 have a CPSW installed.
	8260: The ATM backplane is not installed in the 8260.	Check that the backplane is installed (in the upper part of the 8260).
	Connectivity problem with the backplane (bent pin, badly plugged module. etc.)	Unplug the blade, check the rear pins, as well as the backplane connectors, then plug back the blade carefully.
ATM Media blade has no light on	Module not connected / connectivity problem with the backplane.	<ol style="list-style-type: none"> 1. Unplug the blade, then plug it back carefully. 2. Connect the module to the ATM subsystem. On the 8285 and 8260, use console command SET MODULE x CONNECTED ENABLE.
	Ports not enabled.	From the console terminal, check that the blade status is okay. From the IBM 8260 or IBM 8285 use the console commands, SHOW MODULE n VERBOSE and SHOW PORT n ALL.
	Power budget exceeded.	Check the amount of power available for newly installed modules by using appropriate console command (that is, SHOW POWER all). Add another power supply if necessary or remove one of the previously installed modules.
	The power of that slot has been disabled by the hub management blade.	Use the management console to issue the appropriate commands to enable power to the slot. For the IBM 8260, issue the command SET POWER SLOT n ENABLE.

<i>Table 10 (Page 2 of 5). Typical Problems Seen in Phases 1-3</i>		
Symptom	Possible Causes	Recommended Actions
Command SET MODULE... does not work ("Set rejected").	The blade is physically not well connected, or a pin is bent.	Unplug the blade, check the rear pins, then replug the blade carefully. It is recommended that you use the command SHOW MODULE n VERBOSE to get the exact status of the module.
	The FPGA code of the blade is not compatible with the FPGA of the switch, or with the microcode of the Control-Point (A-CPSW).	Check the prerequisites in the Release Notes or contact your service provider Note: It is recommended that you issue the command SHOW MODULE n VERBOSE to get the exact status of the Module.
ASCII terminal cannot communicate with 8260 or 8285 serial port. Note: RS-232 cabling requirements DIFFER for the 8260 DMM and for the 8260 A-CPSW. Refer carefully to the <i>CPSW User's Guide</i> .	A null-modem is missing or the cable is not a null-nodem cable.	Insert a null-modem between the terminal and the console port, and check that it solves the problem.
	Cable badly plugged, the cable is defective, or the cable does not follow the console port pin-out requirements.	Check the RS-232 cabling requirements in the console port. For the 8260 or the 8285, consult the user's guide. Try with another serial cable.
	The terminal parameters do not match the console ports communications parameters (that is, baud rate, parity, and stop bit).	<ol style="list-style-type: none"> 1. Try the default settings on the terminal (the default 8260/8285 parameters are: 9600 baud, 8 data bits, 1 stop bit, no parity). If it does not work, try different settings (change the baud-rate, the data bits, the parity) until you match the right configuration. 2. If you know that the IP parameters are already set on the 8260/8285, try to open a Telnet session to it. Then, change the RS-232 parameters using the SET TERMINAL... commands. 3. If the ARP-server address was already set correctly in the A-CPSW, but not its IP parameters, you can use a DMM to change the A-CPSW IP address and subnet mask (commands SHOW PORT 9.1 VERBOSE and SET PORT 9.1 IP_ADDRESS/SUBNET_MASK). <p>Once you can PING the A-CPSW, you can Telnet to the A-CPSW and reset the RS-232 parameters from there, using the SET TERMINAL... commands.</p>
Commands SET.... do not work on the 8260 or 8285	The user did not log in with the Administrator's password.	Log out from the session (command LOGOUT), and then enter the administrator's password (factory defaults are 8260 or 8285).
Command SET PORT... does not work (Set rejected).	The corresponding blade is isolated.	From the Terminal Dialog, issue the command SHOW MODULE ALL to check whether the corresponding blade is connected. If it is not connected, enter the command SET MODULE n CONNECTED.
	You are trying to change the parameters of a port that is already enabled.	Disable the port, then enter the command SET PORT... again.

Table 10 (Page 3 of 5). Typical Problems Seen in Phases 1-3

Symptom	Possible Causes	Recommended Actions
Port status is DOWN: Bad FPGA	The FPGA of the media module is not compatible with the FPGA of the ATM Control Point. Check the prerequisites in the Release Notes.	Upgrade the FPGA of the media module.
Module status is DOWN: Hardware Error	The microcode encountered an error when setting the hardware.	<ol style="list-style-type: none"> 1. If the problem persists after a reset, move the module to another slot. 2. If the problem persists, replace the ATM media module.
Port status is DOWN: No Activity It means that there is no light in the receive fiber, or that there is no electrical signal in the receive cable.	Port is not enabled.	Enable port using appropriate management console commands
	The device at the other side of the fiber/cable is powered off.	Power on the remote device.
	A fiber/cable is not well plugged, or is defective.	Replug both sides of the fiber/cable. <ol style="list-style-type: none"> 1. Run wrap tests on the cable. Consult product user guide for procedure. 2. Replace the fiber/cable if necessary.
	For SC-type connectors, the transmit and the receive fibers may be swapped.	Swap the transmit and receive fibers.
	Copper cabling may not be built correctly to swap transmit and receive signalling.	Refer to 8260 or 8285 user guide for the proper pinouts of copper cabling. Note: Pinouts of copper cables that connect UNI links are different from those for copper cables that connect PNNI, NNI or SSI links.
Port status is DOWN: Inconsistency in bandwidth	The bandwidth specifications are different at each end of a connection.	<ol style="list-style-type: none"> 1. Use the SHOW PORT command to display the bandwidth specifications for each port. 2. Use the SET PORT command to ensure that the values match.
Port status is DOWN: Internal error	An internal error is detected on the port.	<ol style="list-style-type: none"> 1. Reset the ATM media module using the RESET MODULE command. 2. If the problem persists, replace the ATM media module.
Port status is DOWN: Invalid VPI-VCI range	The VPI-VCI range entered for a given port is invalid.	Make sure that the range specified for the port is valid, and that the same range is applied at the other end of the connection. Refer to switch user guide for for valid settings.
Port status is DOWN: Connection limit reached	The maximum number of connections (including point-to-point and point-to-multipoint) allowed has been reached, either at the port level or module level. For example, the maximum number allowed for the IBM 8260 is 4064.	Reduce the number of connections.

<i>Table 10 (Page 4 of 5). Typical Problems Seen in Phases 1-3</i>		
Symptom	Possible Causes	Recommended Actions
Port status is DOWN: Duplicate VPC.	When a port is enabled, it is enabled with its associated vpi. The vpi is already in use for this port.	<ol style="list-style-type: none"> 1. Delete the existing VPC that is not required, and disable/enable the port. 2. Disable the port, then re-enable it with a different vpi value.
Port status is DOWN: VPC limit reached	The maximum number of VPCs (64) has been reached.	Delete an existing VPC to allow the creation of a new one.
Port status is Down: Signalling version missing	The signalling version must be specified if the interface does not support ILMI. The signalling version was not specified for a UNI or IISP port or VPCs.	<ol style="list-style-type: none"> 1. Disable the port. 2. Redefine the port specifying the signalling version for the attached device. 3. Re-enable the port.
<p>The status of a UNI port is NOT IN SERVICE.</p> <p>The light/signal is detected on the receive fiber/cable but the remote device is not responding to ILMI.</p>	The peer device or its adapter is defective.	Check the device. Make sure its adapter is OK and the fiber/cable is securely connected to it.
	The transmit fiber/wire is not well plugged or is defective.	<ol style="list-style-type: none"> 1. Replug the cable/fiber. 2. Run port wrap tests. If the test fails the port is bad so try another port. If the wrap is successful replace the cable.
	The peer device does not support ILMI.	Change the UNI port to ILMI-less (only with an 8260/8285 microcode version greater than or equal to 2.0.4).
	A PVP with VPI=0 is/was defined on that port .	Release the PVP, and disable/enable that port.
	The UNI port is defined with the ILMI enabled, but the workstation connected to it has a device driver which does not support the LECS well-known ATM address, and a LECS address is defined in your 8260/8285 .	<p>Check that you have a LECS address configured in your 8260/8285 with the command SHOW LAN_EMUL CONFIGURATION_SERVER. If there should not be any LECS address defined, clear it with the command CLEAR LAN_EMUL CONFIGURATION_SERVER ALL.</p> <p>Take traces of ILMI from the hub (switch) end and from the device end. The 8260 and 8285 has an internal trace that can be activated by the SET TRACE command. See user guide for more details.</p>
<p>The status of a PNNI port is NOT IN SERVICE.</p> <p>The light/signal is detected on the receive fiber/cable but the remote device is not responding to ILMI.</p>	The peer device or its adapter is defective.	Check the device. Make sure its adapter is OK and the fiber/cable is securely connected to it.
	The transmit fiber/wire is not well plugged or is defective.	<ol style="list-style-type: none"> 1. Replug the cable/fiber. 2. Run port wrap tests. If the test fails the port is bad so try another port. If the wrap is successful replace the cable.
	The peer device does not support ILMI.	Change the UNI port to NO-ILMI.
	A PVP with VPI=0 is/was defined on that port .	<p>Release the PVP, and disable/enable that port.</p> <p>Take traces of ILMI from the hub (switch) end and from the device end. The 8260 and 8285 have an internal trace that can be activated by the SET TRACE command. See user guide for more details.</p>

Table 10 (Page 5 of 5). Typical Problems Seen in Phases 1-3

Symptom	Possible Causes	Recommended Actions
Port Status is DOWN: Error Detected	This indicates an ILMI protocol error. This status may be a result of a security violation on the port.	<ol style="list-style-type: none"> 1. Enter the SHOW SECURITY PORT command to check that security is enabled for this port. 2. Enter the SHOW SECURITY LAST_VIOLATION command to check if an address registration has been rejected on the port. 3. Check that the attached device supports ILMI. 4. Perform a wrap test to check the connection to the attached device. 5. Re-enable the port.
Port Status is UP: NO Address Registration	ILMI is up but the connecting device does not accept address registration.	<p>Device address registration needs to be manually down.</p> <p>On 8260 and 8285 with PNNI support this can be done the using the SET REACHABLE_ADDRESS command.</p>

<i>Table 11 (Page 1 of 2). Problems of Phase 3 and 4</i>		
Symptom	Possible Causes	Recommended Actions
PING does not work on IP over ATM between two ARP clients	Cause code 1. ARP Server not properly registered.	<ol style="list-style-type: none"> 1. Verify that the port the ARP server is attached to is OK. 2. Verify ARP server hardware and configuration. 3. Verify media between port and ARP server.
	Cause code 1 The ARP server address entered at the client is incorrect.	Change the ARP server address.
	Clients not properly registered with their switches.	Check phase 3 exit with each client.
	Cause code 3. Hub port may not have enough bandwidth	Reduce the bandwidth or spread it over several hub ports.
	Cause code 3 The ARP server is in another hub and the IISP (NNI) links have not been properly configured.	Review the configuration of the links between the hubs. Note that that they are properly configured for any static routing, reachable addresses and logical links. If there are any network side/user configurations to be done for links, verify that it was done.
	Cause code 3 The VP-tunnel is defective.	Ask your VP tunnel provider to test it.
	Cause code 3. PNNI PVC link does not work.	<ol style="list-style-type: none"> 1. See case study 5 in this chapter regarding addressing on PVC links. 2. If the PVC is on a VOID port. Check where a VPC_Link has has been defined for the port. If so, define a VPC_Link using the VP number of the PVC.
Cause code 31. The IP address of the client is not in the same IP subnet as the server.	Change the IP address or IP subnet mask of the client.	
Configuration bad	Setting error	Verify configuration on client and ARP server: <ul style="list-style-type: none"> • In the device action as ARP server, verify that it is in fact configured as an ARP server, check the ATM address, the IP address and the subnet mask and check the registered clients if possible. • In the ARP clients, check the ARP definition making certain that it is a client, then check for the ARP server, the ATM and IP addresses, the subnet mask of the ARP client and the IP address of the default gateway.

<i>Table 11 (Page 2 of 2). Problems of Phase 3 and 4</i>		
Symptom	Possible Causes	Recommended Actions
PING does not work on IP over ATM between the ARP server and a client.	ARP server or client not registered with their respective switch.	Check the ports of both devices to make sure they are enabled and OKAY. Use console commands of the switch(es) to verify that the address of the client registered properly.
	Bad configuration. Device cannot register with ARP server	Verify configuration on client and ARP server. <ul style="list-style-type: none"> In the device action as ARP server, verify that it is in fact configured as an ARP server, check the ATM address, the IP address and the subnet mask and check the registered clients if possible. In the ARP clients, check the ARP definition making certain that it is a client, then check the ARP server ATM and IP addresses, the subnet mask of the ARP client and IP address of the default gateway.
	Cause code 31. The IP address of the client is not in the same IP subnet as the server.	Change the IP address or IP subnet mask of the client.
TELNET does not work.	TELNET session already active and hub or device may only support one session at a time.	Log off the other session.
	Check that PING works between the device and the hub.	If PING fails, refer above to PING does not work on IP over ATM between two ARP clients.
Network management does not work.	No connection between network management station and hub or device to be managed.	Check that PING works between the NM device and the hub. If it fails refer above to PING does not work on IP over ATM between two ARP clients.
	Network management application is not running.	Check the NM application is running in the ATM device.
	Community names incorrect.	Community names in the NM application and the device should be the same and the access level should allow the NM to receive traps from the device.
TFTP does not work.	Check that PING works between the TFTP file server and the hub.	If PING fails, refer above to PING does not work on IP over ATM between two ARP clients.
	TFTP application not active or properly configured on the server.	Check the TFTP application is operational on the device and correctly configured. For UNIX and AIX servers, make sure the directory on the server allows read/write permission.
	The client is not configured correctly.	Check that address of the TFTP is entered correctly and that file paths and file names entered are correct.
	TFTP times out when PING works.	This may be a problem with the TFTP server code. Some TFTP applications have a configurable parameter that will allow the user to extend the time an application will wait for a response before timing out the session.

5.7 Tools Used during ATM Troubleshooting Case Studies

This is an overview of the investigation tools available for ATM problem troubleshooting. ATM being a connection oriented technology, connectivity is provided by the mean of virtual connections (VCs). Troubleshooting ATM level problems generally turns into troubleshooting virtual connections.

5.7.1 Fully Using End Devices Information

Since the problem is almost always perceived from the end device, one should also immediately try to gather the information available from the device experiencing the problem.

- List and use of VCCs from the end device

Some end devices provide information about the VCC currently running as well as the use of the VCC. This is particularly useful in a LAN Emulation environment since there are several VCCs established for each end device.

- Traffic information

Looking at the frames/cells transmitted on a particular VCC may also be useful.

- Event logging and internal traces if available

8210/MSS provides both an event logging facility as well as a trace facility. Since 8210/MSS is the core of most of the ATM and internetworking services provided by the network, it is often the right point to start investigation. Most of the time, this will allow you to focus on a particular communication layer, and start specific investigations on the suspected layer.

- 8285/8260 internal management interfaces

All the 8260/8285 switches need to be configured using Classical IP or LAN Emulation. Network access options are Telnet and network management. The internal management tools are also powerful troubleshooting tools since they give the reason (UNI clear cause code) for call setup failures.

Network initialization usually starts by configuring internal clients. Looking at the internal client status is the first indicator of the network configuration state. It is really useful when these clients try to connect to a server (ARP, LECS or LES), that is located in a different hub. When this is the case, it can provide interesting information on the status of the ATM path between this hub and the one hosting the target server. Following is the result of a show device issued when the internal clients have been configured while the port and route information was not.

```

8285> show device
8285 Nways ATM Workgroup Switch
Name : 8285
Location :
ITSO Cary

For assistance contact :
ITSO Support Team

Manufacture id: VIM
Part Number: 51H4119 EC Level: E59245
Serial Number: 1436
Boot EEPROM version: v.1.5.0
Flash EEPROM version: v.1.5.1
Flash EEPROM backup version: v.1.3.0
Last Restart : 18:44:16 Wed 9 Jul 97 (Restart Count: 8)

A-8285
-----
ATM address: 39.99.99.99.99.99.00.00.99.99.01.01.99.99.99.99.99.99

> Subnet atm: Connection establishment to ARP cleared
  (Cause code : 0x01. See ATM Forum.)
  IP address: 192.2.1.16. Subnet mask: FF.FF.FF.00
MORE... (<L> to display one more line)
> Subnet lan emulation ethernet/802.3
  Abnormal termination: LECS connection cleared.ATM Forum cause 0x1
  Name :""
--More--
8285>

```

Figure 65. 8285 Show Device Command Result

Note: Clear cause codes are reported in hexadecimal, while most of the literature deals with decimal cause codes.

As highlighted in the previous figure, both the Classical IP client and the LAN Emulation client give indication about SVC clear cause. For generic information about clear cause categories, please refer to the next section. For detailed information about clear cause refer to Appendix C, "UNI 3.0-3.1 Cause Maintenance Error Codes" on page 471.

5.7.1.1 Recommendation about Management through LAN Emulation

When using internal LAN Emulation clients for management purposes, *it is recommended* to create a specific so called Management ELAN that will contain all the internal network devices. The connection between the production network and this management ELAN must be done at layer 3, using a *routed* configuration and not a bridged configuration. When using no specific ELAN for management, or when using a bridged configuration, the internal LAN Emulation clients receive all the broadcasts flowing on the production network. These broadcasts are then sent to the ATM control point and use unnecessary CPU cycles. Using a routed configuration shields the network devices from the broadcasts being sent on the production network.

5.7.2 Nways Campus Manager ATM

Nways Campus manager ATM is a *key component* for ATM problem troubleshooting.

The major benefits are provided by:

1. The ease of use

This is a graphical end user interface, and requires no specific skill to operate, as opposed to any tracing tool.

2. A centralized point of control, and network overall view

This is the only central point in the network. Therefore it allows a fast and effective understanding of a given behavior across several nodes.

3. Summarized and detailed VC information by means of:

- High level and overview information for a given node
- Detailed information for any single VC

These three combined properties often result in fast identification of the problem type or problem location. It makes the network management station a powerful tool for investigating ATM level problems, and allows you to save a lot of troubleshooting time.

Following are the main Nways Campus Manager functions providing VC level information. This information is available for three VC categories:

1. Virtual links

A virtual link is an ATM connection running on a non UNI port.

2. Permanent virtual connections

3. Switched virtual connections

The following examples mainly use switched virtual connection-related functions.

There are mainly three key parameters needed when troubleshooting an ATM network:

- Number of SVC active on a given port
- Summary and detailed information about released VCCs
- Information about VC currently running on a given port

5.7.2.1 Number of SVC Active on a Given Port

It is important to understand macroscopic behavior in a large scale network. Therefore, it may be useful if you observe:

- No SVC

This may indicate a signalling problem, or a physical port problem.

- The number of SVCs changing rapidly

This is of course really subjective and network-dependent, but in some cases, it may indicate instability or retries.

- A large number of SVCs established

This information may indicate when a given port approaches its VC limit.

This is provided by the Switch Interface Configuration panel.

Figure 66. ATM Switch Interface Configuration

The ATM Parameters section provides useful information about the number of SVCs currently running on a given port. It also provides the maximum number of VCs and the number of VPIs, number of VCI bits from which you can compute the VPI and VCI range allowed on this port. Once a wrong behavior is suspected, looking at the call logging activity is the next step in problem investigation.

5.7.2.2 Summary and Detailed Information about Released VCCs

Released VCC information is the information to look at when investigating problems related to VC connection setup:

1. Summary information:
 - Call logging

ATM Call Logging									
Navigation									Help
Node IP Address: 192.168.20.99									Reselect...
Filter Criteria									
Filter Option:		None		Filter Value:				Apply Filter	
SVC Log List									
Index	Interface	Calling Number	Called Number	Creation Time		Clear Time		Cause	
2143827017	1601	0020DA6FA8E0-00	400082100000-03	1997/07/10	17:01:52:00	1997/07/10	17:02:39:00	47	
2143827018	1601	0020DA6FA8E0-03	00A03E000001-00	1997/07/10	17:02:27:00	1997/07/10	17:02:39:00	47	
2143827019	1601	08005A990AB3-01	00A03E000001-00	1997/07/10	17:02:34:00	1997/07/10	17:02:38:00	31	
2143827020	1304	08005A990AB3-01	00A03E000001-00	1997/07/10	17:02:34:00	1997/07/10	17:02:38:00	31	
2143827021	1401	000000000000-00	400082100000-04	1997/07/10	17:02:36:00	1997/07/10	17:02:36:00	100	
2143827022	1401	000000000000-00	400082100000-04	1997/07/10	17:02:31:00	1997/07/10	17:02:31:00	100	
2143827023	1601	08005A990AB3-01	00A03E000001-00	1997/07/10	17:02:26:00	1997/07/10	17:02:30:00	31	
2143827024	1304	08005A990AB3-01	00A03E000001-00	1997/07/10	17:02:26:00	1997/07/10	17:02:30:00	31	
2143827025	1401	000000000000-00	400082100000-04	1997/07/10	17:02:26:00	1997/07/10	17:02:26:00	100	
2143827026	1601	08005A990AB3-01	00A03E000001-00	1997/07/10	17:02:18:00	1997/07/10	17:02:22:00	31	
2143827027	1304	08005A990AB3-01	00A03E000001-00	1997/07/10	17:02:18:00	1997/07/10	17:02:22:00	31	
2143827028	1401	000000000000-00	400082100000-04	1997/07/10	17:02:21:00	1997/07/10	17:02:21:00	100	
2143827029	1401	000000000000-00	400082100000-04	1997/07/10	17:02:16:00	1997/07/10	17:02:16:00	100	
2143827030	1601	08005A990AB3-01	00A03E000001-00	1997/07/10	17:02:10:00	1997/07/10	17:02:14:00	31	
2143827031	1304	08005A990AB3-01	00A03E000001-00	1997/07/10	17:02:10:00	1997/07/10	17:02:14:00	31	
2143904558	1304	08005A990AB3-00	00A03E000001-00	1997/07/09	14:19:03:00	1997/07/09	14:19:03:00	31	
2143904559	1601	08005A990AB3-00	00A03E000001-00	1997/07/09	14:19:03:00	1997/07/09	14:19:03:00	31	
2143904560	1401	000000000000-00	400082100000-04	1997/07/09	14:19:00:00	1997/07/09	14:19:00:00	100	
2143904561	701	002035E10580-00	400082100000-02	1997/07/09	14:18:58:00	1997/07/09	14:18:58:00	1	
2143904562	1201	42000000FF18-02	00A03E000001-00	1997/07/09	14:18:57:00	1997/07/09	14:18:57:00	31	
2143904563	1601	42000000FF18-02	00A03E000001-00	1997/07/09	14:18:57:00	1997/07/09	14:18:57:00	31	
2143904564	1201	42000000FF18-02	400082100000-00	1997/07/09	14:18:57:00	1997/07/09	14:18:57:00	31	

Refresh Details... Clear Stop Query

Reset Close Help

Figure 67. ATM Call Logging

Call logging is a key function for SVC problem troubleshooting since it gives an overview of the SVC activity on a given switch. In this example, we can immediately see that some SVCs are cleared with some suspect cause code. A healthy ATM network should not contain SVCs regularly cleared with cause code 100 or 1 as described in this example. It should in fact mainly contain SVCs being cleared with cause code 16, cause code 31 (Normal), or 47 (Resource unavailable). Clear cause 47 can be generated when an end device is powered off.

These rules should, of course, be used with care, but are excellent indicators of misconfigurations.

When SVCs are regularly cleared with a cause code different from 16, 31 or 47, then one must have a look at why this happens.

2. Detailed information:

- Will be provided by the use of call logging filters

Filters available are:

- Called address
- Calling address

These two filters allow you to focus on a particular station.

- Clear cause

This may allow you to identify commonalities between end devices whose SVCs are cleared with the same cause.

- Interface (physical port)

ATM Call Logging						
Navigation						Help
Node IP Address: 192.168.20.99						Reselect...
Filter Criteria						
Filter Option:	Interface	Filter Value:			701	Apply Filter
SVC Log List						
Index	Interface	Calling Number	Called Number	Creation Time	Clear Time	Cause
2143904741	701	002035E10580-00	400082100000-02	1997/07/09 14:16:03:00	1997/07/09 14:16:03:00	1
2143904749	701	002035E10580-00	400082100000-02	1997/07/09 14:15:58:00	1997/07/09 14:15:58:00	1
2143904751	701	002035E10580-00	400082100000-02	1997/07/09 14:15:53:00	1997/07/09 14:15:53:00	1
2143904759	701	002035E10580-00	400082100000-02	1997/07/09 14:15:48:00	1997/07/09 14:15:48:00	1
2143904763	701	002035E10580-00	400082100000-02	1997/07/09 14:15:43:00	1997/07/09 14:15:43:00	1
2143904769	701	002035E10580-00	400082100000-02	1997/07/09 14:15:38:00	1997/07/09 14:15:38:00	1
2143904773	701	002035E10580-00	400082100000-02	1997/07/09 14:15:33:00	1997/07/09 14:15:33:00	1
2143904775	701	002035E10580-00	400082100000-02	1997/07/09 14:15:28:00	1997/07/09 14:15:28:00	1
2143904783	701	002035E10580-00	400082100000-02	1997/07/09 14:15:23:00	1997/07/09 14:15:23:00	1
2143904787	701	002035E10580-00	400082100000-02	1997/07/09 14:15:18:00	1997/07/09 14:15:18:00	1
2143904793	701	002035E10580-00	400082100000-02	1997/07/09 14:15:13:00	1997/07/09 14:15:13:00	1
2143904797	701	002035E10580-00	400082100000-02	1997/07/09 14:15:08:00	1997/07/09 14:15:08:00	1
2143904805	701	002035E10580-00	400082100000-02	1997/07/09 14:15:03:00	1997/07/09 14:15:03:00	1
2143904811	701	002035E10580-00	400082100000-02	1997/07/09 14:14:58:00	1997/07/09 14:14:58:00	1

Buttons: Refresh, Details, Clear, Stop Query, Reset, Close, Help

Figure 68. Call Logging Filtered by Interface

This allows you to focus on the SVC activity for a given port. In this example, we can see that the device connected to port 7.01 has all its SVCs cleared with cause code 1. Cause code 1 means unallocated or unassigned number. This indicates that a device regularly (every 5 seconds) issues a call to this address. This behavior corresponds to a station trying to reach an ATM address (for instance an ARP server) not registered in this switch. There are three main reasons for this to occur in the case of an ARP server:

- The ARP server ATM address has changed, and some stations have not been reconfigured.
 - The station has a misconfiguration in the ESI/selector part of the ARP server ATM address.
 - The ARP server is not operational.
- Call logging details

Navigation		Help	
Node IP Address:	192.168.20.99		
Interface Index:	1401		
Slot.Port:	14.1		
<i>Calling Number</i>			
Network Prefix:	000000000000000000000000		
User Part:	000000000000 00		
<i>Called Number</i>			
Network Prefix:	390985111111111111111110101		
User Part:	400082100000 04		
<i>Time</i>			
Creation Day:	1997/07/09	Creation Hour:	14:17:00:00
Clear Day:	1997/07/09	Clear Hour:	14:17:00:00
<i>Clear</i>			
Clear Cause:	invalid information element contents		
<i>Parameters</i>			
Backward QoS:	unspecified	Forward QoS:	unspecified
Backward BW:	0	Forward BW:	0
<i>Description</i>			
Close		Help	

Figure 69. ATM Call Details

This gives all information about the call being cleared. The key parameter to look at is the clear cause. This allows you to immediately understand the ATM layer that may be misconfigured inside the network. For instance:

- Cause codes related to route

Cause code 3 indicates no route to destination. This clear cause generally occurs when the local switch is not able to reach the network prefix specified in the called address of the destination. Several case studies illustrate when this may happen and how to fix it. This can also reveal an invalid network prefix part of the called address.

- Cause code related to VPI/VCI not available.

This may indicate that a switch is out of resources. This can also occur on an IISP misconfiguration where both sides have the same signalling configuration (network to network for instance).

- Cause code related to Quality of Service (QoS).

This may happen when running applications that use native ATM calls. This is the case for some voice or video applications, even desktop video conferencing applications. When using this type of application in the network, reserved bandwidth parameters need to be carefully looked at. The network management station will be an essential and unique way to get an overview of the dynamic bandwidth allocation over the network. The dynamic bandwidth allocation may lead to transient problems. In this case it is particularly useful to immediately identify the troubles as being related to QoS.

- Cause codes related to information elements.

This may indicate an incompatibility in the UNI version being run along the SVC path (UNI or IISP ports). This can also be a defective end device as shown in Figure 69 on page 151. The clear cause was 100 (invalid information element contents). This illustrates a particular case where the network management also displayed one invalid information element, since the calling address is equal to 0. When facing this category of clear cause, if a fast consistency checking on UNI and IISP ports does not solve the problem, then lower level investigation tools will be needed.

In this situation, the network management limits have been reached, since the level of information required to further troubleshoot this type of problem is the content of the Call setup, and requires investigation at the signalling level. This level of information will be provided by using the 8260 internal traces, or an ATM network analyzer.

Note: The above comments represent generic advice that points to the first things to look at. Reasons given should be considered as most probable or common failure causes, definitely not as the only possible reasons for the failures.

5.7.2.3 Information about VC Currently Running on a Given Port

1. Summary information:

- The ATM SVC list panel

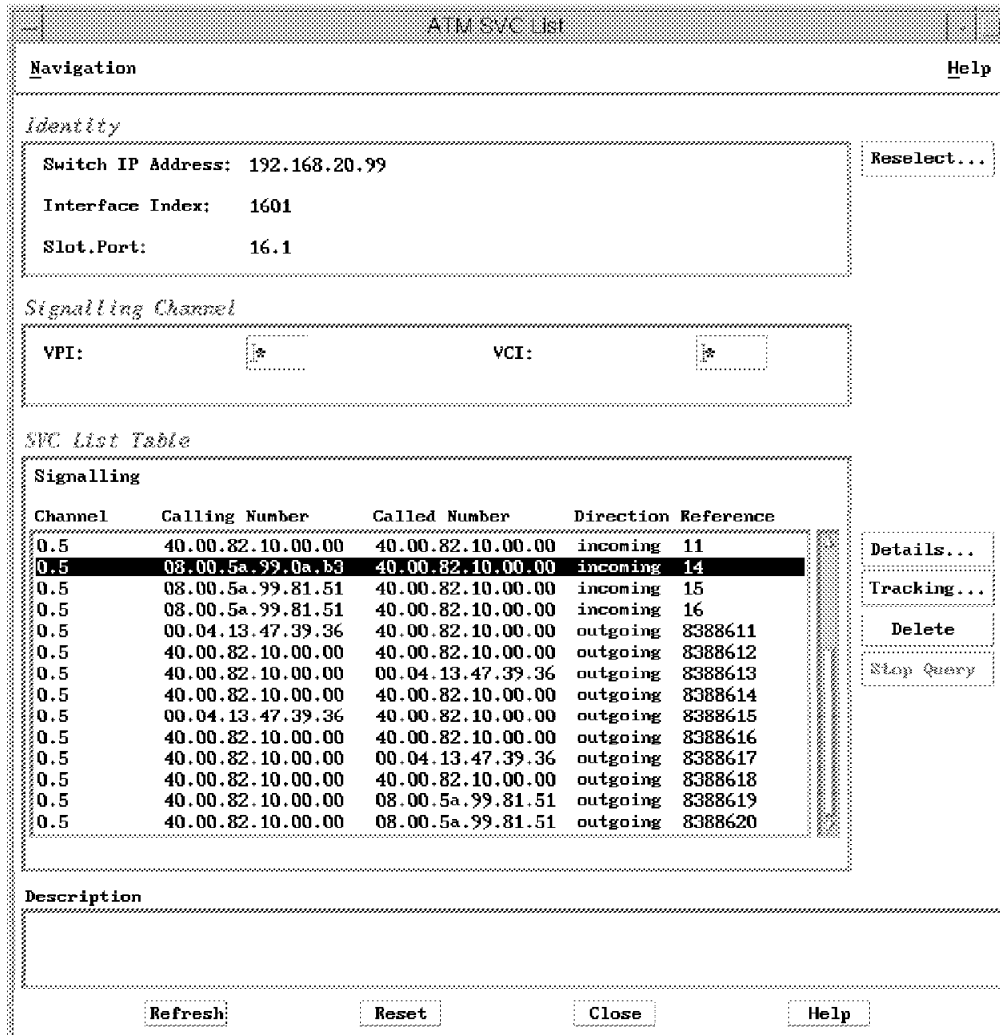


Figure 70. ATM SVC List

This lists all the SVCs running on the port. It immediately allows you to identify calling number and called number. The reference is also interesting for immediate detection of point-to-multipoint VCCs. Point-to-multipoint VCCs will have several entries in the SVC list with the same reference (the call reference) value. From this panel, you can access two more panels for more information about a given VCC. It also allows you to clear a particular VCC in the list. This allows surgically clearing SVCs and does not affect any other connections. A port disable or fiber unplug would have the same effect, but at the port level, which may not be always possible in a production network.

2. Detailed information:

- SVC details panel

ATM SVC Details			
Navigation	Help		
<i>Identity</i>			
Switch IP Address:	192.168.20.99		
Interface Index:	1601		
Slot.Port:	16.1		
<i>Selection</i>			
Signalling Channel:	0.5	Call Reference:	8388625
VPI:	0	VCI:	746
<i>Direction</i>			
SVC Direction:	outgoing		
<i>Calling Number</i>			
ATM Address Network Prefix:	DCC/DFI/AA=0985/11/111111 RD=1111 AREA=01.01		
ATM Address End System:	ESI=00.04.13.47.39.36 SELECTOR=04		
<i>Called Numbers</i>		<i>/Creation</i>	
DCC/DFI/AA=0985/11/111111 RD=1111 AREA=01.01 ESI=40.00.82.10.00.00 SELECTOR=05			
<i>Parameters</i>			
Receive Direction		Transmit Direction	
Type:	Best-Effort	Type:	Best-Effort
Service:	unspecified	Service:	unspecified
Parameters		Parameters	
no parameter		no parameter	
<i>Description</i>			
Refresh		Close	
Help			

Figure 71. ATM SVC Details Panel

- This window allows you to see all the characteristics of a given live VCC
- SVC tracking panel

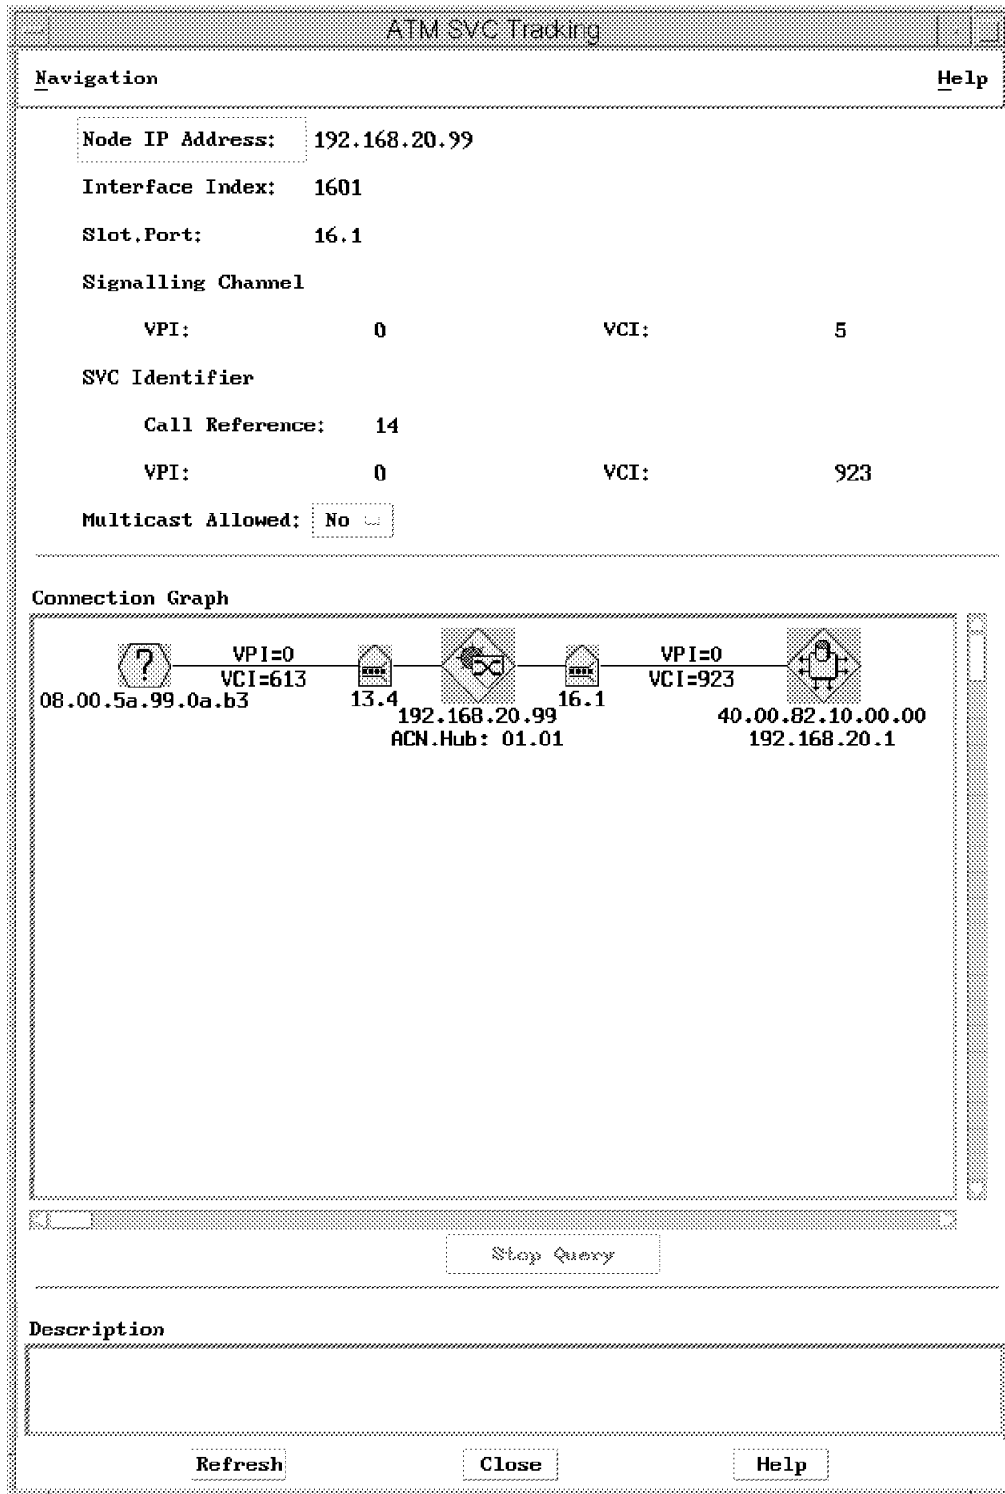


Figure 72. ATM SVC Tracking Panel

This window allows you to see the route used by the VCC, as well as the intermediate VPI/VCI used by the connection. This latter information can be extremely useful when one wants to analyze the data flowing on a particular VCC and the network analyzer is plugged in one of the backbone links inside the network. It is unlikely that you can unplug the analyzer and move it to the port you want to track. This link may have a

large number of VCCs running and it may be difficult or simply impossible to locate the right connection. It becomes obvious and immediate if you correlate with the network management end-to-end view.

As already mentioned, a unique value add of network management stations compared to other debugging tools is the ability to quickly monitor and analyze what happens inside the network regarding connections. As a general rule, one should use the functions described in this paragraph as soon as an ATM level problem is suspected. This may save a considerable amount of time in the troubleshooting work.

5.7.3 8260 Internal Traces and Dumps

The internal 8260 internal traces and dump should be used when the trouble is suspected to be related to a given hub, or when the level of detail needed can't be provided by the higher level tools already mentioned. There are selective traces and dumps available. Since the trace may wrap, it is highly recommended to start the trace only on the components suspected to be in relation to the problem.

5.7.4 An ATM Network Analyzer

Following is an example of the functionality exercised when one uses an ATM network analyzer for pure ATM troubleshooting purposes:

- Capture facilities

Depending on the buffer size of the analyzer and the traffic on a given link, the user is able to capture a reasonable amount of data. The InterWATCH 95000 allows both real-time decoding as well as post capture processing. Using these two methods, one can make sure the capture of a given event or network behavior occurred, and fully process and filter the captured data later on.

- Extensive decoding of all ATM layers

Network analyzers provide decoding from the cell level up to ATM services levels: Classical IP, LAN Emulation, etc.

- VC content analysis, up to application level

In a pure ATM environment, this is a key investigation tool when a protocol problem is suspected.

- Signalling emulation

This is the ability to generate customized ATM calls. This is an interesting functionality since it gives direct access at the VC level , and gets rid of all the upper layer overhead (aging time, caches, etc.) This may also be interesting for signalling stress testing, QoS testing and so on.

An ATM network analyzer mainly deals with VCs. It is therefore important to use other external devices or tools to provide focus on the right VCs. This is mandatory when doing analysis inside backbone links, since the number of VCs running on these links may be so large that it may take a while to find out what you are looking for. As already mentioned, an ATM network analyzer is a perfect complementary tool to the network management station. The network

management station provides summary and network wide information, while the network analyzer provides extensive information on a given point in the network.

5.8 ATM Campus Problems Scenarios

The following scenarios are related to ATM connectivity problems, but do not directly involve LAN emulation, bridging or routing configuration errors.

5.8.1 Case Study 1 - UNI Translation Problem

This scenario illustrates a case where an SVC connection failed due to an incompatibility in UNI message translation between two endstations. There may be many other symptoms when this type of problem occurs in a network. The objective of this scenario is to show a typical approach for this type of problem investigation.

5.8.1.1 Network Diagram

This is a simple configuration involving two endstations.

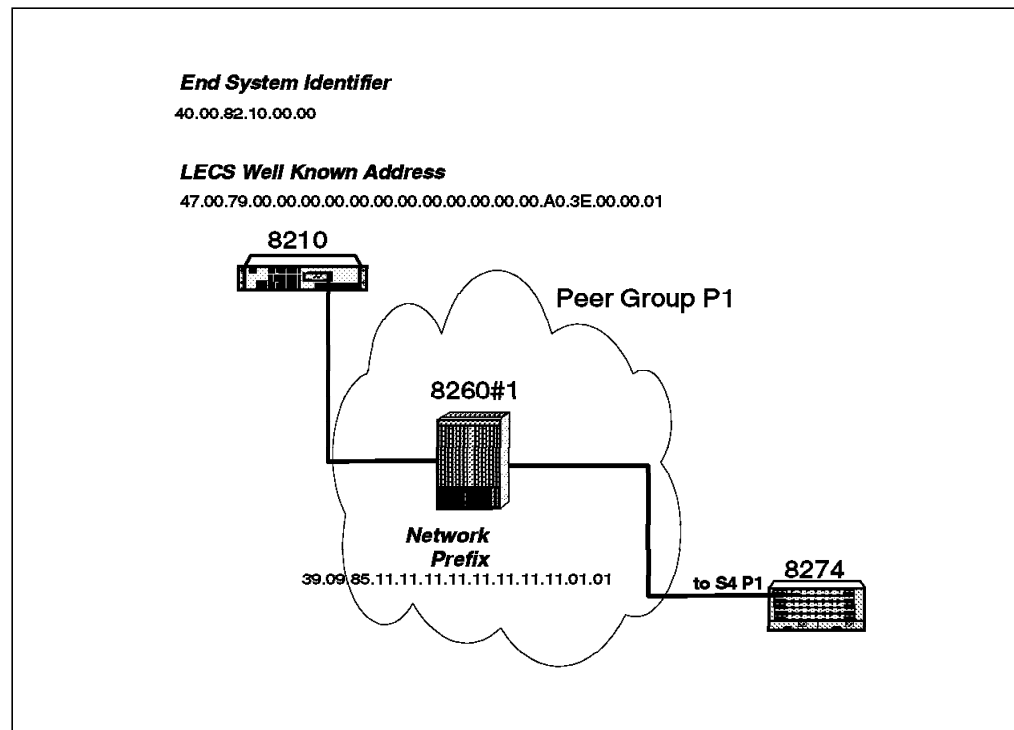


Figure 73. Network Configuration

5.8.1.2 Symptom

Inability of the 8274 LEC to join its emulated LAN.

Checking of all the LAN emulation parameters as recommended in Chapter 7, "ATM Emulated LANs and Logical IP Subnets (LANE 1.0, RFC 1577)" on page 293 did not reveal any misconfiguration. All the LAN emulation parameters are correctly set; the LAN emulation services are up and running. Other LAN emulation clients successfully joined the target ELAN.

5.8.1.3 Methodology

- Investigation starts at the 8274.

Since the 8274 LEC cannot register to the ELAN, let's have a look at some more detailed information related to LAN emulation status. Figure 74 shows the output of the VSS command.

```

/ % vss 4/1 1

Status/Statistic for slot 4 interface 1 Service 1

Service      : LAN Emulation Service 1

LEC status  : Join
ELAN Name    : ETHERNET_ELAN
ELAN Type    : 802.3
LEC ID       : 5
LES address  : 39098511111111111111111010140008210000002 (learned)
BUS address  : 39098511111111111111111010140008210000002
LECS address: 4700790000000000000000000000a03e00000100 (use well-known LECS addr)

BUS
MC Forward VPC/VCC : 0/ 0      MC Send VPC/VCC      : 0/ 0
Echo suppress      :      0

LES
Control Direct VPC/VCC: 0/ 0      Cntl Distribute VPC/VCC: 0/ 0
Control Frames Sent : 1772      Control Frames Rcvd : 1160
LE arps Sent       :      2      LE arps Received     :      0

LECS
Configuration VPC/VCC : 0/ 1
Packets Sent         :      0      Packet Received       :      0

```

Figure 74. Result of the VSS Command on 8274

Figure 74 provides three interesting observations:

- The 8274 cannot reach the LANE operational state. (LEC status is Join.)
 - The LEC learned a LES address and in this case it is the correct one.
 - The LEC is in Join state, and has sent control frames to the LES. This means that it can reach the LES but something may prevent completion of LAN emulation initialization.
- The network management view

The call logging function of Nways Campus Manager ATM revealed an abnormal signalling activity between the ports on which MSS and 8274 were connected. There was a periodic and regular clearing of SVC with a cause 31 (Normal) between these two ports. Even if the clearing cause did not bring much information, this illustrates an unstable state between the 8274 and MSS. This also shows that one of the component entered a retry cycle.
 - Investigation From MSS

The 8274 LEC is not displayed in the list of LECs that joined the ELAN. The next step to perform further problem determination is to look at the MSS Event Logging System. Figure 75 on page 159 is the result of the MSS event logging started on the LES subsystem.

```
LES.087: LES/BUS:'ETHERNET_ELAN'  
:Ctrl Dir estblshd,  
    Calling ATM addr = x39098511111111111111101010020DA6FA8E000  
LES.172: LES/BUS:'ETHERNET_ELAN': adding Proxy Ctrl Dist leaf,  
    LEC ATM addr = x3909851111111111111111101010020DA6FA8E000  
LES.116: LES/BUS:'ETHERNET_ELAN': trmntng LEC:  
    err adding Proxy Ctrl Dist leaf:cause 31  
    LEC ATM addr = x3909851111111111111111101010020DA6FA8E000
```

Figure 75. Event Logging System for LES Subsystem

This output indicates that the LAN emulation initialization failed for this particular LEC. The reason seems to be that the ADD PARTY fails with UNI clear cause code 31. To further investigate, it is required to trace the SVC activity. The easiest way to do this is to start the Event Logging System for the SVC subsystem and filter on VPI=0, VCI=5 to get only the signalling messages. An alternative is to start an 8260 internal trace.

Figure 76 shows the result of the MSS ELS for the SVC subsystem.

```
SVC.015: Enter function leaf: add_leaf, hndl,state,count,57,0,2  
SVC.025: leaf: send succeeded, msg= 09030000 03808000 54588000 09058C05  
SVC.015: Enter function UpdateLeafStateAndCount: old state,new state,count,0,1,2  
SVC.013: Enter function UpdateLeafStateAndCount: conn hndl= 9  
SVC.018: Exit UpdateLeafStateAndCount: state, count,1,3  
SVC.017: Exit add_leaf, rc= 0  
SVC.025: received data = 09038000 03828000 0D548000 03008039  
SVC.002: find_call=17064308  
:hp2.SVC.020: Received Add Party Reject,conn hndl=9,ID=3,state=10:ehp2.  
SVC.015: Enter function leaf: add_party_rej, hndl,state,count,57,1,3
```

Figure 76. Output of the Event Logging System for SVC Subsystem

This clearly highlights a problem at the ATM level since the ADD PARTY message required to add this LEC to the point-to-multipoint VCC (control distribute VCC) is rejected by the switch.

- Using 8260

Commands available from a user interface show a normal status at ATM level. Further investigation requires capturing an 8260 internal trace. The output of this trace is shown in Figure 77.

```
20:57:21 SVC S=09 P=0 TX ADD-PARTY-REJ EPR=8066h CAUSE(1)=31(d)
```

Figure 77. Add Party Clear Cause

The 8260 internal trace shows the ADD party reject with reason 31 (Normal, unspecified).

- Using a protocol analyzer

The protocol analyzer also would show that the ADD PARTY was rejected by the ATM switch.

5.8.1.4 Understanding the Problem

1. The LAN emulation parameters are correct. The investigation shows a problem at the ATM layer.
2. This is probably not an ATM connectivity or ATM routing problem since the Control Direct VCC with the LES is established. However, setup fails for the Control Distribute.

Given these elements, we can surmise that the problem may be related to signalling. The first parameter to check in this case is the UNI version running at both ends as well as potential UNI translation occurring in the path between the LEC and the LES.

To identify the real cause that generated this ADD PARTY reject, it is necessary to look into signalling messages, especially the content of the ADD PARTY message. The AAL parameters Information Element of the ADD PARTY message is not the same for UNI 3.0 and UNI 3.1. The Mode parameter is not present in the 3.1 signalling while it is required in the UNI 3.0 message.

Note: Recent improvement in the switch code for the 8260 and 8285 now allows this configuration to work. The Mode parameter is added automatically as part of UNI translation if the connection is LAN Emulation or Classical IP.

5.8.1.5 Conclusion

In this scenario, suppressing UNI translation will solve the 8274 LEC connection problem. This is done by configuring both UNI ports with the same UNI version and checking that they remain the same once the ports are up and running. Checking is recommended since an auto-detect configuration may have set one port to UNI 3.0 and the other to UNI 3.1. It should be noted that UNI translation problems may also be indicated when an SVC is cleared with clear cause 100 (Invalid Information Element contents).

Figure 78 shows the result of the MSS ELS for LES subsystem captured in the configuration described in this paragraph but running different software versions on the various devices.

```
LES.087: LES/BUS:'ETH1'  
:Ctrl Dir estblshd,  
      Calling ATM addr = x3999999999990000999901020020DA7209D002  
LES.168: LES/BUS:'ETH1':plcng Proxy Ctrl Dist call,  
      LEC ATM addr = x3999999999990000999901020020DA7209D002  
LES.109: LES/BUS:'ETH1':trmntng LEC: Proxy Ctrl Dist call fld: cause 100
```

Figure 78. Output of the Event Logging System for LES Subsystem

Note: As a general rule, UNI translation should be avoided whenever possible.

5.8.2 Case Study 2 - Reachable Address Missing after Migration to PNNI

PNNI offers new facilities, but don't forget the associated specific settings.

5.8.2.1 Network Diagram

This scenario used two peer groups (P1 and P2).

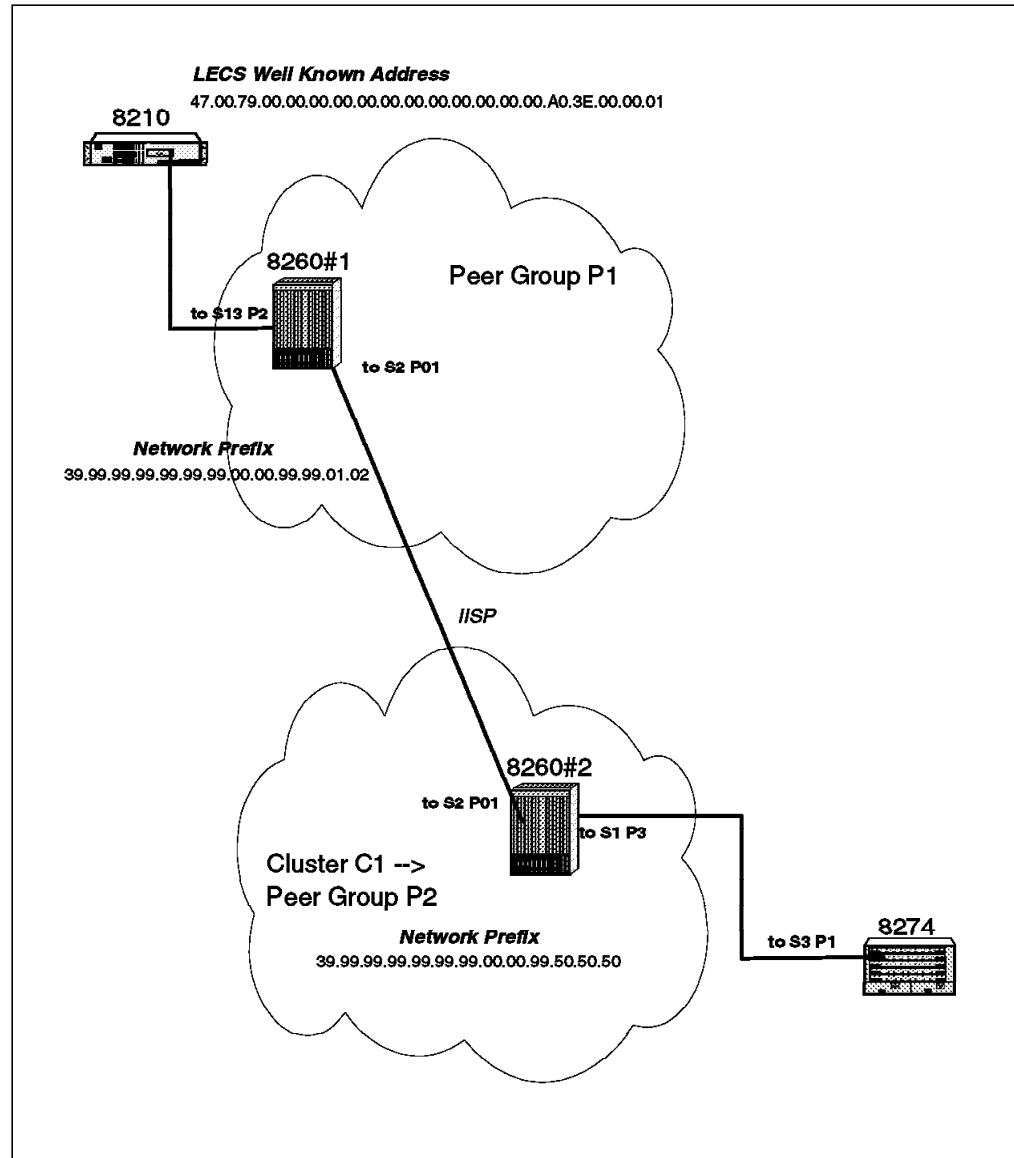


Figure 79. Network Configuration

5.8.2.2 Symptom

After migration of some ATM switches to PNNI, some devices do not reconnect to the emulated LAN. After migration, the user has reconfigured the static routes as they were set in a pre-PNNI configuration.

The problem only affects:

- The devices that use the LECS well-known address.
- Devices using ILMI to get LECS address are able to join the emulated LAN with no problem.

Further investigation shows that when the devices are configured with the real ATM address of the LECS, it solves the problem. Figure 80 shows the example of an 8274 that cannot reestablish LAN emulation connections.

5.8.2.3 Methodology

Investigation can be made from various locations in the network.

From the 8274

```

/ % vas

                                ATM Services
Slot Port Serv      Service      Service
  Num  Num  Num      Description  Type
====  ==  ==  =====
3    1    1    PTOP Bridging Service 1    PTOP 1483
3    1    2    LAN Emulation Service 2    LANE
3    2    1    PTOP Bridging Service 1    PTOP Priv

                                ATM Services
Slot Port Serv VC   Oper
  Num  Num  Num  Typ Status SEL Groups Conn VCI's/Addresses
====  ==  ==  ==  =====
3    1    1    PVC Enabled N/A 3      69
3    1    2    SVC LECS Con 02 2
3    2    1    PVC Disabled N/A 1      100
    
```

Figure 80. Output of the vas Command

The operational status of LECS Con indicates that the 8274 cannot reach the LECS. Note also that no VCs are reported as connected.

From MSS: Start the ELS display for the LECS subsystem. This display does not indicate any activity from this particular client.

From the Network Management: The LAN Emulation Configuration VCC setup seems to fail. The call logging filter applied on interface 103 (Slot 1 Port 3) of 8260#2 reports calls to the LECS well-known address

47.00.79.00.00.00.00.00.00.00.00.00.00.A0.3E.00.00.01

that are cleared with cause code 3 (No Route to Destination).

From the 8260: Starting an internal trace on the 8260 to which the 8274 is connected shows the same result.

Using a Network Analyzer: A network analyzer shows the same result.

Understanding of the Problem: The 8260#2 does not know how to route the call to LECS.

5.8.2.4 Conclusion

Understanding what has changed when migrating to PNNI will shed some light on this problem.

Switch Behavior in a Pre-PNNI Environment: In a pre-PNNI environment the 8260 will not register a foreign address via ILMI. A foreign address is an address whose network prefix is different from the network prefix of the switch itself. The LECS well-known address is a foreign address, and the 8260 does not allow it to register. In order to run a LAN emulation environment the ATM switch needs to be able to process the calls to this LECS WKA. This is done by setting the real LECS ATM address as described in Figure 81.

The new ATM address will be substituted to the well known address. The actual 8260 command to set the LECS address is:

```
set lan_emulation configuration_server active_wka atm_address
```

```
8260ATM#2> set lan_emul configuration_server
Enter ATM address : 39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.00.69.00
```

Figure 81. 8260#2 Set LAN Emulation Configuration Server before Migration

When a LEC gets the LECS address using ILMI, it will get all the addresses configured as LECS addresses, as described in Figure 82. (In this particular case, only one address has been configured.) The LEC then issues a call setup to LECS using the LECS real address.

```
8260ATM#2> show lan_emul configuration_server
Index          ATM address
-----
1 WKA active 39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.00.69.00
```

Figure 82. 8260#2 Show Configuration_server Configuration before Migration

When a LEC uses only LECS well-known address, it issues a call setup with the following called party:

```
47.00.79.00.00.00.00.00.00.00.00.00.00.A0.3E.00.00.01
```

The 8260 will then forward the call to the address declared as **active** in the LECS list. If the LECS is located in another cluster, then the call is forwarded using logical link and static route information (shown in Figure 83 on page 164) to the remote destination.

```

8260ATM#2> show logical_link
Port Vpi Acn Side Mode_Sig Traf Bwidth Status Index
-----
2.01 0 02 user enab 3.0 NRB 0 UP 2

8260ATM#2> show static_route
Index Acn Static route
-----
1 02 3999

```

Figure 83. Links and Routes Configuration for 8260#2 before Migration

Since the call is forwarded with called party unchanged, each switch along the path to destination will use the same method to carry the call up to destination. It can be noticed that there is no need to configure a route to the LECS well-known address.

Switch Behavior in a PNNI Environment: In a PNNI environment, the 8260 allows registration of foreign addresses (*) as illustrated in Figure 84). This means that these addresses are now routed and advertised as any other address in the network.

```

8260ATM#1> show reachable_address all
Port Len Address Active Idx VPI
-----
2.01 104 39.99.99.99.99.99.00.00.99.50.50.50 Y 1 0
4.01 152 39.99.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.00.69 Y Dyn 0
4.01 152 39.99.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.20.69 Y Dyn 0
4.01 152 47.00.79.00.00.00.00.00.00.00.00.00.00.00.00.00.A0.3E.00.00.01 Y Dyn 0 *
6.02 152 39.99.99.99.99.99.99.00.00.99.99.01.02.40.00.82.81.00.66 Y Dyn 0
6.03 152 39.99.99.99.99.99.99.00.00.99.99.01.02.40.00.05.70.00.51 Y Dyn 0
13.02 152 39.99.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.00.72 Y Dyn 0
13.02 152 47.00.79.00.00.00.00.00.00.00.00.00.00.00.00.00.A0.3E.00.00.01 Y Dyn 0 *

```

Figure 84. Foreign Address Registration Example

The 8260 no longer forwards the call to the active LECS. Instead it performs route selection on the LECS WKA. This is reflected in the 8260 by the suppression of the active attribute for configuration server addresses as illustrated in Figure 85.

```

8260ATM#2> show lan_emul configuration_server
Index ATM address
-----
1 39.99.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.00.69.00

```

Figure 85. 8260#2 Configuration_server Configuration after Migration

The LECS will be called at this ATM address only if a LEC uses the ILMI method to get the LECS address. Therefore, it is now necessary to configure routing information for the LECS WKA if the LECS is located in a different peer group. This is done by adding a reachable address to the port used to connect to the LECS as illustrated in Figure 86 on page 165.

```
8260ATM#2> set reachable_address
Enter slot: 2.01
Enter prefix length: 16
Enter reachable address : 47.00
Entry set.
8260ATM#2> show reachable_address all
```

Port	Len	Address	Active	Idx	VPI
2.01	96	39.99.99.99.99.99.99.00.00.99.99.01.	Y	1	-
2.01	16	47.00.	Y	2	-
1.03	152	39.99.99.99.99.99.99.00.00.99.50.50.50.40.00.03.60.00.37	Y	Dyn	0

Figure 86. Route Configuration Needed for 8260#2 after Migration

5.8.3 Case Study 3 - PVC Setup Error

The purpose of this scenario is to explain how to identify PVC problems and describe how to solve the most common PVC misconfigurations on an 8260 running with a PNNI level code.

5.8.3.1 Network Diagram

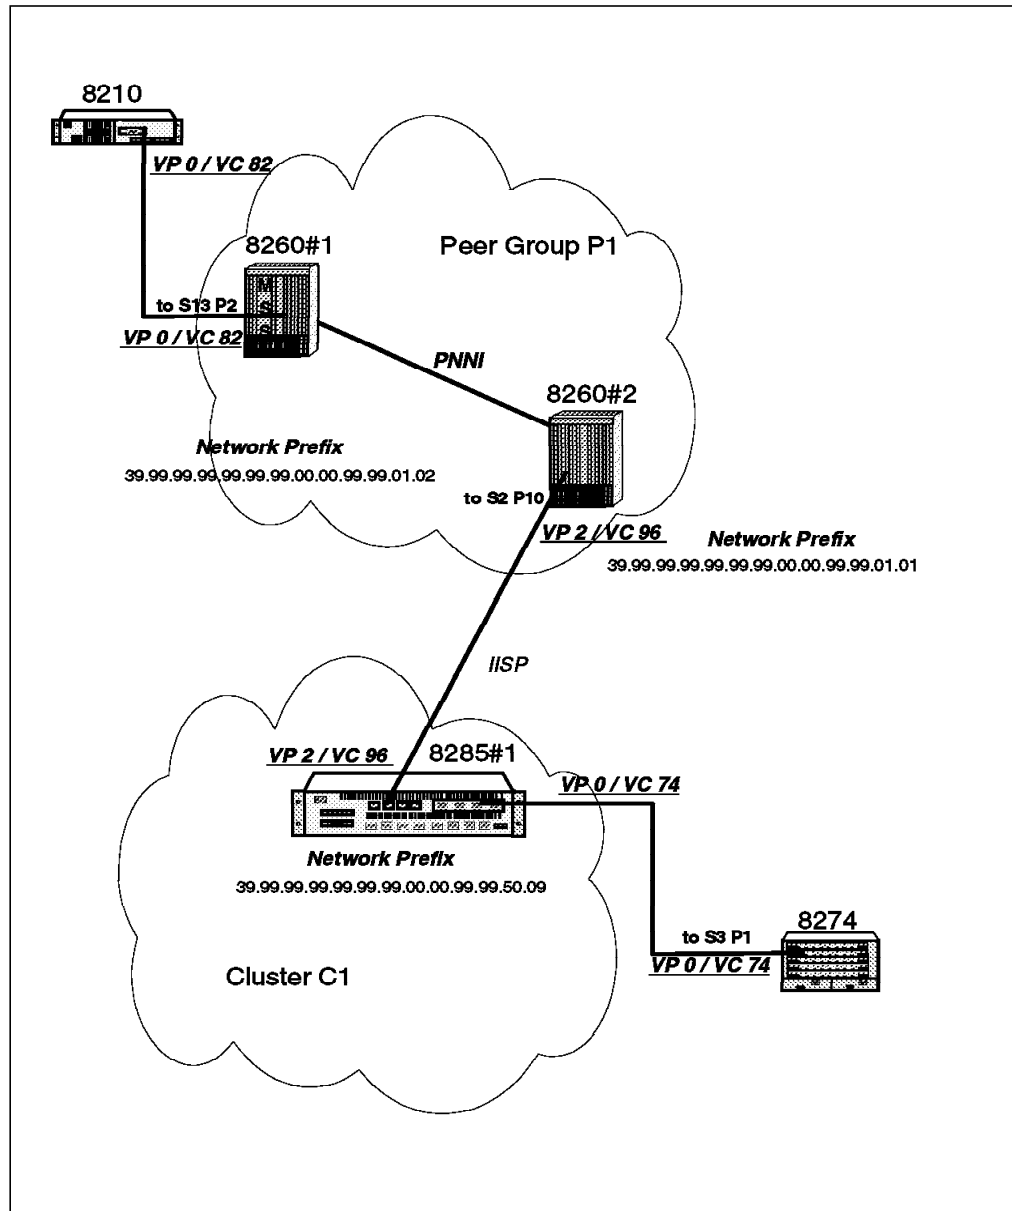


Figure 87. Network Configuration

The 8260 supports soft PVCs. This means that the definition of the PVC is done only at the Ingres and egress points of a group of switches. A group of switches is called a cluster if the switches are running pre-PNNI routing; it is called a peer group if the switches are running PNNI.

A soft PVC cannot be defined across an IISP link. Crossing an IISP link requires that a PVC be defined to one side of the IISP link, and a different PVC be defined starting on the other side of the IISP link.

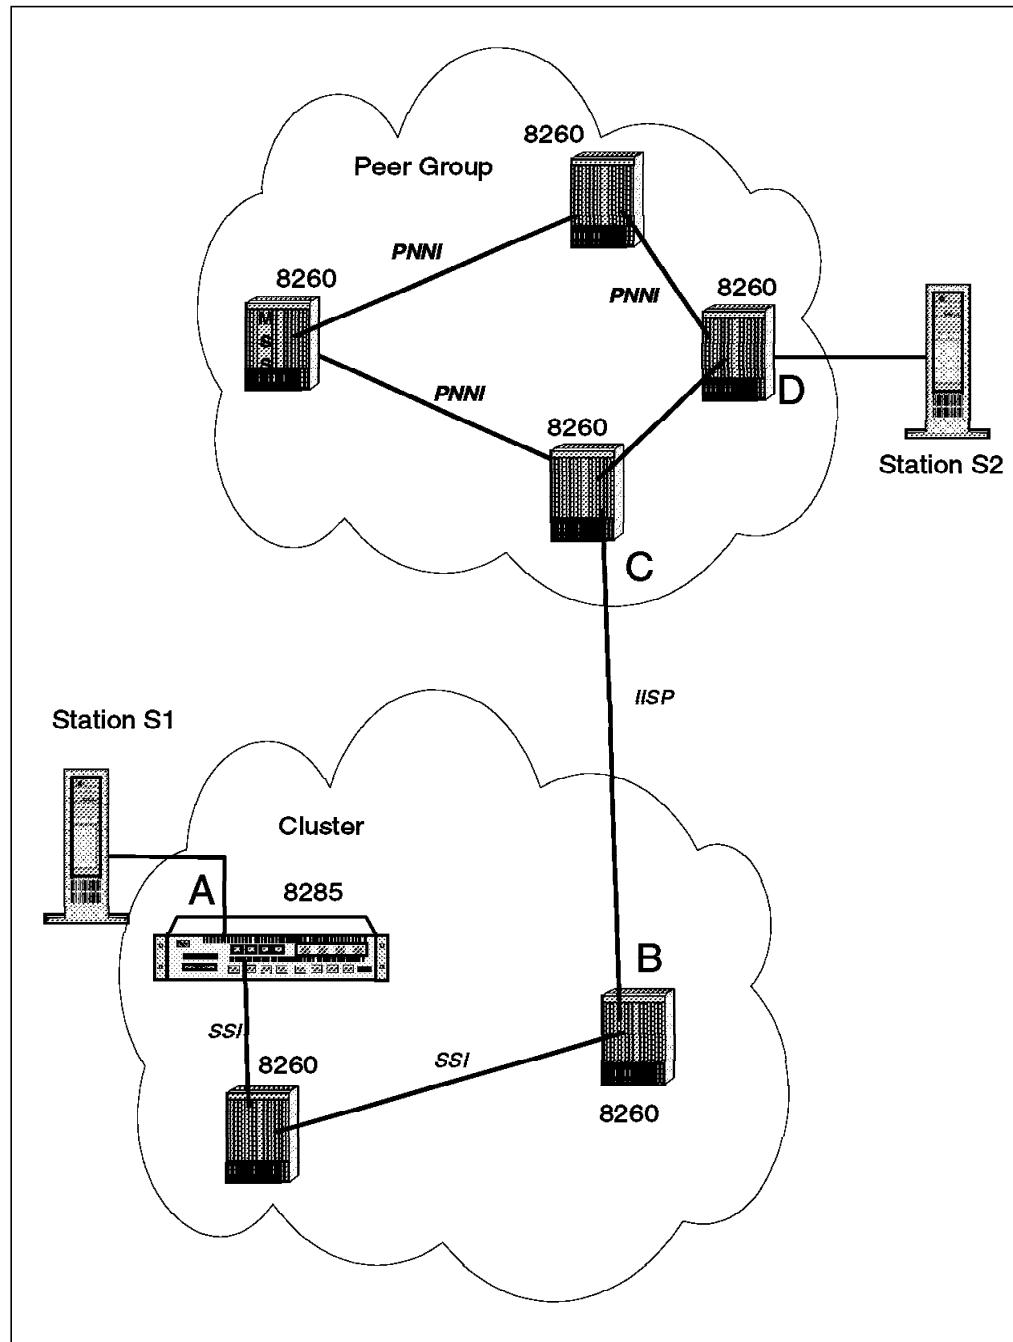


Figure 88. PVC Over IISP Links

As illustrated in this configuration, in order to define a PVC from Station S1 to Station S2, two PVC must be created :

1. One PVC from point A to point B
2. One PVC from point C to point D

Make sure that the Vpi and Vci values match for point B and point C.

5.8.3.2 Symptom

Refer to Figure 87 on page 166. The bridging between the 8274 and the MSS using RFC1483 does not work. Figure 89 displays the 8210 bridging port status.

```
ASRT>list transparent configuration
Filtering database size: 3090
Aging time: 300 seconds
Aging granularity 5 seconds

Port Interface      State  MTU
 1 AT/0:0:82      Enabled 9234
 4 Eth/1          Enabled 1516

ATM Interface+
ASRT>list bridge
Bridge ID (prio/add): 327/40-82-10-14-83-83
Bridge state: Enabled
UB-Encapsulation: Disabled
Bridge type: SR-TB
Number of ports: 4
STP Participation: IEEE802.1d on TB ports and IBM-8209 on SR ports

Port Interface      State  MAC Address      Modes  Maximum
MSDU  Segment  Flags
 1 AT/0:0:82      Up    00-00-00-00-00-00  T      9234      RD
 2 TKR/0          Up    02-00-41-08-80-4E  SR      4544     101      RD
 3 TKR/1          Up    02-00-41-08-C0-4E  SR      4544     401      RD
 4 Eth/1          Up    40-82-10-14-83-00  T      1520      RD

Flags: RE = IBMRT PC behavior Enabled, RD = IBMRT PC behavior Disabled

SR bridge number: 1
SR virtual segment: 272
Adaptive segment: 372
ASRT>?
```

Figure 89. 8210 Bridging Port Status

These results show that for the 8210 side, the bridge port is Enabled, Up and is using VPI:0 VCI:82.

Figure 90 on page 169 displays the 8274 ATM port status.

```

/Interface/ATM % vap
                        ATM Port Table

Slot Port      ATM Port Description      Conn Tran  Media UNI Max  VCI
=====
3   1   ATM PORT                          PVC  STS3c  Multi Pri 1023 10
3   2   ATM PORT                          PVC  STS3c  Multi Pri 1023 10

Slot Port      ATM Network Prefix      End System  Sig Sig  ILMI  ILMI
=====
3   1           N/A                          N/A      N/A N/A  N/A  N/A
3   2           N/A                          N/A      N/A N/A  N/A  N/A
                        Status
-----

Slot Port Tx Seg Sz Rx Seg Sz Tx Buff Sz Rx Buff Sz Oper  SSCOP ILMI
=====
3   1           8192   8192   4600   4600 Enabled Down Down
3   2           8192   8192   4600   4600 Disabled Down Down

/Interface/ATM % vas

                        ATM Services

Slot Port Serv  Service  Service
      Num  Description  Type
=====
3   1   1   PTOP Bridging Service 1   PTOP Priv
3   1   2   rfc1483                PTOP 1483

                        ATM Services

Slot Port Serv VC  Oper
      Num  Typ  Status  SEL Groups  Conn VCI's/Addresses
=====
3   1   1   PVC Enabled  N/A 1   100
3   1   2   PVC Enabled  N/A 2   74

```

Figure 90. 8274 Port Status

From this output, we see that from the 8274 side, the port is configured for PVC; it is Enabled and it is using VPI:0 VCI:74.

5.8.3.3 Methodology

It is now necessary to find out what is the failing communication layer between these two devices. Since the bridging level seems to be correctly defined, a check of the ATM layer is necessary. The check of the ATM layer can be done by gathering information on the frame traffic on the physical interface of each device.

From 8210: Figure 91 on page 170 shows that the VPI:0 VCI:82 is sending frames, but not receiving any data.

```

ATM Interface+list vcc
***** VCCS *****

```

Conn Handle	Conn Type	VPI	VCI	Frames Transmitted	Frames Received	Bytes Transmitted	Bytes Received
6	P-MP	0	761	25	0	2700	0
5	P-P	0	760	3	3	324	324
--More--							
4	P-P	0	759	1	1	44	68
1	SAAL	0	5	137	137	5192	4812
0	N/A	0	82	38	0	2423	0
2	ILMI	0	16	43	43	2346	2748

Figure 91. 8210 ATM Interface Status

From 8274: Figure 92 is the 8274 VC statistics view.

```

/Interface/ATM % vcs

```

ATM Connection Statistics								
Slot	Port	VCI	Rx SDUs	Tx SDUs	Rx Cells	Tx Cells	Rx Octets	Tx Octets
3	1	74	0	18	0	36	0	1728
3	1	100	0	817	0	2012	0	96576

Figure 92. 8274 ATM Statistics

Notice that the VPI:0, VCI:74 is sending frames, but not receiving any. It is now obvious that the ATM layer is not working as expected since both ends are sending frames on their respective VPI/VCI, but neither end ever receives any data. Some investigation must now be done at the PVC level. This will be done by collecting information from the switches inside the network.

Note: For some other type of traffic (video for instance), a quick check of the activity LED on the switch port, whenever possible, will help verify whether there is traffic on the link.

Investigation at 8260 Level: We now take a closer look at the status of the PVC along the path. Detailed status is gathered as described in Figure 93 on page 171.


```

8260ATM#1> show pvc all verbose
-----
PVC: Port 13.02 (id=86,Primary,BE) PTP-PVC VP/VC=0/82
-> Party: (id=0) VP/VC=2/96 STATUS: Failed
   39.99.99.99.99.99.00.00.99.99.01.01.42.00.00.00.02.10.00
Best Effort.
Frame discard : Yes.
Last Active Date : 09:47:56 20 Jun 97 (0 failures)
Status Cause : PVC failure.
   Q93B Cause : 1 (0x1).
    
```

Figure 93. PVC Status: Failed

This PVC has the status Failed with a UNI cause 01. This cause indicates that the remote ATM address is not defined or configured.

The problem in this configuration is that the remote port configured in the PVC does not exist in the remote device. This configuration defines the remote port as being in slot number 2, port number 16 (0x10), instead of in slot number 2, port number 10 as expected.

Understanding PVC addressing will explain why. The address structure used for a port located in a remote hub is represented in Figure 94.




Figure 94. Addressing Rule for a Soft PVC Configuration

- 1** is the network prefix of the remote hub (13 bytes)
- 2** are 4 fixed and mandatory bytes used to build the address.
- 3** is the remote slot (in hexadecimal format).
- 4** is the remote port (in hexadecimal format).
- 5** is the selector and must be set to 0.

Therefore, the correct command required to set a PVC with the remote port number 10 on a module located in slot number 2 in the network configuration described above is:

```
8260ATM#1> set pvc 13.2 86
39.99.99.99.99.99.00.00.99.99.01.01.42.00.00.00.02.0A .00
channel_point_to_point 0.82 2.96 best_effort
PVC set and started.
```

Figure 95. The Correct Command for PVC Setup

After the configuration is changed we can verify the current PVC status:

```
8260ATM#1> show pvc all verbose
-----
PVC: Port 13.02 (id=86,Primary,BE) PTP-PVC VP/VC=0/82
-> Party: (id=0) VP/VC=2/96 STATUS: Active
    39.99.99.99.99.99.00.00.99.99.01.01.42.00.00.00.02.0A.00
Best Effort.
Frame discard : Yes.
Last Active Date : 09:54:14 20 Jun 97 (0 failures)
```

Figure 96. PVC Status: Active

The connection does not work as expected and all the PVCs are reported as Active.

A visual check of the activity LEDs on the ports involved in the desired path may indicate that traffic comes in a particular port but is not forwarded to the expected output port.

If several soft PVCs are involved in the path, check that the VPI/VCI assignments are consistent on both sides of a physical connection between two soft PVCs.

From the Network Management: The network management station may help in solving transient PVC problems.

Two functions can help:

- PVC tracking
This function is used to track a given soft PVC and show the physical path used by the connection inside a given peer group.
- PVC restart
This function may be needed when the PVC has been working, but is no longer active. Restarting the PVC from the network management station may bring it back to operational state.

Using a Network Analyzer: A network analyzer can be useful in troubleshooting a PVC problem, since it gives the ability to check that the network is really configured as expected. It also allows the troubleshooter to identify if the problem is related to the ATM layer or if it is related to upper layer misconfiguration, resulting in a problem such as a difference in frame encapsulation.

Problem Corrected: After resolving the PVC misconfiguration, we get new status results from the physical interfaces.

Figure 97 on page 173 displays the new status from the MSS side.

```
*t 5
ATM Interface+list vcc
```

Conn Handle	Conn Type	VPI	VCI	Frames Transmitted	Frames Received	Bytes Transmitted	Bytes Received
--More--							
4	P-P	0	808	3	3	132	204
1	SAAL	0	5	238	239	6492	6768
0	N/A	0	82	6172	5587	691594	648072
2	ILMI	0	16	257	257	13774	17420

ATM Interface+

Figure 97. 8210 ATM Interface Statistics

Figure 98 displays the result on the 8274 side.

```
/Interface/ATM % vcs
```

ATM Connection Statistics								
Slot	Port	VCI	Rx SDUs	Tx SDUs	Rx Cells	Tx Cells	Rx Octets	Tx Octets
3	1	74	253105	252520	2768885	2777624	132906480	133325952
3	1	100	0	49217	0	122926	0	5900448

Figure 98. 8274 ATM Statistics

5.8.3.4 Conclusion

This scenario illustrates a typical layered investigation. The investigation identified the ATM layer as failing. After correction of the PVC misconfiguration, and subsequent validation of this correction at the hub level and from both sides of the PVC, the ATM layer is now reliable. If the bridging still does not work between the two devices, then further investigation into the bridging configuration needs to be done. Refer to Chapter 8, “ATM Bridging and Routing” on page 417 for details.

Special Note about Defining PVCs on VOID Ports: A PVC may run on a VOID port without any further definition if this VOID port does not have any VPC_Link defined. However, it will not work if the VOID port has a VPC_Link already defined on a different VPI. It should be noted that the link will fail even when the remote port and local port are in the same hub. To avoid this problem, define a VPC_Link with the PVC’s VPI, even if this VPI is VPI=0. Otherwise, the link will fail with cause 0x3.

5.8.4 Case Study 4 -Traffic Reachable Address Problem

This problem can be seen when several servers can be addressed from several peer groups.

5.8.4.1 Network Diagram

This configuration shows two peer groups (P1 and P2)

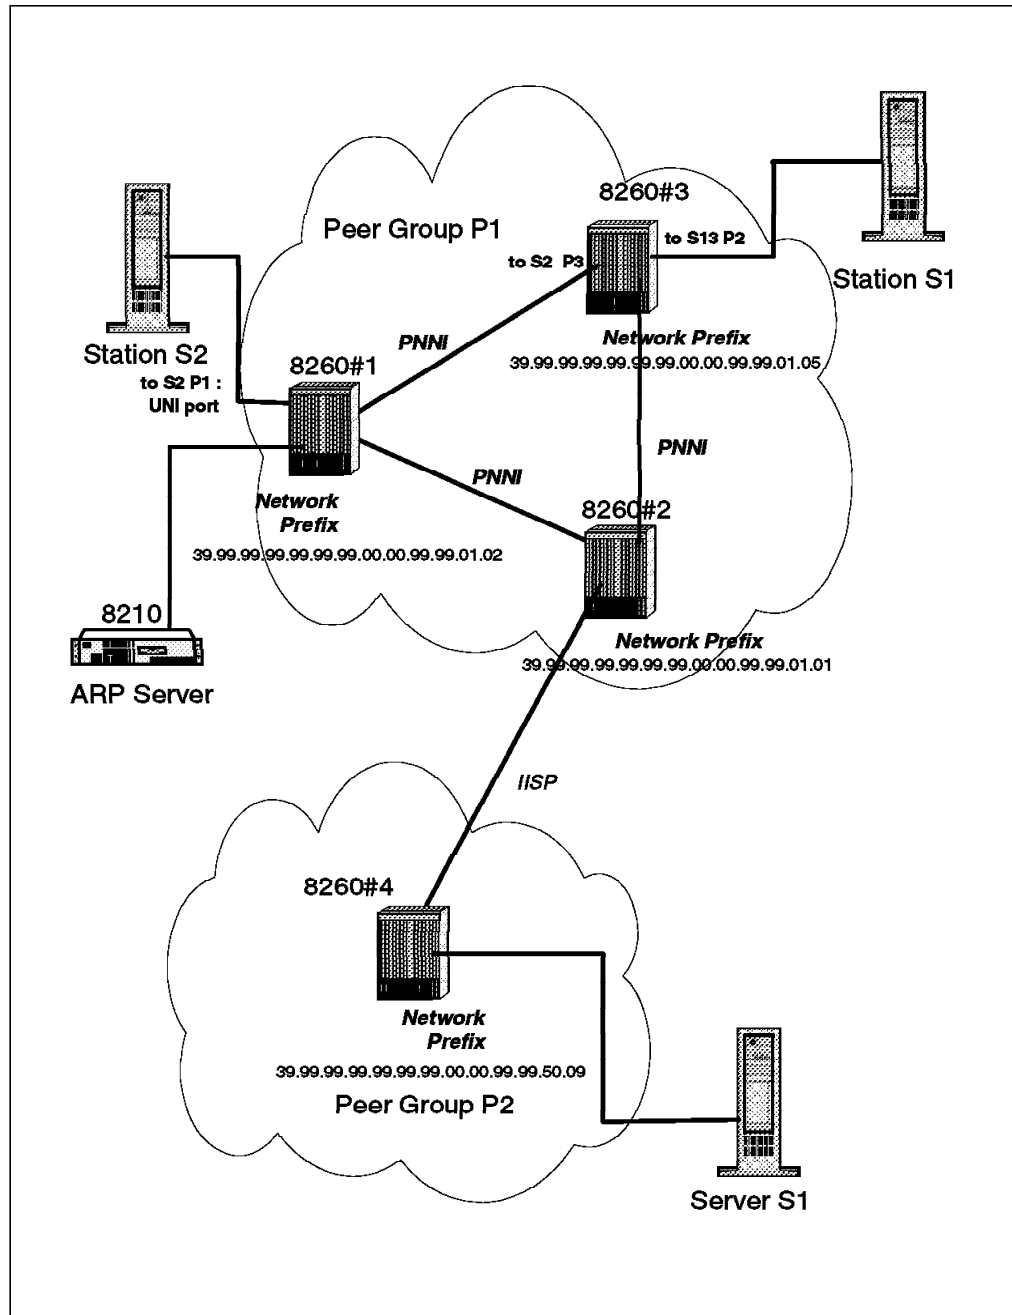


Figure 99. Network Configuration

5.8.4.2 Symptom

Some stations connected to 8260#3 cannot access the server S1 located in peer group P2 during working hours. For example, station S1 is unable to reach server S1 during working hours while it can connect without problem after normal working hours. All these device are running Classical IP and belong to the same subnet.

5.8.4.3 Methodology

Further investigation shows that all the stations connected to 8260#3 have the same problem and verification of all the basic parameters on 8260#3 did not reveal any obvious misconfiguration.

From the Network Management: Call logging filter applied on interface 1302 (slot 13 port 2) of 8260#3 reports SVC being established and immediately cleared with Clear Cause 31 (Normal Unspecified). Since the SVCs are transient, the network management station is not able to track the path followed by the SVC.

From the 8260: Looking at the port where the station is connected shows that there are transient SVCs being set up on top of the VCC to the ARP server.

```
8260ATM#3> show signalling cross_connections port 13.2
In: slot.port vpi.vci type Out: slot.port vpi.vci type Conn Cat
-----
      13.2      0.683  SVC      2.3      0.69  SVC      P2P  UBR
      13.2      0.785  SVC      2.3      0.76  SVC      P2P  UBR
Total number of cross connections = 2
```

Figure 100. Cross Connections of Port 13.2 in Hub 8260#3

Moreover, these transient SVCs are routed to 8260#1 via slot 2 and port 3, while the direct path to peer group P2 is through 8260#2. This single information does not indicate any failure, but understanding why the calls are routed to 8260#1 is the next step in problem determination.

It is now necessary to understand what happens with these calls. To do this we need to explore further into the reachability information within the cluster. Our main purpose is to find the reachability parameters of the destination address as seen by the source node (8260#3).

In order to find what happens with the call, we need to dump the path selection database to get more information about reachability inside the peer group.

The SHOW REACHABLE_ADDRESS ALL command gives the address advertised by a given hub. The dump of the path selection, as shown in Figure 101 on page 176, will give all the reachability information inside the peer group P1, including exterior reachability. The command to dump path selection data is:
dump pnni path_selection

Note: After dumping the path selection data on the 8260, you upload the file via TFTP to a workstation for viewing and analysis. Figure 102 on page 177 shows the TFTP setup at the 8260 for upload of the file to our workstation.

```

Starting PNNI Path Selection Dump at Thu Jun 26 10:25:51 1997

Path selection methods dump
-----
--More--
Spanning trees dump
-----
--More--
Routing tables dump
-----
--More--
Multicast trees dump
-----
--More--
Memory Dump
-----
--More
Allocation Stats Dump
-----
--More--
Graph dump 1
-----
Vertex: 0
      Node Identifier:
      60.A0.39.99.99.99.99.99.00.00.99.99.01.05.40.00.82.60.00.05.00.
--More--
Vertex: 3
      Node Identifier:
      60.A0.39.99.99.99.99.99.00.00.99.99.01.01.40.00.82.60.02.00.00.
--More--
Vertex: 4
      Node Identifier:
      60.A0.39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.60.00.07.00.
--More--
Reachability dump 2
-----
--More--

Vertex: 3 3
      1 prefix(es)
          Prefix 0 :39.99.99.99.99.99.00.00.99.99.50. (96 bits)
--More--

Vertex: 4 4
      1 prefix(es)
          Prefix 0 :39.99.99.99.99.99.00.00.99.99.50. (96 bits)

```

Figure 101. Part of PNNI Path Selection Dump For 8260#3

```
8260ATM#3> dump pnni path_selection
8260ATM#3> show tftp
TFTP Parameters:
Server IP address   : 192.2.1.129.
File Name          : /tmp/pnni_path.dmp.
File type          : Dump.
--More--
8260ATM#3> upload
Upload successful.
8260ATM#3>
```

Figure 102. How to Dump PNNI Path Selection

Understanding reachability information about the address a user wants to call from a node in a peer group is done in two steps when looking at the PNNI path selection dump:

1. Identify the nodes inside the topology. The topology knows about vertices. A vertex is the graphical representation of a physical node. As highlighted in **1** the graph dump section allows you to correlate a physical node to its topological representation which is called a vertex. Vertex 0 is always the local node (from which the dump is taken). From this example we can see that:
 - Vertex 0 is 8260#3.
 - Vertex 3 is 8260#1.
 - Vertex 4 is 8260#2.
2. List all the reachability information relevant to a given address. This is found in the Reachability dump section as highlighted by **2**. This section indicates that the target address (or a part of the target address) is seen as reachable from both vertex 3 **3** and vertex 4 **4**. In this example, the best reachability for the destination address of server S1 is found with a match of 96 bits. Unfortunately, the topology has two reachable network prefixes that match the network prefix of the destination address. Moreover, the length of the prefix is the same: 96 bits. This means that the 8260#3 sees two parallel paths to reach the destination address, vertex 3 and vertex 4 which correspond to 8260#2, and the node where we expected the calls to be routed.

We know that something is wrong in this configuration since no parallel path to peer group P2 is known by the network administrator, but another path is being advertised by 8260#1. A move to 8260#1 is now necessary to understand if this node really advertises reachability for this address and if so, why.

Figure 103 on page 178 gives the reachability information on 8260#1. This confirms that this hub considers this network part as reachable through port 2.1.

```

8260ATM#1> show reachable_address all
Port Len Address Active Idx VPI
-----
 2.01 96 39.99.99.99.99.99.99.00.00.99.99.50. . . . . Y 1 -
14.04 96 39.99.99.99.99.99.99.00.00.99.99.70. . . . . N 2 -
 2.01 152 39.99.99.99.99.99.99.00.00.99.99.01.02.40.00.85.95.11.11 Y Dyn 0 1
 3.02 152 39.99.99.99.99.99.99.00.00.99.99.01.02.08.00.5A.99.86.DC Y Dyn 0
 4.01 152 39.99.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.00.69 Y Dyn 0
 4.01 152 39.99.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.20.69 Y Dyn 0
 4.01 152 47.00.79.00.00.00.00.00.00.00.00.00.00.00.00.A0.3E.00.00.01 Y Dyn 0
 6.03 152 39.99.99.99.99.99.99.00.00.99.99.01.02.40.00.05.70.00.51 Y Dyn 0
13.02 152 39.99.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.00.72 Y Dyn 0
13.02 152 47.00.79.00.00.00.00.00.00.00.00.00.00.00.00.A0.3E.00.00.01 Y Dyn 0
17.02 152 39.99.99.99.99.99.99.00.00.99.99.01.02.12.34.56.78.90.12 Y Dyn 0
    
```

Figure 103. Reachability Information for 8260#1

Looking at port 2.1 gives the following information:

```

8260ATM#1> show port 2.1 verbose

Type Mode Status
-----
2.01: UNI enabled UP

Signalling Version : Auto
> Oper Sig. Version : 3.1
ILMI status : UP
ILMI vci : 0.16
RB Bandwidth : unlimited
Signalling vci : 0.5
Administrative weight: 5040
VPI.VCI range : 3.1023 (2.10 bits)
Connector : RJ45
Media : copper twisted pair
Port speed : 25600 kbps
Remote device is active
    
```

Figure 104. Port 2.1 Configuration

What Really Happens: The port is a UNI port, which is up and running, and has registered an ATM address as illustrated in **1** in Figure 103. Unfortunately, this port was previously used to connect to peer group P2, so reachability information was configured on this port. When the network administrator reconfigured the network, he or she forgot to remove the reachability information of this port. When the ATM station that is is not active and not advertised. However, when the port comes up, the address becomes active, it is advertised throughout the peer group and then the path selection algorithm. Thereafter, we have two paths and because of that some stations reach the expected output port to the destination while some others will be routed to the UNI port and to station S2 that is a dead end. This also explains why some stations did get a good connection at some time and why it failed during normal working hours.

5.8.4.4 Conclusion

This scenario illustrates how to investigate reachability inside a peer group. Dumping the PNNI path selection is straight forward, and it immediately gives information on reachability for a given address from a given node.

5.8.5 Case Study 5 - Reachable Address Missing for Backup LECS

This scenario shows the importance of the way a network is built.

5.8.5.1 Network Diagram

In this network, there are two peer groups (P1 and P2).

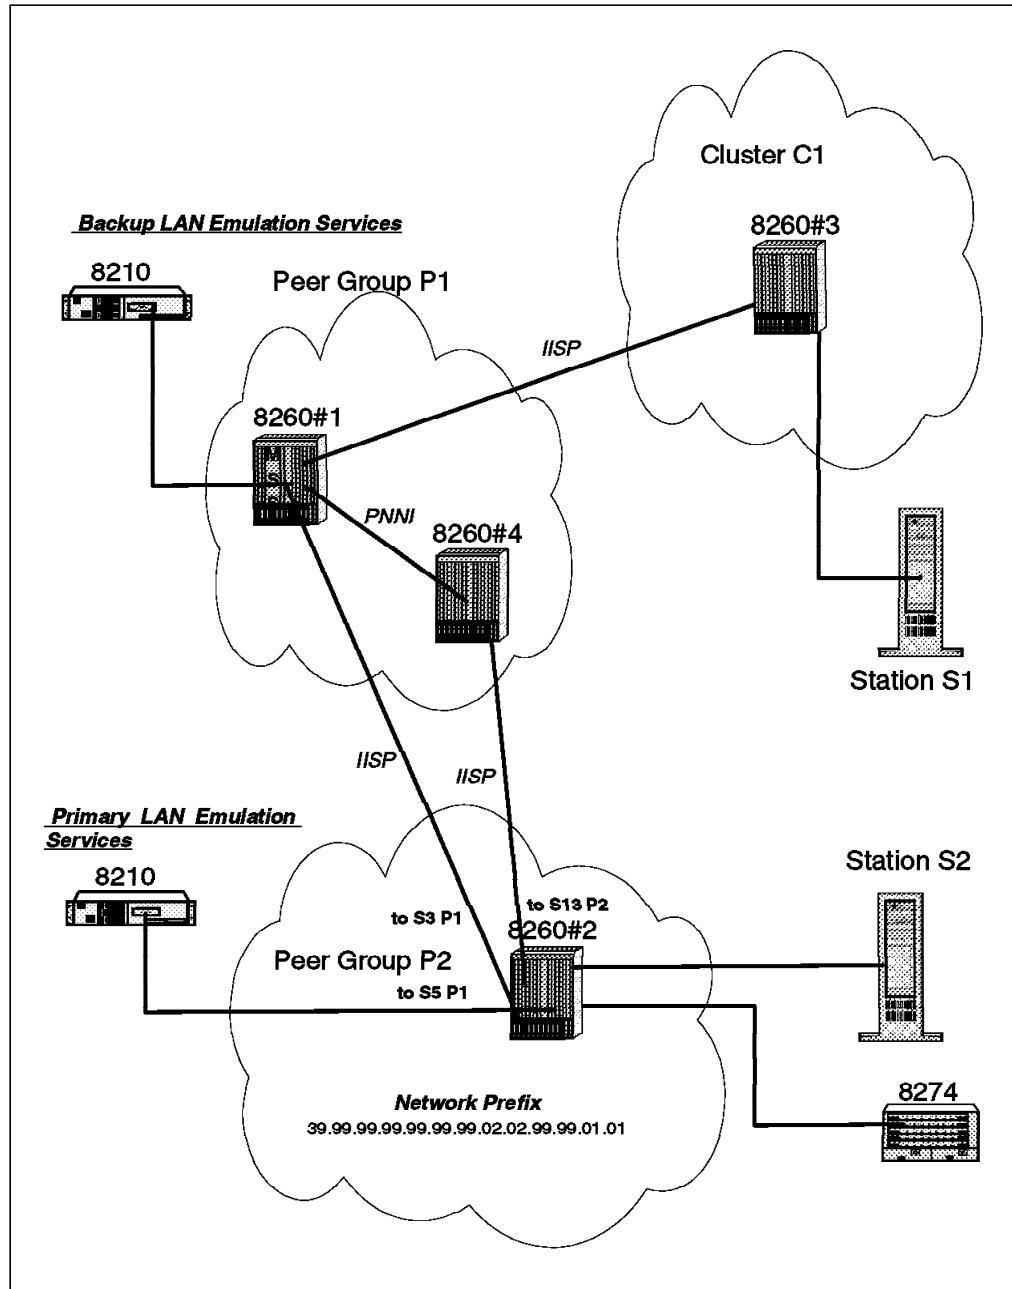


Figure 105. LAN Emulation Redundancy Network Configuration

5.8.5.2 Symptom

LAN Emulation redundancy appears not to be working for some clients in the network.

When the primary LECS fails, some stations can join the backup LECS but others cannot.

5.8.5.3 Methodology

Since some clients are able to reach the backup LECS, this means that it is operational. Investigation must be done at the ATM connectivity level for the failing clients. Our attempts to find the commonality between all the failing clients gives the following results:

- All the clients connected to peer group P2 are unable to connect to the backup LECS.
- The failing clients all use the LECS well-known address.

Figure 106 depicts the reachability information for 8260#2 when the primary LECS is running. The SHOW REACHABLE_ADDRESS command allows us to see all the addresses this switch advertises reachability for to the rest of the peer group members. This command does not give any information on the addresses reachable by other switches inside the peer group.

```
8260ATM#2> show reachable_address all
```

Port	Len	Address	Active	Idx	VPI
3.01	8	39.	Y	2	-
13.02	8	39.	Y	3	-
1.03	152	39.99.99.99.99.99.02.02.99.99.01.01.40.00.03.60.00.37	Y	Dyn	0
1.04	152	39.99.99.99.99.99.02.02.99.99.01.01.40.00.05.80.00.48	Y	Dyn	0
3.02	152	39.99.99.99.99.99.02.02.99.99.01.01.00.20.DA.70.70.80	Y	Dyn	0
5.01	152	39.99.99.99.99.99.02.02.99.99.01.01.40.00.82.10.00.71	Y	Dyn	0
5.01	152	47.00.79.00.00.00.00.00.00.00.00.00.00.00.00.00.A0.3E.00.00.01	Y	Dyn	0
7.01	152	39.99.99.99.99.99.02.02.99.99.01.01.40.00.82.81.00.67	Y	Dyn	0
14.02	152	39.99.99.99.99.99.02.02.99.99.01.01.12.34.56.78.90.12	Y	Dyn	0

Figure 106. Reachability Information with a Running LECS

```
8260ATM#2> show reachable_address all
```

Port	Len	Address	Active	Idx	VPI
3.01	8	39.	Y	2	-
13.02	8	39.	Y	3	-
1.03	152	39.99.99.99.99.99.02.02.99.99.01.01.40.00.03.60.00.37	Y	Dyn	0
1.04	152	39.99.99.99.99.99.02.02.99.99.01.01.40.00.05.80.00.48	Y	Dyn	0
3.02	152	39.99.99.99.99.99.02.02.99.99.01.01.00.20.DA.70.70.80	Y	Dyn	0
7.01	152	39.99.99.99.99.99.02.02.99.99.01.01.40.00.82.81.00.67	Y	Dyn	0
14.02	152	39.99.99.99.99.99.02.02.99.99.01.01.12.34.56.78.90.12	Y	Dyn	0

Figure 107. Reachability Information without LECS

Figure 107 gives the reachability information for this switch when the primary LECS is *not* running. This shows that port 5.1 did not register any address, and therefore the LECS address is no longer a reachable address for this switch. In our configuration, there was only one LECS located in this peer group, which means that now none the switches located in this peer group know how to route a call to the LECS well-known address because it is no longer advertised as reachable by 8260#2. The problem can be fixed by configuring reachability information for the backup LECS address on the IISP ports that provide connectivity to peer group P1 where the backup LECS is located. This is described in Figure 108.

```

8260ATM#2> set reachable_address 13.2 8 47
8260ATM#2> set reachable_address 3.1 8 47
8260ATM#2> show reachable_address all

```

Port	Len	Address	Active	Idx	VPI
3.01	8	39.	Y	2	-
3.01	8	47.	Y	5	-
13.02	8	39.	Y	3	-
13.02	8	47.	Y	4	-
1.03	152	39.99.99.99.99.99.02.02.99.99.01.01.40.00.03.60.00.37	Y	Dyn	0
1.04	152	39.99.99.99.99.99.02.02.99.99.01.01.40.00.05.80.00.48	Y	Dyn	0
3.02	152	39.99.99.99.99.99.02.02.99.99.01.01.00.20.DA.70.70.80	Y	Dyn	0
7.01	152	39.99.99.99.99.99.02.02.99.99.01.01.40.00.82.81.00.67	Y	Dyn	0
14.02	152	39.99.99.99.99.99.02.02.99.99.01.01.12.34.56.78.90.12	Y	Dyn	0

```

8260ATM#2>

```

Figure 108. Setting the Correct Reachability for Backup LECS

The address is configured on two ports to ensure a backup path to the backup LECS. Before using this type of configuration to give reachability to a range of addresses, make sure you have a good understanding of the overall addressing structure and physical connections inside your network.

5.8.5.4 Conclusion

This scenario illustrates how reachable information needs to be added when using IISP connections between switches. Reachability information was only partially configured for this network. In order for a peer group to have proper connectivity with other peer groups, one must make sure that all the network prefixes of the ATM addresses that need to be reached outside of a given peer group are configured as reachable addresses on at least one of the IISP ports connecting to the target remote peer group or cluster.

5.8.6 Case Study 6 - Basic PNNI Connection Problem

No communication between workstations.

5.8.6.1 Network Diagram

There is only one peer group.

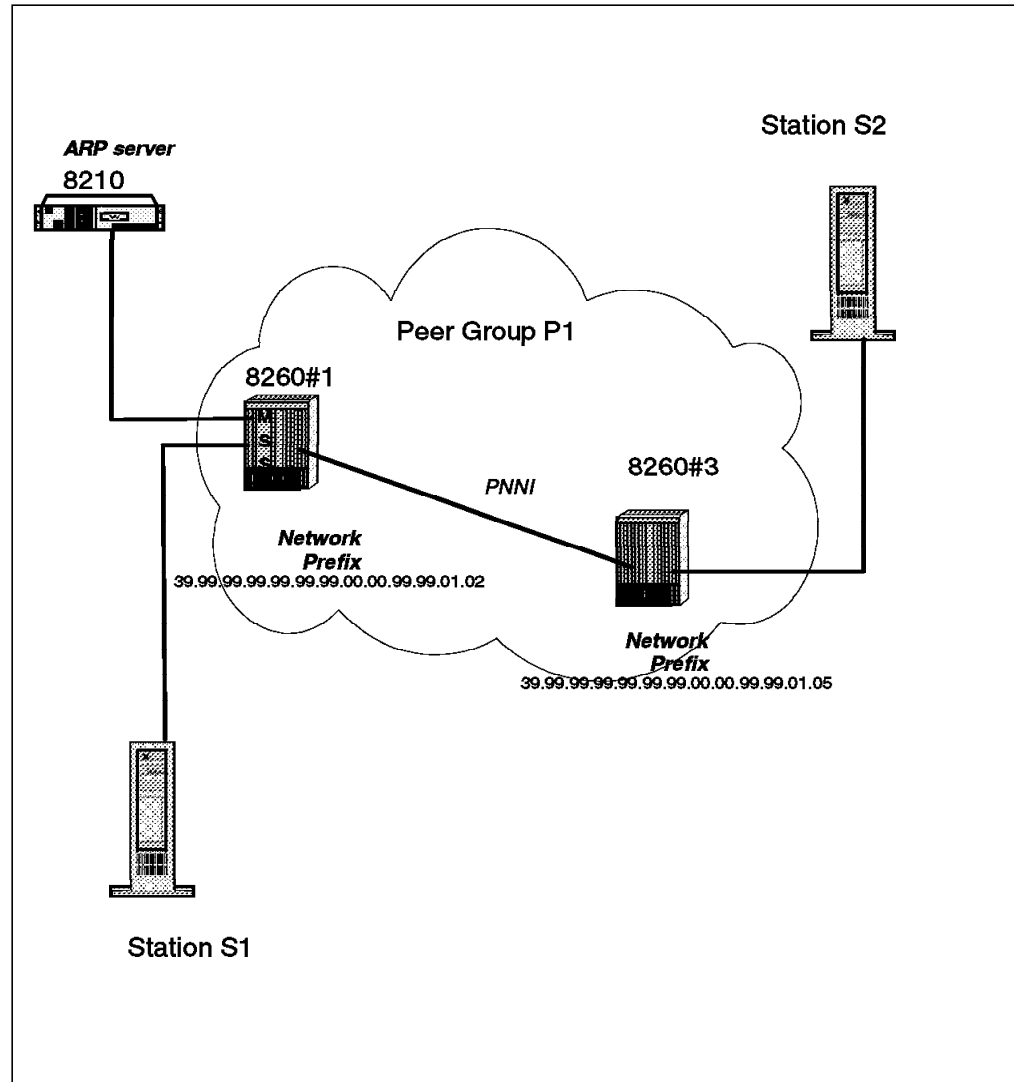


Figure 109. Network Configuration

5.8.6.2 Symptom

Station S2 connected to 8260#3 is not able to reach station S1 connected to 8260#1.

5.8.6.3 Methodology

A quick investigation shows that none of the stations connected to 8260#3 can connect to the ARP server that is connected to 8260#1.

After verification of good physical connectivity between the two hubs, investigation must be done at the ATM layer.

Since this network is made up of nodes running PNNI, checking the PNNI status is the next step in the problem determination. This is done by verifying that the topology of the network seen by each node is consistent with the real topology. The parameters that should be verified are the number of node members in the peer group, the neighboring information for each node, and the PNNI status of a given node.

Figure 110 shows the result of this investigation on 8260#3.

```
8260ATM#3>show pnni peer_group_members
----- Peer Group of Node 0-----
 60.A0.39.99.99.99.99.99.00.00.99.99.01.05.40.00.82.60.00.05.00 connected
 60.A0.39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.60.00.07.00 not cnct.
2 Members.
8260ATM#3>
```

Figure 110. Peer Group Members Seen from 8260#3

Figure 111 shows the result of this investigation on 8260#1.

```
8260ATM#1> show pnni peer_group_members
----- Peer Group of Node 0-----
 58.A0.39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.60.00.07.00 connected
1 Members.
```

Figure 111. Peer Group Members Seen from 8260#1

The output of these commands displays the node identifier. For nodes which do not represent a peer group:

- The first byte of the output is the level indicator which specifies the level of the node's peer group.
- The second byte takes the value 160 (0xA0).
- The remainder of the node ID contains the 20 bytes of the ATM End System Address of the system represented by the node.

Detailed information about the meaning of these fields can be found in the PNNI specification document in Section 5.3.3: Node Identifiers.

From these two commands, it appears that each node has a different view of the PNNI topology. 8260#3 sees two nodes as being members of the peer group, although one of these has the status *not connected*, while 8260#1 only sees itself in the peer group. By looking at the node identifiers, it appears that the ATM addressing is consistent as seen from 8260#3, but 8260#1 has a different Level Indicator. 8260#1 has a level indicator of 58 instead of 60, as seen from 8260#3. Displaying the ATM addressing configuration on each switch gives the following results:

1. From 8260#3:

```
8260ATM#3> show pnni node_0
----- Node 0 -----
ATM addr : 39.99.99.99.99.99.00.00.99.99.01.05.40.00.82.60.00.05.00
Level Identifier : 96 (24 half-bytes and 0 bits)
PGroup Id: 60.39.99.99.99.99.99.00.00.99.99.01
Node Id : 60.A0.39.99.99.99.99.99.00.00.99.99.01.05.40.00.82.60.00.05.00
Unrestricted Transit.
8260ATM#3>
```

Figure 112. Node 0 Configuration for 8260#3

2. From 8260#1

```
8260ATM#1> show pnni node_0
----- Node 0 -----
ATM addr : 39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.60.00.07.00
Level Identifier : 88 (22 half-bytes and 0 bits)
PGroup Id: 58.39.99.99.99.99.99.00.00.99.99
Node Id : 58.A0.39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.60.00.07.00
Unrestricted Transit.
```

Figure 113. Node 0 Configuration for 8260#1

The figures show that the ATM address settings are correct, but the level identifiers are different which translates into different peer group identifiers and explains the inconsistency seen from both nodes. In order to make ATM switches part of the same peer group, both ATM addresses and level identifiers must be consistent in all the nodes.

5.8.6.4 Example of a Three PNNI Nodes Network Valid Configuration

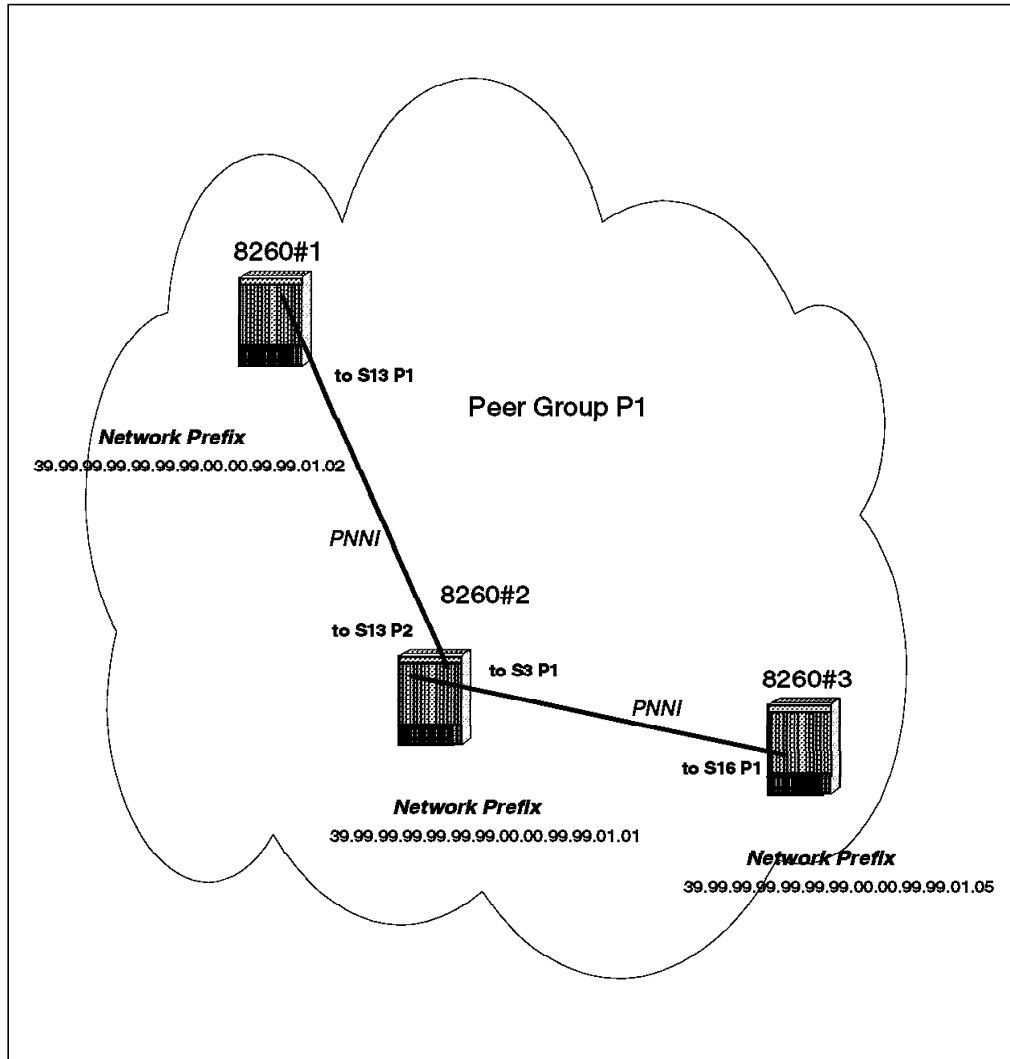


Figure 114. Peer Group Made of Three Nodes

Results from the Various Nodes: The figures that follow are the results of the different commands that can be used to check PNNI configuration from all the nodes involved in a given peer group.

1. Information related to the ATM addressing of the node:

```
8260ATM#1> show pnni node_0
----- Node 0 -----
ATM addr : 39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.60.00.07.00
Level Identifier : 96 (24 half-bytes and 0 bits)
PGroup Id: 60.39.99.99.99.99.99.00.00.99.99.01
Node Id : 60.A0.39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.60.00.07.00
Unrestricted Transit.
```

This indicates that this peer group level identifier is 12 bytes long that is 96 bits (0x60).

2. An example of a good PNNI port status follows. The port status is *UP*.


```
8260ATM#1> show port 13.1 verbose

      Type  Mode   Status
-----
13.01: PNNI enabled UP

ILMI status      : UP
ILMI vci         : 0.16
NNI Bandwidth    : 155000 kbps
RB Bandwidth     : unlimited
Signalling vci   : 0.5
Routing vci      : 0.18
Administrative weight: 5040
VPI.VCI range    : 15.1023 (4.10 bits)
Connector        : TP
Media            : copper twisted pair
Port speed       : 155000 kbps
Remote device is active

Frame format     : SONET STS-3c
Scrambling mode  : frame and cell
Clock mode       : internal
```

3. Each peer group member should be able, when issuing this command, to get the list of all the nodes belonging to the peer group. The expected status for each node in the list is *connected* (Figure 115).

```
8260ATM#1> show pnni peer_group_members
----- Peer Group of Node 0-----
 60.A0.39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.60.00.07.00 connected
 60.A0.39.99.99.99.99.99.00.00.99.99.01.05.40.00.82.60.00.05.00 connected
 60.A0.39.99.99.99.99.99.00.00.99.99.01.01.40.00.82.60.02.00.00 connected
3 Members.
```

Figure 115. Peer Group Members

4. Figure 116, Figure 117 on page 188, and Figure 118 on page 188 show the neighboring information as seen from each of the nodes as well as the ports and VPI used to connect them together.

- Neighboring information from 8260#1:

```
8260ATM#1> show pnni neighbor
----- Neighbors of Node 0-----
60.A0.39.99.99.99.99.99.00.00.99.99.01.01.40.00.82.60.02.00.00: Full
  Port 13.01 vpi=0
8260ATM#1>
```

Figure 116. Neighbors of 8260#1

- Neighboring information from 8260#2:

```
8260ATM#2> show pnni neighbor
----- Neighbors of Node 0-----
60.A0.39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.60.00.07.00: Full
    Port 13.02 vpi=0
60.A0.39.99.99.99.99.99.00.00.99.99.01.05.40.00.82.60.00.05.00: Full
    Port 3.01 vpi=0
8260ATM#2>
```

Figure 117. Neighbors of 8260#2

- Neighboring information from 8260#3:

```
8260ATM#3> show pnni neighbor
----- Neighbors of Node 0-----
60.A0.39.99.99.99.99.99.00.00.99.99.01.01.40.00.82.60.02.00.00: Full
    Port 16.01 vpi=0
8260ATM#3>
```

Figure 118. Neighbors of 8260#3

The **Full** state indicates that the node successfully completed topology information exchange with the neighbor.

5.8.6.5 Conclusion

This scenario illustrates a simple PNNI network and highlights the basic addressing rules to follow.

5.8.7 Case Study 7 - Address Not Reachable between Peer Groups

This scenario represents a problem on a network built with IISP redundancy links.

5.8.7.1 Network Diagram

Peer groups (P1 and P2) are connected through two IISP links.

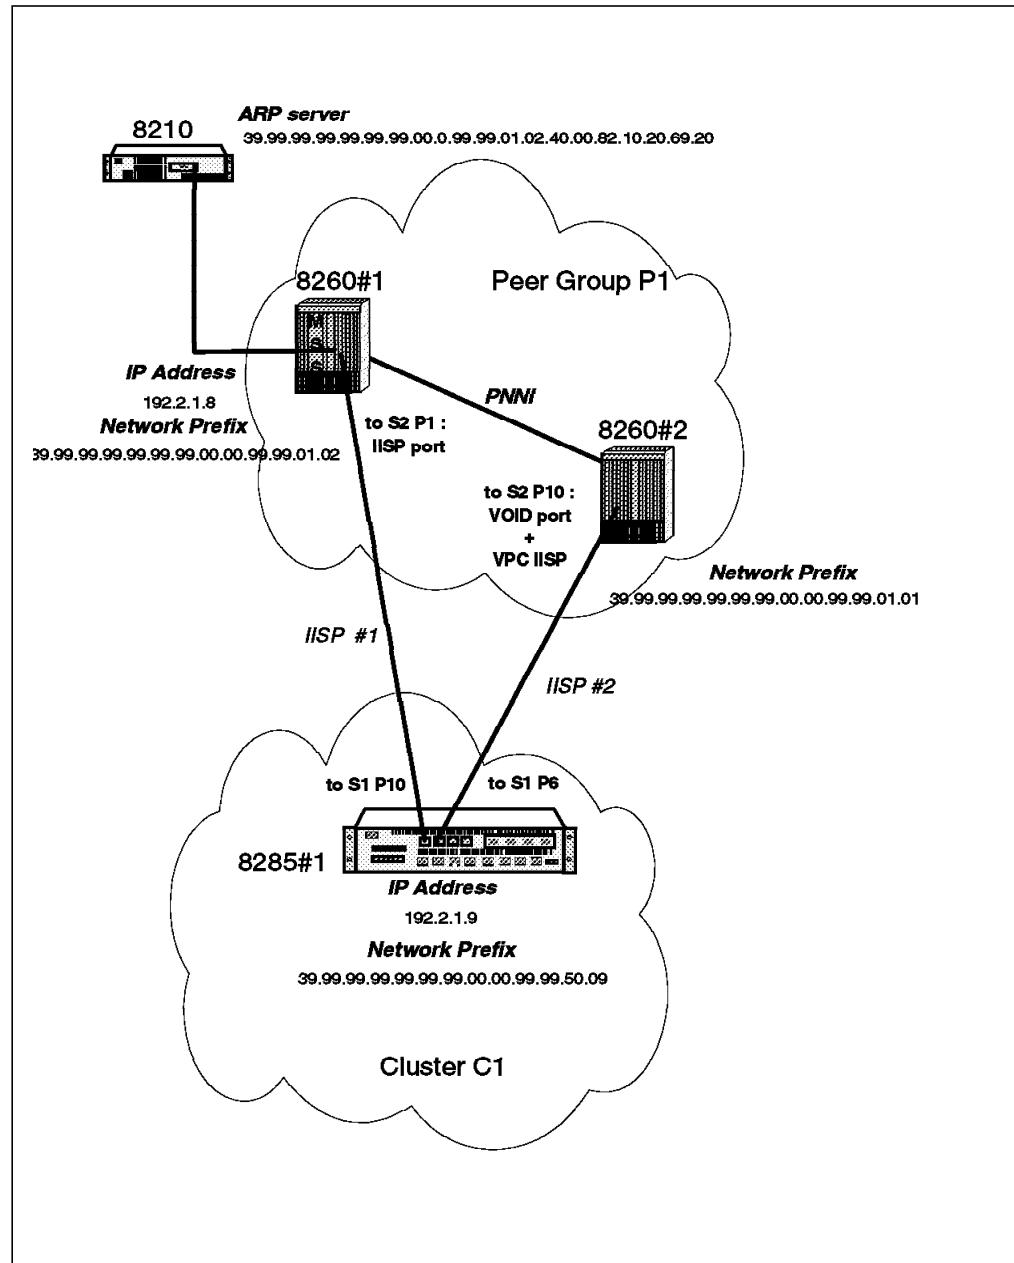


Figure 119. Network Configuration

5.8.7.2 Symptom

After losing the IISP#1 link between the peer group and the cluster, partial and transient connectivity appears in the network. The 8285 internal CIP client is used as a test tool to try to understand the exact behavior of the network. When trying to PING some other stations located in the subnet, the PING does not work. 8260#1 is used as a destination IP address as described in Figure 120.

```
8285#1> PING 192.2.1.8
Starting PING (hit CTRL-C to stop) ...
PING 192.2.1.8: 1 packets sent, 0 received
PING 192.2.1.8: 2 packets sent, 0 received
PING 192.2.1.8: 3 packets sent, 0 received
3 missing responses.
```

Figure 120. Result of First PING from 8285#1

Trying to PING 8285#1 from 8260#1 gives the following results:

IP connectivity is correct.

```
8260ATM#1> PING 192.2.1.9
Starting PING (hit CTRL-C to stop) ...
PING 192.2.1.9: 1 packets sent, 1 received
PING 192.2.1.9: 2 packets sent, 2 received
PING 192.2.1.9: 3 packets sent, 3 received
8260ATM#1>
```

Figure 121. Result of PING from 8260#1

Trying to PING 8260#1 from 8285#1 gives the following results:

IP connectivity is now working.

```
8285#1> PING 192.2.1.8
Starting PING (hit CTRL-C to stop) ...
PING 192.2.1.8: 1 packets sent, 1 received
PING 192.2.1.8: 2 packets sent, 2 received
PING 192.2.1.8: 3 packets sent, 3 received
8285#1>
```

Figure 122. Result of the Second PING from 8285#1

Unfortunately, this PING fails when retried some minutes later. The 8285#1 A is no longer able to reach the IP address of 8260#1.

5.8.7.3 Methodology

Investigation starts at the CIP level. This is done by checking the status of the Classical IP clients experiencing the connectivity problem:

1. From 8285#1

```
8285#1> show device
8285 Nways ATM Workgroup Switch
--More--
A-8285
-----
ATM address: 39.99.99.99.99.99.00.00.99.99.50.09.40.00.82.85.00.09.00

> Subnet atm: Up
IP address: 192.2.1.9. Subnet mask: FF.FF.FF.00
--More--
Default Gateway : OK
-----
IP address: 192.2.1.69

ARP Server:
-----
ATM address: 39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.20.69.20
```

Figure 123. 8285#1 Configuration

This indicates a successful connection to the logical IP subnet.

2. From 8260#1:

```
8260ATM#1> show device
8260 ATM Control Point and Switch Module
--More--
A-CPSW
-----
> Subnet atm: Up
IP address: 192.2.1.8. Subnet mask: FF.FF.FF.00

Default Gateway : OK
IP address: 192.2.1.69

ARP Server:
-----
ATM address: 39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.20.69.20
```

Figure 124. 8260#1 Configuration

This also indicates a successful connection to the logical IP subnet. Moreover, these two stations are configured to register to the same ARP server, and are configured with valid IP addresses (same subnet). No IP routing is involved in this configuration.

Further checking can be done by dumping the ARP table content of the ARP server. Figure 125 on page 192 shows the content of the ARP cache of the ARP server. The ARP server is located in an 8210.

3. From the ARP server:

```

ARP>dump
Network number to dump (0)? 0
Hardware Address      IP Address           Refresh
0/54                 192.2.1.129         17
0/364                192.2.1.72          17
0/49                192.2.1.8          18
0/316              192.2.1.9          19
0/53                 192.2.1.51          19

```

Figure 125. ARP Server Dump

This confirms that the clients did register with the ARP server. Since everything seems correctly configured at the CIP level, let's have a look at the ATM configuration. In this configuration, the two devices are not part of the same peer group.

When the network has parallel links, everything works fine, while a problem appears when loosing IISP#1. This may indicate an ATM connectivity problem between nodes of cluster C1 and nodes of peer group P1 through IISP#2.

The Network Management View: Let's have a look at network management station to see if it gives some more indication. Call logging started on 8260#1 reports calls being rejected with Cause Code 3 (No route to destination). This indicates an ATM routing problem over IISP#2.

ATM Nodes Configuration: Let's look into the configuration of the three nodes involved, and particularly focus on reachability configuration.

1. Port status and link information

a. 8285#1 port 1.6:

```

8285#1> show port 1.6 verbose

      Type Mode      Status
-----
1.06:NNI enabled UP-OKAY

VPI.VCI range      : 3.1023 (2.10 bits)
Connector          : RJ45
Media              : copper twisted pair
Port speed         : 25600 kbps
Remote device is active
IX status          : IX OK
Logical links indexes: 2

```

Port status is normal.

b. 8285#1 port 1.10

```

8285> show port 1.10 verbose

      Type Mode      Status
-----
1.10:NNI enabled UP-OKAY

VPI.VCI range      : 3.1023 (2.10 bits)
Connector          : RJ45
Media              : copper twisted pair
Port speed         : 25600 kbps
Remote device is active
IX status          : IX OK
Logical links indexes: 1

```

Port status is normal.

c. 8285#1 logical links:

```
8285#1> show logical_link all
Port Vpi Acn Side Mode Sig Traf Bwidth Status Index
-----
1.10 0 03 user enab 3.1 NRB 0 UP 1
1.06 0 03 user enab 3.1 NRB 0 UP 2

62 entries empty
```

All logical links are UP.

d. 8260#1 port 2.1:

```
8260ATM#1> show port 2.1 verbose

Type Mode Status
-----
2.01:IISP enabled UP

Signalling Version : 3.1
No ILMI
NNI Bandwidth : 10000 kbps
RB Bandwidth : unlimited
Signalling role : network
Signalling vci : 0.5
Administrative weight: 5040
VPI.VCI range : 3.1023 (2.10 bits)
Connector : RJ45
Media : copper twisted pair
Port speed : 25600 kbps
Remote device is active
```

e. 8260#2 port 2.10:

In order to verify the IISP configuration of the port, two levels need to be checked:

The physical port itself and the vpc_link.

- Physical port status:

```
8260ATM#2> show port 2.10 verbose

Type Mode Status
-----
2.10:VOID enabled UP

No ILMI
VPI.VCI range : 3.1023 (2.10 bits)
Connector : RJ45
Media : copper twisted pair
Port speed : 25600 kbps
Remote device is active
```

- Vpc_link information:

```

8260ATM#2> show vpc_link all verbose

      VPI :Type  Mode      Status
-----
2.10   0:IISP enable UP
Signalling Version   : 3.1
No ILMI
VPC Bandwidth       : 25600 kbps
RB Bandwidth        : unlimited
Signalling role     : network
Signalling vci      : 0.5
Administrative weight: 5040
VPI.VCI range       : 3.1023 (2.10 bits)
VPCI                : 0
    
```

This shows that the nodes have a valid configuration in terms of:

- Signalling version: All set to 3.1.
- Signalling role: One side network, one side user for each IISP.
- VPI.VCI range: Equal on each side of the link.
- VPI used: 0.

Moreover IISP links are UP.

2. Routes configuration

a. 8285 Static routes:

```

8285#1> show static_route
Index Acn Static route
-----
2   03 39
63 empty entries.
    
```

The switch has two logical links and a static route defined to the same virtual ACN number (03). This allows a redundant path to peer group P1 from the 8285#1; everything seems up and running.

b. Reachability information of 8260#1:

```

8260ATM#1> show reachable_address all
Port Len Address                                     Active Idx VPI
-----
2.01  96 39.99.99.99.99.99.00.00.99.99.50. . . . . Y 1 -
3.02 152 39.99.99.99.99.99.99.00.00.99.99.01.02.08.00.5A.99.86.DC Y Dyn 0
4.01 152 39.99.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.00.69 Y Dyn 0
4.01 152 39.99.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.20.69 Y Dyn 0
4.01 152 47.00.79.00.00.00.00.00.00.00.00.00.00.00.00.00.A0.3E.00.00.01 Y Dyn 0
6.02 152 39.99.99.99.99.99.99.00.00.99.99.01.02.40.00.82.81.00.66 Y Dyn 0
6.03 152 39.99.99.99.99.99.99.00.00.99.99.01.02.40.00.05.70.00.51 Y Dyn 0
8260ATM#1>
    
```

This shows that the correct network prefix has been configured, and the address is active.

c. Reachability information for 8260#2:


```

8260ATM#2> show reachable_address all
Port Len Address Active Idx VPI
-----
2.10 96 39.99.99.99.99.99.99.99.00.00.99.99.50. . . . . N 1 -
2.03 152 39.99.99.99.99.99.99.99.00.00.99.99.01.01.40.00.03.50.01.47 Y Dyn 0
3.02 152 39.99.99.99.99.99.99.99.00.00.99.99.01.01.00.20.DA.70.70.80 Y Dyn 0
4.01 152 39.99.99.99.99.99.99.99.00.00.99.99.01.01.40.00.82.10.00.73 Y Dyn 0
4.01 152 47.00.79.00.00.00.00.00.00.00.00.00.00.00.00.A0.3E.00.00.01 Y Dyn 0
5.01 152 39.99.99.99.99.99.99.99.00.00.99.99.01.01.40.00.82.10.00.71 Y Dyn 0
5.01 152 47.00.79.00.00.00.00.00.00.00.00.00.00.00.00.00.A0.3E.00.00.01 Y Dyn 0
7.01 152 39.99.99.99.99.99.99.99.00.00.99.99.01.01.40.00.82.81.00.67 Y Dyn 0
13.01 152 39.99.99.99.99.99.99.99.00.00.99.99.01.01.00.00.C1.10.C8.A0 Y Dyn 0

```

Figure 126. 8260(2) Reachable Addresses

This shows that the correct address has been configured, but the address shows up as inactive (Active : N). This is what causes reachability problems when the other link to cluster C1 is disconnected. Since this address is not active, the 8260#2 does not know how to route the calls whose network prefix match the reachable address defined on port 2.10.

The way to set the correct configuration for this reachable address is described in Figure 127.

```

8260ATM#2>
set reachable_address 2.10 96 39.99.99.99.99.99.99.99.00.00.99.99.50 vpi:0
Entry set.
8260ATM#2> show reachable_address all
Port Len Address Active Idx VPI
-----
2.10 96 39.99.99.99.99.99.99.99.00.00.99.99.50. . . . . Y 1 0
2.01 152 39.99.99.99.99.99.99.99.00.00.99.99.01.01.40.00.03.50.01.48 Y Dyn 0
3.02 152 39.99.99.99.99.99.99.99.00.00.99.99.01.01.00.20.DA.70.70.80 Y Dyn 0
4.01 152 39.99.99.99.99.99.99.99.00.00.99.99.01.01.40.00.82.10.00.73 Y Dyn 0
4.01 152 47.00.79.00.00.00.00.00.00.00.00.00.00.00.00.00.A0.3E.00.00.01 Y Dyn 0
5.01 152 39.99.99.99.99.99.99.99.00.00.99.99.01.01.40.00.82.10.00.71 Y Dyn 0
5.01 152 47.00.79.00.00.00.00.00.00.00.00.00.00.00.00.00.A0.3E.00.00.01 Y Dyn 0
7.01 152 39.99.99.99.99.99.99.99.00.00.99.99.01.01.40.00.82.81.00.67 Y Dyn 0
13.01 152 39.99.99.99.99.99.99.99.00.00.99.99.01.01.00.00.C1.10.C8.A0 Y Dyn 0
8260ATM#2> save all

```

Figure 127. Reachable Address Configuration on a VPC Link

Now, the reachable address is active. This means that it is operational and that reachability information is advertised to all the other members of the peer group this switch belongs to. The switch is now able to route the calls for this reachable address.

This illustrates the major difference in the configuration of the two ports used for the connection to cluster C1.

- Port 2.1 of 8260#1:

This port is configured with a port type of IISP. It *can only* use VPI 0, and the reachable address becomes active as soon as the IISP port is Up. Moreover as shown in Figure 126, there is no VPI value displayed on this reachable address.

- Port 2.10 of 8260#2:

This port is configured with a port type of VOID. On top of this VOID port, an IISP vpc_link needs to be defined. This IISP vpc_link is configured for a given VPI value (0 in this configuration). Definition of the reachable address *requires an explicit VPI value* to be specified. Not doing so will define the address as not active, and will prevent outgoing calls from peer group P1 from reaching cluster C1 through IISP#2.

To summarize:

The VPI value associated with a reachable address is implicit and equals 0 when the port is native IISP. A VPI value needs to be specified when the port is VOID + IISP vpc_link, even if this VPI value is 0.

5.8.7.4 Conclusion

This scenario was built using ATM switches as an example. The same symptoms will be seen from a regular station.

About Symptom and Problem: Be careful when using higher layer protocol commands to investigate SVC problems. Especially for routes troubleshooting, make sure to understand the behavior of all the layers involved. The following parameters need to be considered:

- Using the higher layer protocols adds some overhead to the troubleshooting path since you then need to consider the different mapPING tables being cached in the network.

Most common are for instance:

- ARP cache for Classical IP
 - LE ARP cache
 - IP ARP table for IP over LAN Emulation
 - Address tables in switches for frame forwarding
 - VLAN membership tables
- ATM is connection-oriented; once an SVC is set, it can be used by both ends, and it remains alive for a given amount of time (device dependent).

Back to the symptom described in this scenario, this is what makes this problem appear more complex than it really is. The PING worked sometimes. When the station that had no ATM route to reach its destination could "reuse" the SVC setup when the other station established a connection with the failing station.

So you may consider this problem as transient and station-dependent while it is absolutely solid and network-related. But the environment and the test procedure used can make a problem appear much more complex than it really is.

When using high-layer protocols, try to get rid of timers in the network and influences of any intermediate caching occurring in the network. For instance, this can be done by cleaning ARP tables, unplugging fibers, deleting a logical interface and re-creating it. This may avoid deep and painful investigations on the wrong protocol layer, and avoid focusing investigations on the wrong component in the network.

5.8.8 Case Study 8 - PNNI Path Selection Tracing

The purpose of this scenario is to illustrate how to capture data in order to understand the ATM routing mechanism involved at Call Setup for a *Unicast UBR* call. Understanding the full content of the data captured is out of the scope of this example, and will be seen in the next scenarios.

5.8.8.1 Network Diagram

A network analyzer is used.

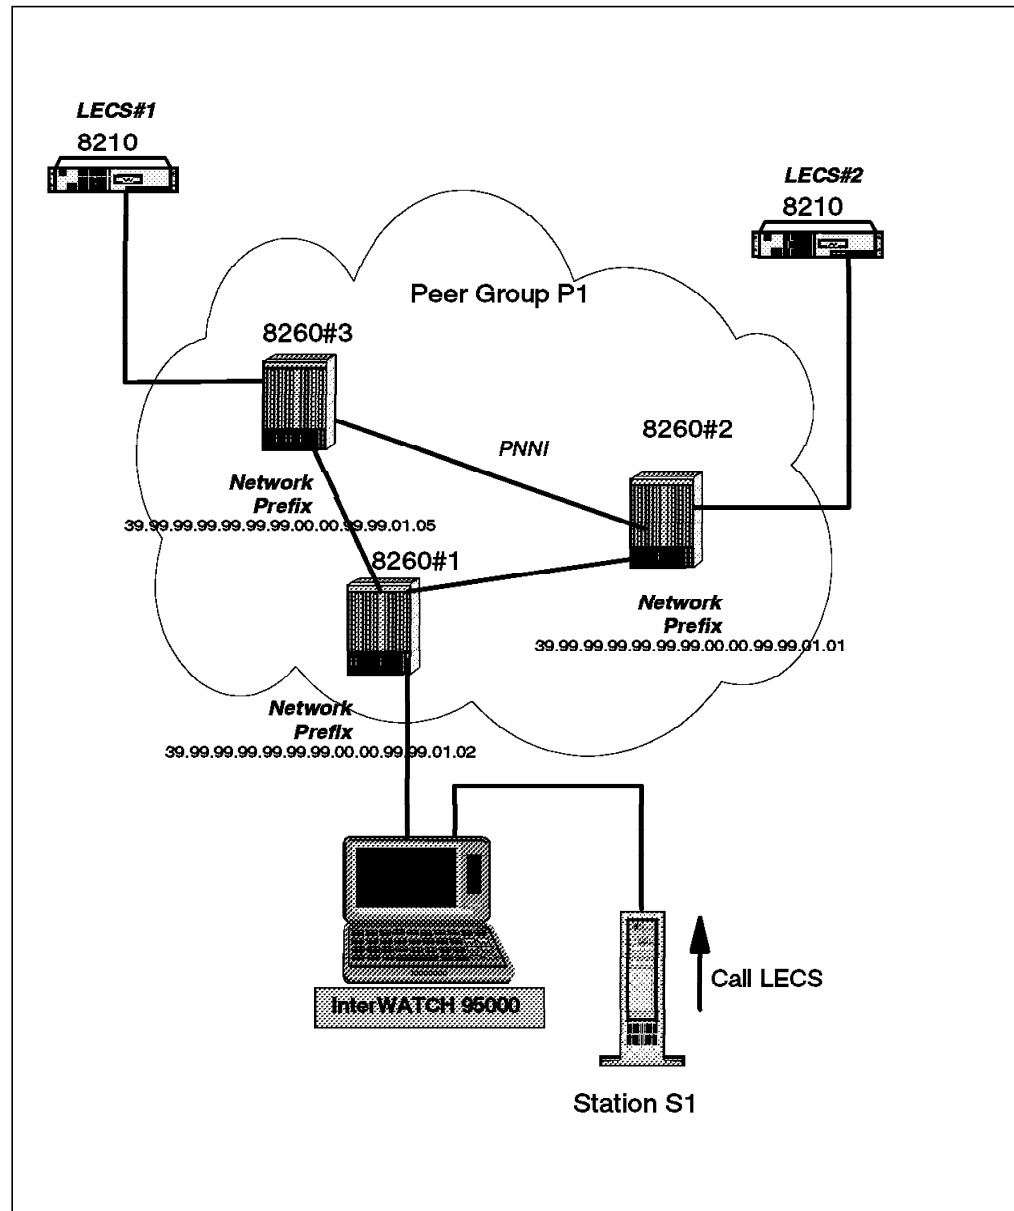


Figure 128. Network Configuration

5.8.8.2 Methodology

In order to understand what happens at the ATM routing level, we have to analyze how a switch makes routing decisions. This information is accessible through the 8260 internal traces and dumps. This example illustrates the data captured from 8260#1 when processing a call to the LECS WKA issued from station S1.

From the 8260 originating the call it is required to start the signalling and PNNI path selection trace as described in Figure 129.

```

8260ATM#1> set trace main_trace on
Main trace is ON
      base trace is off.
      bus trace is off.
signalling messages trace is on.
      ilmi trace is off.
      lec trace is off.
      les trace is off.
      pnni_base trace is off.
      pnni_messages trace is off.
      pnni_neighbor trace is off.
pnni_path_selection trace is on.
      pvc trace is off.
      RFC 1577 trace is off.
      saal trace is off.
      connections trace is off.
      snmp trace is off.

```

Figure 129. How to Start a Selective Trace on 8260

Once the trace is stopped and uploaded, it can be analyzed.

The following figures represent the detailed decoding of the call setup by the Navtel interWATCH 95000. This call is made from station S1 to the LECS WKA. In Figure 130 on page 199, **3** to **9** are the individual information elements of the call setup, and can also be identified in the 8260 internal trace.

```

----- DETAIL -----

PDU Length 144 bytes sliced to 144 bytes
TE Side
AAL-5 CPCS      VCC = (0, 5)
AAL-5 CPCS Overhead:
  Number of padding octets      20
AAL-5 CPCS Trailer:
  User-to-User Indicator        0x00
  Common Part Indicator          0x00
  CPCS Payload Length           116
  Cyclic Redundancy Check       0x6c573ed0

Q.SAAL SSCOP PDU
Pad:                NULLNo more data.
Pad Length:         00----- 0 Octets
Reserved:           --00----
PDU Type:           ----1000  SD(Sequenced Data)
N(S):               0x00002b

UNI-3.1 Signalling Message
PD:                 0x09      Q.2931 (UNI 3.1)      3
Length of CRV:      3 Octets
Flag:               0-----  Origination
CRV:                0x000166
Message Type:       0x05      SETUP                4
Extension Bit:      1-----  Octet group terminated
Spare:              -00-----  Not used
Flag:               ---0----  Instruction not significant
Spare:              ----00--  Not used
Action indicator:   -----00  Clear call
Message Length:     103 Octets

Information Element (IE)
IE Identifier:      0x58      ATM adaptation layer parameters      5
Extension Bit:      1-----  Octet group terminated
Coding Standard:    -00-----  ITU-T standardized
Flag:               ---0----  Instruction not significant
Reserved:           ----0---
IE Instruction:     -----000  Clear call
IE Length:          9 Octets
AAL Type:           0x05      AAL5
Fwd max CPCS-SDU:   1516 Octets
Bk max CPCS-SDU:    1516 Octets
SSCS type:          0x00      Null

Information Element (IE)
IE Identifier:      0x59      ATM traffic descriptor      6
Extension Bit:      1-----  Octet group terminated
Coding Standard:    -00-----  ITU-T standardized
Flag:               ---0----  Instruction not significant
Reserved:           ----0---
IE Instruction:     -----000  Clear call
IE Length:          9 Octets
Fwd pk rate CLP=0+1: 2048
Bk pk rate CLP=0+1: 2048
Best Effort Indctr: 0xbe

```

Figure 130 (Part 1 of 3). NAVTEL InterWATCH 95000 Call Setup Decoding

```

Information Element (IE)
IE Identifier:      0x5e      Broadband bearer capability      7
Extension Bit:     1-----  Octet group terminated
Coding Standard:  -00-----  ITU-T standardized
Flag:             ---0----  Instruction not significant
Reserved:         ----0---
IE Instruction:    -----000  Clear call
IE Length:        3 Octets
Extension Bit:     0-----  Octet group extended
Spare:            -00-----  Not used
Bearer Class:     ---10000  BCOB-X
Extension Bit:     1-----  Octet group terminated
Spare:            -00-----  Not used
Traffic Type:     ---000--  No indication
Timing Requirements: -----00  No indication
Extension Bit:     1-----  Octet group terminated
Suscptblty to Clippng:-00-----  Not susceptible
Spare:            ---000--  Not used
User Plane Connection:-----00  Point-to-point

Information Element (IE)
IE Identifier:      0x5f      Broadband low layer information    8
Extension Bit:     1-----  Octet group terminated
Coding Standard:  -00-----  ITU-T standardized
Flag:             ---0----  Instruction not significant
Reserved:         ----0---
IE Instruction:    -----000  Clear call
IE Length:        9 Octets
Extension Bit:     0-----  Octet group extended
Layer ID:         -11-----  User information layer 3
  Protocol:       ---01011  ISO/IEC TR9577
  Extension Bit:  0-----  Octet group extended
  ISO/IEC TR9577 IPI: -1000000
  Extension Bit:  1-----  Octet group terminated
  IPI (continued): -0-----  0x80 IEEE 802.1 SNAP identifier
  Spare:          --000000  Not used
  Extension Bit:  1-----  Octet group terminated
  SNAP ID:       -00-----
  Spare:         ---00000  Not used
  OUI:           0x00a03e
  PID:           0x0001

Information Element (IE)
IE Identifier:      0x70      Called party number                9
Extension Bit:     1-----  Octet group terminated
Coding Standard:  -00-----  ITU-T standardized
Flag:             ---0----  Instruction not significant
Reserved:         ----0---
IE Instruction:    -----000  Clear call
IE Length:        21 Octets
Extension Bit:     1-----  Octet group terminated
Type of Number:   -000----  Unknown
Address/Numbrng Plan: ----0010  ATM endsystem address
  AFI:            47          ICD ATM Format
  ICD:            0079
  HO-DSP:         0x000000000000000000000000
  ESI:            0x00a03e0000001
  SEL:            0x00

```

Figure 130 (Part 2 of 3). NAVTEL InterWATCH 95000 Call Setup Decoding

Information Element (IE)			
IE Identifier:	0x6c	Calling party number	10
Extension Bit:	1-----	Octet group terminated	
Coding Standard:	-00-----	ITU-T standardized	
Flag:	---0----	Instruction not significant	
Reserved:	----0---		
IE Instruction:	----000	Clear call	
IE Length:	22	Octets	
Extension Bit:	0-----	Octet group extended	
Type of Number:	-000----	Unknown	
Address/Numbering Plan:	----0010	ATM endsystem address	
Extension Bit:	1-----	Octet group terminated	
Presentation Indctr:	-00-----	Presentation allowed	
Spare:	---000--	Not used	
Screening Indctr:	-----00	User-provided, not screened	
AFI:	39	DCC ATM Format	
DCC:	9999		
HO-DSP:	0x99999999000099990102		
ESI:	0x123456789012		
SEL:	0x00		
Information Element (IE)			
IE Identifier:	0x5c	Quality of service parameter	11
Extension Bit:	1-----	Octet group terminated	
Coding Standard:	-00-----	ITU-T standardized	
Flag:	---0----	Instruction not significant	
Reserved:	----0---		
IE Instruction:	----000	Clear call	
IE Length:	2	Octets	
QoS Class Forward:	0x00	QoS class 0 - Unspecified	
QoS Class Backward:	0x00	QoS class 0 - Unspecified	
End of Message			

Figure 130 (Part 3 of 3). NAVTEL InterWATCH 95000 Call Setup Decoding

Figure 131 shows the hexadecimal dump of the call setup as received by the 8260 on its UNI port (interWATCH 95000 capture).

----- HEX DUMP -----			
----f---gX-----	09 03 00 01 66 05 80 00 67 58 80 00 09 05 8c 05	0000	
-----Y-----	ec 81 05 ec 84 00 59 80 00 09 84 00 08 00 85 00	0010	
---[------_---k@	08 00 be 5e 80 00 03 10 80 80 5f 80 00 09 6b 40	0020	
---->--p---G-y-	80 80 00 a0 3e 00 01 70 80 00 15 82 47 00 79 00	0030	
----->-----	00 00 00 00 00 00 00 00 00 00 a0 3e 00 00 01 00	0040	
l-----9-----	6c 80 00 16 02 80 39 99 99 99 99 99 00 00 99	0050	
----4Vx---\-----	99 01 02 12 34 56 78 90 12 00 5c 80 00 02 00 00	0060	
---+-----	08 00 00 2b 00 00 00 00 00 00 00 00 00 00 00	0070	
-----t W>-	00 00 00 00 00 00 00 00 00 00 00 74 6c 57 3e d0	0080	

Figure 131. NAVTEL InterWATCH 95000 Call Setup Binary Dump

5.8.8.3 Understanding the Trace Content

Figure 132 is the result of the trace started in 8260#1 as explained in Figure 129 on page 198.

```

CPT      Sig. buffer (17-2-0-5):                               1      A
0000000019F2D30000000000000000D0007400001302E5D305A6424E00000050 2
0903000166                                                    3
05800067                                                       4
58800009058C05EC8105EC8400                                    5
598000098400080085000800BE                                  6
5E800003108080                                               7
5F8000096B40808000A03E0001                                  8
70800015824700790000000000000000000000A03E00000100        9
6C8000160280399999999999999900009999010212345678901200   10
5C80000200000800002B                                         11

T_PS      Route request primitive received                    B
T_PS      Processing unicast UBR call.
T_PS      Look for best vertex according to restrictive cost.
T_PS      Starting candidate list scan ...
T_PS      Vertex: 4 Cost: 0xffffc66b6                          C
T_PS      Vertex: 3 Cost: 0xffffa6c01
T_PS      End of candidate list scan.
T_PS      Dumping DTL for vertex 3 (1 hops, 54 bytes):        D
T_PS      60.a0.39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.60.00.07.00. 0x0
T_PS      60.a0.39.99.99.99.99.99.00.00.99.99.01.01.40.00.82.60.02.00.00. 0x0

```

Figure 132. Output of 8260 Internal Trace

3 to **11** are labels used to identify the content of the call setup. Each label in Figure 132 has a corresponding label with the full frame decoding in Figure 130 on page 199.

Analyzing an 8260 trace can be done as follows:

1. Locate the call setup in the trace as illustrated in **A**. Of course, using the trace formatter will help for this identification. The objective of this search is to make sure that you found the right call setup message, but the interesting information is the route selection result, and this clearly appears in the trace as illustrated in Figure 131 on page 201.

Refer to Figure 130 on page 199 for the full setup message decoding and content.

2. Analyzing the route selection process as described from step **B** to step **D**:
 - Step **B** indicates that the path selection component is looking for a route.
 - Step **C** shows that two nodes (vertices) provide a route to the destination address.

- Vertex 3 provides a lower cost to reach this address. This is where the call will be routed, and this is done by building the Designated Transit List (DTL). The DTL will be added to the call setup and sent over the network. This illustrates the source routing implementation of PNNI. The first node in the network makes a route selection and builds a DTL, leaving no routing choice to intermediate switches in the path. A DTL is built as follows: the first node of the DTL is the source node (Vertex 0). The last node of the DTL is the destination node (inside the peer group). It may not be the final destination node if the connection goes over an IISP. The DTL also contains all the intermediate nodes when there is more than one hop to reach a destination. In this example, since the target is 1 hop away as mentioned in step **D** the DTL only contains the source and the destination nodes.

5.8.8.4 Conclusion

This was an example of the procedure to follow to capture data related to route selections inside the 8260/8285 ATM switch running PNNI. This was to illustrate and understand the data that can be gathered, and identify the route that will be followed by a given call. The following case study will focus on how to use it.

5.8.9 Case Study 9 - PNNI Path Selection Troubleshooting

This scenario gives information on how to set loading balancing in the PNNI environment.

5.8.9.1 Network Diagram

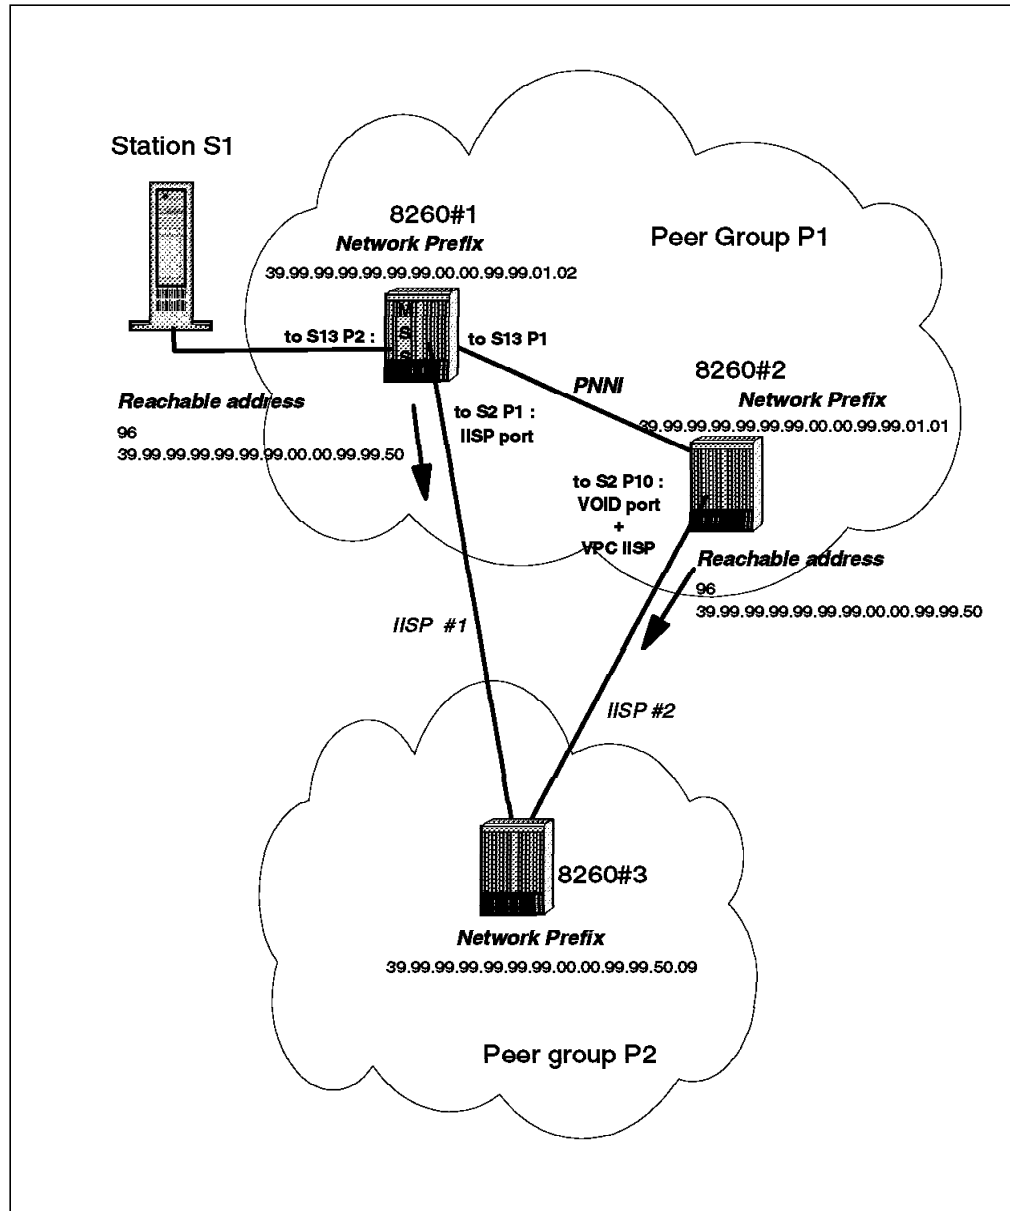


Figure 133. Network Configuration

5.8.9.2 Symptom

The purpose of the network design is to perform load balancing on the two parallel IISP links IISP#1 and IISP#2. The objective is to avoid overloading the PNNI link between 8260#1 and 8260#2 in normal conditions. The expected result is that calls issued from 8260#2 will reach the peer group P2 through IISP#2, and that calls issued from 8260#1 will reach peer group #1 through IISP#1, both nodes using the alternate IISP only as a backup path. By starting node monitoring on both 8260#1 and 8260#2 from the network management station, it

appears that one of the ports, IISP#2, is regularly supporting a heavy load, while the other IISP port reports less activity.

5.8.9.3 Methodology

Investigation can start by verifying the behavior observed by the network management station. For instance, this is done by looking at the cross connection on the port of station S1 connected to 8260#1 that has an SVC established with another station located in the remote peer group P2. Figure 134 shows the cross connection results.

```
8260ATM#1> show signalling cross_connections port 13.2
In: slot.port vpi.vci type Out: slot.port vpi.vci type Conn Cat
-----
      13.2      0.775  SVC      13.1      0.176  SVC      P2P  UBR
      13.2      0.776  SVC      13.1      0.177  SVC      P2P  UBR
      13.2      0.777  SVC      13.1      0.179  SVC      P2P  UBR

Total number of cross connections = 3
```

Figure 134. Cross Connections from Port 13.2

All the SVCs issued from station S1 are routed through the PNNI link to reach peer group P2 through IISP#2 instead of using IISP#1 while it is up and running. This is of course easily illustrated when tracking these SVCs from the network management station.

Understanding that this path is selected from 8260#1 is now required. This can be done by starting an internal trace for both pnni_path_selection and signalling_messages as explained in the previous section, and then restarting the test station S1. This trace gives the following results:

```
CPT      Sig. buffer ( 13-2 -0-5):
0000000019F47EC00000999900B2EA007410201596A464059D424E00000050
090300004B 0580 006658800008058C05EC8105EC84598000098400080085000800BE
5E8000031080805F8000096B40808000A03E0001
7080001582 3999999999999900009999500940008210007213
6C8000160280 3999999999999900009999010212345678901200 5C8000020000004800000E
T_PS Route request primitive received
T_PS Processing unicast UBR call.
T_PS Look for best vertex according to restrictive cost. 1
T_PS Starting candidate list scan ...
T_PS Vertex: 3 Cost: 0xffff47be 2
T_PS Vertex: 0 Cost: 0xffffa3df 3
T_PS End of candidate list scan.
T_PS Dumping DTL for vertex 3 (1 hops, 54 bytes):
T_PS 60.a.0.39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.60.00.07.00. 0x0
T_PS 60.a.0.39.99.99.99.99.99.00.00.99.99.01.01.40.00.82.60.02.00.00. 0x0
```

Figure 135. Output of 8260#1 Internal Trace

The signalling message trace allows you to trace the call setup and identify the calling address, called address and source port as highlighted in Figure 135. Also make sure that this message is really the call setup (0580).

A few lines after the occurrence of the call, the path selection trace gives useful information about the route selected, and why it was selected.

This indicates that the path was selected based on restrictive cost **1**, and that the vertex 3, offers the lower cost based on this criteria. **2** and **3** show the cost of each path to reach the target destination. Dumping the PNNI path selection data will allow you to fully understand what happens in this node. It will particularly allow you to understand the key parameters used for route selection. Figure 136 shows the interesting part of the dump.

```
Starting PNNI Path Selection Dump at Thu Jun 26 10:02:35 1997

Path selection methods dump
-----
      UBR computation:
          Widest Path Modified Dijkstra algorithm           1
          Last computation was: Thu Jun 26 10:01:38 1997

--More--

Graph dump
-----
Vertex: 0
Node Identifier:
60.A0.39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.60.00.07.00.
Vertex: 3                                     2
Node Identifier:
60.A0.39.99.99.99.99.99.00.00.99.99.01.01.40.00.82.60.02.00.00.

Reachability dump
-----
Vertex: 3
1 prefix(es)
  Prefix 0 :39.99.99.99.99.99.99.00.00.99.99.50. (96 bits)

      MaxCR  AvCR  HwmCR  AdmWght  RstCost  AddCost  MctCost
UBR    47169  N/A    47169   5040    0xffff47be  5040    0xffff47be 3
ABR    47169  40093  N/A     5040    0xfffeab21  5040    0xfffeab21
NRT VBR 40093  40093  N/A     5040    0xfffec6c5  5040    0xfffec6c5
RT VBR 40093  40093  N/A     5040    0xfffec6c5  5040    0xfffec6c5
CBR    40093  40093  N/A     5040    0xfffec6c5  5040    0xfffec6c5

Vertex: 0
1 prefix(es)
  Prefix 0 :39.99.99.99.99.99.99.00.00.99.99.50. (96 bits)

      MaxCR  AvCR  HwmCR  AdmWght  RstCost  AddCost  MctCost
UBR    23584  N/A    23584   5040    0xffffa3df  5040    0xffffa3df 4
ABR    23584  20046  N/A     5040    0xffff5591  5040    0xffff5591
NRT VBR 20046  20046  N/A     5040    0xffff6363  5040    0xffff6363
RT VBR 20046  20046  N/A     5040    0xffff6363  5040    0xffff6363
CBR    20046  20046  N/A     5040    0xffff6363  5040    0xffff6363
```

Figure 136. PNNI Path Selection Dump for 8260#1

1 indicates that the algorithm used in this node for UBR is the widest path algorithm. This is the default algorithm used for UBR in 8260/8285. Path selection methods used for a given switch are displayed by using the following command:

```
8260ATM#1> show pnni path_selection
Unspecified bit rate : widest path.
Available bit rate : precomputed path.
```

2 identifies 8260#2 as being vertex 3 in this network topology.

Note: The relation between vertices and nodes is dynamic, and needs to be done each time a node has been restarted.

The reachability dump allows you to see the Resource Availability Information Group (RAIGS) associated with each vertex for this particular destination network prefix. Since this is a UBR call, the corresponding information for each vertex is located in lines **3** and **4**. From this information, it is possible to understand why the calls are routed to 8260#2:

1. The path selection algorithm used is the widest path, so the administrative weight configured for these links is not used.
 2. The cost used for path selection is the restrictive cost (RstCost), and this cost appears to be smaller for vertex 3: 8260#2 (**3**). The following is an overview of the parameters that are used for restrictive cost calculation:
 - The restricted cost is obtained by the following rule: $RstCost = 0xffffffff - 0xMaxCr$ where $0xMaxCr$ is the Maximum Cell Rate in hexadecimal notation. $MaxCr$ is initialized as follows:
 - For PNNI or IISP ports:
It is the value of the NNI bandwidth parameter set for the port.
 - For VOID port + VPC_Link:
It is the value of the VPC bandwidth parameter set for the VPC link.
- $MaxCr$ is then updated according to an algorithm based on the number of calls running on the resource (port or VPC_link). This value is in cells/second while the bandwidth specified in the configuration is in kbps.

Viewing the configuration of both IISP links allows you to highlight the parameter settings for the calls being routed through IISP#2.

1. Parameters for IISP#1:

```

8260ATM#1> show port 2.1 verbose

      Type Mode      Status
-----
2.01: IISP enabled  UP

Signalling Version : 3.1
No ILMI
NNI Bandwidth      : 10000 kbps
RB Bandwidth       : unlimited
Signalling role    : network
Signalling vci     : 0.5
Administrative weight: 5040
VPI.VCI range      : 3.1023 (2.10 bits)
Connector          : RJ45
Media              : copper twisted pair
Port speed         : 25600 kbps
Remote device is active

```

2. Parameters for IISP#2:

```

8260ATM#2> show vpc_link 2.10 0 verbose

      VPI : Type Mode      Status
-----
2.10   0: IISP enable  UP
Signalling Version : 3.1
No ILMI
VPC Bandwidth      : 20000 kbps
RB Bandwidth       : unlimited
Signalling role    : network
Signalling vci     : 0.5
Administrative weight: 5040
VPI.VCI range      : 3.1023 (2.10 bits)
VPCI              : 0

```

The bandwidth specified was 10000 kbps for IISP#1, while it was set to 20000 kbps for IISP#2.

Note: 47169Cells/s ($0\text{x}B841$ in hexa) = $20000\text{ kbps} / (53\text{ bytes} \times 8\text{ bits})$. The calculation of the associated restrictive cost gives $0\text{xfffffff}-0\text{xb}841=0\text{xffff}47\text{be}$.

To perform the load balancing expected by the network designer, the two nodes can be configured with the shortest path algorithm for route selection.

Figure 137 illustrates how to proceed.

```

8260ATM#1> set pnni path_selection ubr:shortest_path
To activate issue COMMIT after your last 'set pnni...' entry.
To cancel all changes since previous COMMIT, issue UNCOMMIT.
8260ATM#1> commit
Enter parameter: pnni
COMMIT successfully executed.
To save new configuration issue SAVE.
8260ATM#1> save all
8260ATM#1> show pnni path_selection
Unspecified bit rate : shortest path.
Available bit rate : precomputed path.

```

Figure 137. Changing UBR Path Selection Method

Since the default value for administrative weight on all links is 5040, there should be no other tuning required to achieve load balancing based on the location of the calling station.

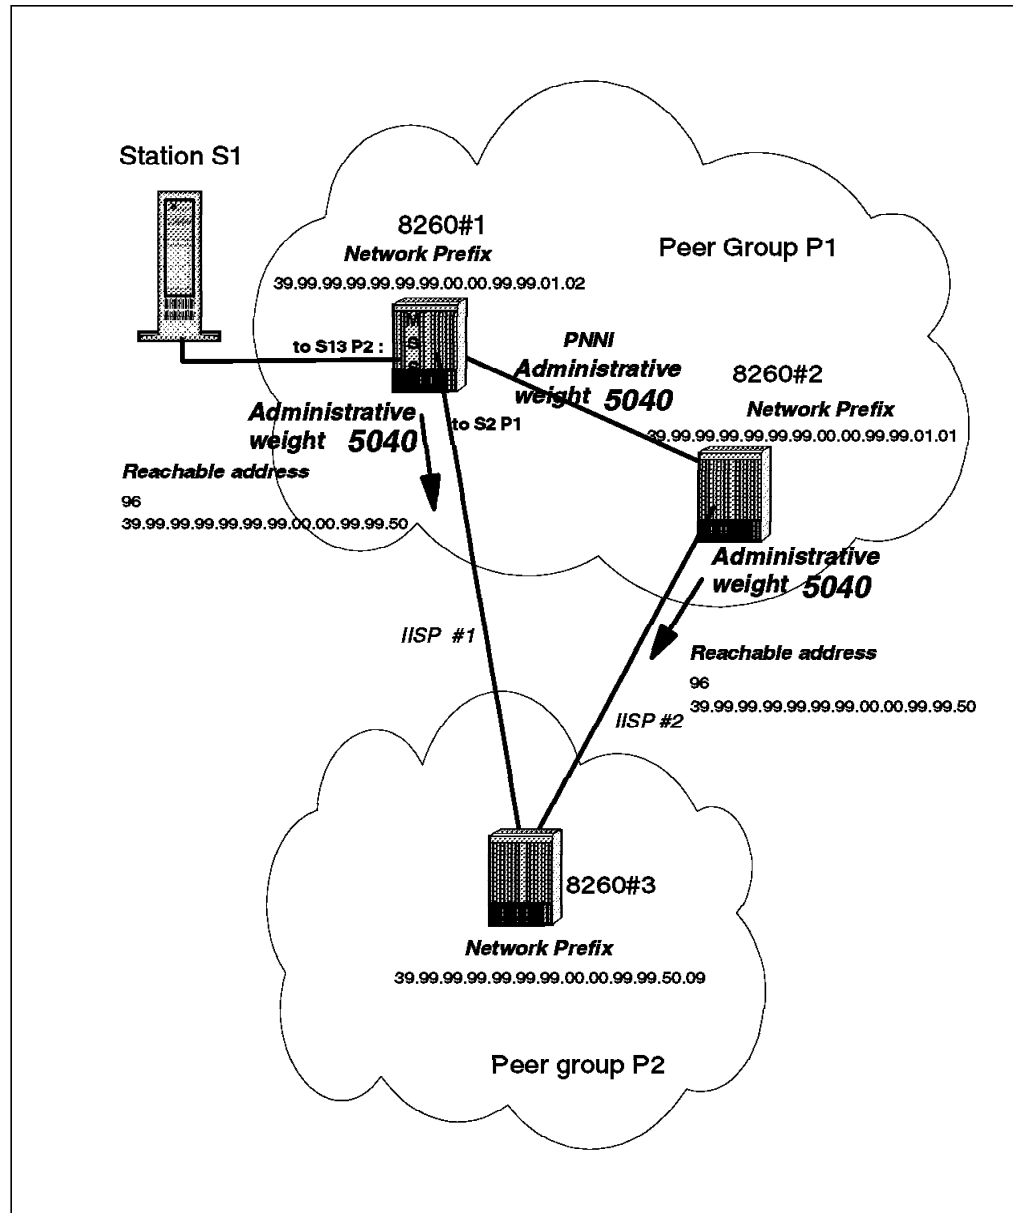


Figure 138. Administrative Weight Configuration on the Links

Using this path selection method, the local IISP link (zero hop) will always be preferred to the remote IISP link (one hop) since the cost to reach peer group P2 will be composed of both the IISP administrative weight and the PNNI link administrative weight.

After the hub reconfiguration illustrated in Figure 137 on page 208, the display of the cross connections for the port to which station S1 is connected gives the following results:

```

8260ATM#1> show signalling cross_connections port 13.2
In: slot.port vpi.vci  type  Out: slot.port vpi.vci  type      Conn  Cat
-----
      13.2      0.785  SVC      2.1      0.76  SVC      P2P  UBR
      13.2      0.786  SVC      2.1      0.78  SVC      P2P  UBR
      13.2      0.787  SVC      2.1      0.79  SVC      P2P  UBR

Total number of cross connections = 3

```

The calls are routed through the local IISP link. The route selection now behaves as expected: all the UBR calls issued from 8260#1 will use IISP#1 if it is available, while all the UBR calls issued from 8260#2 will use IISP#2 if available. Each IISP can, however, be used by both switches if one fails. This configuration only applies to outgoing calls, that is, calls going from peer group P1 to peer group P2. Some more configuration needs to be done in 8260#3 if the network administrator wants the same behavior for incoming calls.

Figure 139 shows the capture of the internal trace when restarting station S1 with this new configuration.

```

47005 06]26 10:25:30 CPT      Sig. buffer (13-2-0-5):
0000000019F4E600000000000000B4120074008007D4A4640647424E00000050
090300004B0580006658800008058C05EC8105EC84598000098400080085000800BE
5E8000031080805F8000096B40808000A03E0001
70800015823999999999999900009999500940008210007213
6C80001602803999999999999999000099990102123456789012005C8000020000004800002E
47010 06]26 10:25:30 T_PS Route request primitive received
47011 06]26 10:25:30 T_PS Processing unicast UBR call.
47012 06]26 10:25:30 T_PS Look for best vertex according to additive cost. 1
47013 06]26 10:25:30 T_PS Starting candidate list scan ...
47014 06]26 10:25:30 T_PS Vertex: 0 AddCost: 5040
47015 06]26 10:25:30 T_PS Vertex: 3 AddCost: 10080 2
47016 06]26 10:25:30 T_PS End of candidate list scan.
47017 06]26 10:25:30 T_PS Dumping DTL for vertex 0 (0 hops, 27 bytes):
47018 06]26 10:25:30 T_PS
60.a0.39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.60.00.07.00. 0x0

```

Figure 139. Path Selection Trace when Using Shortest Path

The trace shows:

- 1** the cost used for path selection is now the additive cost (AddCost).
- 2** as expected from configuration, vertex 0 (the local node) offers the best cost (5040), while vertex 3 offers a cost of 10080. 10080 is the cost of the remote IISP link + the cost of the PNNI link. Figure 140 on page 211 is the associated dump of the PNNI path selection, and allows you to understand the cost calculation for AddCost based on the topological database content.


```

Starting PNNI Path Selection Dump at Thu Jun 26 10:25:51 1997

Path selection methods dump
-----
      UBR computation:
          Modified Bellman-Ford shortest path algorithm
          Last computation was: Thu Jun 26 10:25:50 1997
Spanning trees dump
-----
      UBR spanning tree:
          3->(cost: 0xffffa6c01, dist: 1, AddCost: 5040, PortId: 0xa1000000)
          0->(cost: 0x0, dist: 0, AddCost: 0, PortId: 0x0)
--More--

Graph dump
-----
Vertex: 0
Node Identifier:
60.A0.39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.60.00.07.00.
Vertex: 3
Node Identifier:
60.A0.39.99.99.99.99.99.00.00.99.99.01.01.40.00.82.60.02.00.00.
Reachability dump
-----

Vertex: 0
1 prefix(es)
Prefix 0 :39.99.99.99.99.99.99.99.00.00.99.99.50. (96 bits)

      MaxCR  AvCR  HwmCR  AdmWght  RstCost      AddCost  MctCost
UBR    47169  N/A    47169   5040    0xffff47be   5040    0xffff47be
ABR    47169  40093  N/A     5040    0xfffeab21   5040    0xfffeab21
NRT VBR 40093  40093  N/A     5040    0xfffec6c5   5040    0xfffec6c5
RT VBR 40093  40093  N/A     5040    0xfffec6c5   5040    0xfffec6c5
CBR    40093  40093  N/A     5040    0xfffec6c5   5040    0xfffec6c5

Vertex: 3
1 prefix(es)
Prefix 0 :39.99.99.99.99.99.99.99.00.00.99.99.50. (96 bits)

      MaxCR  AvCR  HwmCR  AdmWght  RstCost      AddCost  MctCost
UBR    47169  N/A    47169   5040    0xffff47be   5040    0xffff47be
ABR    47169  40093  N/A     5040    0xfffeab21   5040    0xfffeab21
NRT VBR 40093  40093  N/A     5040    0xfffec6c5   5040    0xfffec6c5
RT VBR 40093  40093  N/A     5040    0xfffec6c5   5040    0xfffec6c5
CBR    40093  40093  N/A     5040    0xfffec6c5   5040    0xfffec6c5

```

Figure 140. PNNI Path Selection Dump

This dump shows the following information:

- 1** The shortest path is now used.
- 2** In the spanning tree dump section, 3-> indicates reachability information for vertex 3.

Vertex 3 (8260#2) has an AddCost of 5040. This is the cost of the PNNI link **3**.

4 shows that both IISP links also have additive cost equal to their administrative weight (AdmWgth).

To summarize: The path selection for UBR can be done based on two algorithms:

1. Shortest path: this is mainly static tuning since the main parameter used for path selection is based on the administrative weight parameter. This was the tuning requested in this configuration.
2. Widest path: This offers dynamic path selection based on the dynamic updates of link costs in the network. Initial values are based on the NNI bandwidth or VPC bandwidth parameters.

Note: The criteria used for updating the cost is the number of connections and not the traffic flowing on a given connection.

5.8.9.4 About Network Prefix Length

From a topology standpoint, the routing to a given address is done based on the longest match on the network prefix that matches the network prefix of the destination address. In the previous scenario there were parallel links to a remote peer group, and the nodes 8260#1 and 8260#2 had identical reachable addresses configured to reach peer group P2. Identical reachable addresses mean that both network prefix and network prefix length are identical. This is the only configuration where the topology will consider having two parallel paths and then make a decision based on cost. This cost will be a restrictive cost for widest path or additive cost for shortest path. In the previous scenario, the best match for the destination address was on 96 bits and it was advertised by both 8260#2 and 8260#1. Any difference in the length of the network prefix used to configure a reachable address immediately gives priority to the longest prefix.

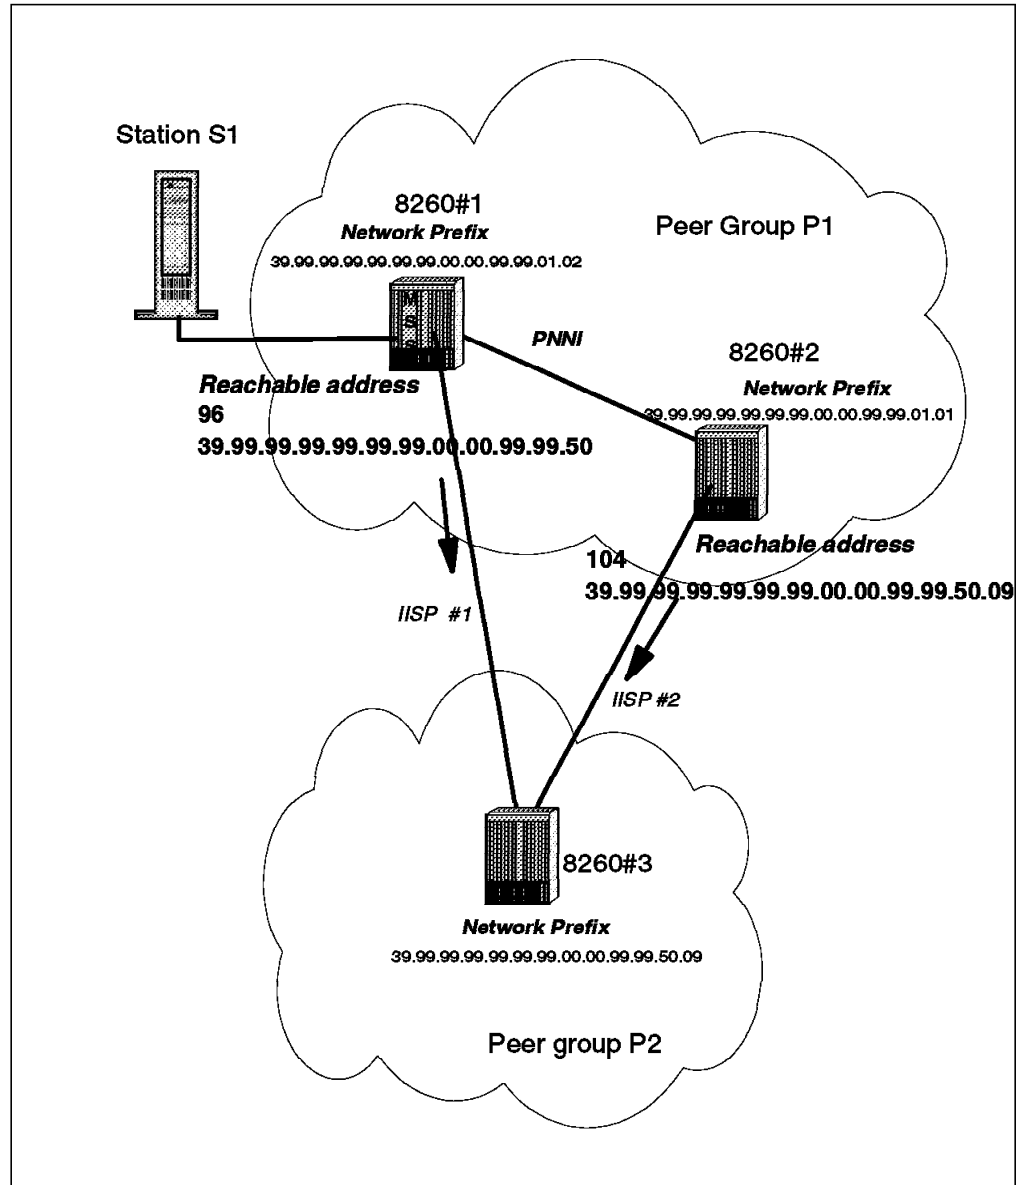


Figure 141. Reachability with Different Network Prefix Length

Figure 142 illustrates the result of the path selection in a configuration using different prefix lengths as shown in Figure 141.

```

T_PS Route request primitive received
T_PS Processing unicast UBR call.
T_PS Look for best vertex according to restrictive cost.
T_PS Starting candidate list scan ...
T_PS Vertex: 3 Cost: 0xffff1426
T_PS End of candidate list scan.
T_PS First access to vertex 3. Computing the DTL.
T_PS Dumping DTL for vertex 3 (1 hops, 54 bytes):
T_PS 60.a0.39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.60.00.07.00. 0x0
T_PS 60.a0.39.99.99.99.99.99.00.00.99.99.01.01.40.00.82.60.02.00.00. 0x0
    
```

Figure 142. Path Selection Trace when Using Different Prefix Lengths

Configuring the reachability with different prefix lengths as described in the previous figure immediately turns IISP#2 into *THE* path to reach peer group P2. IISP#1 will be used from peer group P1 only if IISP#2 is down and the path selection will not be based on any cost since the topology considers the destination address as reachable only through one link as highlighted by **1** in Figure 142 on page 213.

Note: This illustrates the fact that the length of the network prefix is a key parameter when configuring reachability information. It also shows the impact of using different prefix lengths when configuring reachability information on parallel links to a given destination.

In other words, prefix length takes precedence over any link cost consideration.

5.8.10 About Duplicate Addresses

The following figure depicts a network where both an 8210 and an MSS are connected to the same hub.

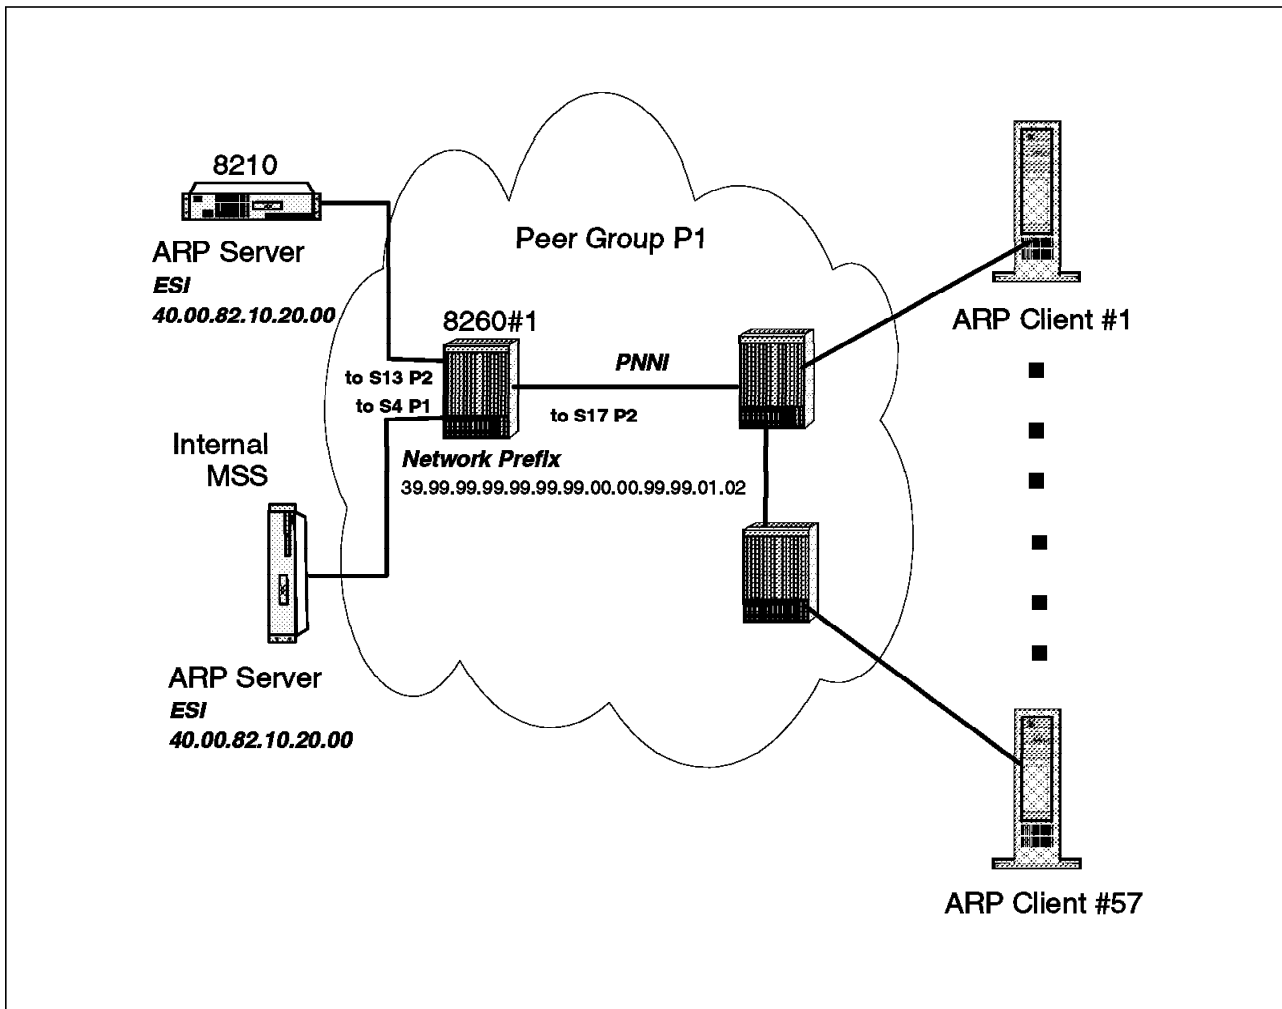


Figure 143. Multiple 8210 Connected to the Same Hub

The MSS and 8210 are configured to provide redundant ARP Server functionality. This means that both the MSS and the 8210 will try to register the same ATM address with the switch.

This was not allowed in an 8260 not running PNNI, but it is allowed in a switch running PNNI. It is the default configuration as illustrated by the following command:

- Extract of show device command:

```
8260ATM#1> show device
8260 ATM Control Point and Switch Module
Name : 8260ATM
Location :
Cary

For assistance contact :
Advanced Technical Support Group

Manufacture id: VIME
Part Number: 51H3659 EC Level: E28028
Boot EEPROM version: v.3.0.0
Flash EEPROM version: v.3.0.0
Flash EEPROM backup version: v.2.5.1
--More--
Device configured for PNNI port capability. No LES can start.
Dynamic RAM size is 16 MB. Migration: off. Diagnostics: enabled.
Device defined as primary.
Duplicate ATM addresses are allowed.
8260ATM#1>
```

Displaying the reachable address will give a status on ARP server address registration.

```
8260ATM#1> show reachable_address all
Port Len Address Active Idx VPI
-----
3.02 152 39.99.99.99.99.99.00.00.99.99.01.02.08.00.5A.99.86.DC Y Dyn 0
4.01 152 39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.00.69 Y Dyn 0
4.01 152 39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.20.00 Y Dyn 0 1
6.02 152 39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.81.00.66 Y Dyn 0
6.03 152 39.99.99.99.99.99.00.00.99.99.01.02.40.00.05.70.00.51 Y Dyn 0
13.02 152 39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.00.72 Y Dyn 0
13.02 152 39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.20.00 Y Dyn 0 2
17.02 152 39.99.99.99.99.99.00.00.99.99.01.02.00.00.C1.10.C8.A0 Y Dyn 0
8260ATM#1>
```

Both ARP Server ATM addresses are registered to the switch. One of these addresses is registered to port 4.1, **1** while the other is registered to port 13.2. **2**.

Following is the configuration of the ATM UNI ports on which the duplicate addresses are registered.

- Port 4.1 configuration:

```
8260ATM#1> show port 4.1 verbose
```

Type	Mode	Status	Daughter Card Description
4.01:	UNI	enabled UP	ATM MSS Server

```

Signalling Version : Auto
> Oper Sig. Version : 3.0
ILMI status       : UP
ILMI vci          : 0.16
RB Bandwidth      : unlimited
Signalling vci    : 0.5
Administrative weight: 5040
VPI.VCI range     : 15.1023 (4.10 bits)
Connector         : NONE
Media             : none
Port speed        : 155000 kbps
Remote device is active

```

```
ATM MSS Server Card Informations:
```

```

P/N:13H6775 EC level:051125 Manufacture: RALO
Model Number      : 5300
Boot Code Version : v.1.0.0
Operational Code Version : v.1.0.0
IP address        : 192.2.1.69
IP Subnet Mask    : 0.0.0.0

```

Physical Port:	Status	Network	Speed	Connector	MAC Address
01	OKAY	ATM	N/A	Backplane	002035990132

- Port 13.2 configuration:

```
8260ATM#1> show port 13.2 verbose
```

Type	Mode	Status
13.02:	UNI	enabled UP

```

Signalling Version : 3.1
> Oper Sig. Version : 3.1
ILMI status       : UP
ILMI vci          : 0.16
RB Bandwidth      : unlimited
Signalling vci    : 0.5
Administrative weight: 5040
VPI.VCI range     : 15.1023 (4.10 bits)
Connector         : SC DUPLEX
Media             : multimode fiber
Port speed        : 155000 kbps
Remote device is active

```

```

Frame format      : SONET STS-3c
Scrambling mode   : frame and cell
Clock mode        : internal

```

- Path selection used:

```

8260ATM#1> show pnni path_selection
Unspecified bit rate : shortest path.
Available bit rate   : precomputed path.
8260ATM#1>

```

The switch is configured to use the shortest path and the administrative weight of both UNI ports is the default value that is **5040**.

The following figures depict the connections being established to an ARP server after the switch 8260#1 has been migrated to PNNI. This illustrates the case where 57 stations try to register almost simultaneously to the ARP server. These ARP clients reach ARP servers of 8260#1 through the port 17.2: a PNNI link. Listing the cross connections established from this port indicates to which UNI ports the calls have been routed.

```
8260ATM#1> show signalling cross_connections port 17.2
In: slot.port vpi.vci type Out: slot.port vpi.vci type Conn Cat
-----
```

In: slot.port	vpi.vci	type	Out: slot.port	vpi.vci	type	Conn	Cat
17.2	0.49	SVC	13.2	0.448	SVC	P2P	UBR
17.2	0.50	SVC	13.2	0.451	SVC	P2P	UBR
17.2	0.51	SVC	13.2	0.452	SVC	P2P	UBR
17.2	0.52	SVC	13.2	0.453	SVC	P2P	UBR
17.2	0.53	SVC	13.2	0.454	SVC	P2P	UBR
17.2	0.54	SVC	13.2	0.455	SVC	P2P	UBR
17.2	0.55	SVC	13.2	0.456	SVC	P2P	UBR
17.2	0.56	SVC	13.2	0.457	SVC	P2P	UBR
17.2	0.57	SVC	13.2	0.458	SVC	P2P	UBR
17.2	0.58	SVC	13.2	0.459	SVC	P2P	UBR
17.2	0.59	SVC	13.2	0.460	SVC	P2P	UBR
17.2	0.60	SVC	13.2	0.461	SVC	P2P	UBR
17.2	0.61	SVC	13.2	0.462	SVC	P2P	UBR
17.2	0.62	SVC	13.2	0.463	SVC	P2P	UBR
17.2	0.63	SVC	13.2	0.464	SVC	P2P	UBR
17.2	0.64	SVC	13.2	0.465	SVC	P2P	UBR
17.2	0.65	SVC	13.2	0.466	SVC	P2P	UBR
17.2	0.66	SVC	13.2	0.467	SVC	P2P	UBR
17.2	0.67	SVC	13.2	0.468	SVC	P2P	UBR
17.2	0.68	SVC	4.1	0.105	SVC	P2P	UBR
17.2	0.69	SVC	13.2	0.469	SVC	P2P	UBR
17.2	0.70	SVC	13.2	0.543	SVC	P2P	UBR
17.2	0.71	SVC	13.2	0.544	SVC	P2P	UBR

Figure 144. Dynamic Load Balancing of Calls to a Duplicate Address (1 of 2)

17.2	0.72	SVC	4.1	0.411	SVC	P2P	UBR
17.2	0.73	SVC	13.2	0.545	SVC	P2P	UBR
17.2	0.74	SVC	4.1	0.413	SVC	P2P	UBR
17.2	0.75	SVC	13.2	0.547	SVC	P2P	UBR
17.2	0.76	SVC	4.1	0.420	SVC	P2P	UBR
17.2	0.77	SVC	13.2	0.549	SVC	P2P	UBR
17.2	0.78	SVC	4.1	0.424	SVC	P2P	UBR
17.2	0.79	SVC	13.2	0.551	SVC	P2P	UBR
17.2	0.80	SVC	4.1	0.425	SVC	P2P	UBR
17.2	0.81	SVC	13.2	0.552	SVC	P2P	UBR
17.2	0.82	SVC	4.1	0.426	SVC	P2P	UBR
17.2	0.83	SVC	13.2	0.553	SVC	P2P	UBR
17.2	0.84	SVC	4.1	0.427	SVC	P2P	UBR
17.2	0.85	SVC	13.2	0.554	SVC	P2P	UBR
17.2	0.86	SVC	4.1	0.428	SVC	P2P	UBR
17.2	0.87	SVC	13.2	0.555	SVC	P2P	UBR
17.2	0.88	SVC	13.2	0.556	SVC	P2P	UBR
17.2	0.89	SVC	13.2	0.557	SVC	P2P	UBR
17.2	0.90	SVC	4.1	0.432	SVC	P2P	UBR
17.2	0.91	SVC	13.2	0.558	SVC	P2P	UBR
17.2	0.92	SVC	4.1	0.433	SVC	P2P	UBR
17.2	0.93	SVC	13.2	0.559	SVC	P2P	UBR
17.2	0.94	SVC	4.1	0.435	SVC	P2P	UBR
17.2	0.95	SVC	13.2	0.560	SVC	P2P	UBR
17.2	0.96	SVC	4.1	0.436	SVC	P2P	UBR
17.2	0.97	SVC	13.2	0.561	SVC	P2P	UBR
17.2	0.98	SVC	4.1	0.437	SVC	P2P	UBR
17.2	0.99	SVC	4.1	0.438	SVC	P2P	UBR
17.2	0.100	SVC	4.1	0.439	SVC	P2P	UBR
17.2	0.101	SVC	13.2	0.562	SVC	P2P	UBR
17.2	0.102	SVC	4.1	0.440	SVC	P2P	UBR
17.2	0.103	SVC	13.2	0.563	SVC	P2P	UBR
17.2	0.104	SVC	4.1	0.441	SVC	P2P	UBR
17.2	0.105	SVC	13.2	0.564	SVC	P2P	UBR

Total number of cross connections = 57
8260ATM#1>

Figure 145. Dynamic Load Balancing of Calls to a Duplicate Address (2 of 2)

All these calls are calls to the ARP server ATM address.

This clearly highlights the dynamic behavior of the 8260. The switch starts to route the incoming calls (input 17.2) to 13.2 and then dynamically performs load balancing on both addresses registered to ports 13.2 and 4.1.

From a user (ARP client) standpoint, this leads to some of the stations being registered to the primary ARP server, while others are connected to the backup ARP server. It has in fact created two isolated logical IP subnets, one in each active ARP server. This will lead to a symptom where you get partial connectivity inside the logical IP subnet. The only easy way to identify such a problem is to dump the contents of both ARP server tables to see that they are both active at the same time.

How to Fix It: The duplicate address parameter seen when issuing a show device is device-related. This means that a given switch will let a duplicate address be registered. *This parameter applies to a single switch, and not to the network.* Disabling this parameter on the hub where the redundant ARP servers are connected will lead to switch behavior identical to a switch not running PNNI.

The second address registration will fail at ILMI while the first is alive. This is performed using the following command:

```
8260ATM#1> set device duplicate_atm_addresses:

Possible completions:
    allowed
    not_allowed

8260ATM#1> set device duplicate_atm_addresses:not_allowed
This call will reset the ATM subsystem.
Are you sure ? (Y/N) Y
```

This is then verified by a issuing a show device command:

```
8260ATM#1> show device
8260 ATM Control Point and Switch Module
--More--
Device configured for PNNI port capability. No LES can start.
Dynamic RAM size is 16 MB. Migration: off. Diagnostics: enabled.
Device defined as primary.
Duplicate ATM addresses are not allowed.
```

Figure 146. 8260 Configuration For Redundant ARP Servers

Other Workable Configurations: We want the calls to be routed to the primary ARP server, so what about using different administrative weight on the two UNI ports (4.1 and 13.2)?

Setting a lower administrative weight for the port hosting the primary ARP server may have the expected effect from a pure ATM standpoint; the calls will all be routed to that port when the address is alive. However, in the case of a temporary failure of the primary ARP server, some stations may register with the backup ARP server (since the primary is temporarily not accessible). When it comes up again, it is the preferred path (lower cost) and the new calls are routed to the primary ARP server. This has the same result as the initial configuration: the Logical IP Subnet is once again split between the two ARP Servers. **Therefore using different administrative weights on the UNI ports is not suitable for this configuration**

Redundant ARP Servers: Redundant ARP server support has been improved in MSS/8210 Release 1.1. One of these improvements is a RED Channel (monitoring VCC) being established between the primary ARP server and the redundant ARP server. It prevents the redundant ARP server from registering its ATM address while the primary ARP server is alive, and it resets the Redundant ARP Server's ATM address register when the primary ARP server recovers. The default 8260 configuration of duplicate ATM addresses may then be convenient in such an environment.

This illustrates a case where the ATM configuration needs to be made according to upper layer behavior. In this particular case, the ARP server redundancy behavior has changed between MSS/8210 Release 1.0 and MSS/8210 Release 1.1. It also highlights how important it is to fully understand the new functions coming with a given level of code, as well as the impacts on the overall

configuration. Understanding the changes before entering a migration may result in a safe migration from one level to the other. Not doing so, and discovering the problem at network restart generally leads to time and energy being lost in painful investigations. Reading the release notes, or any documentation coming with any new release of a product often turns out to be the winning approach.

Note: The call balancing depicted in Figure 145 on page 218 may also happen when the duplicate address being called is the LECS well-known Address. Therefore the recommendations about LECS configuration consistency given later in this section are also applicable to this particular configuration. Connecting both LECSs to the same hub is probably not the most reliable configuration one can expect. The most common configuration for LAN emulation redundancy may look like the following, where LECS are connected to different hubs.

5.8.10.1 Duplicate Addresses in the Network

Redundant LAN emulation services is often required to build a reliable network. This requires several LECSs to be installed in the network. Several LECSs in a network will lead to duplicate LECS well-known addresses advertised by one or several nodes. The following figure illustrates such a network.

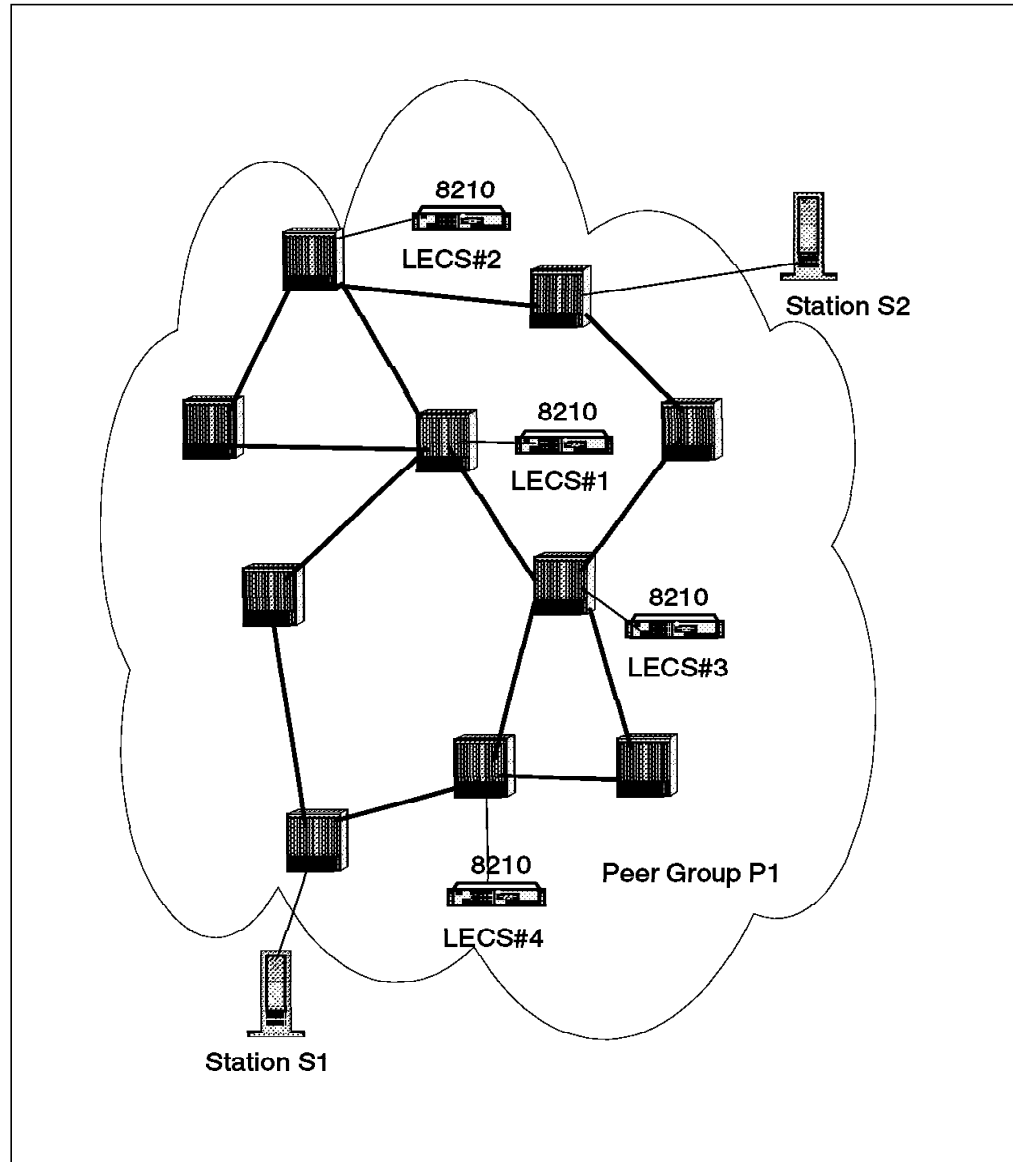


Figure 147. A Network with Multiple LECS

A natural question coming immediately when looking at this network configuration is: which LECS will be called by any given station? The answer to this question is given by understanding how the LECS WKA are processed in the topology database. Since the LECS registers the LECS WKA, the topology advertises the LECS WKA with a length of 152 bits (full address) as illustrated in the following extract of a PNNI path selection dump.

```
Vertex: 5
1 prefix(es)
Prefix 0 :
47.00.79.00.00.00.00.00.00.00.00.00.00.00.00.00.A0.3E.00.00.01. (152 bits)
```

Each node containing a LECS advertises this same address. The path selection for the calls with the LECS WKA as destination address will follow the UBR path selection rules as explained in the previous section. Now the match of the destination address is 152 bits while it was only 96 bits of the network prefix in the previous examples. Based on the algorithm selected for path selection and

the actual status of the routing parameters, this can lead to a station calling a LECS, for instance LECS#1 at a given time and the same station calling another LECS, for example LECS#3, a few seconds later. This will happen if the dynamic topology information modified the cost and created a lower cost path to LECS#3. This may not occur if the shortest path algorithm is used in all the switches; however, if the preferred LECS is not reachable, another LECS must be accessed for the LAN emulation services to be reliable. *This means that any station in a peer group may call any LECS in its peer group.* For this reason, one should not consider a given LECS as primary or secondary, but really consider the set of LECS as a distributed resource inside the peer group. This implies that all the LECS in a given peer group *must be* loaded with *the exact same configuration*. Not doing this will lead to unpredictable LAN emulation configurations. There is no way to prevent a group of stations attached to different hubs from calling a specific LECS located inside the same peer group, and allow, at the same time, this LECS to be called by a different group of stations.

Running Different LAN Emulation Domains in the Same Network: In large networks, or for security reasons, a network administrator may be interested in segmenting the emulated LAN into different domains. A LAN emulation domain can be defined as the set of resources (ELANs) a given LECS or set of LECSs has control over. These domains have completely different configurations rules, and one must prevent a given LEC from accessing the wrong LECS. This can only be achieved as illustrated in the following figure.

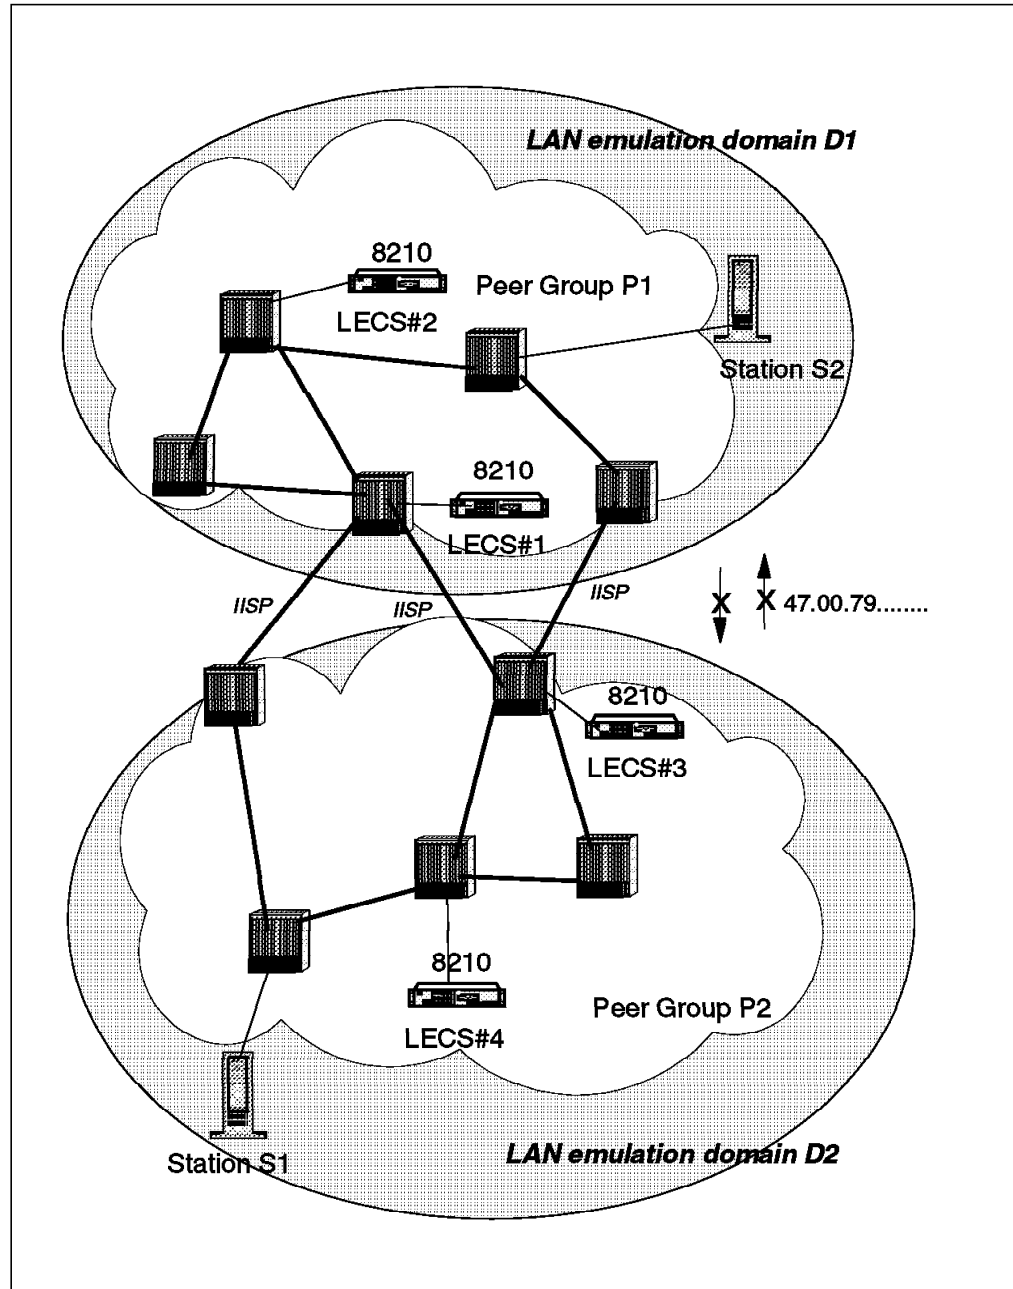


Figure 148. Creating LAN Emulation Domains

In the case of several domains, one must be sure that the clients of Domain D1 do not access LECS of domain D2 and vice versa. The only way to achieve this is to run at least one peer group per LAN emulation domain and not configure reachability information for the LECS WKA or part of the LECS WKA network prefix from one cluster to the other. This will ensure that station S1 reaches LECS#3 or LECS#4 but *never reaches* LECS#1 or LECS#2 and vice versa for Station S2.

5.8.10.2 Conclusion

PNNI provides a lot of flexibility in the way ATM addresses are handled in the network. As a result, it is really key to realize how powerful and simple it can be. It is important to design and implement the network according to the new functionalities offered by PNNI.

5.8.11 Case Study 10 - Reachability Configuration

This scenario will investigate activity reported when some LEDs seem abnormal.

5.8.11.1 Network Diagram

Three peer groups form the network.

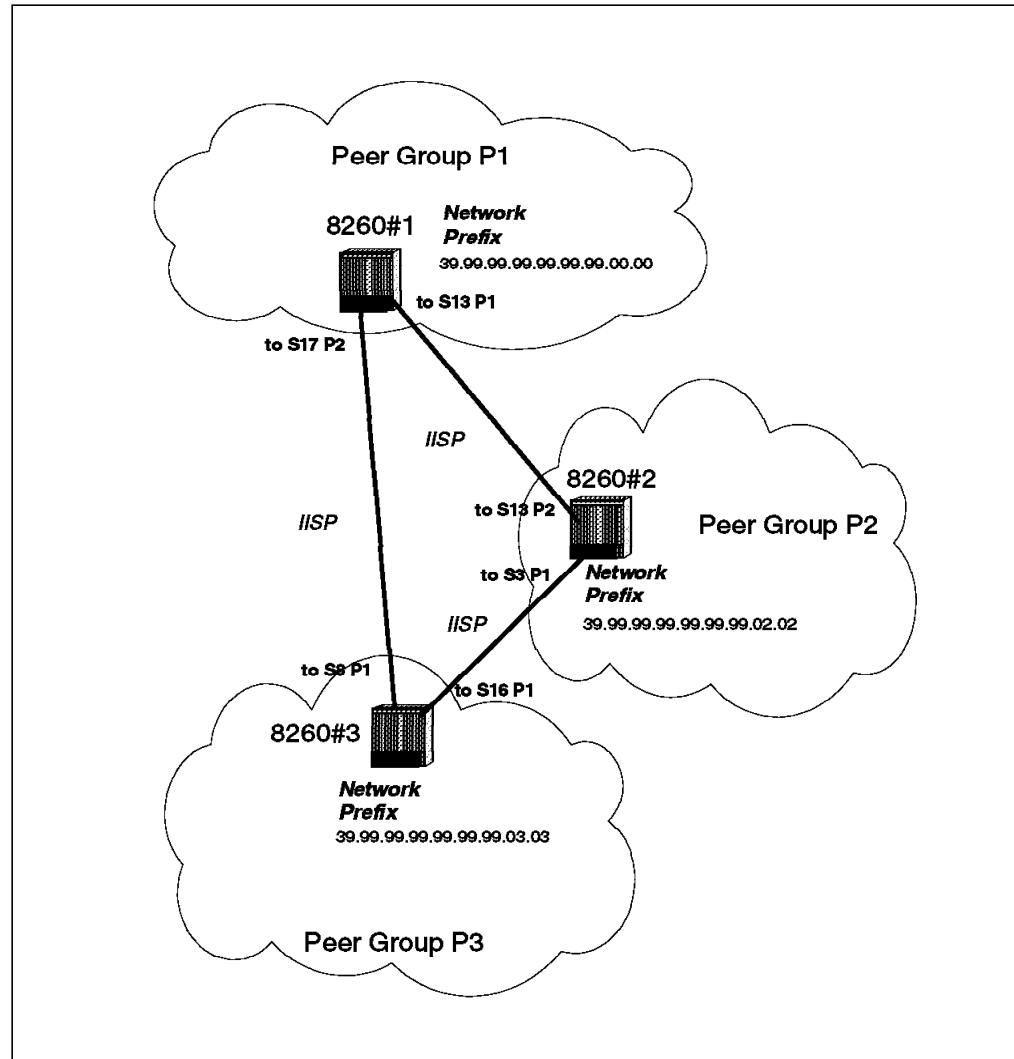


Figure 149. Network Configuration

5.8.11.2 Symptom

The activity LED of some ports is on but not blinking for several seconds during off hours.

5.8.11.3 Methodology

By looking at the ports whose activity LED is on, we find out that this occurs for IISP ports. Connecting to the switch and displaying port status reveals that everything is UP and running. Next step is to look at the SVC activity for one of these ports. This is done by issuing the following command on port 13.1, for instance:

```

8260ATM#1> show signalling cross_connections port 13.1
In: slot.port vpi.vci type Out: slot.port vpi.vci type Conn Cat
-----
--More--
13.1 0.230 SVC 17.2 0.786 SVC P2P UBR
13.1 0.231 SVC 17.2 0.787 SVC P2P UBR
13.1 0.232 SVC 17.2 0.788 SVC P2P UBR
13.1 0.239 SVC 17.2 0.795 SVC P2P UBR
13.1 0.240 SVC 17.2 0.796 SVC P2P UBR
13.1 0.241 SVC 17.2 0.797 SVC P2P UBR
13.1 0.242 SVC 17.2 0.798 SVC P2P UBR
13.1 0.243 SVC 17.2 0.799 SVC P2P UBR
13.1 0.250 SVC 17.2 0.806 SVC P2P UBR
13.1 0.251 SVC 17.2 0.807 SVC P2P UBR
13.1 0.252 SVC 17.2 0.808 SVC P2P UBR
13.1 0.253 SVC 17.2 0.809 SVC P2P UBR
13.1 0.254 SVC 17.2 0.810 SVC P2P UBR
13.1 0.261 SVC 17.2 0.817 SVC P2P UBR
13.1 0.262 SVC 17.2 0.818 SVC P2P UBR
13.1 0.263 SVC 17.2 0.819 SVC P2P UBR
13.1 0.264 SVC 17.2 0.820 SVC P2P UBR
13.1 0.265 SVC 17.2 0.821 SVC P2P UBR
13.1 0.272 SVC 17.2 0.828 SVC P2P UBR
13.1 0.273 SVC 17.2 0.829 SVC P2P UBR
13.1 0.274 SVC 17.2 0.830 SVC P2P UBR
13.1 0.275 SVC 17.2 0.831 SVC P2P UBR
13.1 0.276 SVC 17.2 0.832 SVC P2P UBR
13.1 0.283 SVC 17.2 0.839 SVC P2P UBR
13.1 0.284 SVC 17.2 0.840 SVC P2P UBR
13.1 0.285 SVC 17.2 0.841 SVC P2P UBR
13.1 0.286 SVC 17.2 0.842 SVC P2P UBR
13.1 0.287 SVC 17.2 0.843 SVC P2P UBR
13.1 0.294 SVC 17.2 0.850 SVC P2P UBR
13.1 0.295 SVC 17.2 0.851 SVC P2P UBR
13.1 0.296 SVC 17.2 0.852 SVC P2P UBR
13.1 0.297 SVC 17.2 0.853 SVC P2P UBR
13.1 0.298 SVC 17.2 0.854 SVC P2P UBR
13.1 0.305 SVC 17.2 0.861 SVC P2P UBR
13.1 0.306 SVC 17.2 0.862 SVC P2P UBR
13.1 0.307 SVC 17.2 0.863 SVC P2P UBR
13.1 0.308 SVC 17.2 0.864 SVC P2P UBR
--More--
Total number of cross connections = 435
8260ATM#1>
8260ATM#1>

```

Figure 150. Cross Connections from Port 13.1

From this output, we can see that:

1. There are quite a lot of connections established on this IISP port. The investigation takes place during off hours and the network is usually not loaded at this time.
2. These connections are almost all directed to the same output port, namely 17.2. This port is another IISP port.
3. Most ports have adjacent VCI values. This means that they have been allocated in a short period of time.

Immediately reissuing the same command gives the following results:

3. Reachability information for 8260#3:

```

8260ATM#3> show reachable_address all
Port Len Address Active Idx VPI
-----
 8.01  8 39. . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . Y  3  -
16.01  8 39. . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . Y  2  -
 8.02 152 39.99.99.99.99.99.99.99.03.03.99.99.01.05.02.00.82.72.00.0A Y Dyn  0
13.01 152 39.99.99.99.99.99.99.99.03.03.99.99.01.05.40.00.82.10.00.70 Y Dyn  0
13.01 152 39.99.99.99.99.99.99.99.03.03.99.99.01.05.40.00.82.10.20.70 Y Dyn  0
13.01 152 47.00.79.00.00.00.00.00.00.00.00.00.00.00.00.00.A0.3E.00.00.01 Y Dyn  0
15.02 152 39.99.99.99.99.99.99.99.03.03.99.99.01.05.40.00.03.40.00.13 Y Dyn  0
    
```

The reachability configuration configured on the IISP ports can be described as follows:

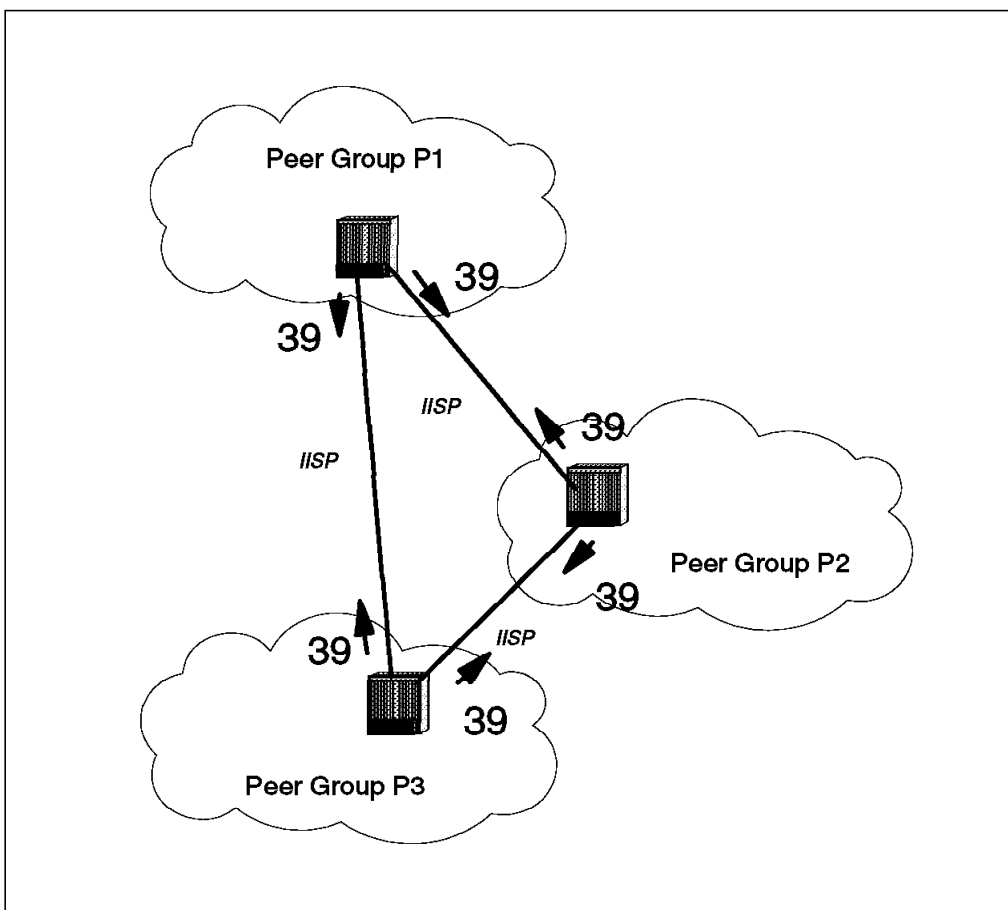


Figure 152. Reachability Configuration

Understanding of the Problem: Drawing the reachability information configured on each switch for this network topology clearly highlights a routing loop in the network.

When This Configuration Generates a Routing Loop: Key parameters of the addressing configuration:

- The three peer groups have a level identifier of 96 bits.
- The reachable address is 8 bits long.

Each time a switch receives a call setup for a destination address that does not match its peer group network prefix, but matches for instance one of the prefixes advertised as reachable on IISP links, it will forward this call through the IISP having reachability for this prefix. In this scenario, all the DCC ATM format addresses whose network prefixes do not match the network prefix and length of a PNNI peer group will circulate through the network. The reason is that they match the forwarding criteria without having any valid destination peer group to stop in. This call goes from switch to switch along the various IISPs. If the forwarding configuration creates a loop then the call will be processed along the loop path until one of the switches reaches VCC exhaustion. The call is then released with cause value: No VPCI/VCI available.

This explains the activity LED going on from time to time as well as the burst of SVC activity seen on the IISP links.

How Does It Happen?: This happens because the ATM addressing has been reorganized in the network; new network prefixes were assigned to the switches, and some stations had logical interfaces configured for Classical IP still trying to connect to the configured ARP server ATM address. Unfortunately, since the network prefixes have changed, this address does not exist anymore. Each time a station tries to call this address, the call setup enters the loop process.

The following is an example of reachability configuration that prevents the calls from circulating forever. For all these switches, the reachability information is configured so that the preferred IISP used to reach a neighboring peer group is the one that directly connects to this peer group. This is done by entering the longer prefix match on this IISP: in this configuration 72 bits. The other IISP is configured as a backup path if needed; the network prefix match is 64 bits.

Note: 72 bits and 64 bits were the minimum length that allowed you to differentiate a remote peer group's network prefix and provide preferred/alternate paths since the three peer groups have a common network prefix composed of the first 56 bits.

Calls that do not match the local peer group network prefix or any foreign address being registered inside this peer group are candidates to be forwarded through IISP.

Let's consider several cases:

1. The network prefix of the destination address in the call setup does not match any of the 64 first bytes configured as reachable addresses. The call is rejected with the cause: No route to destination.
2. The network prefix of the destination address in the call setup matches the first 72 bytes of the network prefix being reachable on a given IISP. The call will be routed to the neighboring peer group through the IISP that connects these two peer groups:
 - The network prefix of the destination address matches the network prefix of the remote peer group. The call is processed inside the peer group.
 - The network prefix of the destination address does not match the network prefix of the peer group. The peer group does not know how to reach this address and clears the call with the cause: No route to destination.
3. The network prefix of the destination address in the call setup only matches the first 64 bytes of the network prefix being reachable on a given IISP. The

call will be routed through one of the two IISPs based on routing parameters as explained in the previous scenarios. It will then be rejected by any of those with the cause: No route to destination.

Reachability Configuration after Switches Reconfiguration: Follow the information that is found on a good network setup.

1. Reachability configuration from 8260#1:

```
8260ATM#1> show reachable_address all
Port Len Address Active Idx VPI
-----
13.01 72 39.99.99.99.99.99.02.02. . . . . Y 3 -
13.01 64 39.99.99.99.99.99.03. . . . . Y 5 -
17.02 64 39.99.99.99.99.99.02. . . . . Y 4 -
17.02 72 39.99.99.99.99.99.03.03. . . . . Y 6 -
3.02 152 39.99.99.99.99.99.00.00.00.99.99.01.02.08.00.5A.99.86.DC Y Dyn 0
6.03 152 39.99.99.99.99.99.00.00.00.99.99.01.02.40.00.05.70.00.51 Y Dyn 0
13.02 152 39.99.99.99.99.99.00.00.00.99.99.01.02.40.00.82.10.00.69 Y Dyn 0
13.02 152 39.99.99.99.99.99.00.00.00.99.99.01.02.40.00.82.10.20.69 Y Dyn 0
13.02 152 47.00.79.00.00.00.00.00.00.00.00.00.00.00.00.A0.3E.00.00.01 Y Dyn 0
17.01 152 39.99.99.99.99.99.00.00.00.99.99.01.02.00.20.DA.72.09.D0 Y Dyn 0
```

2. Reachability configuration from 8260#2:

```
8260ATM#2> show reachable_address all
Port Len Address Active Idx VPI
-----
3.01 64 39.99.99.99.99.99.00. . . . . Y 3 -
3.01 72 39.99.99.99.99.99.03.03. . . . . Y 4 -
13.02 72 39.99.99.99.99.99.00.00. . . . . Y 2 -
13.02 64 39.99.99.99.99.99.03. . . . . Y 5 -
1.03 152 39.99.99.99.99.99.02.02.99.99.01.01.40.00.03.60.00.37 Y Dyn 0
1.04 152 39.99.99.99.99.99.02.02.99.99.01.01.40.00.05.80.00.48 Y Dyn 0
3.02 152 39.99.99.99.99.99.02.02.99.99.01.01.00.20.DA.70.70.80 Y Dyn 0
5.01 152 39.99.99.99.99.99.02.02.99.99.01.01.40.00.82.10.00.71 Y Dyn 0
5.01 152 47.00.79.00.00.00.00.00.00.00.00.00.00.00.00.A0.3E.00.00.01 Y Dyn 0
7.01 152 39.99.99.99.99.99.02.02.99.99.01.01.40.00.82.81.00.67 Y Dyn 0
14.02 152 39.99.99.99.99.99.02.02.99.99.01.01.12.34.56.78.90.12 Y Dyn 0
8260ATM#2>
```

3. Reachability configuration from 8260#3:

```
8260ATM#3> show reachable_address all
Port Len Address Active Idx VPI
-----
8.01 72 39.99.99.99.99.99.00.00. . . . . Y 2 -
8.01 64 39.99.99.99.99.99.02. . . . . Y 5 -
16.01 64 39.99.99.99.99.99.00. . . . . Y 3 -
16.01 72 39.99.99.99.99.99.02.02. . . . . Y 4 -
8.02 152 39.99.99.99.99.99.03.03.99.99.01.05.02.00.82.72.00.0A Y Dyn 0
13.01 152 39.99.99.99.99.99.03.03.99.99.01.05.40.00.82.10.00.70 Y Dyn 0
13.01 152 39.99.99.99.99.99.03.03.99.99.01.05.40.00.82.10.20.70 Y Dyn 0
13.01 152 47.00.79.00.00.00.00.00.00.00.00.00.00.00.00.A0.3E.00.00.01 Y Dyn 0
15.02 152 39.99.99.99.99.99.03.03.99.99.01.05.40.00.03.40.00.13 Y Dyn 0
8260ATM#3>
```

5.8.11.4 Configuring Reachability in a Peer Group with Two Outgoing IISP Links

Following are some examples of typical network configurations where a peer group is connected to the rest of the network using two IISPs. The purpose of this section is to answer the following question: How should I configure reachability information on a given IISP link?

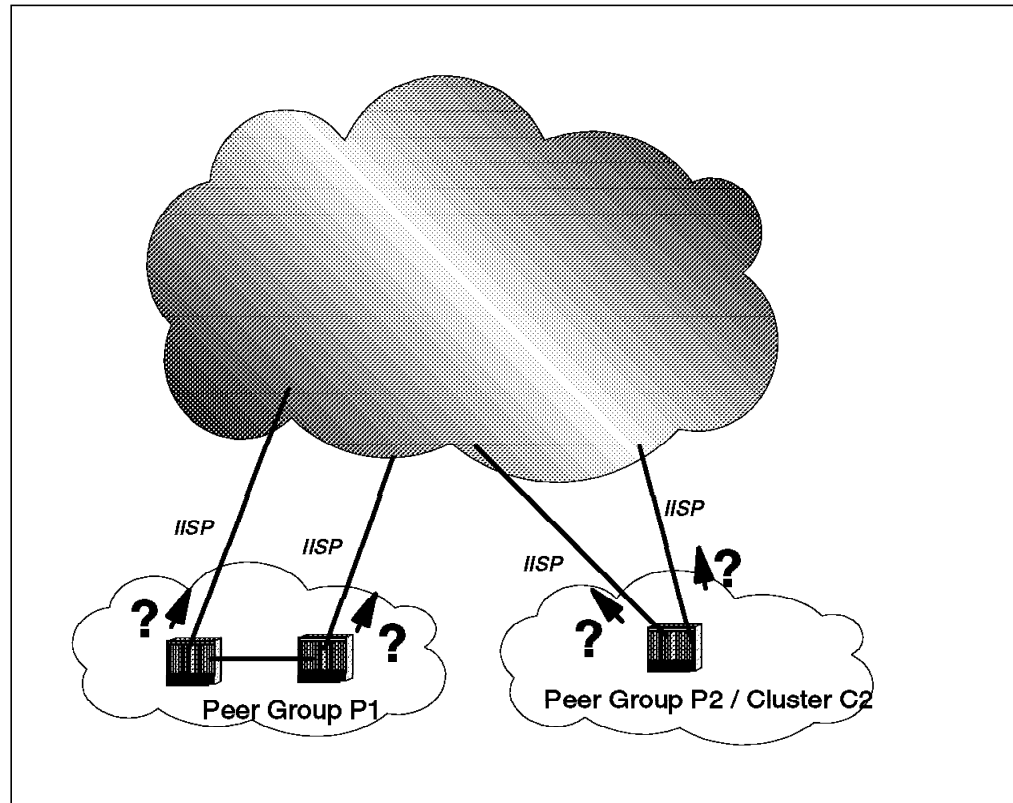


Figure 153. How to Configure Reachability Information on Parallel IISPs

The answer to this question is highly dependent on the two following parameters:

1. What is the network topology these links are connected to? The way reachability information should be configured is highly dependent on the topology being run on the other end of the IISP link.
2. What do we expect from the IISP? The answer to this question will give some guidelines for reachability configuration.
 - If load balancing is expected for outgoing calls, then a reachable address needs to be specified with equal length on both of the outgoing IISP links.
 - If load balancing is to be based on the number of connections running on a given IISP interface, then all the switches of the peer group must be configured to use the widest path algorithm.
 - If load balancing is to be more static and based for instance on the location of the calling station (only for IISP connected to two different switches), then all the switches must be configured to use the shortest path, with adequate configuration of the administrative weight on all links.

- If one link is to be used as primary, and the other as backup for a given network prefix, then there are mainly two solutions:
 - a. Use the same prefix length but different administrative weights on these links. This has several limitations:
 - The administrative weight implies shortest path being configured in all the switches of the peer group.
 - Administrative weight is configured for a given IISP and cannot be specified for a particular network prefix in the case of several network prefixes being accessible through the same IISP. This solution could not have been used in the triangle network topology shown in Figure 138 on page 209.
 - b. Use different prefix lengths to reach a given network prefix. This is the easiest configuration that allows implicit specification of one path as the preferred path, and the other one as the backup. This was the option selected in the triangle topology case.

The following section illustrates how the topology of the remote end influences the configuration of reachability information.

1. The peer group is connected to a backbone peer group.

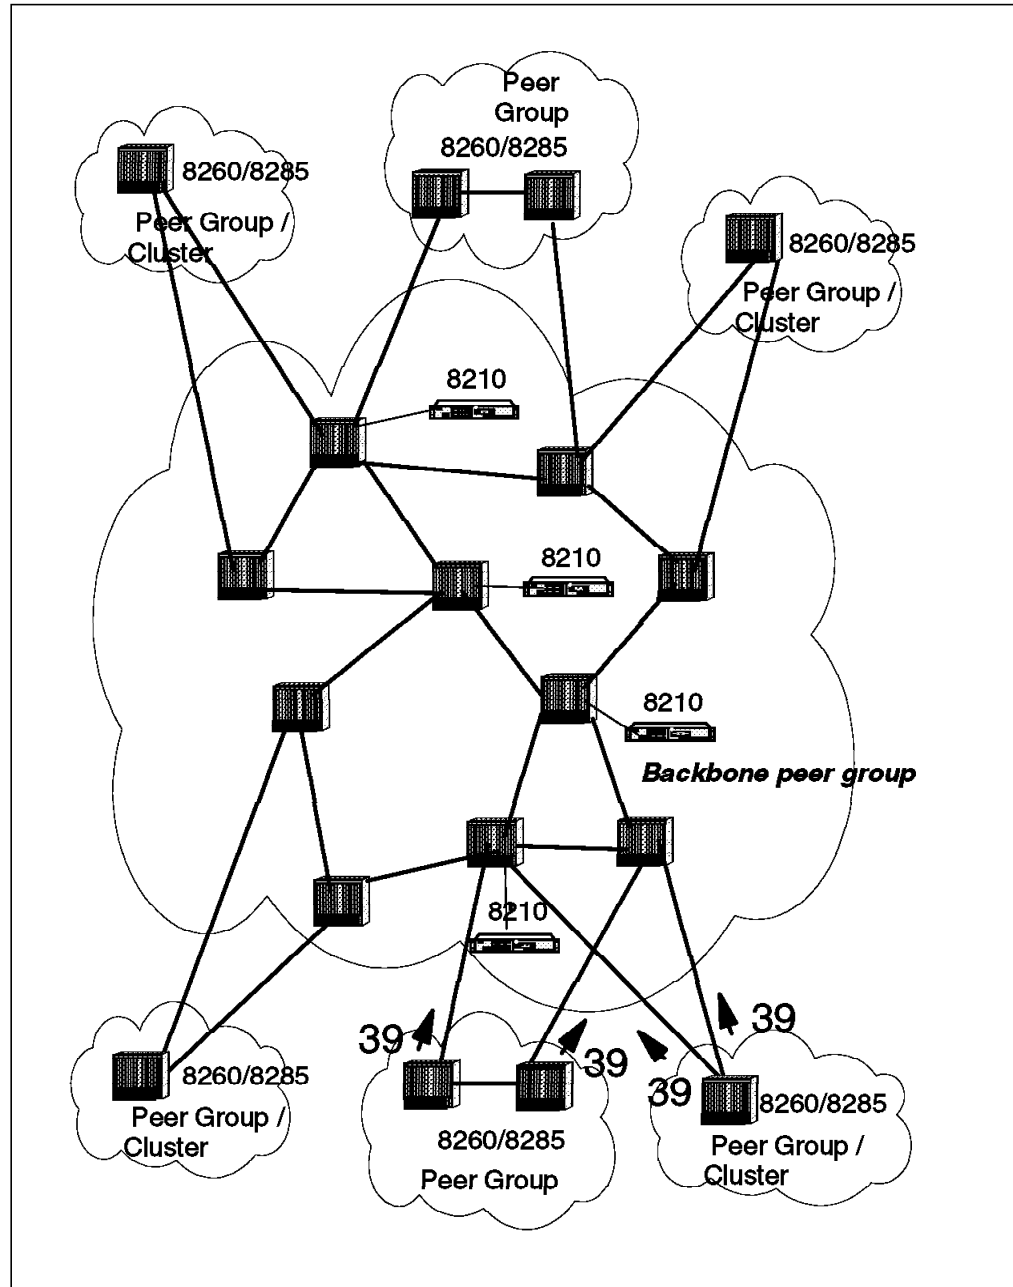


Figure 154. Peer Group Connected to a Backbone Peer Group

In this configuration, reachability information can be simply configured by specifying one byte of the network prefix being reachable through the IISP. In this example, we give reachability to all the addresses using the DCC ATM format (39.xx.) through the IISP link. No other reachability information for DCC ATM format addresses has been configured in other switches of the peer group. When a call setup is received, if the network prefix of a called DCC ATM format address does not match the network prefix of the peer group, then the call is routed to one of the two IISPs since the best match found will be the prefix advertised by the IISP, and the prefix length match is 8 bits long. This offers several advantages over a configuration where all individual network prefixes would be configured on each link:

- a. Simple: This type of configuration is easy to do, and is not really subject to misconfiguration by operator.
- b. Effective: There are fewer network prefixes advertised in the topology.
- c. Scalable: When connecting a new cluster or peer group to the backbone peer group, no reconfiguration needs to be done since this type of addressing can be seen as a default gateway concept for DCC ATM format destination addresses.

2. The peer group is connected to a peer group with the same type of configuration.

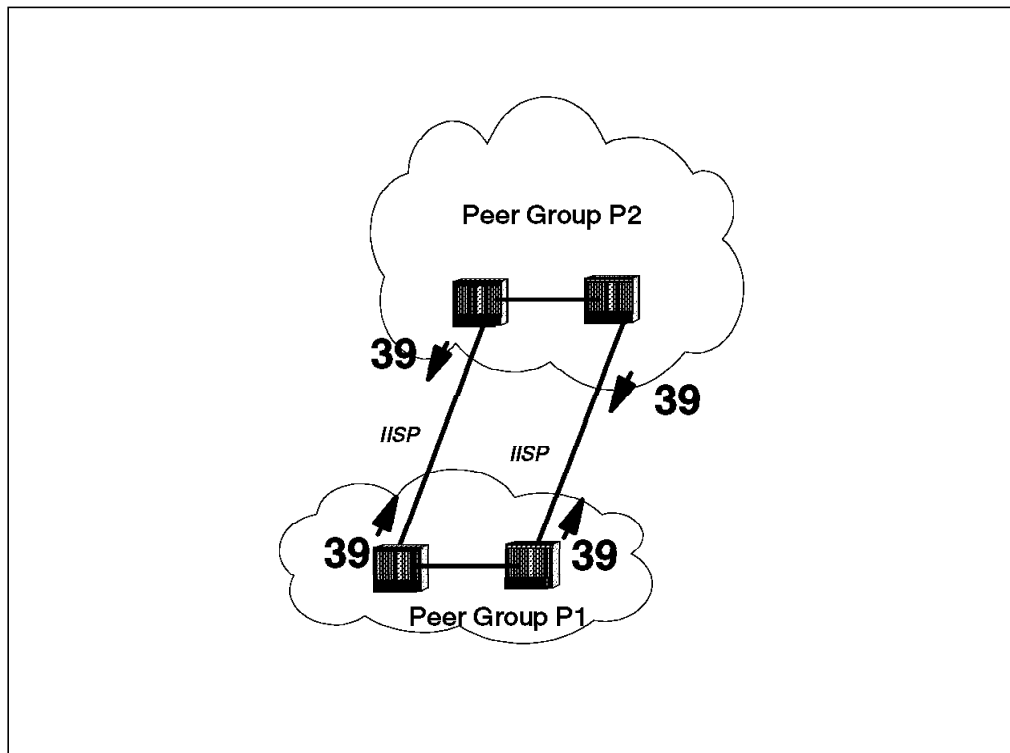


Figure 155. Peer Group Connected to a Similar Peer Group

In this configuration, if the remote peer group implements the same addressing scheme, then the network offers a routing loop for a range of addresses as already explained. In this case, the reachability configuration needs to be reconsidered.

3. The peer group is connected to two different isolated peer groups.

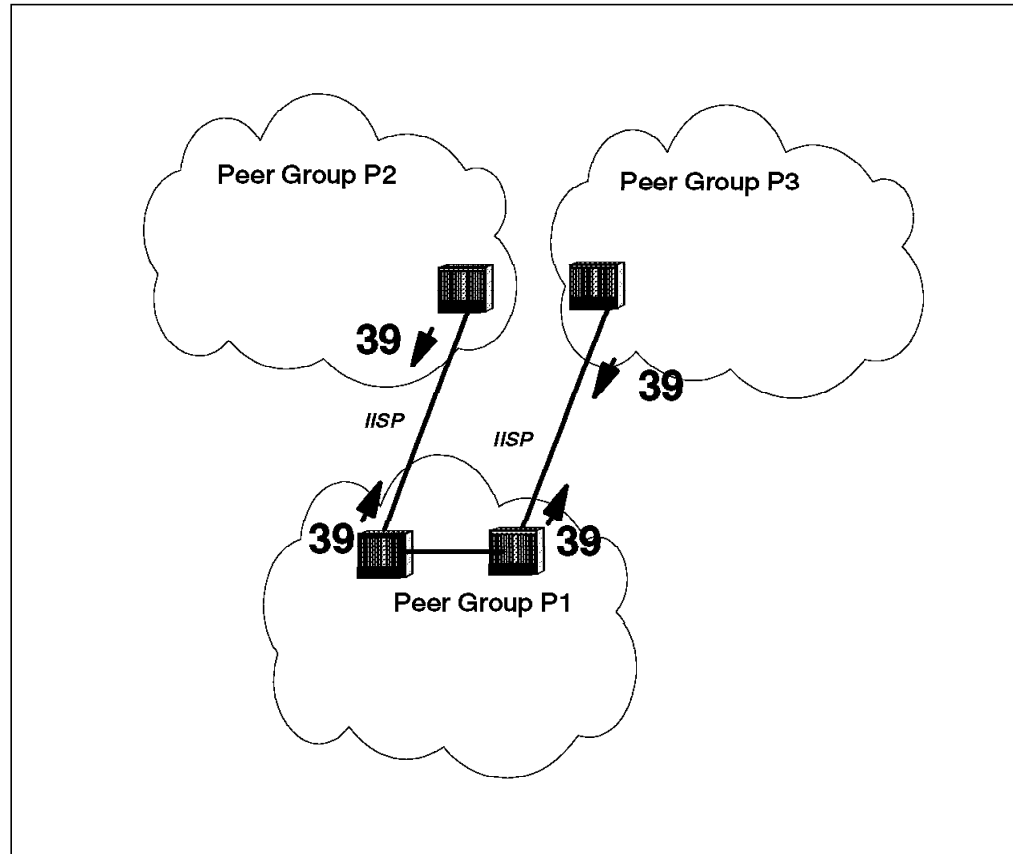


Figure 156. Peer Group Connected to Two Isolated Peer Groups

In this configuration, the reachability configuration does not present any loop. It is in fact incomplete since the routing has no way of knowing which IISP to use when receiving a call inside peer group P1 whose destination is either peer group P2 or peer group P3. This will lead to one of the peer groups not being reachable from peer group P1 *except* in the particular case where *crankback* is enabled in the switches of peer group P1. However, this is not the objective of crankback and will reduce signalling performances since statistically, half of the calls will complete thanks to crankback. In this configuration, the calls are forwarded through one of the IISPs. If the destination address is not in this remote peer group, then the call is retried on the other IISP link where it will succeed.

5.8.11.5 Conclusion

These misconfigurations may not be visible from the user point of view. Since this has only transient effects, however, in a large network it may put undesired overload on the signalling. This will not bring the network down, at least not if the network is made of 8260s or 8285s, but may reduce signalling capabilities and lead to valid calls being rejected from time to time. This chapter also illustrates the fact that simple reachability configurations can be the best solutions in some network configurations while they can be the worst in other configurations. Therefore, it is recommended to disable crankback when doing validation of the reachability configuration in the network. It may then be enabled after test completion.

Chapter 6. LAN Switches in Campus Networks

This chapter provides guidelines for troubleshooting some of the common problems that can specifically arise when using LAN switches in campus networks and illustrates these through the use of some problem scenarios.

6.1 What Makes a Healthy Switched Campus Environment

LAN switch networks can be considered healthy when their network environment is stable; devices are able to communicate and the switches are responding to the demands (traffic) of the users communicating across them. The use of LAN switches in a network can often improve the performance of the entire network; however if used in a poorly designed network they can often affect the healthy state of the whole network. Three factors that are important to consider when designing healthy LAN switched network environments are:

- Switch configuration

As switches become more function-rich, their configuration becomes more complex. Considerable care should be taken when building your switch network environment to avoid communication problems in the longer term.

- Broadcast control

A network using LAN switches can control broadcasts on the network, through the concept of virtual LANs. A virtual LAN need not be constrained to a physical area, but may be spread out to selected users in different locations or even across a WAN. VLANs therefore make it possible to build large switched networks in which the VLANs control the spread of broadcasts.

The design of the VLANs in your network can have significant effects on the spread of broadcasts and hence on the overall performance of your network. Monitoring broadcasts on your network segments is essential to verify their healthy state.

Workstations are usually assigned to VLANs by analyzing the frames sent by a workstation. It is important to understand exactly how devices are assigned to VLANs for your particular switches when designing your network; otherwise, devices may be assigned to incorrect VLANs. This could cause significant problems when devices try to communicate.

- High-speed connections

Switches provide point-to-point inter-connections between ports. Inter-switch links and links to server devices may, however, be shared by a number of devices connected via a switch. Any shared switch link should, if possible, have enough bandwidth to cope with all the traffic requirements of the communicating devices sharing the link. It is therefore recommended, to maximize the performance of your switched network environment, that all shared links use high-speed connections, such as ATM or Fast Ethernet. This will prevent a shared link from becoming a bottleneck in your network environment.

6.2 Rules for Using LAN Switches

In today's networking environment, with its ever increasing demand for more bandwidth, LAN switches have become commonplace. In their simplest form switches provide the same functionality as bridges and interlinking hubs, but are often cheaper, contain more ports and provide greater throughput. They have therefore enabled greater segmentation of traditional legacy LAN environments and have enabled increases in bandwidth without the need to replace network adapters or cabling at each workstation. To achieve the greatest increase in bandwidth, workstations may be attached directly to the switch. In this case the workstation receives the dedicated bandwidth of an entire network segment.

Switches interconnect LAN segments (or devices) together. Each switch port is of a particular LAN type and as such must follow the rules associated with that LAN type. In the case of some switch ports they may be configured to run certain services over a single physical port using virtual ports, for example Classical IP and LANE 1.0 on ATM. These ports are bound by the rules covering the services they support.

Since switches usually provide bridging functionality, as you might expect, they typically support all the common bridging techniques discussed in 2.5, "Bridges and Bridging Methodologies" on page 27. This may include source route bridging (SRB), transparent bridging (TB), source route transparent bridging (SRTB) and translational bridging (SRT). Not all bridging types are supported in all switching LAN environments. Also, like bridges, they often support the use of traffic filters used to prevent certain traffic from one port crossing over to another port. Capabilities vary from switch to switch but usually include MAC address filtering and may possibly include protocol filtering.

Most switches allow different types of LANs to be interconnected. Frames will be converted from one frame type to the other to allow devices attached to different types of LANs to communicate.

Some switch-to switch or switch-to-workstation connections support running in full-duplex mode. This effectively doubles the bandwidth of the interconnection. Both the switches and the workstation involved in the connection must support full-duplex mode. Special drivers are normally required on workstations.

Some more advanced functions and techniques used by some switches are considered in the following sections.

6.2.1 Interconnecting Switches

In summary, switches generally allow three types of connection:

- Connection to another switch

When interconnecting switches great care should be taken to avoid the introduction of loops in your network. This can result in unpredictable traffic flow and broadcast storms in your network. Loops can be avoided by using the spanning tree algorithm. This is normally disabled by default on switches.

- Connection to a concentrator or hub

The switch port may be attached to a hub port or in some cases, involving token-ring, to a hub ring in/ring out port.

- Connection to a workstation or another LAN device

Connecting workstations to switch ports can sometimes cause problems when the spanning tree algorithm is enabled on the switch port. If the spanning tree is enabled, the switch port needs to detect, when it becomes active, if the new connection has introduced a loop in the network. It will enter the normal spanning tree pre-forwarding state and during this time the port will be blocked, preventing the workstation from accessing the network. There will therefore be a delay while the switch uses the spanning tree algorithm. This can cause problems for fast machines that think the network is down while the port is initializing. The spanning tree should be disabled for workstation ports if problems are experienced.

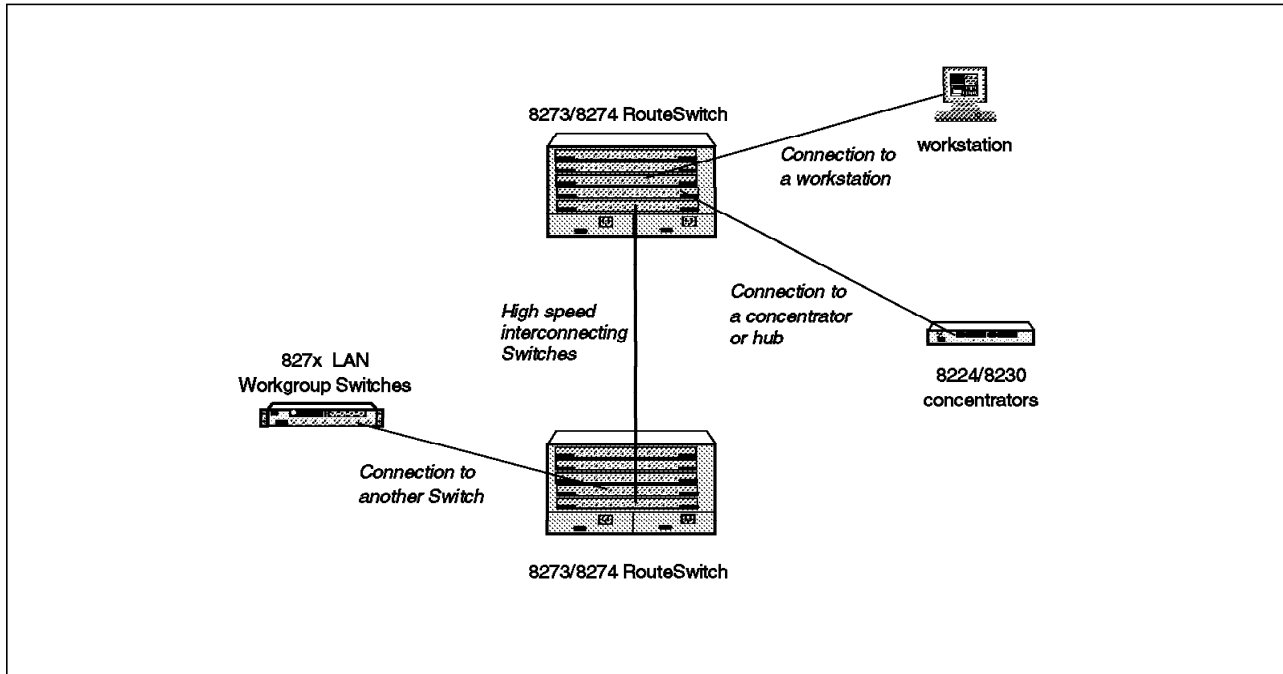


Figure 157. LAN Switch Connections

Figure 157 illustrates the different connections usually supported by a switch.

Depending on the capabilities of the switch, switch connections may involve the use of a straight-through or a cross-over cable. Straight-through cables have the same pin connections for both ends of each cable. Cross-over cables cross the transmit and receive pins to allow the transmit pin on one device to be connected to the receive pin on the device it is connected to. The pin sequence used is different for different LAN types.

Typically switch and hub ports are crossed internally, so when connecting a switch to another switch or to a hub a cross-over cable may be required. Some switches and some hubs have enough intelligence to work out whether a port needs to be internally crossed or not and hence switch-to-switch or switch-to-hub connections are automatically configured when using a standard straight-through cable. To find out what kinds of cable are required when interconnecting your switch to other devices refer to the product documentation you received with your switch.

As well as interconnecting switches using standard switch ports, some switches provide proprietary high-speed interconnection techniques. Three of these used by IBM switches include:

- Etherpipes and tokenpipes

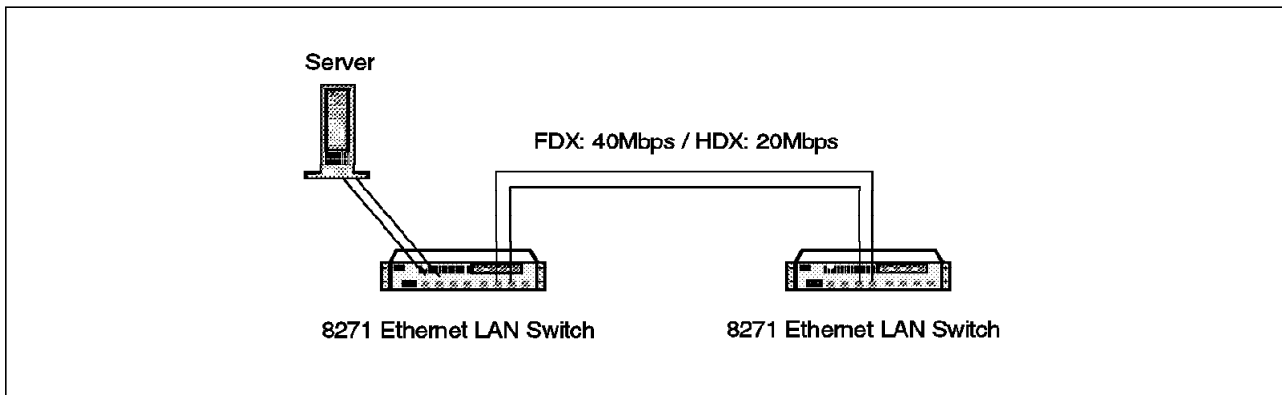


Figure 158. Etherpipes and Tokenpipes

Etherpipes and tokenpipes are used with the IBM Nways 8271 Ethernet LAN Switch and IBM Nways 8272 Token-Ring LAN Switch products respectively. Etherpipes are high bandwidth interconnections between two IBM Nways 8271 Ethernet LAN Switches and tokenpipes are high-bandwidth interconnections between two IBM Nways 8272 Token-Ring LAN Switches.

Etherpipes and tokenpipes are multi-link inter-switch connections. A maximum of four switch ports on two switches can be interconnected and the bandwidth is concatenated into a single pipe. Each port provides a full-duplex link between switches. Hence a four-port Etherpipe provides an 80 Mbps inter-switch connection and a four-port tokenpipe provides a 128 Mbps inter-switch connection.

- Trunking used by the IBM RouteSwitch products




Figure 159. Trunking Used by the RouteSwitch

The IBM Nways 8273 Ethernet RouteSwitch and IBM Nways 8274 LAN RouteSwitch products support high-bandwidth switch-to-switch trunking. VLAN trunking provides the capability to transport multiple VLAN groups through high-speed RouteSwitch-to-RouteSwitch links. These switches currently support VLAN trunking using FDDI or ATM links.

When using trunking, the frames are encapsulated within a proprietary frame, implicitly identifying them with their VLAN group number to provide

the data separation and to prevent interaction with non-trunked stations. ATM trunks can be configured as PVCs or SVCs.

For FDDI links, group multiplexing over switch-to-switch links is also possible using IEEE 802.10. 802.10 is a standard for multiplexing trunks over FDDI networks. Frames are first translated then tagged with a header containing the group information.

- Point-to-point-bridging (PToP) used by the IBM RouteSwitch products

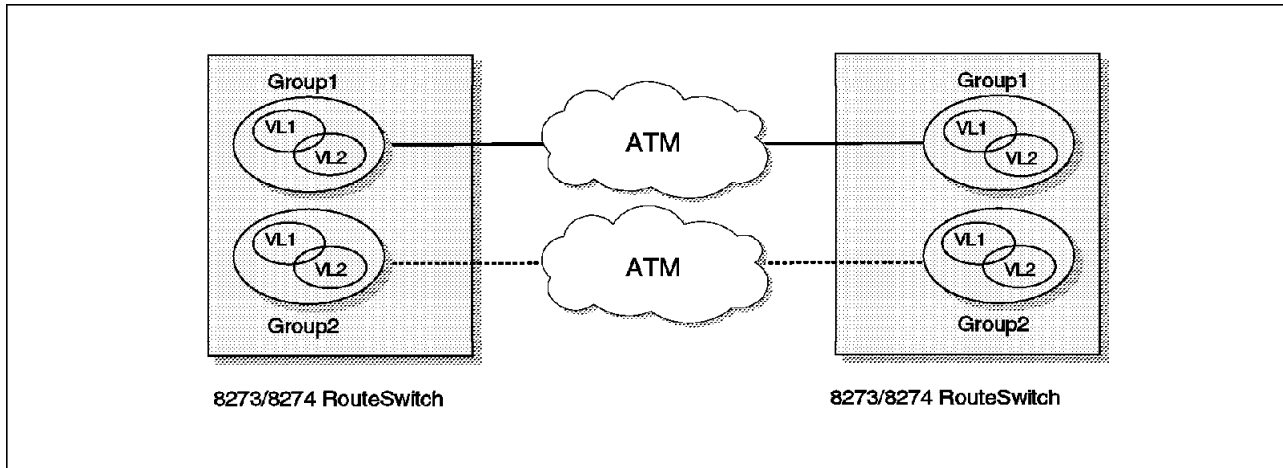


Figure 160. PToP Bridging Used by the RouteSwitch

The IBM Nways 8273 Ethernet RouteSwitch and IBM Nways 8274 LAN RouteSwitch products also support high bandwidth switch-to-switch connections using point-to-point bridging. PToP bridging allows two groups to be connected together over an ATM network. The ATM connection supports the use of SVCs and PVCs. Two types of encapsulation are supported: RFC 1483 using LLC encapsulation and a proprietary encapsulation.

6.2.2 Virtual LANs (VLANs)

Virtual LANs are logical LANs rather than physical collections of devices. They segment the physical LAN into different logical regions. This can be used to further increase the available bandwidth of devices attached to a switch or on a segment attached to a switch, by minimizing the broadcast traffic sent to all devices. Devices should be organized into regions with the network resources they utilize.

Most LAN switches on the market support the creation of port-based VLANs. These are where different ports on a switch are segmented into logical groups. Data traffic is then prevented from traversing the switch between these groups. Each group of ports forms a logically separate LAN. These groups of ports are sometimes referred to as virtual switches.




Figure 161. Port-Based VLANs or VLAN Groups

Figure 161 illustrates a port-based VLAN or VLAN group.

Some more advanced LAN switches allow the creation of VLANs through the use of policies. The assignment of network devices to their respective VLAN is performed by monitoring traffic from the port to which the device is attached and based on policies assigning stations to the correct VLAN. These allow the creation of VLANs independent of the physical location of devices. Network workstations can move physical location but will still be connected to the logical LAN containing the network resources they utilize. Stations may often be a member of multiple VLANs at the same time.

VLANs may also be composed of different LAN types. A single VLAN may consist of token-ring, Ethernet, Fast Ethernet, FDDI and ATM devices. VLANs therefore make it possible to build large switched networks in which the VLANs control the spread of broadcasts. Resources may be centralized, allowing easier management while still optimizing the network bandwidth available.

The creation of VLANs in LAN switches is not standardized; all implementations are proprietary. As an example of the way VLANs are used and the problems that can occur with them we consider the IBM Nways 8273 Ethernet RouteSwitch and IBM Nways 8274 LAN RouteSwitch products. These products allow the creation of port and policy-based VLANs. The policies they currently support are:

- Port rules
- MAC address rules
- Protocol rules
- Network address rules
- User-defined rules

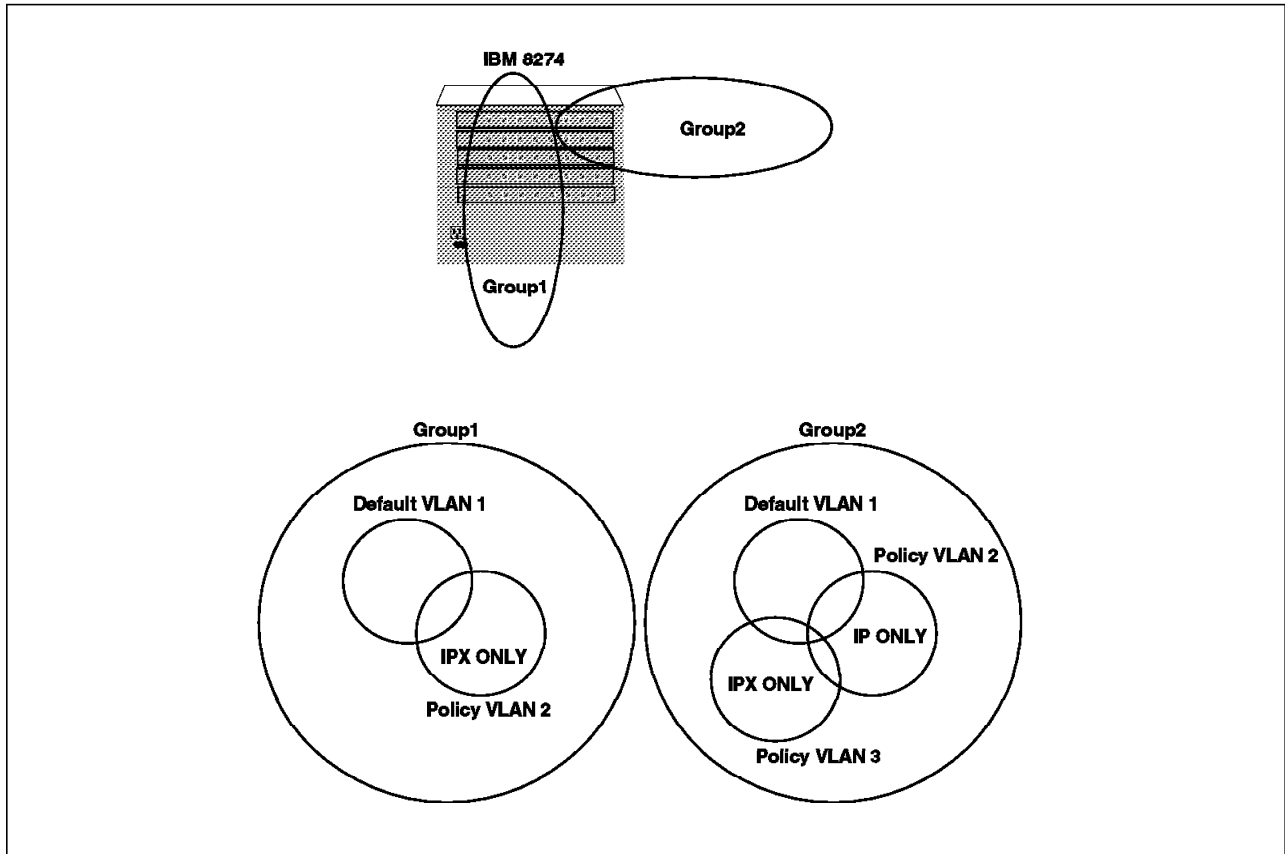


Figure 162. Policy and Port VLANs

Figure 162 illustrates port and policy VLANs. The IBM Nways 8274 LAN RouteSwitch may be segmented into port-based groups which form a broadcast domain and then each group may be segmented into policy-based VLANs.

In order to avoid problems in creating VLANs it is important to understand how ports and workstations are assigned to a particular VLAN by the policies defined in the switch.

The following example summarizes the process used by the IBM Nways 8274 LAN RouteSwitch to assign stations to VLANs:

1. Initially all ports in a defined group (a port-based VLAN) are members of the default VLAN.
2. A frame is received on a port
3. The switch hardware looks up the source address to ensure it has been learned.
4. If the source address is not recognized (not in Content Addressable Memory (CAM)), the frame is forwarded to the switch Management Processor Module (MPM) for processing.
5. The MPM examines the entire frame to determine which VLAN(s) the source address qualifies for membership in and adds the source address and switch port to the VLAN membership list. It also removes the source address from the default (unknown) VLAN.
6. Broadcasts and multicasts are always sent to the MPM for processing.
7. Simultaneously the switch forwards the frame.

8. If the destination address is recognized, the frame will be compared against the filtering database to determine if the destination device belongs to a VLAN associated with the port that received the frame from the source device. If so, the switch will do any necessary translations required and will forward the frame to the correct port; otherwise the frame will be dropped.
9. If the destination address is not recognized but the source address is known, the frame is forwarded to all ports in common VLANs with the source address.
10. If the destination address is not recognized and the source address is unknown, the frame is forwarded to all ports within the same group.

<i>Table 12. Unicast Processing</i>		
MAC Address	Known Destination	Unknown Destination
Known Source Address	Frame is switched to destination port if source and destination addresses have at least one VLAN in common. Frame is not analyzed by the MPM and no additional VLAN assignment is made.	Frame is forwarded out of all ports which have at least one VLAN in common with the source address (except optimized device switching ports). Frame is not analyzed by the MPM and no additional VLAN assignment is made.
Unknown Source Address	The frame is completely analyzed by MPM and VLAN assignment is updated. Then the frame is switched to the destination port if the source and destination addresses have at least one VLAN in common.	Frame is flooded out all ports which are in the same group (except optimized device switching ports). The frame is completely analyzed by MPM and VLAN assignment is updated.

<i>Table 13. Broadcast and Multicast Processing</i>	
MAC Address	Processing
Known Source Address	Frame is forwarded out of all ports which have at least one VLAN in common with the source MAC address. The frame is completely analyzed by MPM and VLAN assignment is updated.
Unknown Source Address	Frame is flooded out of all ports that are in the same group as the source port. The frame is completely analyzed by MPM and VLAN assignment is updated.

Note: Optimized device switching ports are optimized for the connection of a single workstation. Since the switch knows a single device, possibly a server, is connected to the port, frames are not flooded out of the port.

6.2.3 Address Tables and Aging Timers

Switches normally determine how to switch frames by learning MAC addresses and building up address tables for each port on the switch. The timers that affect these tables are of special importance when stations are moved or links are temporarily broken.

If a device is moved from one port to another a switch will take time to learn the change. The most important factors affecting the time this takes are the aging times for addresses in the switch. If too high, the switch will take a long time to discover moves. If too short, the switch will not be efficient since it will have to keep learning where devices are attached after a short period of non-transmission.

Timer values are usually configurable in most switches. Different switches will also provide different timers depending on the capabilities of the switch. For the IBM Nways 8274 LAN RouteSwitch the important timers controlling address aging are shown in Table 14.

Timer Name	Description	Value
Aging Time	The timeout period in seconds for aging out dynamically learned forwarding information.	10-1000000 (Default 300) secs
Auto-Tracker VLAN Aging Time	The length of time in seconds to remember which VLAN a station belonged to even after the station has been aged out of the bridge filtering database.	10-1000000 (Default 1200) secs

If a switch port is connected to a LAN that has more devices attached than the maximum number of allowable entries in the switch port's address tables, problems may occur. If a switch port table becomes full, addresses are discarded before they age, or time out. This may cause performance problems as the switch must keep re-learning the ports to which stations are attached. It is therefore highly recommended that the maximum number of devices on a segment attached to a switch port be kept lower than the maximum entries in the switch port table. This can be achieved by segmenting the LAN attached to the switch port.

6.2.4 Operating Modes and Performance

Many factors can affect the performance of a switch. It typically depends on the functions running in your switch and the amount of frame processing that the switch must perform. If you are using policies to filter traffic between ports, all traffic must be analyzed prior to transmission, to determine if it meets the defined criteria. If you are using VLANs to control the spread of traffic, then frames again must be analyzed depending on the VLAN policies defined in your switch. (For more information see 6.2.2, "Virtual LANs (VLANs)" on page 241.)

One other factor that may adversely affect network performance is the speed adaption capabilities of the switch. Switches are often used to connect servers at a higher access rate than clients. This requires the switch to perform speed adaption. This is handled by buffering within the switch itself, but the buffers have limited capabilities. Protocols must adapt their flow to the capabilities of the switch. For most protocols this is performed automatically; however some protocols, for example RIPL, only have limited flow control mechanisms. Therefore moving servers to high-speed ports when using these protocols may actually decrease server performance.

The switch may also run in a number of modes. Different switches may support different types of data switching modes. Commonly, switches will support one, two or three of the following modes:

- Cut-through (normal switch mode)

The switch starts to forward a frame through its output port prior to receiving the entire frame at its input port. This allows very fast frame forwarding but does mean corrupted frames are forwarded.

- Store and forward (normal bridge mode)

The switch stores the entire frame before forwarding it to allow the switch to verify the frame is complete and without errors by validating the frame check sequence at the end of the frame. This is slower for forwarding frames but ensures corrupt data is discarded sooner.

- Adaptive cut-through

Frames are monitored for errors while in cut-through mode. If excessive errors are detected, the port will automatically change to store-and-forward operation.

6.3 Problem Determination Guidelines

Figure 163 summarizes some of the most common problems found when using LAN switches in campus networks and illustrates the problem determination procedure we used to detect and solve these problems.

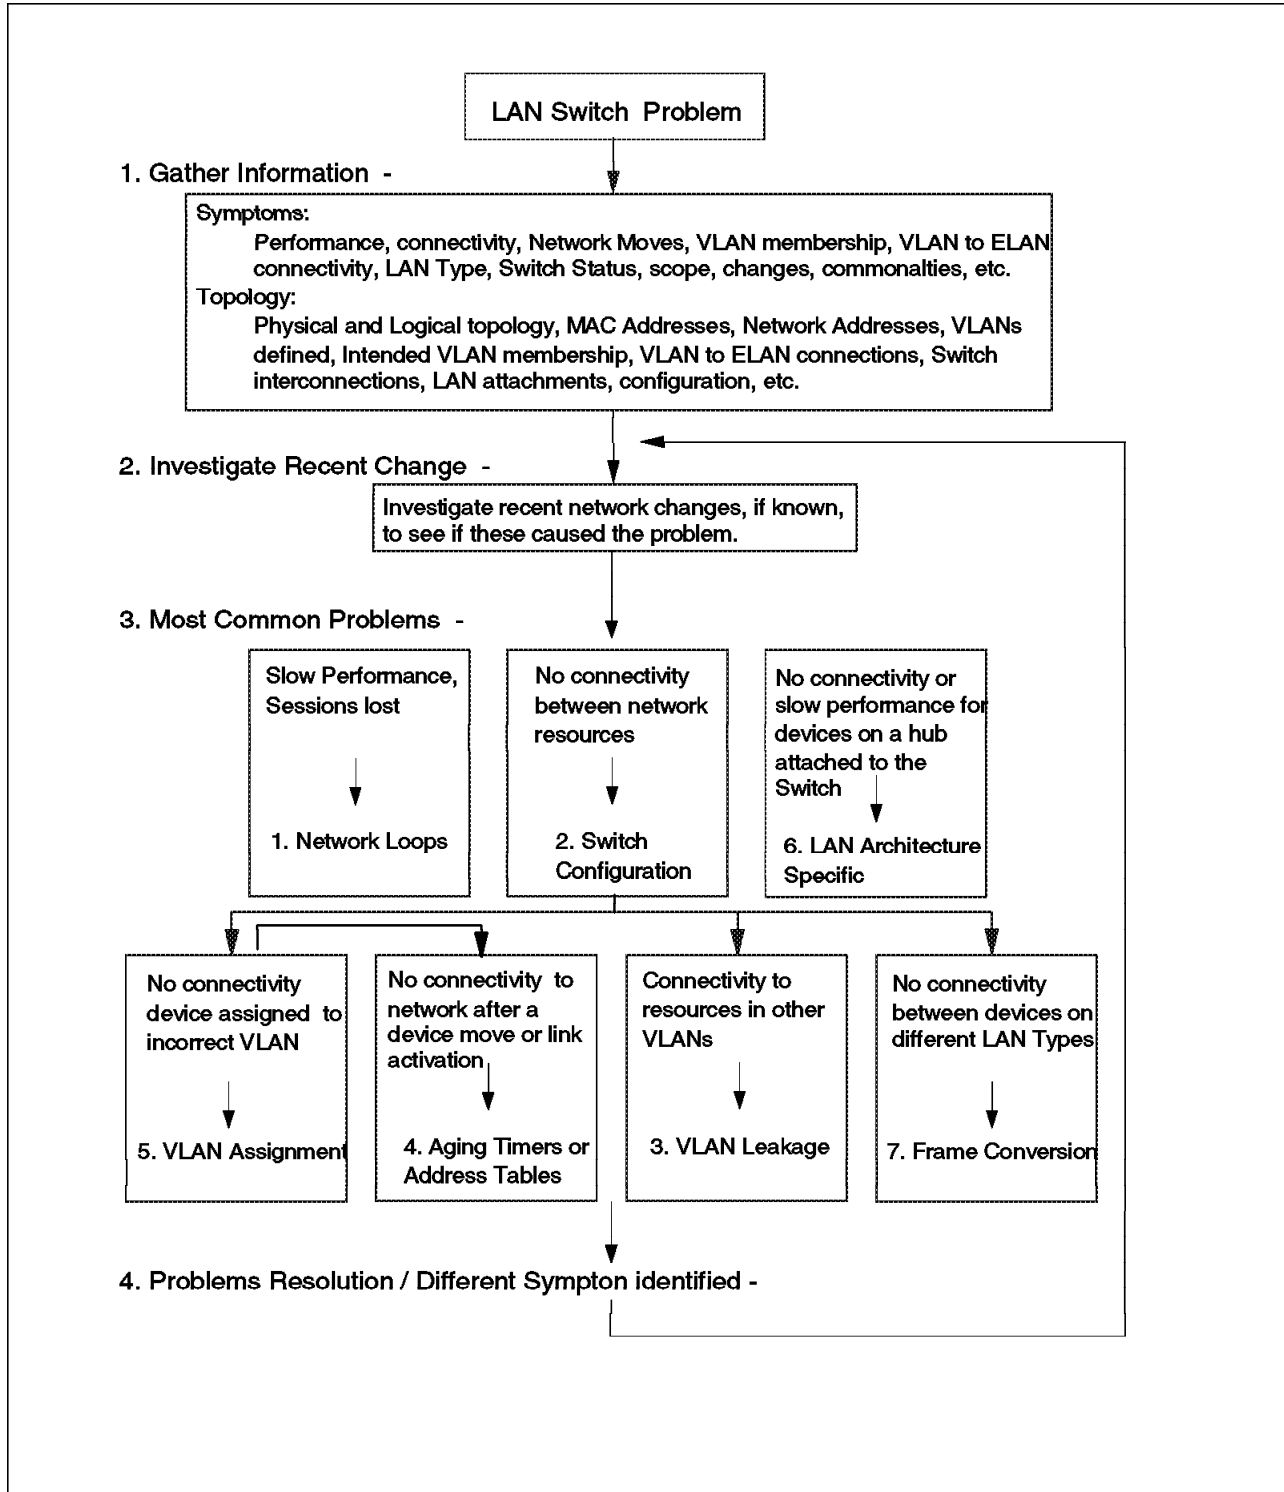


Figure 163. Common Switch Problems and Our Problem Determination Methodology

The four sections can be described as follows:

1. Gather information.

As mentioned in Chapter 4, “Problem Determination Guidelines” on page 91 the key to effective problem determination is to gather as much information regarding a problem’s symptoms and the network topology as possible. Some of the specific questions you may need to ask in a switched LAN environment are:

- Are you using VLANs? What VLANs are defined? What policies are used? What devices should be on each VLAN? Does traffic on one VLAN affect devices on another VLAN?
- Have devices been moved or new links activated? Has the problem only started happening since the move or activation?
- What devices are affected by the problem? Where are they connected to the switch? What are their MAC and network addresses (IP/SNA/NetBIOS/IPX)?
- Are all or most of the devices affected on a single LAN segment attached to one port of the switch?
- Are you using VLANs with emulated LANs? Do you have connectivity problems between these logical environments?
- How is your network interconnected? Do you use other switches, bridges or routers in your network to interconnect segments? Are you experiencing performance problems?

2. Investigate recent changes.

As mentioned in Chapter 4, “Problem Determination Guidelines” on page 91, most problems occur when you have changed something in your network. Effective change control and management can allow changes to be easily detected and backed out quickly in the event of a problem.

3. Interrogate the LAN switches based on the problem symptom experienced.

Some of the most common problems involving LAN switches involve:

- Network loops

If switches are interconnected to other switches or via bridges and routers, loops can be formed in your network. This often results in some network segments being flooded with traffic and users often experience very bad performance. Different LAN types handle network loops in different ways but most use the spanning tree algorithm to prevent traffic flooding. (For more information see 2.5.1.1, “Spanning Tree Protocol (STP)” on page 28.)

Spanning tree problems are not unique to LAN switches. LAN bridges or routers implementing bridge functions can also result in network loops. For more information on problems that can occur with network loops and the use of the spanning tree algorithm see 8.5.2, “Case Study 2: Spanning Tree Loop” on page 450.

- LAN switch configuration

As LAN switches implement more and more functions, their configuration becomes more complex. Problems can occur when switches are configured incorrectly. This often results in devices not receiving the connectivity they expect from the network.

Extreme care should be taken when defining VLANs, both port/group based and policy-based. Incorrectly defining groups and VLAN policies may result in devices being attached to an incorrect VLAN and could mean a workstation cannot get access to the network resources it requires.

Switch configuration problems are discussed further in 6.5.3, “Case Study (1-1): Simple Configuration Problem - Ports in Wrong Groups” on page 258.

- VLAN assignment

VLAN policies are defined correctly but MAC devices, which should be members of a VLAN, are not accessible by other MAC devices connected to the VLAN. This is typically because the source MAC device sends a unicast frame to a destination MAC device unknown by the switch and the device is not connected to a port on a VLAN in common with the source address.

For more information on the problems that can occur with VLAN assignments see 6.5.4, “Case Study (1-2): VLAN Assignment - Devices in Wrong VLANs” on page 265.

- LAN architecture specific

Some LAN types may experience specific problems due to the use of switches in the network. This may affect devices attached directly to the switch but may also affect all devices attached to a downstream hub or concentrator. One example of a LAN type that can experience problems in using LAN switches is token-ring. Many of the problems are not related to the switch itself but are due to incompatibilities in the way devices implement the LAN architecture. An example of this is the way devices use the ARI and FCI bits in token-ring frames.

Token-ring frames contain a status byte at the end of each frame (see 2.3.1.1, “Frame Format” on page 18). This includes the ARI (address recognized) and FCI (frame copied) bits. These are often used by a sending station. The sender sends data in a frame. The frame goes around the ring and is either picked up by the destination station or where the destination station is on a different ring by a bridge (usually a source route bridge). After it is picked up by the destination or bridge, the frame continues on the ring until it reaches the sending station that removes it from the ring. The destination station or source-route bridge often sets the ARI/FCI bits to indicate the frame has been received and copied. IEEE 802.5 (TR standard) states that these bits should not be used by the sender to validate that the data was received; however, many sending station drivers use them this way. The sending station can tell the device is on a different ring by the presence of the RIF field. However the RIF field is not used by transparent switches, which forward frames using address tables. The switch can therefore either leave the ARI/FCI bits alone or set the bits to 1.

- Never setting the bits may cause problems with layer 2 or 3 protocols that use these bits. Data may be retransmitted or sessions re-established even though the data was sent and received correctly.
- Always setting the bits may cause problems since the receiver station may assume that another station with the same address resides downstream on the same ring. It will then generate a soft error reporting MAC frame.

Problems can usually be minimized by tailoring whether a switch sets or doesn't set the ARI/FCI bits on a per port basis. The IBM Nways 8272 Token-Ring LAN Switch, IBM Nways 8273 Ethernet RouteSwitch and IBM Nways 8274 LAN RouteSwitch products all support this. To actually fix this problem the drivers on the token-ring workstations should be replaced with ones that fully support the IEEE 802.5 standard. The latest drivers for IBM network adapters are available from <http://www.networking.ibm.com/>

- Frame translation

Many problems can occur when frames are converted from one LAN type to another. For example, different LAN types support different MAC addressing standards. Some fields, present in the frames of one LAN type, are not present in the frames used by other LAN types (for example, RIF); different encapsulation types are available on different LAN types, etc. No standard way of translating frames exists. Switches and bridges implement proprietary techniques to translate different frame types. For more information on frame conversion see 2.5.4, "Source-Route Translational Bridging (SR-TB)" on page 32.

- Aging timers and address tables

Switches usually determine how to switch traffic by learning MAC addresses of devices accessible via their ports. Problems can occur when a device moves from one port to another port or when links to other parts of the network become active. If a new link becomes active, the switch may have to reconfigure all of its links according to the spanning tree algorithm. This can result in paths from one device to another being disabled and new paths being enabled. The switch must re-learn which MAC addresses are accessible via which ports. Convergence of the spanning tree algorithm to determine which inter-switch links should be blocked and which enabled can take some time.

After a device move or after the spanning tree has converged it may still take the switch some time to relearn which MAC addresses are accessible through which ports. The most important parameters that govern this time are the aging time for a switch and if VLANs are used, the VLAN aging time. Refer to 6.2.3, "Address Tables and Aging Timers" on page 245 for more information.

For more information on the problems that can occur with aging times see 6.6, "Case Studies for Address Aging Time" on page 270.

- VLAN leakage

Since the assignment of devices to policy VLANs is usually done dynamically through learning unknown source addresses, some frames may be forwarded, under certain conditions, to ports that belong to a different VLAN. This is called VLAN leakage and can cause significant problems. For example, this may lead to a device that is only a member of a NetBIOS VLAN and directly connected to a switch port, receiving IP ARP broadcasts. In summary, VLAN leakage may occur:

- If a known device belongs to two VLANs, its broadcast, multicast and unknown destination frames are forwarded to both VLANs.
- If devices that are assigned to different VLANs join an additional common VLAN, the broadcast domain is expanded to the common

VLAN. This is discussed in 6.7.3, “Case Study (3): VLAN Leakage Problem” on page 284.

- The very first frame of all devices is forwarded to all ports within the group regardless of VLAN membership.
 - If a known MAC address is timed out and removed from the switch’s memory, the next broadcast from such a device is forwarded to all ports within the group regardless of VLAN membership.
 - Unicasts are forwarded to all ports within a group if source and destination addresses have timed out of the switch’s memory.
4. Resolve problem or investigate a different symptom identified during the problem determination process.

The following sections describe our methodology for fixing problems with some of the most common problems mentioned.

6.3.1 LAN Switch Configuration Problem Methodology

Numerous problems can cause devices to fail to communicate. One of the most common is the configuration of the LAN switch itself. To investigate any connectivity problems between devices it is usually best to investigate the LAN switches in your network to check they have been configured correctly. Start with the switch closest to the devices experiencing the problem.

Check the switches’ LEDs to ensure the device and its modules are operating correctly. Refer to the product documentation that came with your switch to verify their correct status and to rectify any problems you discover.

Most switches support a command interface and possibly a network management application that can verify the status of modules and ports. Ensure the ports with devices experiencing a connectivity problem are enabled and active.

Finally make sure the ports are assigned to the correct VLAN groups you have defined in your network. VLAN groups are broadcast domains and ports in one broadcast domain can not normally communicate with ports in another broadcast domain.

6.3.2 VLAN Assignment Problem Methodology

LAN switches assign devices to policy VLANs by analyzing certain traffic coming from the device. This process is described in more detail for the IBM Nways 8273 Ethernet RouteSwitch and IBM Nways 8274 LAN RouteSwitch products in 6.2.2, “Virtual LANs (VLANs)” on page 241. Problems can therefore occur when some devices have communicated via the LAN switch and been assigned to their respective VLANs but other devices have not. Problems may also occur if a device is idle for a period of time. The switch may age its address out of its VLAN tables possibly making it inaccessible to other devices on a VLAN.

To check for VLAN assignment problems you need to understand what VLANs should exist and what devices are assigned to them. You then need to interrogate your network to discover which ports are assigned to which VLANs and which MAC devices are assigned to each VLAN. You may usually use commands on a LAN switch or a network management station to verify VLAN policies, port and MAC assignments. By knowing this information and by

understanding the way stations are assigned to VLANs you can usually identify whether a problem exists.

6.3.3 Aging Timers and Address Tables Problem Methodology

Moving stations between switches or ports may cause problems. It takes switches time to reconfigure their address table to detect the move. Make sure that the port a moved station is moved to is in the same group and VLAN as the port the moved station was moved from.

If you experience communication problems after workstations move ports, you may need to wait a few minutes for the switch to learn about the move. If this causes significant problems you may wish to try reducing the aging time or VLAN aging time of your switch to shorten the time it takes the switch to age out its address tables. This may reduce the performance of the switch and in some cases may cause devices that have not communicated for a short time to be removed from their VLANs. Great care should therefore be taken when modifying the timers.

To diagnose aging timer problems check the MAC address assignment to ports and to VLANs in your switch. Make sure that MAC addresses in the switch tables are assigned to the correct ports and correct VLANs. If any are assigned to the wrong port, then it is probably because the machine has moved and the switch must relearn its location.

Network traces can also show you the traffic flowing from a device. By knowing what traffic the switch uses to update its address tables you can determine why the switch has not updated its tables.

Finally looking at the settings for the aging time and VLAN aging time will show you how often the switch ages out its address table.

6.3.4 VLAN Leakage Problem Methodology

If MAC addresses and ports are assigned to multiple VLANs, it's possible that VLAN leakage may occur. Check the defined VLAN policies in your switches and also look at the frames being sent from your stations. If a station sends multiple protocol frames which match different VLAN policies, the MAC address and port of this station will be assigned to multiple VLANs.

Verify the VLAN membership of ports. When a port belongs to multiple VLANs and is attached to a downstream hub, VLAN leakage may occur when a new station is connected to the downstream hub. The switch will not recognize the new source address and will forward broadcasts, multicasts and unknown destination frames, to all ports that are in the same group.

Verify the MAC address assignment in your switch and determine whether there are common VLANs for devices that should only be assigned to different VLANs. The broadcast domain will have been expanded to the common VLAN.

6.3.5 Diagnosing LAN Switch Problems

The following chart lists generic LAN problems with their probable causes. follow the recommended actions to solve your problem.

<i>Table 15. Diagnosing Problems Concerning Connectivity between Devices</i>		
Symptom	Probable Cause	Recommended Actions
Abnormal LEDs on LAN switch	Hardware problem	Refer to the product-specific documentation available for your switch.
No LED lights on LAN switch	No power to device	Check power cable and source and refer to the product-specific documentation available for your switch.
Abnormal module/slot status	Hardware problem or switch initialization problem	Refer to the product-specific documentation available for your switch.
Abnormal port status or port not enabled	Hardware problem	Refer to the product-specific documentation available for your switch.
	LAN switch configuration	Refer to the product-specific documentation and enable the port.
Ports assigned to different VLAN groups	LAN switch configuration	<ol style="list-style-type: none"> 1. Refer to your network documentation to discover which ports should be in which groups. 2. Configure your LAN switch so that ports are in the same VLAN group or check device interconnecting VLAN groups.
Ports assigned to same VLAN group but different policy or multicast VLANs	VLAN assignment or VLAN timers	<ol style="list-style-type: none"> 1. Check what ports and devices are assigned to each VLAN. 2. Update VLAN policies to add a port policy for server devices. 3. Set up an automated procedure to send data from server devices every few minutes. This will ensure they are assigned and do not drop from their VLANs. 4. Increase aging time and/or VLAN aging time, if changing policies or automated procedure is not possible.
Ports should be assigned to different VLAN but traffic is flowing between them	VLAN leakage	<ol style="list-style-type: none"> 1. Check what ports and devices are assigned to each VLAN. 2. Determine whether ports or devices share a common VLAN. 3. Update VLAN policies to rectify problem.

6.4 Hints and Tips in IBM Switch Environments

This section proposes a procedure to troubleshoot a LAN network in an environment using switches.

6.4.1 Hardware Check

The hardware has to be checked first.

- Management Processor Module (MPM)

Check the MPM power supply LEDs to verify that the power supply is operating correctly. If one power supply is inserted, then the *PS1* LED (just *PS* on the IBM Nways 8273 Ethernet RouteSwitch) should be on steady green. If two power supplies are installed, then both the *PS1* and *PS2* LEDs should be on solid green.

All modules are subjected to extensive power-on diagnostics during the Power-On Self-Test cycle. While the diagnostics are running, the MPM *OK2* LED will blink amber. When diagnostics are complete and all modules are operating correctly, then the *OK1* LED on all modules should be on solid green and the *OK2* should be blinking green.

- Network Switch Module (NSM)

LEDs on network switch modules vary by the network interface type and by a module's application. However, two LEDs are common to all switching modules. These LEDs, *OK1* and *OK2*, provide information on the hardware and software status of the module.

The ports which are attached to devices should be active. The network switch module includes one row of LEDs for each port.

- *STA* (status): On green continuously when a good cable connection exists. Flashes Green slowly when the port has been disabled.
- *LINK* (link status): On green when the module has a valid physical link. Under normal conditions, this LED should always be on when a cable is connected.
- *ACT* (activity): On green when the port is transmitting or receiving packets/cells.

6.4.2 Login to the User Interface

Use a terminal emulation software package to access a session with the switch. For the IBM Nways 8273 Ethernet RouteSwitch and IBM Nways 8274 LAN RouteSwitch set your terminal emulation software parameters to the following:

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

After the system has booted, which will take about a minute, you will see a login prompt:

```

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

*****

IBM Corporation LAN RouteSwitch - Copyright (c) 1994, 1995, 1996
      System Name:      A8274
      System Location:  ITS0
      Primary MPM
Command      Main Menu
-----
File         Manage system files
Summary      Display summary info for VLANs, bridge, interfaces, etc.
VLAN         VLAN management
Networking   Configure/view network parameters such as routing, etc.
Interface    Enter the Physical Interface Configuration/Parameter Sub-menu
Security     Configure system security parameters
System       View/set system-specific parameters
Services     View/set service parameters
Switch       Enter Any to Any Switching Menu
Help         Help on specific commands
Exit/Logout  Log out of this session
?           Display the current menu contents

/ %
    
```

Figure 164. RouteSwitch Main Console

The default system prompt is / %. The / indicates that you are at the main menu. The % is a prompt character.

6.4.3 Console Commands

When you log on to the IBM Nways 8274 LAN RouteSwitch the first screen you see is a menu containing various submenus for different commands and services. All commands may be issued from any level of this tree of commands. If you do not know what commands are available, you may type ? at any command menu to see the list of commands.

The following tables contain a listing of the menu structure and some useful user interface commands on the IBM Nways 8273 Ethernet RouteSwitch and IBM Nways 8274 LAN RouteSwitch.

Table 16 (Page 1 of 2). The RouteSwitch Submenu Structure

Submenu	Description
Main	Enter the main menu
VLAN	View list of active VLANs on an interface
->at	Enter the Auto-Tracker sub-menu
->br	Enter the Bridge Configuration/Parameter submenu
Interface	Enter the Interface Configuration menu
->atm	Enter the ATM Management submenu
->fddi	Enter the FDDI Management submenu
->tok	Enter the Token-Ring Management submenu

Submenu	Description
Networking	Configure /View network parameters
->ip	Enter IP networking command submenu
->ipx	Enter IPX networking command submenu
Services	View /Set service parameters
System	View /Set system-specific parameters
Summary	Display summary info for VLANs, bridge, interface, etc.
Switch	Enter any-to-any switching port translations
Exit	Log out of this session

Command	Description	Reference
VLAN Submenu		
gp	View the list of groups currently defined	6.5.3, "Case Study (1-1): Simple Configuration Problem - Ports in Wrong Groups" on page 258
via	View ports assigned to the selected group	6.5.3, "Case Study (1-1): Simple Configuration Problem - Ports in Wrong Groups" on page 258
addvp	Add ports to a group	6.5.3, "Case Study (1-1): Simple Configuration Problem - Ports in Wrong Groups" on page 258
modvp	Modify existing VPORT configuration information	6.7.3, "Case Study (3): VLAN Leakage Problem" on page 284
Auto-Tracker Submenu		
viatrl	View an Auto-Tracker rule configuration	6.5.4, "Case Study (1-2): VLAN Assignment - Devices in Wrong VLANs" on page 265
modatvl	Modify definition of an Auto-Tracker VLAN	6.5.4, "Case Study (1-2): VLAN Assignment - Devices in Wrong VLANs" on page 265
cratvl	Create an Auto-Tracker VLAN	6.7.3, "Case Study (3): VLAN Leakage Problem" on page 284
vivl	View list of active VLANs on an interface	6.5.4, "Case Study (1-2): VLAN Assignment - Devices in Wrong VLANs" on page 265
vimcrl	View a multicast VLAN configuration	6.5.4, "Case Study (1-2): VLAN Assignment - Devices in Wrong VLANs" on page 265
fwtvl	View VLAN assignment of learned MAC addresses	6.5.4, "Case Study (1-2): VLAN Assignment - Devices in Wrong VLANs" on page 265
Bridge Submenu		
stc	Configure spanning tree parameters on selected group	6.5.4, "Case Study (1-2): VLAN Assignment - Devices in Wrong VLANs" on page 265
macinfo	Locate learned bridge MAC addresses in this chassis	6.6.3, "Case Study (2-1): Aging Timer Problem without VLAN Polices" on page 271
System Submenu		
slot	View Slot table information	6.5.3, "Case Study (1-1): Simple Configuration Problem - Ports in Wrong Groups" on page 258

6.5 Case Studies for Configuration and VLAN Assignment Problems

This section contains various case studies involving configuration problems and VLAN assignment problems, using IBM products to demonstrate our problem determination guidelines.

6.5.1 Network Environment

The test environment used for the following case studies is described below:

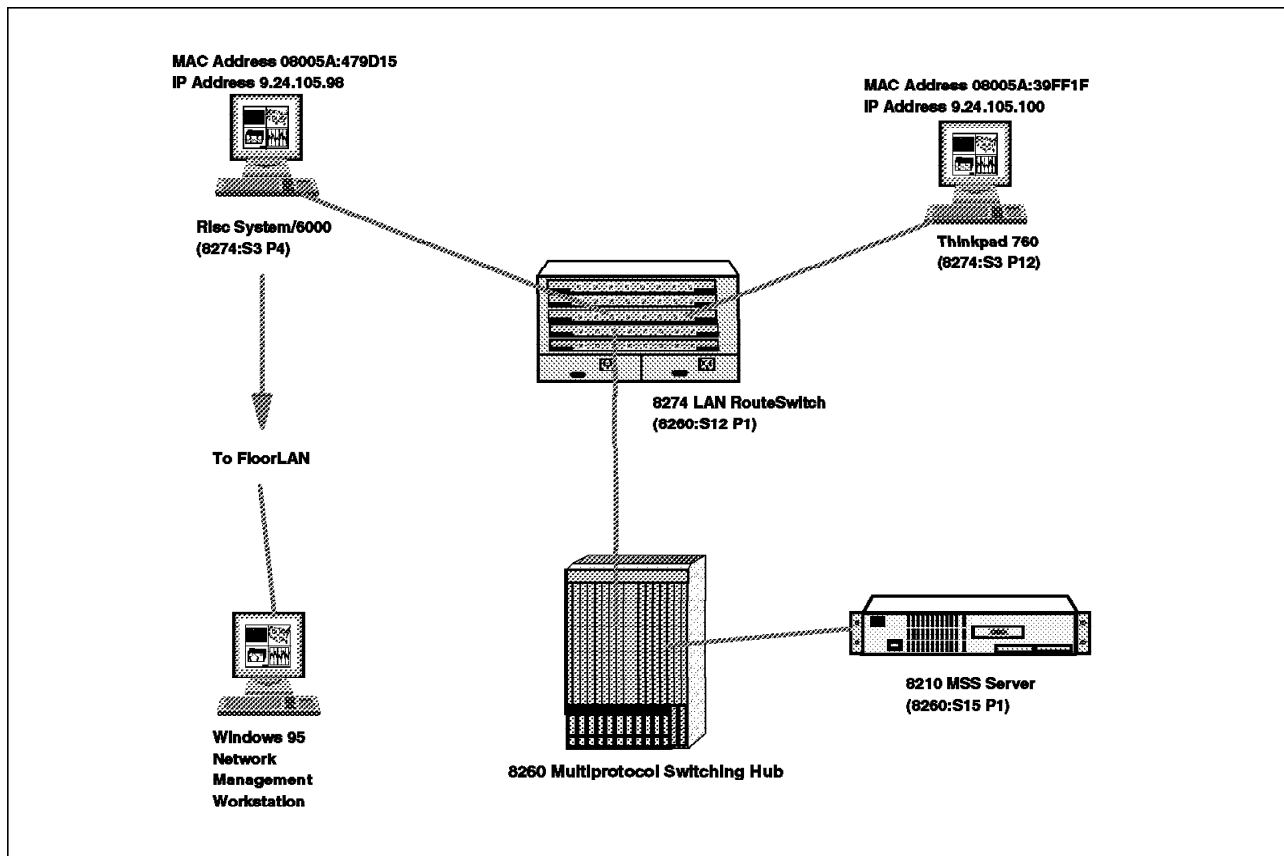


Figure 165. LAN Switch Environment (Physical View)

6.5.2 Symptom: Clients Cannot Communicate with Each Other

Many problems can cause clients to fail to communicate. This section considers some of the most common problems.

6.5.3 Case Study (1-1): Simple Configuration Problem - Ports in Wrong Groups

The logical configuration of the IBM Nways 8274 LAN RouteSwitch is shown in Figure 166.

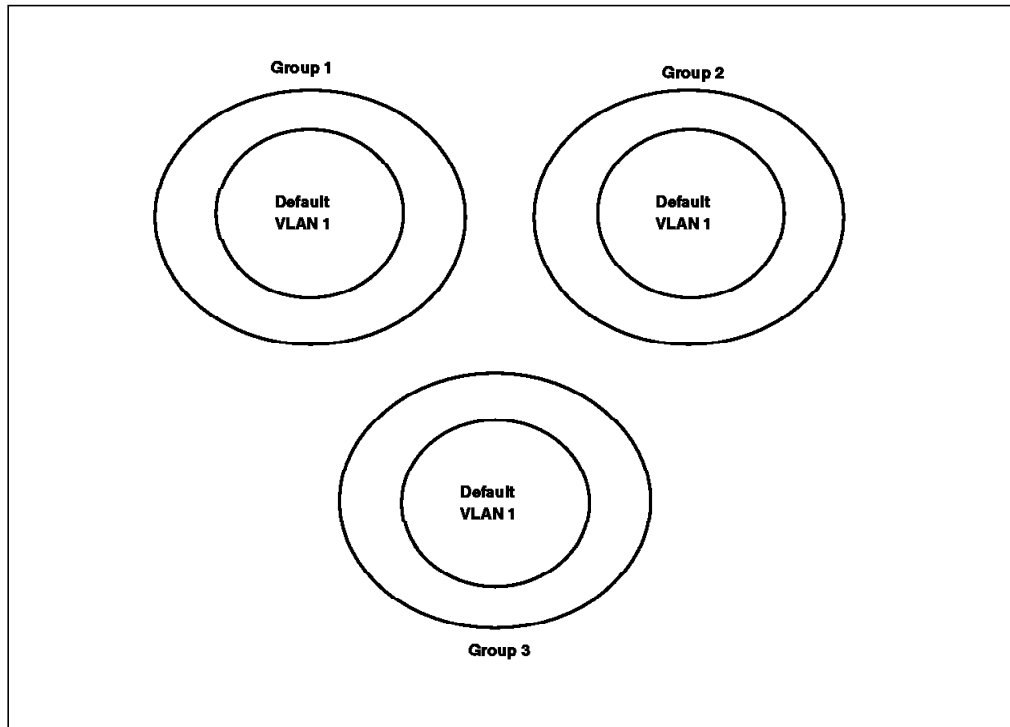


Figure 166. LAN Switch Environment (Logical View Case 1-1)

The logical switch environment consists of:

- Three groups (port VLANs)
- One default VLAN for each group

In this example the ThinkPad 760 could not ping the RISC System/6000.

6.5.3.1 Methodology

To find out why, we looked at the obvious status and configuration items discussed in 6.3.1, “LAN Switch Configuration Problem Methodology” on page 251.

First, we checked the IBM Nways 8274 LAN RouteSwitch LEDs to determine the status of the device and its modules. Figure 167 on page 259 shows a representation of the front panel of the IBM Nways 8274 LAN RouteSwitch. It is captured from the Windows 95 version of the RouteManager software.

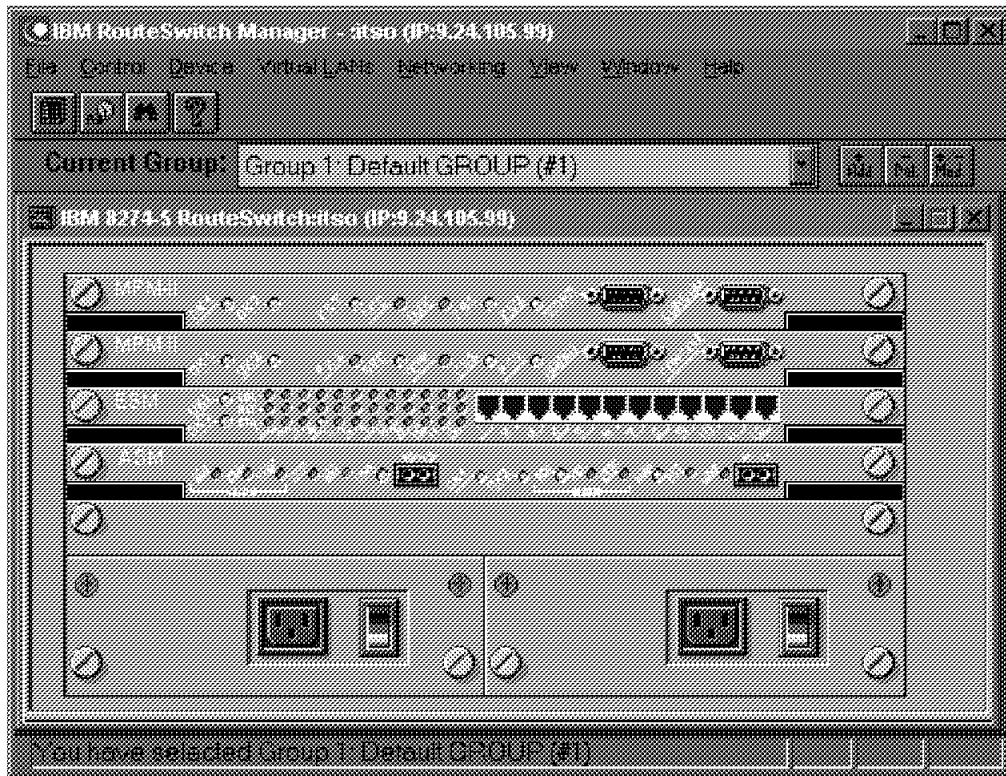


Figure 167. LAN Switch LEDs

LEDs can provide a useful first check to determine whether a hardware failure has occurred. The troubleshooting section of the product guide contains valuable information on the status of a device's LEDs. All the LEDs looked normal.

To further investigate the status of the switch, use its internal commands. The status of the switch can also be interrogated via the management application. This is not always available so we chose to investigate using the internal commands. Check the status of slots and ports via the following commands.

```

/ % slot
      Module-Type Adm-Status HW   Board   Mfg   Firmware-Version
Slot Part-Number Oper-Status Rev   Serial # Date   Base-MAC-Address
-----
1*   MPM-II      Enabled  1.00   64999721 12/17/96 2.1.3
      5012013 Operational
                                     00:20:da:75:8e:d0
2     MPM-II      Enabled  1.00   64999795 12/21/96 2.1.3
      5012013 Redundant
                                     00:20:da:75:e1:c0
3     Ether/12   Enabled  4.04   64385471 11/09/96 2.1.3
      5011206 Operational
                                     00:20:da:73:f1:b0
4     HSM2       Enabled  3.01   63565266 08/26/96 2.1.3
      5011006 Operational
                                     00:20:da:6f:a8:e0
                                     00:20:da:6f:a8:f0
4-1   ATM 2Meg    1.01   63158891 08/27/96
      5011624
4-2   ATM 2Meg    1.01   63158842 08/27/96
      5011624
5     Empty
/ %

```

Figure 168. 8274 Command to View the Status of Its Modules

The IBM Nways 8274 LAN RouteSwitch contained a primary MPM (slot 1), a backup MPM (slot 2), an Ethernet switching module ESM (slot 3) and a high-speed switching module (Slot 4). The slots showed they were all enabled and in the case of the primary MPM, ESM and HSM, operational. The backup MPM is in redundant mode (since it is a backup). Other modes are possible; for more information refer to the product guide that is shipped with your IBM Nways 8274 LAN RouteSwitch.

```

/ % via

                                GROUP Interface Attachments For All Interfaces

GROUP:
Slot/Intf      Description                Service/
-----      -----                -
Instance      Protocol      Admin
-----      -----                -----
3.1 :*                Rtr / 2      CIP      Enabled
2.1 :*                Rtr / 3      IP       Enabled
2.1 :*                Rtr / 3      IPX      Enabled
1:3/1 Virtual port (#1)   Brg / 1      Tns      Enabled
2:3/2 Virtual port (#2)   Brg / 1      Tns      Enabled
1:3/3 Virtual port (#3)   Brg / 1      Tns      Enabled
1:3/4 Virtual port (#4)   Brg / 1      Tns      Enabled
1:3/5 Virtual port (#5)   Brg / 1      Tns      Enabled
1:3/6 Virtual port (#6)   Brg / 1      Tns      Enabled
1:3/7 Virtual port (#7)   Brg / 1      Tns      Enabled
2:3/8 Virtual port (#8)   Brg / 1      Tns      Enabled
1:3/9 Virtual port (#9)   Brg / 1      Tns      Enabled
1:3/10 Virtual port (#10)  Brg / 1      Tns      Enabled
1:3/11 Virtual port (#11)  Brg / 1      Tns      Enabled
2:3/12 Virtual port (#12)  Brg / 1      Tns      Enabled
1:4/1 Virtual port (#18)  Lne / 1      Tns      Enabled
2:4/1 Virtual port (#14)  Lne / 2      Tns      Enabled
3:4/1 Virtual port (#17)  CIP / 1      Tns      Enabled
1:4/2 Virtual port (#15)  Brg / 1      Tns      Enabled
1:5/1 Virtual port (#13)  Brg / 1      Tns      DETACHED

2:5/2 Virtual port (#16)  Brg / 1      Tns      DETACHED

/ %

```

Figure 169. 8274 Command to View the Assignment of Ports to Groups

The command above shows the status of all ports and the groups to which they are attached. Ports 5/1 and 5/2 are from a two-port FDDI switching module that has been removed. The switch continues to remember the configuration but the ports show detached.

Port 3/4 contains the RISC System/6000; the port is enabled and set up for transparent bridging (service=Brg and Protocol=Tns). The port is associated with group 1 **1**.

Port 3/12 contains the ThinkPad; the port is enabled and set up for transparent bridging (service=Brg and Protocol=Tns). The port is associated with group 2 **2**.

The two ports are therefore in different VLAN groups. A VLAN group is a logical broadcast domain. Groups can be attached via routers and external bridges but not normally interlinked internally. To view the definitions for the groups on the IBM Nways 8274 LAN RouteSwitch use the following command.

```

/ % gp
Group
ID          Group Description      Network Address  Proto/
(:VLAN ID) (IP Subnet Mask) Encaps
=====
  1 Default GROUP (#1)
  2 New GROUP (#2)                192.168.7.5     IP /
                                   (ff.ff.ff.00)  FDDI
                                   70000007       IPX /
                                   (0020da:758ed4) FLLC
  3 for CIP                        192.168.20.133 CIP /
                                   (ff.ff.ff.80)  1483
/ %

```

Figure 170. 8274 Command to Show the Defined Groups

The RISC System/6000 and the ThinkPad are in the same IP subnetwork and so could be connected via routers. You can also confirm that the IBM Nways 8274 LAN RouteSwitch is not set up to route between groups by using the *via* command, shown in the routeSwitch commands table 256. The IBM Nways 8274 LAN RouteSwitch has three router interfaces, shown at the top of the command list: An IP router, a CIP router and an IPX router. The IP router is associated with group 2 but group 1 has no such router. The two groups are therefore not router connected by the IBM Nways 8274 LAN RouteSwitch. They could still be connected by an external-bridge or router.

Note

Externally bridging different groups within the same IBM Nways 8274 LAN RouteSwitch is not recommended, especially when bridging between LAN Emulation LECs on the same physical interface.

To fix the problem change the IBM Nways 8274 LAN RouteSwitch definition and add port 3/12 to group 1.

```

/ % addvp 1 3/12
3/12 - This interface has already been assigned to GROUP 2 -
(New GROUP (#2)).
Do you wish to remove it from that GROUP and assign it (with
new configuration values) to this VLAN (n)? y

Slot 3 Port 12 Configuration:
Description :
Bridge Mode: {Auto-Switch (a),
              Optimized Device Switching(o),
              Spanning Tree Bridge(b)}          (b) :
Flood Limit (bytes / second)                    (64000) :
Output format type: {Default(IP-Eth II; IPX-802.3) (d),
                    Ethernet II (e),
                    SNAP(s),
                    LLC(1)}                    (d) :
Ethernet_802.2 Pass Through: {Yes | No }        (y) :
Admin Status { disable (d), enable (e) }        (e) :
slot/port 3/12 is not currently being mirrored
Mirroring enabled { no (n), yes (y) }           (n) :

Adding port 3/12 to GROUP 1...
/ %

```

Figure 171. 8274 Command to Change Group Assignment of a Port

Then confirm that the ports are in the same group.

```

/ % via

GROUP Interface Attachments For All Interfaces

GROUP:
Slot/Intf      Description          Service/
              Instance          Protocol  Admin
              =====          =====
3.1 :*         Rtr / 2          CIP       Enabled
2.1 :*         Rtr / 3          IP        Enabled
2.1 :*         Rtr / 3          IPX       Enabled
1:3/1 Virtual port (#1) Brg / 1     Tns       Enabled
2:3/2 Virtual port (#2) Brg / 1     Tns       Enabled
1:3/3 Virtual port (#3) Brg / 1     Tns       Enabled
1:3/4 Virtual port (#4) Brg / 1     Tns       Enabled
1:3/5 Virtual port (#5) Brg / 1     Tns       Enabled
1:3/6 Virtual port (#6) Brg / 1     Tns       Enabled
1:3/7 Virtual port (#7) Brg / 1     Tns       Enabled
2:3/8 Virtual port (#8) Brg / 1     Tns       Enabled
1:3/9 Virtual port (#9) Brg / 1     Tns       Enabled
1:3/10 Virtual port (#10) Brg / 1     Tns       Enabled
1:3/11 Virtual port (#11) Brg / 1     Tns       Enabled
1:3/12 Virtual port (#12) Brg / 1     Tns       Enabled
1:4/1 Virtual port (#18) Lne / 1     Tns       Enabled
2:4/1 Virtual port (#14) Lne / 2     Tns       Enabled
3:4/1 Virtual port (#17) CIP / 1     Tns       Enabled
1:4/2 Virtual port (#15) Brg / 1     Tns       Enabled
1:5/1 Virtual port (#13) Brg / 1     Tns       DETACHED

2:5/2 Virtual port (#16) Brg / 1     Tns       DETACHED

/ %

```

Figure 172. 8274 Command to View the Assignment of Ports to Groups

Now that the ports are attached to the same VLAN group confirm connectivity by pinging the ThinkPad from the RISC System/6000.

6.5.3.2 Conclusion

As the configuration of switches becomes more and more complex it becomes significantly easier to make configuration mistakes. Checking the device, slot and port status can provide valuable information regarding the particular configuration mistake and make correcting the problem much faster and easier.

6.5.4 Case Study (1-2): VLAN Assignment - Devices in Wrong VLANs

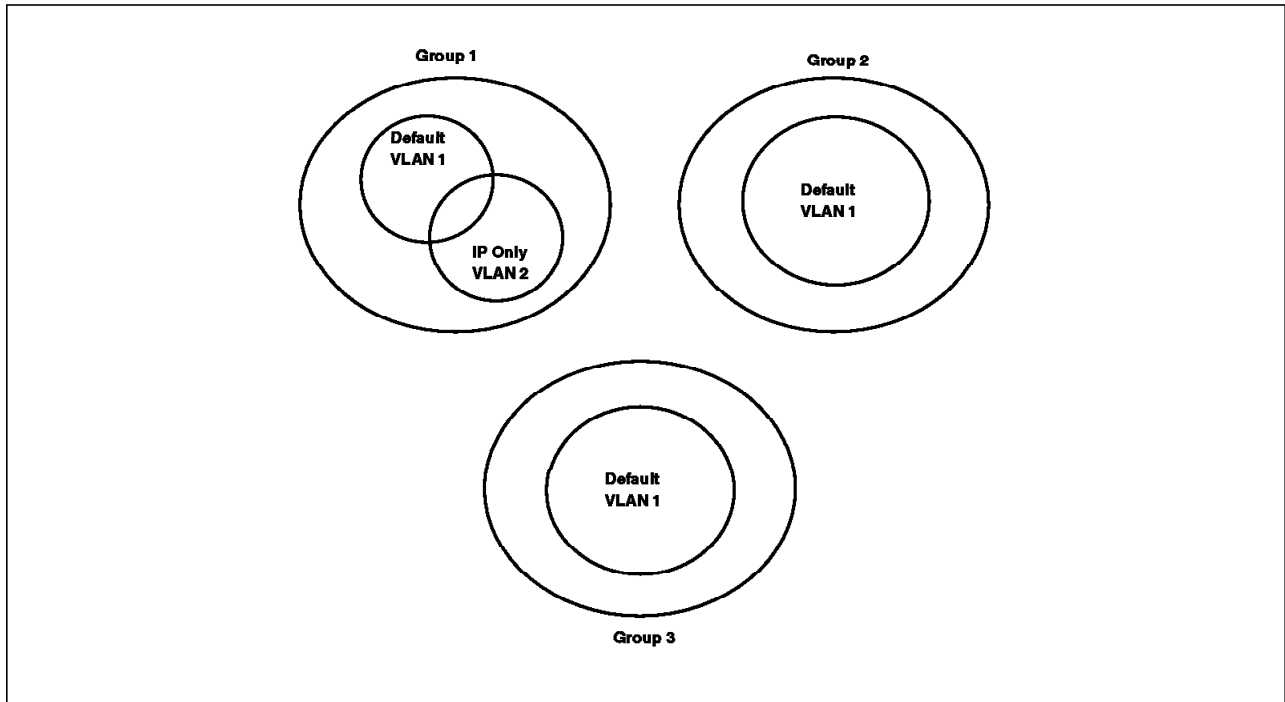


Figure 173. LAN Switch Environment (Logical View Case 1-2)

The logical switch environment consists of:

- Three groups (port VLANs)
- One default VLAN for each group
- One IP Only VLAN in group 1

In this scenario the ThinkPad could not ping the RISC System/6000 but the switch was configured correctly and both ports were in the same VLAN group.

6.5.4.1 Methodology

We first followed the methodology described in 6.3.1, “LAN Switch Configuration Problem Methodology” on page 251 to check that the switch was configured correctly. We found all the LEDs were normal; all slots were working normally, all ports were enabled, and the RISC System/6000 port and ThinkPad ports were both in group 1.

We then used the methodology described in 6.3.2, “VLAN Assignment Problem Methodology” on page 251 to further investigate the problem.

First we checked what policy VLANs and multicast VLANs were defined on the switch using the following two commands.

```

/ % viatr1
VLAN  VLAN  Rule Rule          Rule  Rule
Group:Id  Num  Type          Status Definition
-----
      1: 2    1  PROTOCOL RULE Enabled Protocol = IP
/ %
    
```

Figure 174. 8274 Command to Verify Policy VLANs Defined

```

/ % vimcr1
There are no Multicast VLANs configured
/ %
    
```

Figure 175. 8274 Command to Verify Multicast VLANs Defined

This confirmed that only one VLAN (other than the default) had been defined and that it used a protocol VLAN rule for IP only.

To investigate the VLAN assignment of ports we used the following command.

```

/ % viv1
          Virtual Interface VLAN Membership
Slot/Intf/Service/Instance  Group  Member of VLAN#
-----
1 /1 /Rtr /2      3      1
1 /1 /Rtr /3      2      1
3 /1 /Brg /1      1      1
3 /2 /Brg /1      2      1
3 /3 /Brg /1      1      1
3 /4 /Brg /1      1      1      1
3 /5 /Brg /1      1      1
3 /6 /Brg /1      1      1
3 /7 /Brg /1      1      1
3 /8 /Brg /1      2      1
3 /9 /Brg /1      1      1
3 /10 /Brg /1     1      1
3 /11 /Brg /1     1      1
3 /12 /Brg /1     1      1 2      2
4 /1 /Lne /1      1      1
4 /1 /Lne /2      2      1
4 /1 /CIP /1     3      1
4 /2 /Brg /1      1      1
5 /1 /Brg /1      1      1
5 /2 /Brg /1      2      1
/ %
    
```

Figure 176. 8274 Command to View the VLAN Assignment of Ports

The command shows that port 3/4 (with the RISC System/6000 attached) is only assigned to default VLAN 1 of group 1 **1** while port 3/12 (with the ThinkPad attached) is assigned to default VLAN 1 and IP-only VLAN 2 of group 1 **2**. This indicated that the devices may have a problem. In this environment we assumed both devices should be members of the IP-only VLAN.

Then we checked what MAC devices were associated with the defined VLANs.


```
/ % fwtvl 1
Enter Slot/Interface (return for all ports)
 : Total number of MAC addresses learned for Group 1: 2
Maximum number of entries to display [20] :
  MAC Address      Slot/Intf/Service/Instance  AT VLAN Membership
  -----
08005A:39FF1F     3/ 12/ Brg/ 1                2
/ %
```

Figure 177. 8274 Command to Verify the MAC Addresses Learned for a Group

For this command you must specify the group for which you want to see the MAC to VLAN assignments. You can see from the command above that only the ThinkPad's MAC address has been learned by the switch and it is connected to VLAN 2.

This is probably because the RISC System/6000 has not yet sent any data via the switch or it has been idle for some time and has been aged out of the switch address and VLAN tables. IP ARP broadcasts sent by the ThinkPad to discover the address of the RISC System/6000 will only be sent to ports on common VLANs with the ThinkPad's MAC address, which in the case above is VLAN 2. Since the port with the RISC System/6000 connected is not a member of VLAN 2, it will not receive any of these broadcasts and the two devices will not be able to communicate.

It may be worth checking the spanning tree parameters to discover what the aging timers are for aging out MAC addresses and decreasing these if necessary. The default aging timers are usually adequate. To check the spanning tree parameters use the command below. Since the IBM Nways 8274 LAN RouteSwitch implements separate spanning trees for each group, in its latest release, you must specify the group whose spanning tree you wish to view or change.

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

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

/ %
```

Figure 178. 8274 Command to Verify the Spanning Tree Parameters

You can see from this command that the aging time is 300 seconds **1** and the Auto-Tracker VLAN aging time is 1200 seconds **2**. These are the defaults and probably do not need to be changed.

If your network has network server devices that remain idle for some time, for example LAN printers, it is probably best to assign the ports to which these devices attach directly to their corresponding VLAN. A VLAN can have a number

of policies associated with it. In this case choose to change VLAN 2 to be IP-only or port 3/4 **1**.

```

/ % modatvl 1:2
VLAN 1: 2 is defined as:
  1. Description = IP Only VLAN
  2. Admin Status = Enabled
  3. Rule Definition
      Rule Num   Rule Type   Rule Status
      1          Protocol Rule Enabled
Available options:
  1. Set VLAN Admin Status
  2. Set VLAN Description
  3. Add more rules
  4. Delete a rule
  5. Set rule Admin Status
  6. Quit
Option = 3
Select rule type:
  1. Port Rule
  2. MAC Address Rule
  3. Protocol Rule
  4. Network Address Rule
  5. User Defined Rule
Enter rule type (1): 1
Set Rule Admin Status to [(e)nable/(d)isable] (d): e
Enter the list of ports in Slot/Interface format: 3/4 1
  1. Description = IP Only VLAN
  2. Admin Status = Enabled
  3. Rule Definition
      Rule Num   Rule Type   Rule Status
      1          Protocol Rule Enabled
      2          Port Rule   Enabled
Available options:
  1. Set VLAN Admin Status
  2. Set VLAN Description
  3. Add more rules
  4. Delete a rule
  5. Set rule Admin Status
  6. Quit
Option = 6
/ %

```

Figure 179. 8274 Command to Modify the Policies of a VLAN

Then re-check the VLAN assignment of ports.

```

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

```

Figure 180. 8274 Command to View the VLAN Assignment of Ports

Ports 3/4 and 3/12 are now both in VLANs 1 and 2. The ThinkPad now pings the RISC System/6000 successfully.

6.5.4.2 Conclusion

VLAN assignment is learned through monitoring certain frames arriving at the switch. If these frames do not arrive or have not arrived for a considerable period of time, then devices may not be attached to their correct VLANs. If using policy VLANs it is always worth checking to establish whether the devices are assigned to the correct VLANs or not and if required, create policies to ensure server machines are always attached to the correct VLANs.

6.6 Case Studies for Address Aging Time

Moving stations between switches or ports can cause problems. It takes switches time to reconfigure themselves to detect the move. This section describes the factors affecting the time switches take to reconfigure themselves.

6.6.1 Network Environment

In our network environment the endstations are connected to a downstream hub and a downstream switch. Most switches will recognize when a workstation, connected directly to the switch, is disconnected and moved. They will immediately flush the address tables for the device's port and will therefore immediately recognize the station has moved.

The test environment used for the following case studies is shown below:

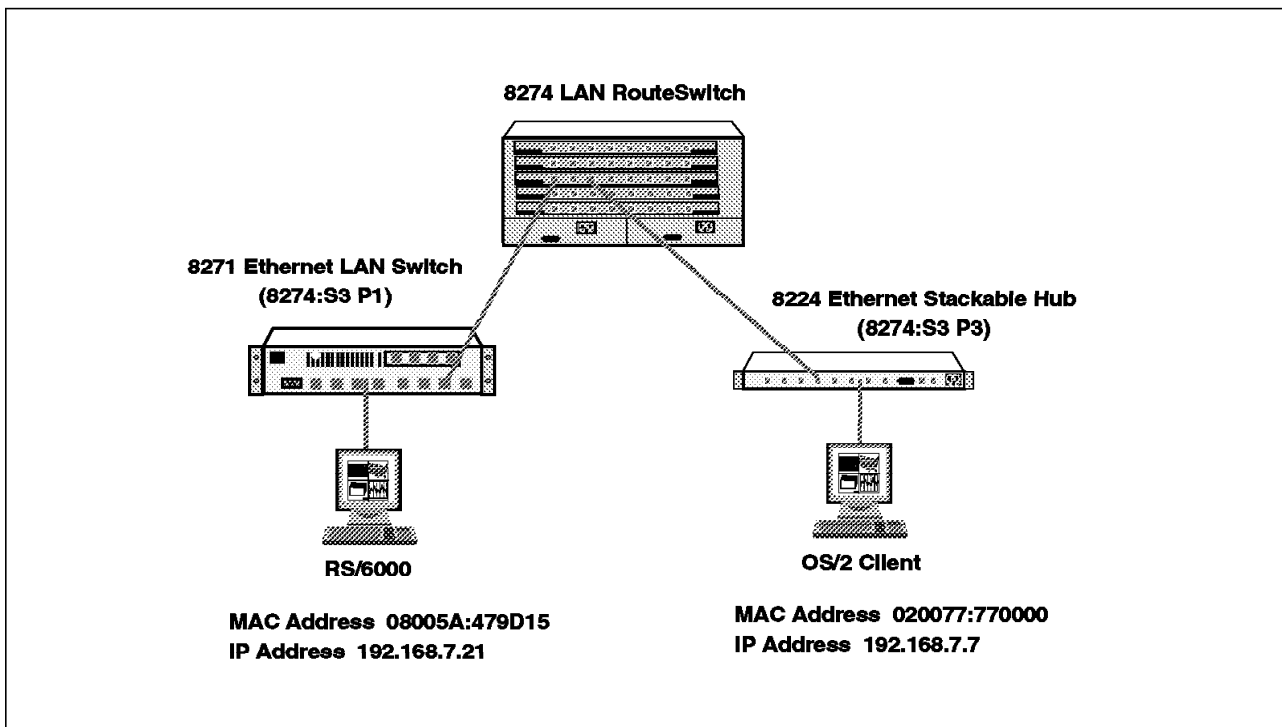


Figure 181. LAN Switch Environment for Aging Timer Test

6.6.2 Symptom

Workstations can't communicate via the switch for several minutes after being moved.

6.6.3 Case Study (2-1): Aging Timer Problem without VLAN Polices

The IBM Nways 8274 LAN RouteSwitch had no groups or VLAN policies, so all ports belong to the same broadcast domain. After the RISC System/6000 was moved from the IBM Nways 8271 Ethernet LAN Switch to the IBM Nways 8274 LAN RouteSwitch (slot3/port2), the OS/2 client couldn't ping to the RISC System/6000.

6.6.3.1 Methodology

First follow the methodology described in 6.3.1, "LAN Switch Configuration Problem Methodology" on page 251 to check that the switch was configured correctly. We found all the LEDs were normal, all slots were working normally, all ports were enabled, and that all ports were in group 1.

We also used the procedure described in 6.3.2, "VLAN Assignment Problem Methodology" on page 251 to confirm both ports were attached to the same VLAN (VLAN 1).

Then we continued to investigate the problem using the procedure described in 6.3.3, "Aging Timers and Address Tables Problem Methodology" on page 252.

First investigate the MAC address table of ports by using the following command:

```
/ % macinfo
Enter MAC address ([XXYYZZ:AABBCC] or return for none) :
Enter Slot Number (1-5) : 3
Total number of MAC addresses learned for this slot: 2
Maximum number of entries to display [20] :
          Non-Canonical   Group CAM   Last   Exp   ATM
S1/If/Service/In  MAC Address  MAC Address T  ID   Indx S  Seen  Timer VCI
-----
  3/ 1/   Brg/ 1 08005A:479D15 10005A:E2B9A8 E   1 17E6 T   34   300   0
  3/ 3/   Brg/ 1 020077:770000 4000EE:EE0000 E   1 17E4 T   76   300   0
```

Figure 182. 8274 Command to Display the MAC address Information For a Slot

In the above screen the RISC System/6000 was still assigned to slot3/port1, which was the port with the IBM Nways 8271 Ethernet LAN Switch attached. However the RISC System/6000 had been moved to slot3/port2. The IBM Nways 8274 LAN RouteSwitch was therefore forwarding ping packets, which were sent from the OS/2 client, to slot3/port1 based on this address table. The two clients were therefore failing to communicate.

After 1 minute the two started to communicate. So we checked the MAC address table of ports again.

```

/ % macinfo
Enter MAC address ([XXYYZZ:AABBCC] or return for none) :
Enter Slot Number (1-5) : 3
Total number of MAC addresses learned for this slot: 2
Maximum number of entries to display [20] :

```

S1/If/Service/In	MAC Address	Non-Canonical MAC Address	Group T ID	CAM Indx	Last S Seen	Exp Timer	ATM VCI
3/ 2/ Brg/ 1	08005A:479D15	10005A:E2B9A8	E	1 17E4	T 11	300	0
3/ 3/ Brg/ 1	020077:770000	4000EE:EE0000	E	1 17E6	T 0	300	0

Figure 183. 8274 Command to Display the MAC Address Information for a Slot

The RISC System/6000 has now been moved to the address table for port 2. The aging time for the table is set to 300 seconds (shown in the Exp Timer column). It would take the switch 300 seconds to detect the move if only unicast frames were being sent, so we concluded that the RISC System/6000 must be sending some other broadcast frames to the IBM Nways 8274 LAN RouteSwitch

The following is the network trace from slot3/port2:

Number	Abs Time	Destination	Source	Interpretation
203	17:42:05.98787	192.168.7.255	192.168.7.21	RIP (TCP/IP) Response ID=IP Entries=5
266	17:43:05.98739	192.168.7.255	192.168.7.21	RIP (TCP/IP) Response ID=IP Entries=5

Figure 184. Network Trace from the Risc System/6000 Port on the 8274 Switch

The routing information protocol (RIP) was enabled on the RISC System/6000. Therefore the RISC System/6000 sent a RIP broadcast packet every 60 seconds. The IBM Nways 8274 LAN RouteSwitch analyzed the broadcast packet, discovered the move and updated its address tables.

We confirmed that the "routed" process, which uses the RIP protocol, was running in the RISC System/6000 by using the following command.

```

# ps -ef | grep routed
root 9340 5870 0 17:30:09 - 0:00 /usr/sbin/routed -g
root 10010 30974 2 17:44:26 pts/0 0:00 grep routed

```

Figure 185. RS/6000 Command to Display Active Processes Running

6.6.3.2 Conclusion

Some stations send broadcast frames periodically. For example Novell servers send RIP and SAP frames and IP gateways may send RIP frames. These allow the IBM Nways 8274 LAN RouteSwitch to detect station moves and update its address tables. If these are not sent, it may take the switch the time specified by the aging time parameter to learn of the workstation move.

6.6.4 Case Study (2-2): Aging Timer Problem with VLAN Polices

A second VLAN was defined in the IBM Nways 8274 LAN RouteSwitch based on the following VLAN policies:

1. Network address rules (IP address: 192.168.7.0)
2. Port rules (Slot3/Port2, Slot3/Port3)

After the OS/2 client is moved from the IBM Nways 8224 Ethernet Stackable Hub to the IBM Nways 8274 LAN RouteSwitch (slot3/port2), the RISC System/6000 can't ping it for some time.

6.6.4.1 Methodology

First follow the methodology described in 6.3.1, "LAN Switch Configuration Problem Methodology" on page 251 to check that the switch was configured correctly. We found all the LEDs were normal, all slots were working normally, all ports were enabled, and that all ports were in group 1.

We then used the procedure described in 6.3.2, "VLAN Assignment Problem Methodology" on page 251 to investigate VLAN membership of the ports and devices.

We checked that the VLAN polices were defined as the information above describes.

```
/ % viatr1
VLAN  VLAN  Rule Rule      Rule      Rule
Group:Id  Num  Type      Status  Definition
-----
   1:  2      1  NET ADDR RULE  Enabled  IP Addr = 192.168.7.0
                                           IP Mask = 255.255.255.0
                                           2  PORT RULE    Enabled  3/2/Brg/1
                                           3  PORT RULE    Enabled  3/3/Brg/1
```

Figure 186. 8274 Command to View VLANs Defined

The policies were correctly defined so we checked the port VLAN assignment.

```

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

```

Figure 187. 8274 Command to View the VLAN Assignment of Ports

The RISC System/6000 port and the OS/2 client port were both in group 1 and VLAN 2.

After more than 2 minutes, the RISC System/6000 still couldn't ping the OS/2 client so we investigated further using the procedure described in 6.3.3, "Aging Timers and Address Tables Problem Methodology" on page 252. We first checked the port MAC address table by using the following command:

```

/ % macinfo
Enter MAC address ([XXYYZZ:AABBCC] or return for none) :
Enter Slot Number (1-5) : 3
Total number of MAC addresses learned for this slot: 2
Maximum number of entries to display [20] :
                Non-Canonical   Group CAM   Last   Exp   ATM
S1/If/Service/In  MAC Address  MAC Address T  ID   Indx S  Seen  Timer VCI
-----
3/ 1/   Brg/ 1 08005A:479D15 10005A:E2B9A8 E   1 17E6 T   10   300   0
3/ 3/   Brg/ 1 020077:770000 4000EE:EE0000 E   1 17E4 T   297  300   0

```

Figure 188. 8274 Command to Display the MAC Address Information for a Port

In the above screen the OS/2 client was assigned to slot3/port3, the IBM Nways 8274 LAN RouteSwitch was therefore forwarding the ping packet sent from the RISC System/6000 to slot3/port3. The Last Seen time for slot3/port3 was shown as 297 seconds, close to the aging time of 300 seconds. After it expired, the RISC System/6000 could ping the OS/2 client. We confirmed that this was changed using the following command:


```
/ % macinfo
Enter MAC address ([XXYYZZ:AABBCC] or return for none) :
Enter Slot Number (1-5) : 3
Total number of MAC addresses learned for this slot: 2
Maximum number of entries to display [20] :
S1/If/Service/In  MAC Address      Non-Canonical  Group CAM   Last   Exp  ATM
                  MAC Address      MAC Address    T  ID   Indx S  Seen  Timer VCI
-----
3/ 1/   Brg/ 1 08005A:479D15 10005A:E2B9A8 E   1 17E4 T   13   300   0
3/ 2/   Brg/ 1 020077:770000 4000EE:EE0000 E   1 17E6 T    2   300   0
```

Figure 189. 8274 Command to Display the MAC Address Information for a Port

6.6.4.2 Conclusion

The clients cannot communicate until the aging time expires since the OS/2 client does not send any frames onto the network after its move.

The following diagram shows the data flow in the aging process.

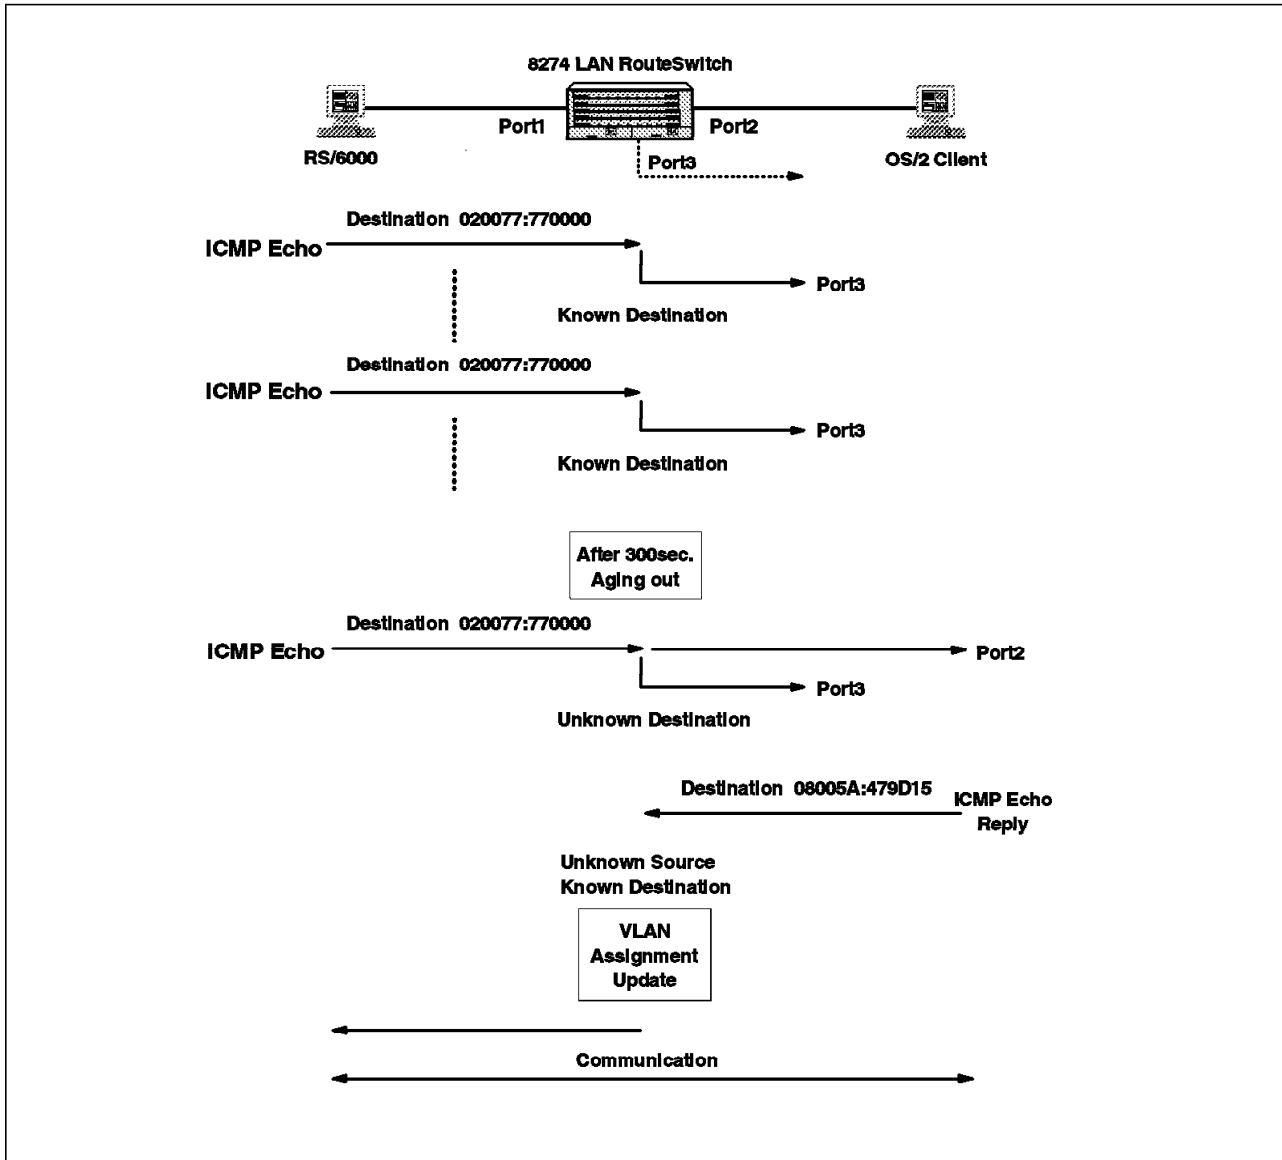


Figure 190. Data Flow in the Aging Timeout

After aging out, the ICMP echo frame is forwarded out of all ports that have at least one VLAN in common with the RISC System/6000. In this diagram, the RISC System/6000 has the OS/2 client MAC address entry in his ARP table. It therefore always sends unicast ICMP ECHO requests. If the RISC System/6000 doesn't have an entry for the OS/2 client in its ARP cache, it will send an IP ARP broadcast which is forwarded out of all ports that have at least one VLAN in common with the RISC System/6000. The stations would then be able to communicate immediately.

6.6.5 Case Study (2-3): Auto-Tracker VLAN Aging Timer Problem

We changed the IBM Nways 8274 LAN RouteSwitch definition for VLAN 2 so that it was just based on an IP network address (192.168.7.0) rule as shown below. After initialization the RISC System/6000 and OS/2 client were assigned to VLAN 2. After some time the RISC System/6000 couldn't ping the OS/2 client.

```
/ % viatr1
VLAN  VLAN  Rule Rule          Rule  Rule
Group:Id  Num  Type          Status Definition
-----
      1: 2      1 NET ADDR RULE Enabled IP Addr = 192.168.7.0
```

Figure 191. 8274 Command to View VLANs Defined

6.6.5.1 Methodology

The following diagram describes how the VLAN assignment of the MAC addresses of these devices and their ports change over time.

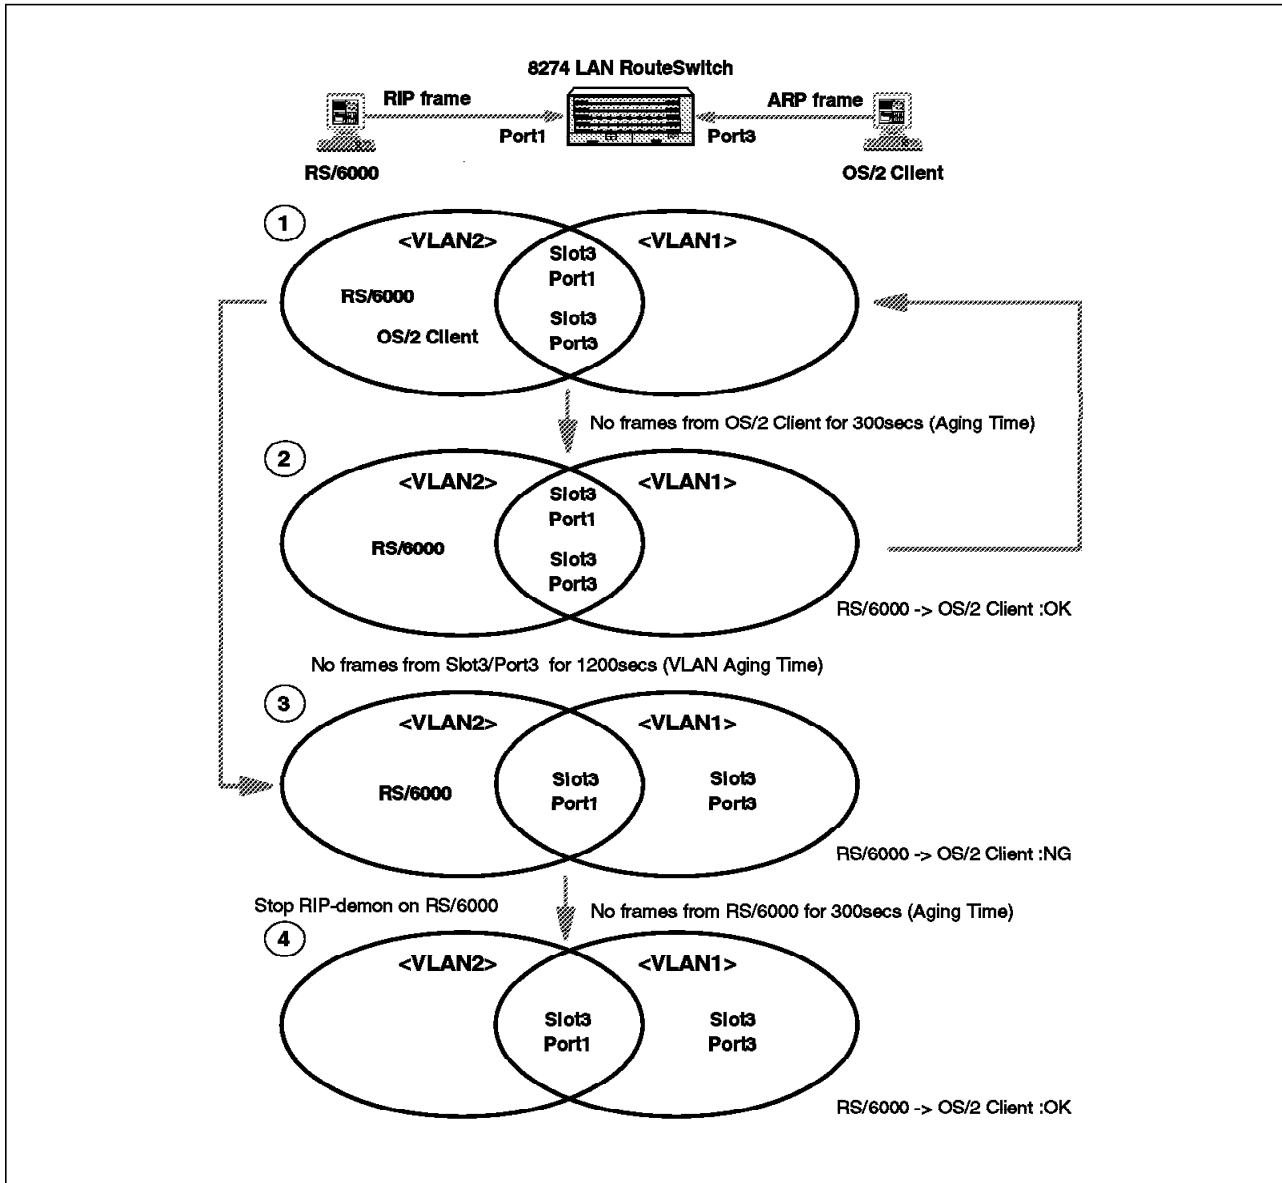


Figure 192. Auto-Tracker VLAN Aging Timer

To demonstrate why the RISC System/6000 and OS/2 client can't communicate we investigate each stage of the diagram above.

1. Initialization:

Initially all ports and MAC devices belong to the default VLAN 1. The RIP frame of RISC System/6000 is received at the IBM Nways 8274 LAN RouteSwitch port. The frame is completely analyzed by the MPM. As this frame matches the defined policy (network address rule) for VLAN 2, the following will occur:

- The MAC address of RISC System/6000 is removed from the default VLAN 1 and moved to VLAN 2. Therefore, the RISC System/6000 is no longer a member of the default VLAN 1.
- The port to which the RISC System/6000 is connected (in this case slot3/port1) is added to VLAN 2 but also remains in VLAN 1, because all ports are always members of the default VLAN 1.

- VLAN 2 is activated because slot3/port1 is the first port assigned to that VLAN.

The OS/2 client sends an IP ARP broadcast to communicate with the RISC System/6000. The frame is also analyzed by the MPM. The MAC address of the OS/2 client is moved to VLAN 2 and slot3/port3 is added to VLAN2. The actual MAC addresses and port assignments can be checked by the following commands:

- vivl - Display the VLAN membership of ports
- fwtvl - Display the VLAN membership of MAC addresses

```

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

Figure 193. 8274 Command to Display the VLAN Membership of Ports

```

/ % fwtvl
Enter Slot/Interface (return for all ports) :
Total number of MAC addresses learned for Group 1: 2
Maximum number of entries to display [20] :
MAC Address           Slot/Intf/Service/Instance   AT VLAN Membership
-----
08005A:479D15         3/ 1/ Brg/ 1           2
020077:770000         3/ 3/ Brg/ 1           2
    
```

Figure 194. 8274 Command to Display the VLAN Membership of MAC Addresses

We can see that the ports and MAC addresses are assigned to VLAN 2.

2. No frames are sent from the OS/2 client for 300 seconds (aging time).

All learned MAC addresses and their VLAN membership information is stored in the Content Addressable Memory (CAM). If the IBM Nways 8274 LAN RouteSwitch doesn't receive a subsequent frame from the OS/2 client within the aging time (default: 300 secs.), the OS/2 client's MAC address and VLAN membership flag are removed from the CAM, but port3/3 remains in VLAN 2 as seen below. In this situation the OS/2 client's MAC address becomes an unknown address again. The unicast frames that are sent from the RISC System/6000 are forwarded out of all ports (3/1, 3/3) that have at

least one VLAN (VLAN 2) in common with the RISC System/6000. The OS/2 client will therefore receive the frames.

```

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

Figure 195. 8274 Command to Display the VLAN Membership of Ports

Ports 3/1 and 3/3 are both in VLAN 1 and VLAN 2.

```

/ % fwtvl
Enter Slot/Interface (return for all ports) :
Total number of MAC addresses learned for Group 1: 1
Maximum number of entries to display [20] :
MAC Address      Slot/Intf/Service/Instance  AT VLAN Membership
-----
08005A:479D15   3/ 1/ Brg/ 1 2
    
```

Figure 196. 8274 Command to Display the VLAN Membership of MAC Addresses

The OS/2 client's MAC address has been removed from VLAN 2.

3. No frames from slot3/port3 for 1200 seconds.

The port timeout or VLAN aging time is different from the aging time and set to 1200 seconds as the default value. If the OS/2 client is connected to a dedicated port at the IBM Nways 8274 LAN RouteSwitch and doesn't send a frame out that matches the VLAN 2 policy within a time period of 1200 seconds, the OS/2 client's port assignment is deleted from VLAN 2. Thus, the OS/2 client's port is no longer in a common VLAN with the RISC System/6000. The frames from the RISC System/6000 will not be forwarded to slot3/port3.

```

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

```

Figure 197. 8274 Command to Display the VLAN Membership of Ports

```

/ % fwtvl
Enter Slot/Interface (return for all ports) :
Total number of MAC addresses learned for Group 1: 1
Maximum number of entries to display [20] :
MAC Address      Slot/Intf/Service/Instance  AT VLAN Membership
-----
08005A:479D15    3/ 1/ Brg/ 1 2

```

Figure 198. 8274 Command to Display the VLAN Membership of MAC Addresses

4. No frames from the RISC System/6000 for 300 seconds.

If we now stop all frames from the RISC System/6000, including RIP frames, for 300 seconds, the IBM Nways 8274 LAN RouteSwitch removes the RISC System/6000's MAC address from VLAN 2 and the MAC address of the RISC System/6000 also becomes an unknown address. Frames from the RISC System/6000 will now be flooded out all ports, which are in the same group and the RISC System/6000 can communicate with the OS/2 client again.

```

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

```

Figure 199. 8274 Command to Display the VLAN Membership of Ports

```

/ % fwtvl
Enter Slot/Interface (return for all ports) :
Total number of MAC addresses learned for Group 1: 0

```

Figure 200. 8274 Command to Display the VLAN Membership of MAC Addresses

6.6.5.2 Conclusion

The above case study describing the IBM Nways 8274 LAN RouteSwitch behavior may cause communication problems. To stop these problems from happening you could configure your devices to send some frames periodically. The IBM Nways 8274 LAN RouteSwitch will then assign workstations to VLANs by monitoring these frames. Silent devices connected to a dedicated port at the IBM Nways 8274 LAN RouteSwitch, can be configured to always join a specific VLAN by adding a port rule to the VLAN policy.

6.7 Case Studies for VLAN Leakage

VLAN leakage occurs when two devices in different VLANs receive data from each other. This section contains the case study to simulate VLAN leakage.

6.7.1 Network Environment

The test environment used for the following case study is described below.

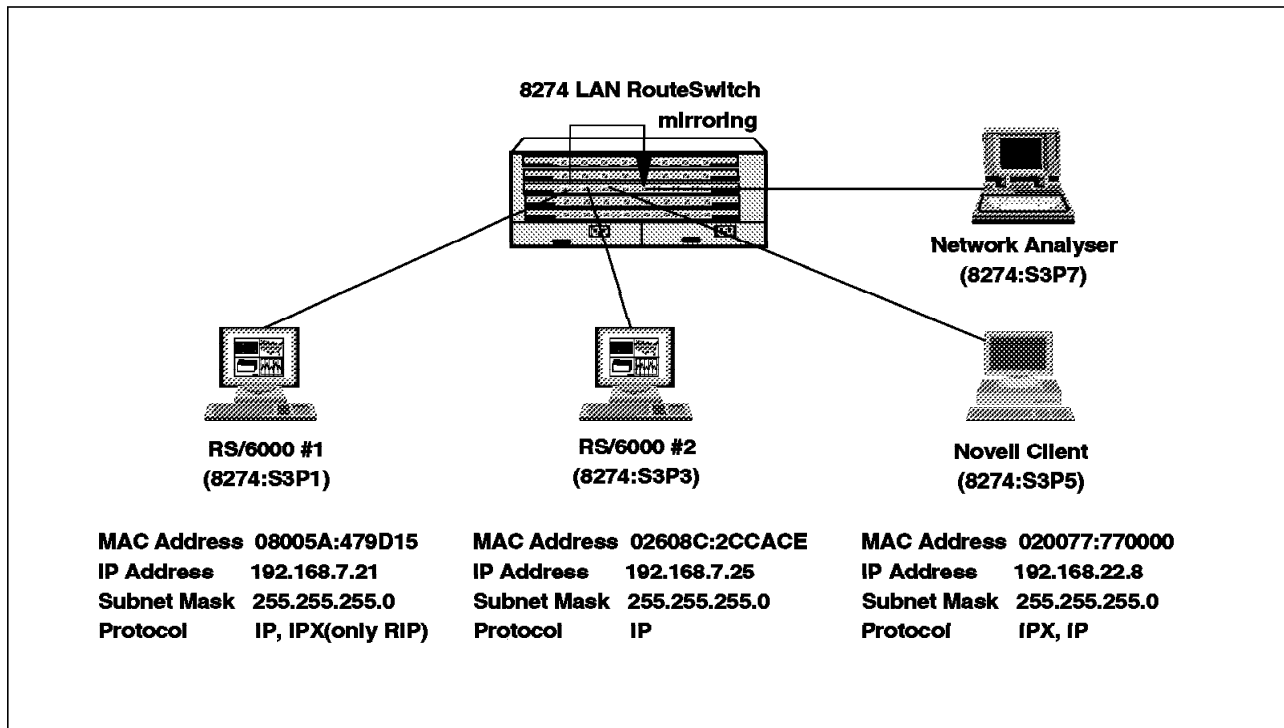


Figure 201. LAN Switch Environment for VLAN Leakage Test

On the IBM Nways 8274 LAN RouteSwitch you can set up port mirroring for any pair of 10 Mbps Ethernet ports within the same switch chassis. We enabled port mirroring of port slot3/port1 to port slot3/port7. When you enable port mirroring, the active/mirrored port (slot3/port1) transmits and receives its network traffic normally, and the mirroring port (slot3/port7) receives a copy of all the transmitted and received frames from the active port. You can then connect an RMON probe or a network analysis device to the mirroring port to see all the data from the mirrored port without disrupting the mirrored port. Port mirroring can be set up using the `addvp` or `modvp` commands as follows:

```

/ % modvp 1 3/1
Modify local port 1 (Virtual port (#1)) ? (y) :

Slot 3 Port 1 Configuration:
Description (Virtual port (#1)) :
Bridge Mode: {Auto-Switch (a),
              Optimized Device Switching(o),
              Spanning Tree Bridge(b)}          (b) :
Flood Limit (bytes / second)          (64000) :
Output format type: {Default(IP-Eth II; IPX-802.3) (d),
                    Ethernet II (e),
                    SNAP(s),
                    LLC(1)}                  (d) :
Ethernet_802.2 Pass Through: {Yes ,, No }    (y) :
Admin Status { disable (d), enable (e) }     (e) :
slot/port 3/1 is not currently being mirrored
Mirroring enabled { no (n), yes (y) }        ( ) : y
Mirroring vport slot/port ? ( ) : 3/7

```

Figure 202. 8274 Command to Modify a Port Setup

6.7.2 Symptom

If the RISC System/6000 #1 and the Novell client are assigned to different VLANs, but join an additional common VLAN, the broadcast domain is expanded to the common VLAN and the two can communicate.

6.7.3 Case Study (3): VLAN Leakage Problem

In this environment all ports of the IBM Nways 8274 LAN RouteSwitch belonged to the default group 1. At first the IBM Nways 8274 LAN RouteSwitch had two VLANs defined with the following rules:

- VLAN 2: Network address-based VLAN (IP 192.168.7.0)
- VLAN 3: Network address-based VLAN (IP 192.168.22.0)

Then we defined an additional VLAN with the following rule:

- VLAN 4: Protocol-Based VLAN (IPX)

We confirmed that the IBM Nways 8274 LAN RouteSwitch flooded the IP frames from the Novell client to the RISC System/6000 and attached a different IP subnetwork by using a network analyzer.

6.7.3.1 Methodology

First we followed the methodology described in 6.3.1, “LAN Switch Configuration Problem Methodology” on page 251 to check that the switch was configured correctly. We found all the LEDs were normal, all slots were working normally, all ports were enabled, and that all ports were in group 1.

We then used the procedure described in 6.3.4, “VLAN Leakage Problem Methodology” on page 252 to investigate VLAN membership of the ports and devices.

We checked that the VLAN policies were defined as described above by using the following command:

```

/ % viatr1
VLAN  VLAN  Rule Rule      Rule      Rule
Group:Id  Num  Type          Status    Definition
-----
   1:  2      1  NET ADDR RULE Enabled  IP Addr = 192.168.7.0
                                     IP Mask = 255.255.255.0
   1:  3      1  NET ADDR RULE Enabled  IP Addr = 192.168.22.0
                                     IP Mask = 255.255.255.0
    
```

Figure 203. 8274 Command to View VLANs Defined

If every station has sent at least one IP frame to the IBM Nways 8274 LAN RouteSwitch, all devices will be known by the IBM Nways 8274 LAN RouteSwitch and assigned to the appropriate VLAN as shown in Figure 204:

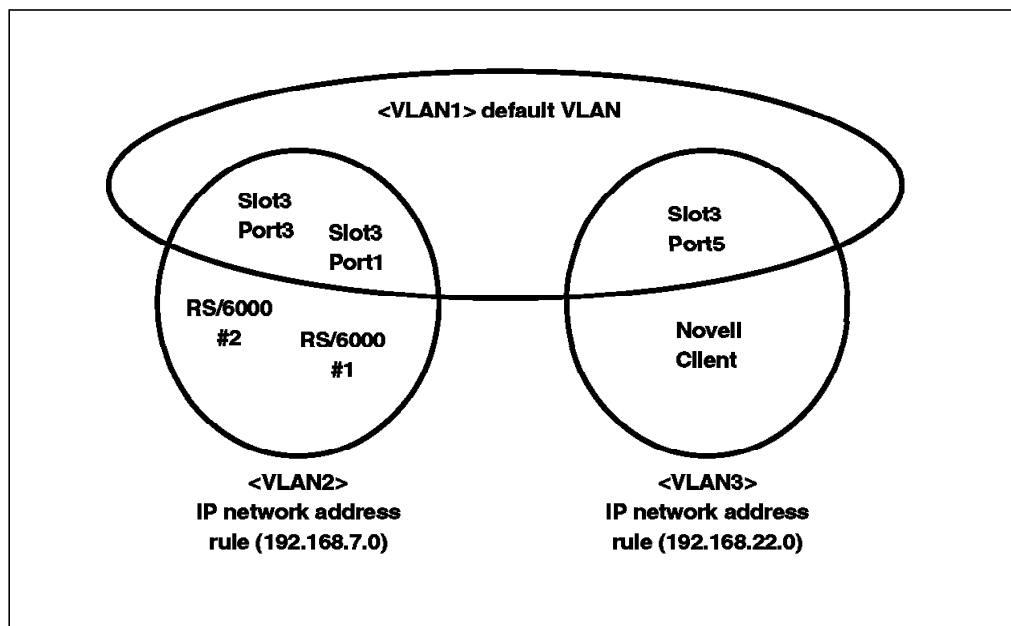


Figure 204. Logical View: Two IP Network-Based VLANs

If the RISC System/6000 #1 sends an IP ARP broadcast, this broadcast will be forwarded to all ports which are in the same circle as the RISC System/6000 #1, that is VLAN 2 only. In this case, the ARP frames will show up at slot3/port1 and slot3/port3. In the same way, a broadcast from the Novell client will only be forwarded to VLAN 3. The actual ports and MAC address assignments can be checked by the following commands:

```

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

```

Figure 205. 8274 Command to View the VLAN Assignment of Ports

```

/ % fwtvl
Enter Slot/Interface (return for all ports) :
Total number of MAC addresses learned for Group 1: 3
Maximum number of entries to display [20] :
  MAC Address      Slot/Intf/Service/Instance   AT VLAN Membership
-----
08005A:479D15     3/ 1/ Brg/ 1 2
020077:770000     3/ 5/ Brg/ 1 3
02608C:2CCACE     3/ 3/ Brg/ 1 2

```

Figure 206. 8274 Command to Display the VLAN Membership of MAC Addresses

Then we created an additional VLAN rule based on IPX protocol as follows:

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

Figure 207. 8274 Command to Create a VLAN Rule

We checked that the VLAN policies reflected the additional VLAN rule.

```
/ % viatr1
VLAN  VLAN  Rule Rule      Rule  Rule
Group:Id  Num  Type          Status Definition
-----
 1: 2      1  NET ADDR RULE Enabled IP Addr = 192.168.7.0
                               IP Mask = 255.255.255.0
 1: 3      1  NET ADDR RULE Enabled IP Addr = 192.168.22.0
                               IP Mask = 255.255.255.0
 1: 4      1  PROTOCOL RULE Enabled Protocol = IPX
```

Figure 208. 8274 Command to View VLANs Defined

If every station has sent at least one IP/IPX frame to the IBM Nways 8274 LAN RouteSwitch, all devices will be known by the IBM Nways 8274 LAN RouteSwitch and assigned to the appropriate VLAN as shown in Figure 209 on page 288:

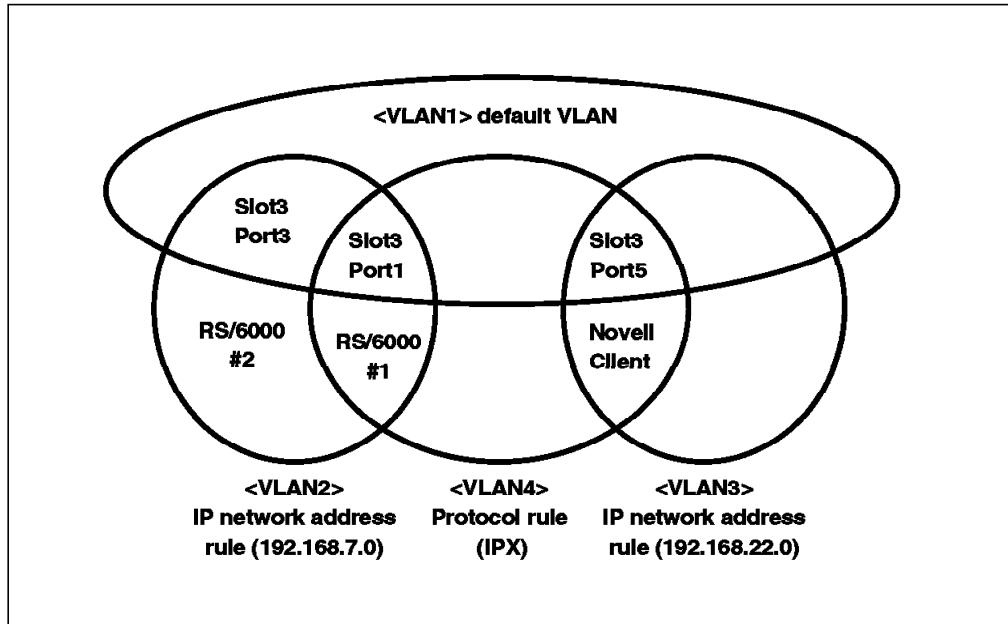


Figure 209. Logical View: Two IP Network and IPX-Based VLANs

The actual ports and MAC address assignments can be checked with the following commands:

```

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

Figure 210. 8274 Command to View the VLAN Assignment of Ports

```
/ % fwtv1
Enter Slot/Interface (return for all ports) :
Total number of MAC addresses learned for Group 1: 3
Maximum number of entries to display [20] :
  MAC Address          Slot/Intf/Service/Instance  AT VLAN Membership
  -----
08005A:479D15         3/ 1/ Brg/ 1                2 4
020077:770000         3/ 5/ Brg/ 1                3 4
02608C:2CCACE         3/ 3/ Brg/ 1                2
```

Figure 211. 8274 Command to Display the VLAN Membership of MAC Addresses

The following traces were captured at slot3/port1 in the IBM Nways 8274 LAN RouteSwitch:

Number	DeltaTime	Destination	Source	Interpretation
38	2.5 sec	IEEE 802.1D Bridges BPDU BPDU_Type=Configuration	00-20-DA-73-F1-B0	1
39	2.5 sec	IEEE 802.1D Bridges BPDU BPDU_Type=Configuration	00-20-DA-73-F1-B0	
40	2.5 sec	IEEE 802.1D Bridges BPDU BPDU_Type=Configuration	00-20-DA-73-F1-B0	
41	2.5 sec	IEEE 802.1D Bridges BPDU BPDU_Type=Configuration	00-20-DA-73-F1-B0	
42	2.3 sec	IEEE 802.1D Bridges BPDU BPDU_Type=Configuration	00-20-DA-73-F1-B0	
43	1.3 sec	00000002.FFFFFFFF RIP (Novell) Response #Nets=1	00000002.08005A479D15	2
44	135.5 ms	00000000.FFFFFFFF SAP Nearest Service Query ServerType=File Server (SLIST source)	00000000.020077770000	3
45	1.0 sec	IEEE 802.1D Bridges BPDU BPDU_Type=Configuration	00-20-DA-73-F1-B0	
46	381.5 ms	00000000.FFFFFFFF SAP Nearest Service Query ServerType=File Server (SLIST source)	00000000.020077770000	
47	1.4 sec	00000000.FFFFFFFF SAP Nearest Service Query ServerType=File Server (SLIST source)	00000000.020077770000	
48	712.8 ms	IEEE 802.1D Bridges BPDU BPDU_Type=Configuration	00-20-DA-73-F1-B0	
49	694.0 ms	00000000.FFFFFFFF SAP General Service Query ServerType=File Server (SLIST source)	00000000.020077770000	
50	1.8 sec	IEEE 802.1D Bridges BPDU BPDU_Type=Configuration	00-20-DA-73-F1-B0	
51	600.4 ms	00000000.FFFFFFFF SAP Nearest Service Query ServerType=File Server (SLIST source)	00000000.020077770000	
52	1.4 sec	00000000.FFFFFFFF SAP Nearest Service Query ServerType=File Server (SLIST source)	00000000.020077770000	
53	494.8 ms	IEEE 802.1D Bridges BPDU BPDU_Type=Configuration	00-20-DA-73-F1-B0	
54	911.5 ms	00000000.FFFFFFFF SAP Nearest Service Query ServerType=File Server (SLIST source)	00000000.020077770000	
55	1.4 sec	00000000.FFFFFFFF SAP General Service Query ServerType=File Server (SLIST source)	00000000.020077770000	
56	182.9 ms	IEEE 802.1D Bridges BPDU BPDU_Type=Configuration	00-20-DA-73-F1-B0	
57	2.2 sec	00000000.FFFFFFFF SAP Nearest Service Query ServerType=File Server (SLIST source)	00000000.020077770000	
58	275.9 ms	IEEE 802.1D Bridges BPDU BPDU_Type=Configuration	00-20-DA-73-F1-B0	
59	269.4 ms	192.168.7.255 RIP (TCP/IP) Response ID=IP Entries=5	192.168.7.21	4
60	693.7 ms	192.168.22.5 ARP REQUEST Desired_Prtcl=192.168.22.5 Protocol=DOD Internet Protocol	192.168.22.8	5
61	166.1 ms	00000000.FFFFFFFF SAP Nearest Service Query ServerType=File Server (SLIST source)	00000000.020077770000	
62	813.6 ms	192.168.22.5 ARP REQUEST Desired_Prtcl=192.168.22.5 Protocol=DOD Internet Protocol	192.168.22.8	
63	557.8 ms	IEEE 802.1D Bridges BPDU BPDU_Type=Configuration	00-20-DA-73-F1-B0	
64	34.9 ms	00000000.FFFFFFFF SAP Nearest Service Query ServerType=File Server (SLIST source)	00000000.020077770000	
65	407.3 ms	192.168.22.5 ARP REQUEST Desired_Prtcl=192.168.22.5 Protocol=DOD Internet Protocol	192.168.22.8	
66	1.0 sec	192.168.22.5 ARP REQUEST Desired_Prtcl=192.168.22.5 Protocol=DOD Internet Protocol	192.168.22.8	

Figure 212. Network Trace from Port Slot3/Port1

From the trace above we can see:

- **1** The IBM Nways 8274 LAN RouteSwitch was configured to use the spanning tree protocol, as defined by the IEEE 802.1 D standard, and BPDU frames were being forwarded from the IBM Nways 8274 LAN RouteSwitch.
- **2** After we created the IPX protocol VLAN rule, an IPX-RIP frame was sent by the RISC System/6000 #1. This broadcast frame was completely analyzed by MPM and the VLAN assignment was updated in the CAM. Then the MAC address of RISC System/6000 #1 belonged to VLAN 4 but also remained in VLAN 2 and slot3/port1 was added to VLAN 4.
- **3** An IPX-SAP frame was sent by the Novell client. The MAC address of the Novell client was also assigned to VLAN 4 and this broadcast frame was flooded out of slot3/port1 and slot3/port5, which are members of VLAN 4.
- **4** The RISC System/6000 #1 sent the IP RIP frame periodically; therefore the MAC address of RISC System/6000 #1 remains in VLAN 2.
- **5** We issued a ping from the Novell client to generate an IP ARP broadcast. Its MAC address belonged to VLAN 3 and VLAN 4, so the broadcast frame was flooded out slot3/port1 and slot3/port5. The Novell client sent this frame to VLAN 3 (192.168.22.0), but the RISC System/6000 #1 that was assigned to VLAN 2 (192.168.7.0) received this frame. VLAN leakage has occurred.

6.7.3.2 Conclusion

If the source device is a member of multiple VLANs, some VLAN leakage may occur during the frame flooding process as shown in Figure 213. VLAN leakage may occur only among VLANs in the same group. The frames do not leak between groups.

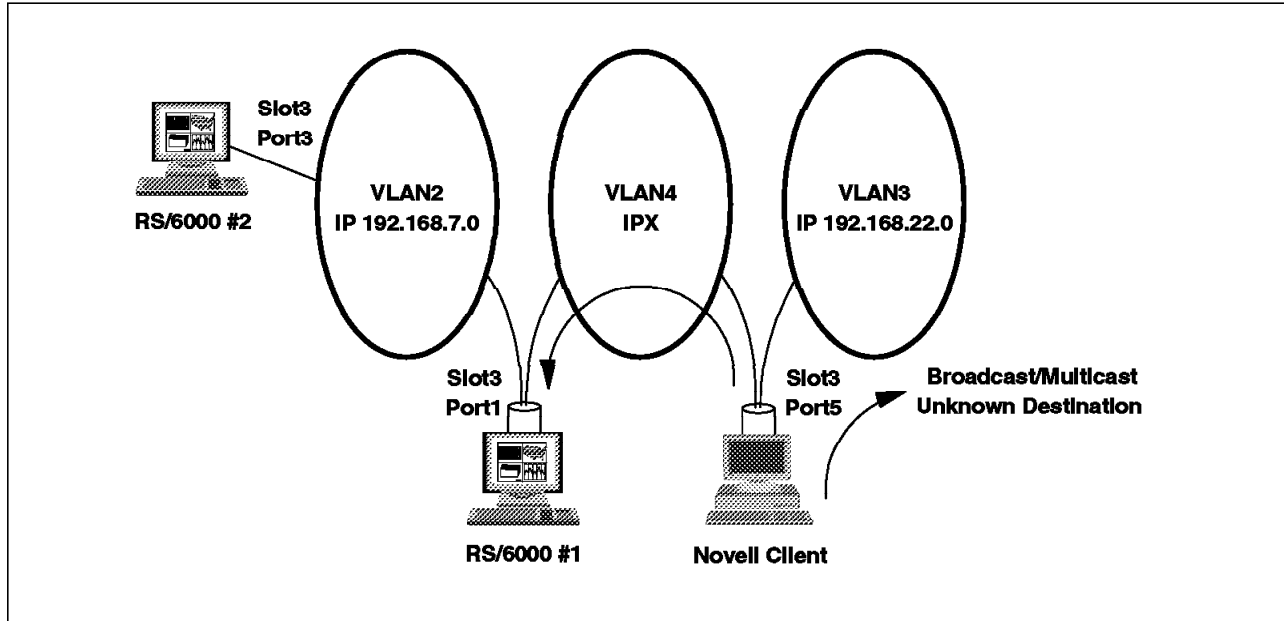


Figure 213. VLAN Leakage

Chapter 7. ATM Emulated LANs and Logical IP Subnets (LANE 1.0, RFC 1577)

This chapter provides guidelines for problem determination in an ATM Forum-Compliant LAN Emulation (LANE 1.0) or Classical IP (RFC 1577) environment. It also uses some typical problem scenarios to illustrate the use of these guidelines in diagnosing problems within IBM networks.

7.1 What Makes a Healthy LANE 1.0 or Classical IP Environment

LANE 1.0 provides a way for Ethernet/IEEE 802.3 and token-ring/IEEE 802.5 LAN MAC-layer protocols, and their higher-layer protocols and applications, to transparently communicate with devices across an ATM network. Clients implement a LAN emulation protocol that enables the client to transfer data over the ATM network and provides LAN services to the higher layer protocols. Classical IP provides a way to transport IP data over ATM networks.

Both are dependent on the ATM network to provide the connectivity between devices. The healthy state of these networks is, therefore, reliant on the healthy state of the underlying ATM network. However the customization and use of these technologies can have a significant impact on the state of the ATM networks they use. The healthy condition of ATM networks is further discussed in Chapter 5, "Starting Problem Isolation in an ATM Network" on page 109.

LANE 1.0 implements an address resolution and broadcast function for emulated LAN traffic. If excessively used, this can significantly impact network utilization and as a result cause data loss and transmission delays. This can also have a considerable impact on the servers used to provide these functions which will impact data transmission.

Classical IP also uses a method of address resolution but does not permit broadcasts. Constant address polling will impact the ability of devices to transmit data and will therefore impact performance.

To confirm your network is performing to the best of its abilities it is therefore necessary to monitor both address resolution and, for LANE 1.0, broadcasts on your network.

IBM has implemented a number of proprietary functions in its IBM Nways Multiprotocol Switched Services (MSS) server product that fully support the ATM Forum-Compliant LAN Emulation (LANE 1.0) standard and extend it to limit the impact of significant numbers of address resolution polls and broadcasts in a LANE 1.0 environment. For more information refer to 7.2.3, "IBM MSS Server Performance Extensions" on page 301.

7.1.1 Address Resolution

The most important items to monitor for both LANE 1.0 and Classical IP are:

- The number of address resolution polls
- The number of unsuccessful address resolution polls

Measuring these items over a period of time (possibly hourly or daily) will give a good indication of the state of the address resolution process.

In a LANE 1.0 or Classical IP environment using SVCs this is usually possible by using a network management machine and collecting Management Information Protocol (MIB) statistics from the central address resolution server. In a Classical IP environment using PVCs this will not be possible since there is no central server. Classical IP environments using PVCs impose a considerable overhead for the network administrator, in defining and keeping track of all the PVC links. Therefore they are usually very small. The address resolution process is performed by each client and hence distributed throughout the network. The impact on each individual client and its network connection is considerably reduced. Our ability to monitor this process is therefore not usually a problem. If it becomes necessary to monitor, we may be able to collect adapter statistics from each client or from critical clients.

Many factors affect the number of address resolution polls made in your network, for example the size of the network, the size of the device's address cache, and the time before a cache ages out. The best way to determine an acceptable level is to measure the number of polls over a period of time to enable you to baseline your network.

In normal working conditions where all clients are communicating, address resolution failures will only occur when a client attempts to connect to a device that does not exist. This is an attempt to connect to an unknown MAC address with LANE 1.0 or an unknown TCP/IP address in Classical IP. Most network protocols will send out a broadcast to resolve the MAC address of a client before attempting to connect to it. The number of unsuccessful address resolution attempts will therefore be very low. To determine an acceptable level it is also best to baseline your network, but as a rough guide this value should be somewhere between 10 and 500 hundred per day.

7.1.2 Broadcasts

Excessive numbers of broadcasts in your network can give you a good indication that something is wrong. They may be caused by a badly functioning device or set of devices or may simply be an indication of high utilization. The number of broadcasts that is acceptable or not acceptable is very difficult to judge. The best method of evaluating this is again to baseline your network.

7.2 Rules of ATM Forum-Compliant LAN Emulation (LANE 1.0) and Classical IP (RFC 1577) Networks

The ATM LAN Emulation SubWorking Group (LANE SWG) completed the architecture model and specification for ATM LAN Emulation late in 1994. The specification was adopted by the ATM Forum in March 1995. Prior to this, IBM as well as other vendors, was already researching and developing technologies in this field. This has led to a number of proprietary specifications of LAN Emulation. Most of these proprietary versions have very similar concepts and design rules to the ATM Forum Compliant version. IBM's intention has always been that our proprietary version products will be phased-out in favor of (or upgraded to) products conforming to the ATM Forum's LAN emulation specifications as those specifications become available.

RFC 1577 is a much earlier standard defined in IETF RFC 1577 titled Classical IP and ARP over ATM and commonly referred to as Classical IP. IBM has been delivering this functionality for the RISC System/6000 since March of 1994.

Most networking vendors use LAN Emulation or Classical IP as the encapsulation scheme for data traffic over ATM networks; however, many other solutions exist. The ATM Forum has defined another specification for Multiprotocol over ATM (MPOA). MPOA is intended to address multiple protocol suites and leverage the Quality of Service (QoS) benefits of ATM. The ATM Forum has also extended the LANE 1.0 standard and created LANE 2.0. These new standards are not currently in wide-spread use and will therefore not be examined in this book.

This section summarizes some of the key rules for ATM Forum-Compliant LAN Emulation (LANE 1.0) and Classical IP (RFC 1577) networks.

7.2.1 ATM Forum-Compliant LAN Emulation (LANE 1.0) Initialization Rules

The main components of an emulated LAN are as follows:

- The ATM network
Provides the underlying connectivity between all the components on the emulated LAN.
- LAN Emulation Configuration Server (LECS)
Assigns individual LE clients to different emulated LANs using its own policies, configuration databases and information supplied by the LE clients. Clients may or may not support the use of the LECS. Each LECS may assign clients to many emulated LANs. Networks can be configured to support a backup LECS.
- LAN Emulation Server (LES)
Provides a facility for registering and resolving MAC addresses and route descriptors to ATM addresses. It responds directly to queries or forwards a query. Each emulated LAN can have only one LES. It may be configured to use a backup LES in case of failure.
- Broadcast and Unknown Server (BUS)
Handles data sent by an LE client to the broadcast MAC address ('FFFFFFFF'), functional addresses, all multicast traffic, route explorer frames and initial unicast frames that are sent by an LE client before the target ATM address is resolved. Frames are serialized and can not be interleaved. Each emulated LAN can have only one BUS server. It may be configured to use a backup BUS in case of failure.
- LAN Emulation Client (LE client or LEC)
An LE entity is, for example a user workstation, bridge or switch, which performs data forwarding, address resolution and control functions. Bridges and switches that connect an emulated LAN to a legacy LAN are known as Proxy LE clients (Proxy LECs)
- Connections
Different connection VCCs are used for control information and data.

The different control and data VCCs are shown in the diagram below:




Figure 214. ATM LAN Emulation Connection VCCs

The following table shows the number of VCCs required by an emulated LAN.

Table 18. Number of VCCs Required by Emulated LAN	
VCC Type	Number Required
Configuration Direct VCC	One per LEC using the LECS
Control Direct VCC	One per LEC attached to each emulated LAN
Control Distribute VCC	One per emulated LAN
Multicast Send VCC	One per LEC attached to each emulated LAN
Multicast Forward VCC	One per emulated LAN
Data Direct VCC	One for each LEC-to-LEC connection
Note: The configuration direct VCCs are normally dropped by LECs after configuration is complete.	

The ATM network must be capable of supporting all the control and data VCCs required by all emulated LANs.

7.2.1.1 LE Service Operation

The operation of LAN emulation is defined by the LAN emulation user-to-network interface. This is summarized below:

1. Initialization

Figure 215 below summarizes the initialization process.

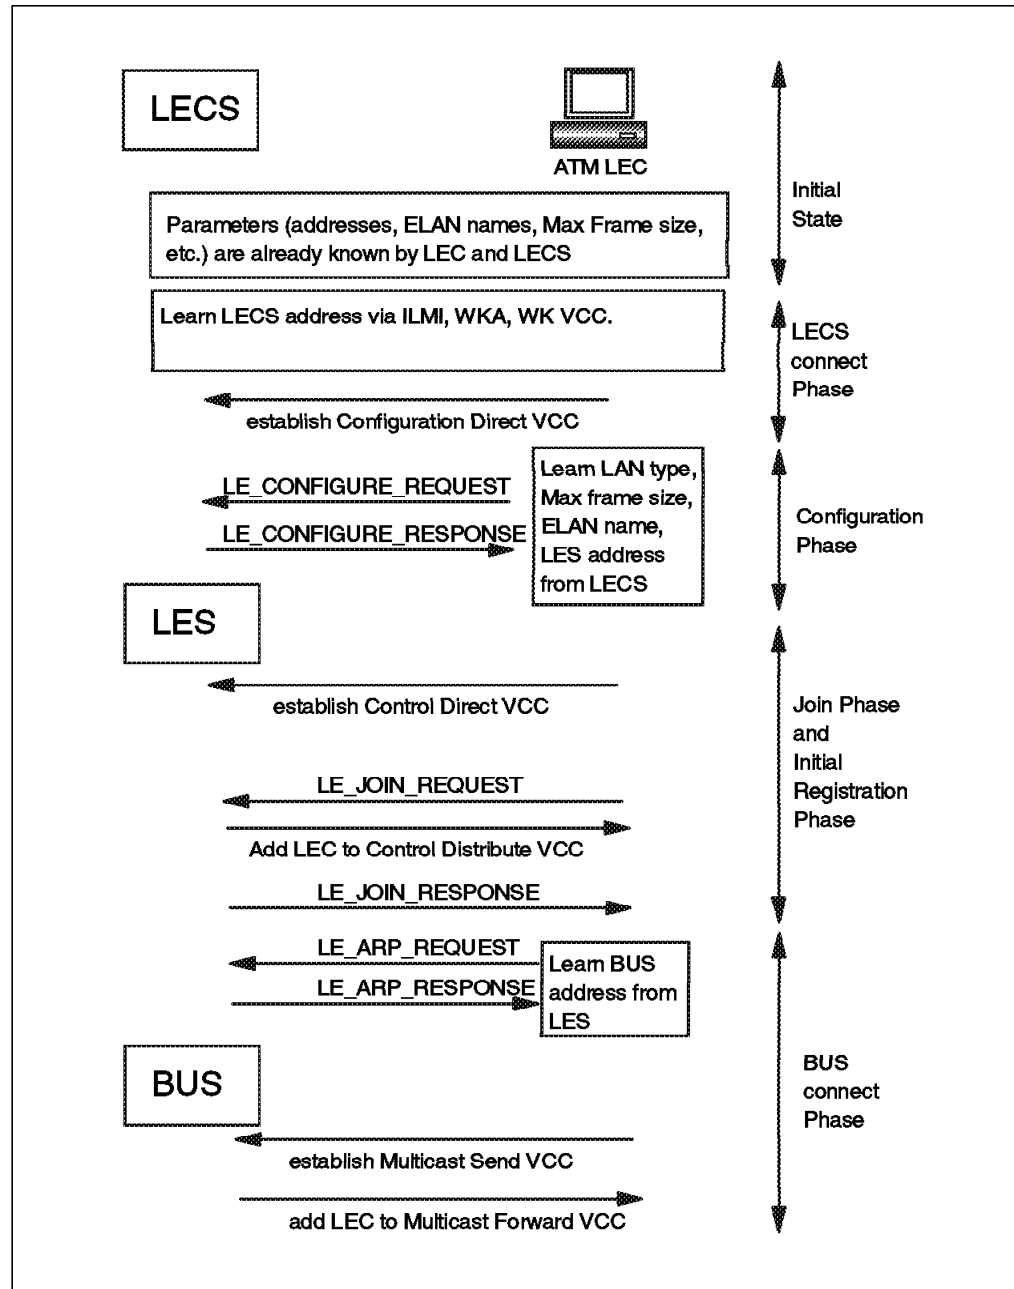


Figure 215. LE Initialization Process

The initial state of an LE client is an implementation issue, but certain parameters are subject to range constraints. Parameters have minimum, maximum and default values. If any parameter falls outside of the range, the LE client may fail to function. The ATM Forum defined parameters for workstations and servers are shown in D.2, "ATM Forum LAN Emulation

Client Parameters” on page 480 and D.1, “ATM Forum LAN Emulation Server Parameters” on page 479.

During initialization an LE client may or may not choose to use a LECS. If the LE client chooses to use a LECS it may discover the address of the LECS in any one of three ways:

a. ILMI

The hub to which the LE client is attached must be configured to provide the LECS address during ILMI via a Get or GetNext request.

b. Using the well-known address (WKA) of the LECS
(X'4700790000000000000000000000A03E00000100')

Normally the hub, to which the LE client is attached, must be configured to map the address of the LECS to the WKA of the LECS. Some hubs support the registration of foreign addresses. In this case the LECS may be set up to use the actual WKA and the ATM network must route calls to this address across the network to the correct hub.

c. Using the Well-Known VCC (WK VCC) of the LECS (VPI=0 VCI=17)

Normally the hub, to which the LE client is attached, must be configured with a PVC to the LECS. The PVC must use the WK VCC of the LECS.

Once the LE client knows the ATM address of the LECS it establishes a configuration direct VCC to its address. This is known as the LECS connect phase.

The client obtains various configuration parameters from the LECS. These include the ATM address of its LES, LAN type, maximum frame size, and the emulated LAN name. It does this by issuing an LE_CONFIGURE_REQUEST frame to and receiving an LE_CONFIGURE_RESPONSE frame from the LECS. If the LE client does not receive an LE_CONFIGURE_RESPONSE frame within a specified configurable time period, it can retry until the configured retries are exhausted, at which time the configuration fails. This is known as the configuration phase.

Note

If an LE client chooses not to use a LECS, it must be pre-configured with the ATM address of the LES.

Once the LE client is configured it establishes a control direct VCC to the LES and attempts to join the emulated LAN. It does this by issuing an LE_JOIN_REQUEST frame to the LES, containing its configured parameters. The LES validates the request by checking for duplicate MAC and ATM addresses, maximum frame size and LAN type. If the conditions are met for a successful join, the LES adds the LE client to its control distribute VCC and sends an LE_JOIN_RESPONSE frame to the LE client. If no LE_JOIN_RESPONSE is received by the LE client within a configurable time period (default 120 seconds), the client fails to join the emulated LAN. This is known as the join phase.

During registration an LE client may register one MAC address with the LES. This is known as the initial registration phase.

After the client has successfully joined the emulated LAN it discovers the address of the BUS by sending an LE_ARP_REQUEST frame to the LES to resolve the broadcast MAC address. The LES responds with the BUS address in an LE_ARP_RESPONSE frame. The LE client then establishes a

multicast send VCC and the BUS adds the client to its multicast forward VCC. If either of these fail, the LE client will be removed from the emulated LAN. If these complete, the LE client is now a member of the emulated LAN. This final stage is known as the BUS connect phase.

2. Data transfer

The diagram below summarizes the data transfer process.

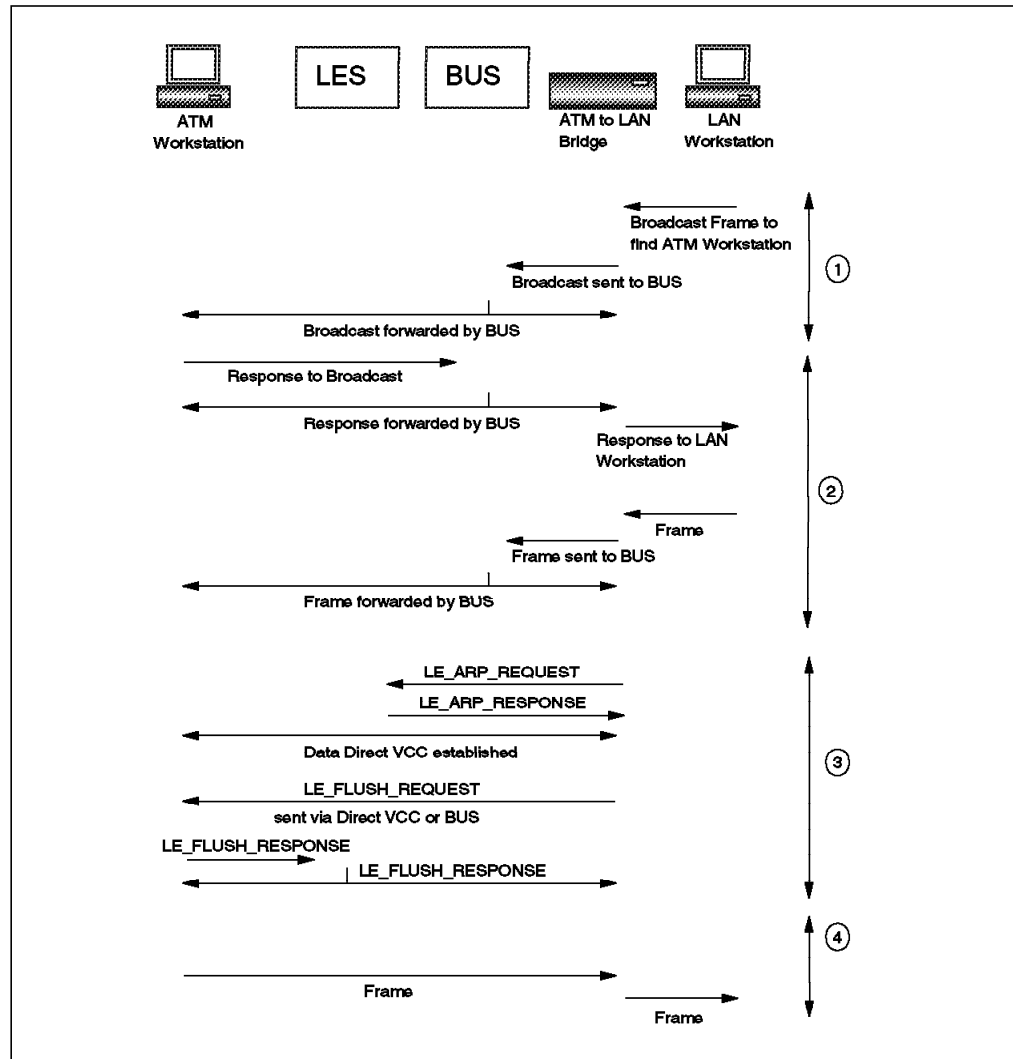


Figure 216. LE Data Transfer Process

Figure 216 shows four stages of data transfer between two LE clients. These are:

- Stage 1 - Broadcast via BUS

Most of today's networking protocols establish the MAC address of a station they wish to communicate with by sending out a broadcast frame on the network. In Figure 216 the LAN workstation sends out a broadcast frame which is picked up by the ATM-to-LAN bridge (Proxy LE client) and forwarded via the BUS to all other ATM workstations and bridges (LE clients).

- Stage 2 - Unicast data sent via BUS

LE clients may then use the BUS to send unicast frames. The destination station sends a response to the original LAN workstation via the BUS to the ATM-to-LAN bridge which forwards the response on to the LAN workstation. The LAN workstation may then send more unicast frames through the ATM-to-LAN bridge which forwards frames via the BUS.

- Stage 3 - LE clients set up a direct connection

LE clients may continue to send frames via the BUS until they have completed setting up a direct connection. LE clients continue to send data while at the same time an LE client tries to determine the ATM address of the destination LE client. The ATM-to-LAN bridge tries to find out the ATM address of the ATM workstation. It does this by sending an LE_ARP_REQUEST frame, containing the MAC address of the ATM workstation, to the LES. The LES responds with the ATM address of the workstation via an LE_ARP_RESPONSE frame. The ATM-to-LAN bridge then establishes a data direct VCC with the ATM workstation.

If the LES does not know the ATM address for a workstation's MAC address sent in an LE_ARP_REQUEST frame, it distributes the LE_ARP_REQUEST to all LE clients and proxy LE clients. If a proxy LE client knows the MAC address is a station connected to one of the LAN segments to which it is attached, it responds to the LES with its ATM address. The LES can then respond to the original station.

Once the data direct VCC is established the ATM-to-LAN bridge sends an LE_FLUSH_REQUEST frame via the BUS to the ATM workstation. Once the LE_FLUSH_REQUEST is sent it holds or discards any frames destined for the workstation. The ATM workstation will then respond with an LE_FLUSH_RESPONSE which it sends via the LES to the ATM to LAN bridge.

- Stage 4 - LE clients send data via direct connection

Once the LE_FLUSH_RESPONSE is received, all subsequent frames sent between the ATM workstation and bridge will be sent via the data direct VCC. The bridge will then forward the frames to the LAN-attached workstations. The flush process ensures that all frames are delivered in order.

At any time LE clients may register or unregister MAC-ATM address pairs with the LES by sending LE_REGISTER_REQUEST frames or LE_UNREGISTER_REQUEST frames to the LES and by receiving LE_REGISTER_RESPONSE frames or LE_UNREGISTER_RESPONSE frames from the LES. They may also advertise changes in remote address bindings using the LE_NARP_REQUEST and LE_NARP_RESPONSE frames and indicate topology changes using the LE_TOPOLOGY_REQUEST and LE_TOPOLOGY_RESPONSE frames.

The format of the ATM Forum LAN Emulation frames is shown in D.3, "Configuration Frame Format" on page 484.

7.2.2 Differences between IBM-Compliant and ATM Forum-Compliant LAN Emulation

The main differences between IBM and ATM Forum-Compliant LAN Emulation (LANE 1.0), are with the BUS and LECS. The BUS that is identified within the ATM Forum specification as a separate entity, is, of course present within IBM-compliant LAN emulation but fully integrated within the LES. Due to this the ATM Forum emulated LAN requires more VCCs.

IBM-compliant LAN emulation uses a point-to-point default VCC between each LE client and the LES, a point-to-multipoint general multicast VCC between the LES and every LE client and a point-to-multipoint bridge multicast VCC between the LES and every proxy LE client.

In IBM-Compliant LAN emulation there is no LECS. The ATM Forum LECS concept may simplify the management of emulated LANs significantly. Learning server addresses dynamically relieves LE clients from hard-coding server addresses, simplifies LES backup and lays the foundation for LES distribution.

7.2.3 IBM MSS Server Performance Extensions

The IBM Nways Multiprotocol Switched Services (MSS) server provides a multiprotocol networking solution for the ATM environment. It can be a LECS, LES, BUS or ATMARP server and provides multiprotocol bridging and routing functions between emulated LANs and logical IP subnets.

IBM has implemented a number of proprietary functions in its MSS server product which fully support the ATM Forum-Compliant LAN Emulation (LANE 1.0) standard and extend it to provide additional functionality. This section summarizes some of the main extensions that can be used to improve or monitor network performance. All of these functions are optional to activate.

- Intelligent LES (ILES)

The ILES splits the control distribute VCC into two point-to-multipoint VCCs, a proxy control distribute VCC and a non-proxy control distribute VCC. These can be used to limit the address resolution polls, sent out by the LES, to only proxy LE clients. This reduces the level of traffic sent to non-proxy LE clients.

- Intelligent BUS (IBUS)

The IBUS splits the multicast forward VCC in to two point-to-multipoint VCCs, a proxy multicast forward VCC and a non-proxy multicast forward VCC. It then only sends unicast frames destined for a proxy LE client on the proxy multicast forward VCC and vice versa. This reduces the traffic sent to LE clients.

- Broadcast manager (BCM)

The BCM learns information from the BUS and converts broadcast frames to unicast frames when it has learned the frames' intended destination. It supports IP, IPX and NetBIOS. This significantly reduces the number of broadcasts sent needlessly to all workstations and hence reduces the level of traffic flowing to LE clients.

- Source route manager (SRM)

Token-ring endstations normally use source route bridging to find the most efficient route through a bridged network. SRM is an extension to the BCM function of the MSS server, and whenever possible it will convert all route explorer and spanning tree explorer frames to specifically routed frames.

- BUS monitor

The BUS monitor allows network administrators to monitor the usage of the BUS.

- Quality of Service (QoS) for LE clients (Version 1.1 only)

ATM Forum-Compliant LAN Emulation (LANE 1.0) does not define the use of ATM's QoS features for emulated LANs. It assumes the use of best-effort data direct VCCs between LE clients. This extension enables the use of these QoS parameters for data direct VCCs.

- SuperVLAN (Version 1.1 only)

SuperVLAN support provides a number of extensions to improve the performance of bridged traffic between emulated LANs.

- Next-Hop Resolution Protocol (NHRP) (Version 1.1 only)

The MSS server now supports the use of the NHRP protocol to improve performance of routed traffic.

For more information on these and other functions of the IBM Nways Multiprotocol Switched Services (MSS) server refer to *Understanding and Using the MSS Server* publication SG24-4915.

7.2.4 LECS and IBM MSS Server LES/BUS Redundancy

Fault tolerance is a major requirement in today's networks. The LANE 1.0 standard provides a mechanism to define a redundant LECS. A backup LECS may be configured and installed in your network and the network may be configured to provide this LECS address via ILMI. LE clients attempt to set up a configuration direct VCC to their LECS but if this fails, will re-request the next configured LECS address and attempt to set up a connection to it. The IBM Nways Multiprotocol Switched Services (MSS) server only supports a single LECS instance per MSS. Multiple LECS must be implemented using multiple MSS servers.

Note

When using backup LECS servers it is important to keep both the primary and backup servers synchronized. If not, clients may be assigned to different emulated LANs depending on the LECS server they use at initialization.

The LANE 1.0 standard does not however provide the ability to support a redundancy mechanism for the LES/BUS function. The IBM Nways Multiprotocol Switched Services (MSS) server implements a proprietary method for implementing a backup LES/BUS for any emulated LAN. When the LES/BUS is configured, the network administrator has the option to enable its redundant operation. If enabled, the LES/BUS must be defined as a primary or backup server for its particular emulated LAN. During initialization the primary LES/BUS establishes a redundancy VCC to the backup LES/BUS. The presence of this VCC indicates that the primary LES/BUS is operational. If the primary LES/BUS can't establish the redundancy VCC, for any reason, it retries every 5 seconds. The backup LES/BUS will become active if the primary fails. To ensure all clients are connected to a single LES/BUS, the backup LES/BUS will release all VCCs when the primary becomes active again. LE clients will only connect to the backup LES/BUS if they learn its address from the LECS. If the LES/BUS fails clients will be dropped from their emulated LAN and will attempt to rejoin by contacting the LECS. The LECS will return the LES/BUS address as follows:

- If the LES is in the same MSS as the LECS, and is configured as primary and is active, the LECS returns the local LES ATM address.

- If the LES is in the same MSS as the LECS, and is configured as backup but no redundancy VCC has been established, the LECS again returns the local LES address.
- If the LES is in the same MSS as the LECS, is configured as backup, and a redundancy VCC is established, the LECS returns the ATM address of the primary LES.
- If the LES is not in the same MSS as the LECS, the LECS maintains a short-memory of the LES ATM address it has last provided. If the LE client attempts to rejoin an emulated LAN within a timeout period (5 minutes) for the cached entry, the LECS will return the alternate LES address and update its memory.

7.2.5 Classical IP (RFC 1577) Rules

Classical IP defines the operation of IP and the Address Resolution Protocol (ARP) over ATM. Clients that communicate together via Classical IP form an IP only emulated LAN, known as a Logical IP Subnet (LIS). Clients in a LIS must all use the same IP network subnet number and address mask. Clients on different LISs can only communicate via a router. All clients must be directly connected to the ATM network. The connection between clients may use either PVC or SVC connections. An ATMARP server is only required when using SVC connections. The clients and ATMARP server (if required) use the ATMARP and InATMARP protocols to resolve IP addresses to ATM addresses and vice versa. ATMARP is the same protocol as the IP ARP protocol with extensions needed to support ARP in a unicast server ATM environment. InATMARP is the same protocol as the original IP InARP protocol (or RARP protocol) but is applied to ATM networks.

Classical IP (RFC 1577) encodes ATMARP and InATMARP packets in AAL-5 PDUs using LLC/SNAP encapsulation.

Note

Classical IP (RFC 1577) does not support IP broadcast or multicast traffic.

7.2.5.1 Classical IP with SVC Connections

Figure 217 on page 304 describes the use of Classical IP using SVC connections.

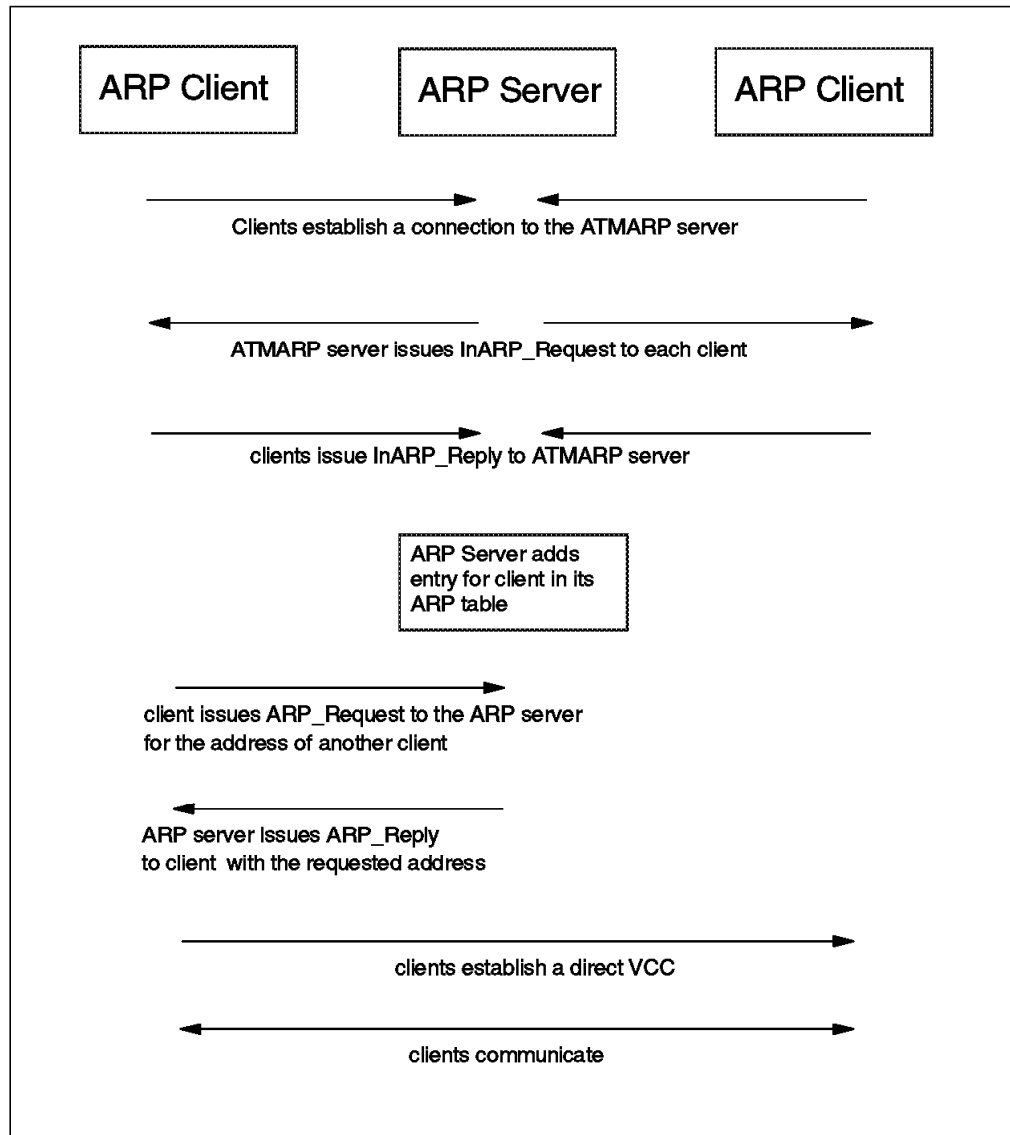


Figure 217. Classical IP (RFC 1577) Using SVC Connections

To support Classical IP using SVC connections, a single ATMARP server must be located within each LIS. This server is responsible for resolving the ATMARP requests of all the clients within the LIS.

Each client must be configured with the ATM address of the ATMARP server. They all establish a point-to-point VCC to the ATMARP server. Once established the ATMARP server issues an InARP_REQUEST frame to the client. The client replies with its address in an InARP_REPLY frame. The ATMARP server builds its ARP table based on these replies. Clients may then request the address of another client from the ATMARP server by sending an ARP_Request and receiving an ARP_Reply. The client then sets up a direct point-to-point VCC with the destination client. Frames are then sent over the direct VCC.

If a client requests an unknown address from the ATMARP servers the ARP server will send out an ARP_NAK frame as a reply. The client can therefore determine the difference between an ATMARP server failure and an unknown address. Clients must re-register their addresses with the ATMARP server every 20 minutes. Client ARP entries are valid for a maximum of 15 minutes and

ATMARP server ARP entries are valid for a maximum of 20 minutes. Prior to aging out an entry an ATMARP server sends an InARP_REQUEST to the client.

7.2.5.2 Classical IP with PVC Connections

To support Classical IP using PVC connections, there is no need for an ATMARP server. Every client must be connected to every other client. Each client must have a mechanism for determining what PVCs it has and which are being used with LLC/SNAP encapsulation. At initialization each client issues an InARP_REQUEST frame on all of its PVCs. The clients at the other end reply with their ATM address in an InARP_REPLY frame. Each client then updates its ARP table with all addresses. To send data a client will check its ARP table and send data via the corresponding PVC.

7.3 Problem Determination Guidelines

Figure 218 summarizes the process we used for solving problems that can occur with ATM LAN emulation and Classical IP (RFC 1577). We identified four stages in the diagnosis process.

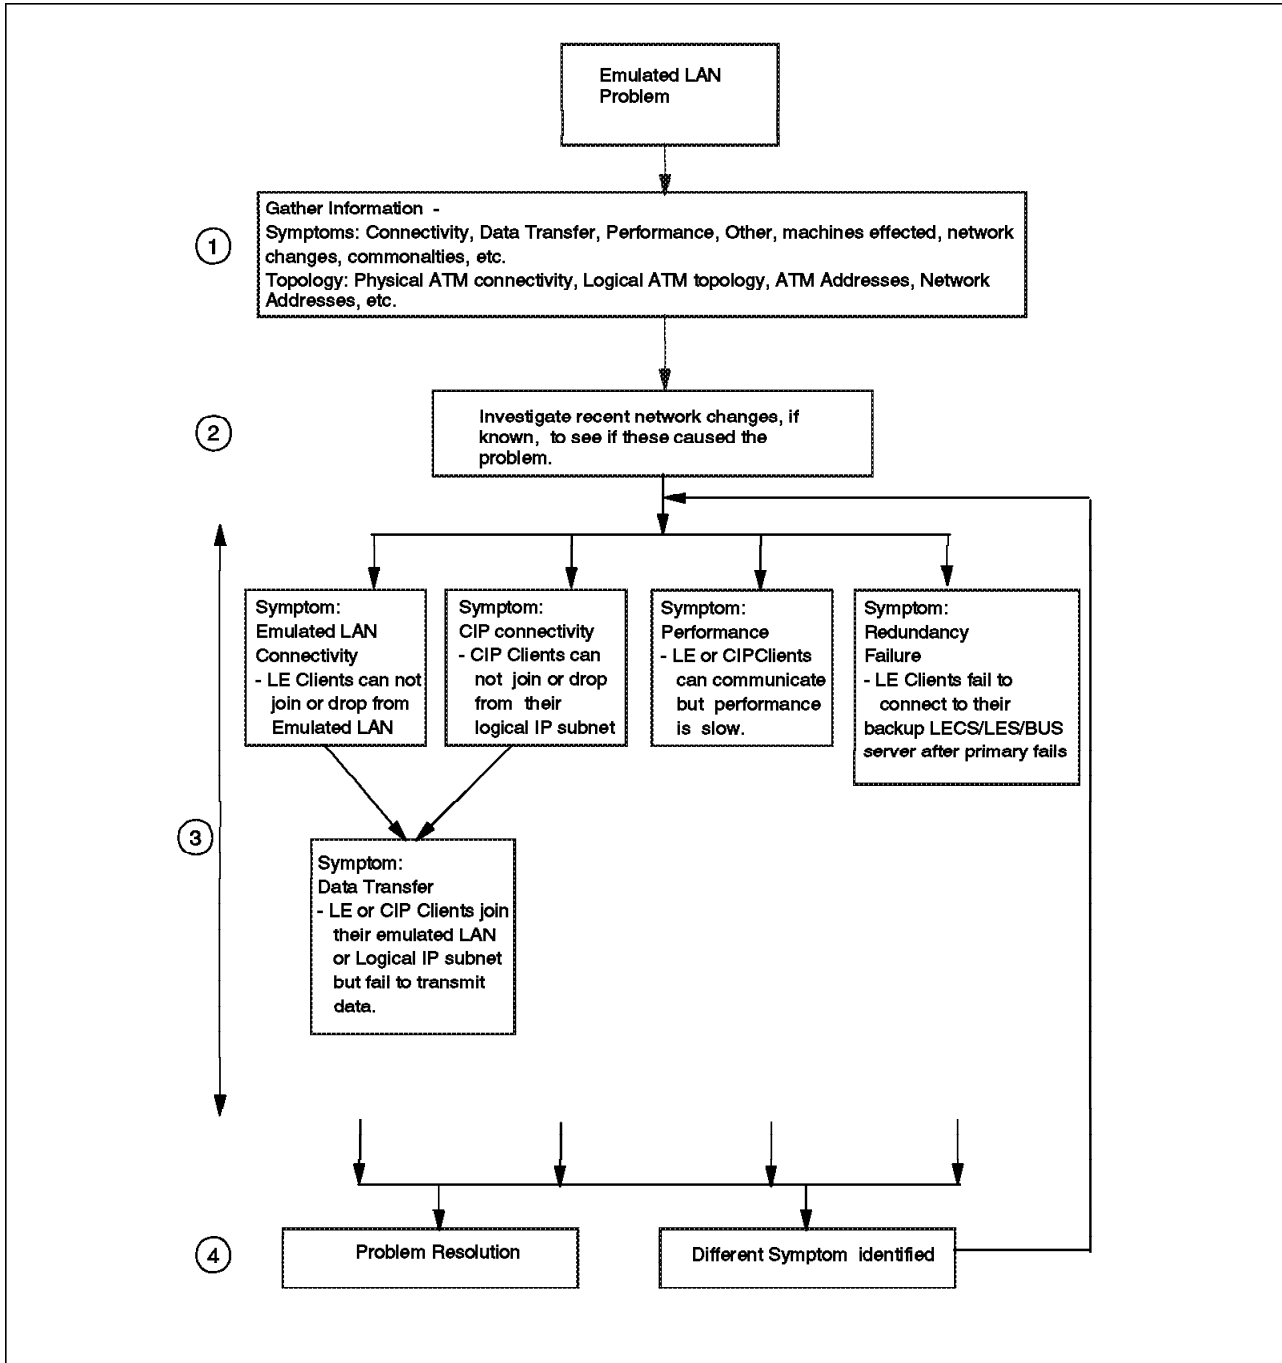


Figure 218. Problem Determination Methodology

The four stages can be described as follows:

1. Gather information.

As mentioned in Chapter 4, "Problem Determination Guidelines" on page 91, the key to effective problem determination is to gather as much information as possible regarding the problem.

Some of the specific questions you may need to ask in a LANE 1.0 or Classical IP environment are shown below. Users of the network are the best source of information. By talking to the users try and answer the following questions.

First, gather as much information as possible regarding the symptoms of the problem you are experiencing.

- What devices are affected by the problem? Are all of the devices attached to the same emulated LAN or logical IP subnet? Are they all affected or only a subset of them? Where are the devices connected to the ATM network? What are their ATM, MAC and network addresses (IP/SNA/NetBIOS/IPX)?
- What do the effected devices have in common? Are they on the same hub or connected to a certain portion of the network? Are they all using a common device, for example, a common LECS, LES, BUS or ATMARP server, or a common file, print or application server?
- What happens when the problem occurs? Do the effected devices always fail to get their required connectivity or do they fail only sometimes? Do they always connect to their servers but then at some point lose their connectivity? Can they connect but receive lower performance than usual?
- What has changed in the network since the problem started happening? Have devices been added or reconfigured?
- When does the problem occur? Does it only happen at a particular time of the day or when users are doing something specific, such as using a particular application or doing a large file transfer?

Secondly, gather as much information as possible regarding the physical and logical topology of the network. Some of the key questions to answer are:

- How is your ATM network physically connected? How are the hubs and switches interconnected? What other LAN segments/networks are connected to the ATM network?
- How is your network logically configured? What emulated LANs or logical IP subnets have been defined and how are these logically interconnected between each other and to other LAN segments/networks?
- What are the ATM, MAC and network addresses (IP/SNA/NetBIOS/IPX) of the key network components? These includes hubs, bridges, routers, switches, LECS, LES, BUS and ATMARP servers.
- Where are the key LANE 1.0 or Classical IP network devices physically connected to your network?

Good quality network documentation should provide the answers to these questions. If your documentation is not up-to-date or not maintained, a network management station configured for your networking environment should provide many of the answers.

2. Investigate recent network changes.

Most problems occur when you have changed something in your network. Effective change control and management can make sure one change does not impact the rest of your network. It also allows you to easily identify changes that may have introduced problems and allow you to back out of these changes quickly and hence resolve the problem. You can then examine why the changes caused the problem.

3. Interrogate the network based on the problem symptoms.

If the change causing the problem cannot be identified or is unknown, you must interrogate the network to discover the problem. The key information you need to discover depends on the problem you are experiencing. We classified the most common problems that can occur into five main types:

- Emulated LAN connectivity problems

Emulated LAN connectivity problems include all cases where a number of devices on an emulated LAN cannot communicate with each other because one or more of them are no longer members of the emulated LAN. This may be because they failed to join or have been dropped from it after a period of time.

- Classic IP connectivity problems

Classic IP connectivity problems include all cases where a number of devices on a logical IP subnet can not communicate with each other because one or more of them are no longer members of the logical IP subnet. This may be because they failed to join or have been dropped from it after a period of time.

- Data transfer problems

Data transfer problems include all cases where a number of devices on an emulated LAN or logical IP subnet cannot communicate with each other but all devices are members of the emulated LAN or logical IP subnet. This may be due to problems with higher layer protocols, data encapsulation, etc.

- Performance problems

These include problems where LE clients from an emulated LAN can communicate but performance is lower than expected.

- LE redundancy failures

These are special cases where redundancy features have been installed and configured but do not work correctly at a time of failure.

It is often relatively easy to discover whether the problem involves devices on an emulated LAN or logical IP subnet and whether the devices can or cannot communicate. It is more difficult to determine, from information gathered from users, whether or not all devices are connected to their emulated LANs or logical IP subnets or not. To determine this you often need to perform some simple tests from the devices experiencing the problem. Some examples of tests you can perform are:

- Ping another client.
- Log on to any LAN server.
- Log in to any Novell server.
- Connect to any Win95 or Win NT server.

- Start host emulation and connect to the host.
- Interrogate LES/BUS to determine clients connected to the emulated LAN.

Note: Only some LES/BUS/ATMARP products will support commands to determine what devices are linked to them. This function is supported by the IBM MSS server.

These tests should help to confirm whether the devices are attached to their emulated LANs or logical IP subnets and hence help you to determine the area in which the problem lies.

The methodology used to diagnose and fix problems, for the different types of problems listed above, may often vary. The different methodologies we used in this redbook are considered in the sections below.

4. Resolve problem or investigate a different symptom.

By following the problem methodologies described below, you will usually resolve the problem. In some cases you may identify that the problem is actually related to a different problem area, in which case you may need to repeat the entire problem determination process to investigate this new area to resolve the problem.

7.3.1 Emulated LAN Connectivity Problem Methodology

Numerous problems can occur when connecting to an emulated LAN. The problem diagnosis methodology used to discover problems is as follows:

- Stage 1 - Confirm that the essential emulated LAN devices do not have a hardware or power problem.
- Stage 2 - Confirm that the ATM network is working correctly and all necessary devices are connected.
- Stage 3 - Confirm that you have not made a simple mistake with the configuration of your LE environment and that the LE service servers are working correctly.
- Stage 4 - Interrogate the LE service servers, the VCCs established and the LE data flows.
- Stage 5 - Investigate emulated LAN timers and advanced parameters.

7.3.1.1 Stage 1 - Hardware and Power Problems

The essential devices in an emulated LAN are as follows:

- LE service components:
 - LAN Emulation Configuration Server (LECS)
 - LAN Emulation Server (LES)
 - Broadcast and unknown server (BUS)
- LAN Emulation Clients (LECs), for example user workstations or bridges/switches connecting ATM to legacy LAN environments (proxy LECs).

If you are physically close to the LE service components making up your emulated LAN, check the devices' front panel LEDs to ensure the devices do not have a hardware problem and are powered up correctly.

If possible, use internal commands to discover the status of the physical devices providing the LE services to your emulated LAN.

Refer to the troubleshooting section of your product-specific documentation to confirm the products are not experiencing hardware problems and that the LEDs confirm power to the devices.

7.3.1.2 Stage 2 - ATM Problems

Emulated LANs are dependent on the underlying ATM network to provide connectivity between all the devices attached to the emulated LAN. Numerous problems can occur causing the devices to lose this connectivity. These problems may involve the hubs and switches making up the ATM network or problems with the devices attached to the emulated LAN. The essential LE devices required for an emulated LAN are listed in 7.3.1.1, "Stage 1 - Hardware and Power Problems."

If a problem is restricted to machines connected in a particular region of your campus network, this may indicate a problem with the ATM network. Internal hub commands or network traces may be necessary to confirm your ATM network is working correctly. The status of your ATM ports and modules, inter-hub links and the ESI addresses registered with your ATM switches will provide a useful guide in verifying the correct working of your ATM network.

Refer to Chapter 5, “Starting Problem Isolation in an ATM Network” on page 109 for more information on how to make sure your ATM network is configured correctly, is working as designed, is providing the connectivity between all attached devices and that all the essential LE devices (listed in 7.3.1.1, “Stage 1 - Hardware and Power Problems” above) are connected to the network.

7.3.1.3 Stage 3 - Simple Configuration and LE Service Status Problems

The most common causes of all problems connecting to an emulated LAN, once all the emulated LAN devices are connected to the ATM network, are simple configuration mistakes and failures in the LE service components.

Many of the devices in the network require configuring to ensure LE clients connect to the correct emulated LANs. Many problems can be solved by checking a few parameters and the status of a few components and correcting any errors you find. It is often difficult to discover the source of the configuration error. This section is intended as a guide to identify the most important configuration parameters and components to check, to enable faster problem determination.

The following is a list of key parameters and components to verify on the different LE devices:

- In each ATM hub with LE clients attached

Check to see if all hubs, or all the hubs that have LE clients experiencing problems attached, have been configured correctly to enable each LE client, configured to use a LECS, to determine the address of its LECS. LE clients can determine the address of their LECS in a number of ways, as described in 7.2.1.1, “LE Service Operation” on page 297. To discover how to configure the hubs used in your network refer to the documentation that came with each hub.

- LECS

The LECS is a key component in the assignment of LE clients to their emulated LANs. Check its status to ensure it is working correctly. This can normally be done through the use of internal commands.

For each emulated LAN you define within your LECS, you must define the LES ATM address, the emulated LAN name, the LAN type (token-ring or Ethernet), the assignment policy priorities and a number of assignment policies which determine which LE clients will be assigned to which emulated LANs. The policies can cover:

- ATM address prefix
- ELAN name
- ELAN type
- MAC address
- Maximum frame size
- Route descriptor (source route bridging proxies only)

Double-check the LECS status, parameters, policies and defined priorities to ensure that LE clients will be assigned to the correct emulated LANs. For information regarding how to check these values refer to the documentation that came with your LECS.

- LES/BUS

The ATM Forum has defined six parameters for LES servers. These are listed in D.1, "ATM Forum LAN Emulation Server Parameters" on page 479. The most important ones to check to validate connectivity are:

- LAN server's ATM address
- LAN type
- Maximum frame size
- BUS ATM address

For each of the parameters above ensure that the values are consistent with the definitions in the LECS, BUS and LE clients. For information on how to check these parameters refer to the documentation that came with your LECS, LES, BUS and LE clients. It is also essential to check the status of the LES and BUS. LE clients will fail to communicate without both the LES and BUS being available.

- LEC and proxy LEC configuration

If LE clients are configured incorrectly and one or more of their parameters are out-of-bounds from the defined maximum and minimum values, the LE client will never attempt to connect to the LECS. This often results in a message on the client during boot up. The message usually mentions the parameter that is configured incorrectly.

The ATM Forum has defined 28 parameters for LE clients. These are listed in D.2, "ATM Forum LAN Emulation Client Parameters" on page 480. The most important ones to check to validate connectivity are dependent on the policies you have defined in your LECS. Depending on the policies you should check:

- LE client ATM address
- LAN type
- Maximum data frame size (which must match the size in the LES)
- ELAN name
- MAC address
- Route descriptors (if applicable)
- LES address (if used)

These should match the definitions in the LECS and LES. For information on how to check these parameters refer to the documentation that came with your LECS, LES and LE client.

7.3.1.4 Stage 4 - Interrogate LE Service, VCCs and Data Flows

In order to confirm whether an LE client or clients have connected successfully with the LE service servers you must monitor each stage of the LE initialization process while the problem is happening to discover which stage is causing the problem. The LE initialization process is summarized in 7.2.1.1, "LE Service Operation" on page 297.

Some problems are intermittent and hence it is very difficult to locate the problem unless it reoccurs. If a problem is intermittent, it is often related to broken or ill fitting cable connections. Schedule a time when the network can be down and check or replace the cables and ATM ports connecting the essential

LE service components to the ATM network. Problems with these devices will affect all devices on the emulated LAN. By being pro-active in changing them you may fix the problem and will not have to wait for the problem to reoccur.

To truly diagnose the problem, you must wait for the problem to reoccur. You will then need to monitor each LE service component, establish the clients it serves, monitor the ATM network to determine the VCCs that have been established and possibly examine the LE protocol data units (PDU) flowing from clients to each LE service component.

Some manufacturer's implementations of the key LE service components, the LECS, LES and BUS, will provide commands or logs, which allow a network troubleshooter to analyze statistics regarding each component's operation. These can be used to determine which clients are connected to each of the implemented services and can indicate why any failed to connect. This is usually relatively quick and easy to determine. The IBM MSS server offers these facilities. This will allow you to determine which component the LE clients are failing to connect to and hence at which stage of the LE initialization process the problem lies.

The VCCs established can also provide a useful guide to determine the source of a problem. To determine the VCCs established from and to LE clients in your network, you may be able to use a network management station (if available) or you may be able to interrogate your ATM switch. The list of VCCs will identify which LE clients have connected and may also identify problems in your ATM network, for example, whether the maximum allowable number of VCCs in your ATM switch has been exceeded. Refer to Chapter 5, "Starting Problem Isolation in an ATM Network" on page 109 for more information on discovering whether you have exceeded the maximum number of VCCs available in your network. It is essential to know the ATM addresses of the LE service components to determine whether the LE clients have established SVCs to these addresses.

The connection to the LECS will usually be dropped after configuration, so you may need to use a fast refresh rate on your network management station to see the VCC setup before it is disconnected. If a client uses a LECS and has a connection to its LES, then it has already communicated successfully with the LECS server.

If you have no other way to discover which LE service servers LE clients are connected to or what VCCs they have established, you may be able to use internal trace tools within your ATM switch or a network analyzer to examine the ATM SVC call setups as LE clients set up their VCCs to the LE service servers.

Trace the ATM network while you power on one of the LE clients. Analyze the trace and search for the ATM address of the LE client. You should be able to find all the SVC calls to the ATM switch as the LE client attempts to connect to its LECS, LES and BUS.

Finally, to determine what is happening when clients attempt to connect examine the LE data flows between a client and each LE service server. The status of the data flows is shown in each LE control frame. For a list of status codes refer to D.8, "Control Frame Status Values" on page 491.

It may be possible to examine the LE data flows using the facilities of your LE service servers or you may need to use a network analyzer. This will allow you

to discover exactly the point in the LE initialization process where the problem occurs.

Using traces can be very time consuming and difficult to set up. They should only be used as a last resort. For even a small network a trace can result in a very large amount of data. For a network of any size it can be virtually impossible to locate the specific SVC setups or LE PDUs you are interested in. The use of filters can make a significant difference.

If you use a network analyzer, connect it to the network between an LE client (or proxy LE client) that is not functioning and the ATM switch. This will enable you to only monitor traffic to and from one machine. Finding the problem with one machine will often allow you to fix the problem for all your clients.

7.3.1.5 Stage 5 - Investigate LE Timers and Advanced Parameters

In small emulated LAN environments most problems will be fixed by following the previous four stages of the methodology. The defaults used for the numerous other parameters do not usually cause problems. In large emulated LAN environments some of these other parameters will need to be tuned or your network may need to be redesigned to eliminate problems. If the parameters are not tuned or they are exceeded, some clients may fail to connect to their emulated LANs or may be dropped from their emulated LANs after a period of time.

The most important parameters to check in the various LE service components are as follows:

- LECS
 - Maximum number of configuration direct VCCs to a LECS (default 128)

This is the maximum number of simultaneous connections the LECS can support. If this is exceeded, the LECS will release all VCCs that have not been used for the number of seconds specified by the VCC idle time. Ensure this parameter is great enough to handle all the simultaneous requests to access the LECS. If you discover that connections to the LECS are failing, you may need to increase this parameter.
 - LECS VCC idle time (default 60 seconds)

The length of time a configuration direct VCC can remain idle before being released by the LECS when the maximum number of configuration direct VCCs is exceeded. You may also need to reduce this timer if you discover that connections to the LECS are failing.
 - Validate best effort PCR (default: No)

This parameter specifies whether to validate the Peak Cell Rate of Best Effort VCCs. If this parameter is enabled, the requested VCCs will be rejected if their peak cell rate is greater than the minimum data rate of the ATM interface. If this is enabled and configuration direct VCCs are being rejected, you may need to disable this value.
 - Maximum configuration direct VCC reserved bandwidth (default: 0)

The maximum reserved bandwidth the LECS will accept for a configuration direct VCC. The LECS will reject calls for a maximum reserved bandwidth higher than this value. If this is non-zero that is enabled and configuration direct VCCs are being rejected, you may need to increase this parameter.

- LES

- Control timeout (default: 120 seconds)

The period used to timeout request/response control frame transactions. The LE client must be sent an LE_JOIN_RESPONSE frame before the control timer expires, after sending an LE_JOIN_REQUEST frame. If the LES does not process the REQUEST frame in time, the join fails. If you experience join failures, you may need to increase this timer.

- Validate best effort PCR (default: No)

This parameter specifies whether to validate PCR for best effort VCCs to the LES. Connections will be rejected if the PCR is greater than this maximum. If this is enabled and control direct VCCs are being rejected, you may need to disable this parameter.

- Maximum control direct VCC reserved bandwidth (default: 0)

The maximum reserved bandwidth that will be expected for control direct VCCs. If this is non-zero, that is enabled, and control direct VCCs are being rejected, you may need to increase this parameter.

- BUS

- Validate best effort PCR (default: No)

This parameter specifies whether to validate PCR for best effort VCCs to the BUS. Connections will be rejected if the PCR is greater than this maximum. If this is enabled and multicast send VCCs are being rejected you may need to disable this parameter.

- Maximum multicast send VCC reserved bandwidth

The maximum reserved bandwidth that will be excepted for multicast send VCCs. If this is non-zero that is enabled and multicast send VCCs are being rejected, you may need to increase this parameter.

Check your LECS, LES and BUS statistics to determine whether the defaults for the parameters above need to be changed.

7.3.1.6 LE Connectivity Problems and Recommended Actions

You may discover the problem within any of the stages described above. Table 19 on page 316 describes many of the common problems and their solutions, based on the information you gather above.

Table 19 (Page 1 of 4). Diagnosing Problems Connecting to Emulated LANs

Symptom	Probable Cause	Recommended Actions
Abnormal LEDs on essential emulated LAN or ATM device	Hardware problem	Refer to the product-specific documentation available for the device.
No LED lights on essential emulated LAN or ATM device	No power to device	Check power cable and source and refer to the product-specific documentation available for the device.
Only LE clients in one region of your network are affected	ATM connectivity problem	Refer to Chapter 5, "Starting Problem Isolation in an ATM Network" on page 109 to diagnose ATM network problems.
LE service server (LECS/LES/BUS) not connected to ATM network	ATM connectivity problem	Refer to Chapter 5, "Starting Problem Isolation in an ATM Network" on page 109 to diagnose ATM network problems.
Clients fail to connect to one of their LE service servers and the number of VCCs in your network exceeds the maximum for your switch	ATM maximum capacity exceeded	Refer to Chapter 5, "Starting Problem Isolation in an ATM Network" on page 109 to diagnose whether you have exceeded the maximum number of VCCs allowed in your network.

<i>Table 19 (Page 2 of 4). Diagnosing Problems Connecting to Emulated LANs</i>		
Symptom	Probable Cause	Recommended Actions
Client(s) fail to set up SVC to LECS. (LECS is connected to ATM network.)	Adapter driver problem	Check to see if you are using the latest adapter drivers. Check with the network adapter card's manufacturer to discover if you are using the latest adapter drivers and whether they have any known problems.
	LECS address not discovered or unreachable	Check the hub to which an LE client is attached to determine if it has been set up correctly, to enable the LE client to discover the address of its LECS (see 7.2.1.1, "LE Service Operation" on page 297).
	Client does not support the use of LECS	Refer to the documentation available with your client driver to find out if the client does or does not support the use of a LECS. Some clients must be hard coded with the address of their LES and do not support the use of a LECS.
	LE client configuration	If you see a message mentioning that one or more of an LE client's parameters is out-of-bounds during boot up of the client, check the client's configuration to ensure all parameters are correct. See 7.3.1.3, "Stage 3 - Simple Configuration and LE Service Status Problems" on page 311 for more information.
	LECS failure or driver problem	<ol style="list-style-type: none"> 1. Restart the LECS and check whether it is working correctly. Refer to the documentation that came with your LECS to perform self tests on the LECS to check it is working correctly. 2. Check the microcode used in your LECS to determine if it is at the latest level. Contact your LECS manufacturer to check whether they know of any problems in their LECS microcode. 3. Finally trace the LE PDUs between the LE client and LECS to discover whether the client or the LECS is not working correctly. Report the problem to the manufacturer of your LECS or network adapter.
	VCCs to the LECS are being rejected	<ol style="list-style-type: none"> 1. Check whether the maximum number of VCCs has been exceeded and possibly increase the maximum value and/or decrease the VCC idle time used in your LECS or redesign your network to reduce the number of clients that simultaneously attempt to connect to the LECS. 2. Check to determine whether the LECS validates maximum PCR value for VCCs and if so disable this to see if it is causing problems. 3. Check to see if the LECS has a maximum reserved bandwidth for VCCs and if so disable this maximum to see if it is causing problems.

Table 19 (Page 3 of 4). Diagnosing Problems Connecting to Emulated LANs

Symptom	Probable Cause	Recommended Actions
Clients connect to LECS but fail to connect successfully to LES. (LES is connected to the ATM network.)	Adapter driver problem	As stated in previous section.
	LECS and/or LE client configuration	<p>The LECS may be rejecting a client(s) based on the LE client's configuration and/or the configuration of the LECS policies. It may also be assigning clients to a LES which does not exist.</p> <p>Check the LECS policies and the LE client's definitions to determine if both have been defined correctly. Refer to 7.3.1.3, "Stage 3 - Simple Configuration and LE Service Status Problems" on page 311 to understand the most common parameters to check.</p>
	LES and/or LE clients configuration	<p>The join to the LES has failed due to the client being configured incorrectly. Make sure the configuration of your LES and LE client is consistent.</p> <p>The join to the LES may have been rejected due to a duplicate MAC or ATM address definition on the LE client. It may also have been rejected due to differences between the LES and LE client's LAN type or maximum frame size.</p> <p>Refer to 7.3.1.3, "Stage 3 - Simple Configuration and LE Service Status Problems" on page 311 to understand the most common parameters to check.</p>
	LES failure	<ol style="list-style-type: none"> 1. Restart the LES and check whether it is working correctly. Refer to the documentation that came with your LES to perform self-tests, to check it is working correctly. 2. Check the microcode used in your LES to determine if it is using the latest drivers. Contact your network adapter manufacturer to check whether they know of any problems in their LES microcode. 3. Finally trace the LE PDUs between the LE client and LES to discover whether the client or the LES is not working correctly. Report the problem to the manufacturer of your LES or network adapter.
	VCCs to the LES are being rejected	<ol style="list-style-type: none"> 1. If joins to the LES are timing out, increase the control timer value to see if this fixes the problem. 2. Check to determine whether the LES validates maximum PCR value for VCCs and if so, disable this to see if it is causing problems. 3. Check to see if the LES has a maximum reserved bandwidth for VCCs and if so, disable this maximum to see if it is causing problems.

<i>Table 19 (Page 4 of 4). Diagnosing Problems Connecting to Emulated LANs</i>		
Symptom	Probable Cause	Recommended Actions
Clients connect to LES but fail to connect to BUS. (BUS is connected to the ATM network.)	Adapter driver problem	As stated in previous section.
	LES not providing the correct BUS address	<ol style="list-style-type: none"> 1. The LES may be incorrectly configured with the address of the BUS. Refer to the documentation that came with your LES to make sure the BUS address has been defined correctly. 2. Trace the LE PDUs between a client and the LES to determine if the LES is not providing the address of the BUS or whether it is providing an incorrect address. Report the problem to the manufacturer of your LES.
	BUS failure	<ol style="list-style-type: none"> 1. Restart the BUS and check whether it is working correctly. Refer to the documentation that came with your BUS to perform self-tests on the BUS to ensure it is working correctly. 2. Check the microcode used on the BUS to determine if it is using the latest drivers. Contact your network adapter manufacturer to check whether they know of any problems in using a BUS. 3. Finally, trace the LE PDUs between the LE client and BUS to discover whether the client or the BUS is not working correctly. Report the problem to the manufacturer of your BUS or network adapter.
	VCCs to the BUS are being rejected	<ol style="list-style-type: none"> 1. Check to determine whether the BUS validates maximum PCR value for VCCs and if so disable this to see if it is causing problems. 2. Check to see if the BUS has a maximum reserved bandwidth for VCCs and if so disable this maximum to see if it is causing problems.
<p>Note:</p> <p>You can access information and the latest network drivers for IBM networking equipment and adapters from IBM's networking home page :http://www.networking.ibm.com/.</p>		

Many problems can be avoided by utilizing a redundant LECS in your network and/or a redundant LES/BUS for each emulated LAN. IBMs MSS server supports configuration as a redundant LECS and/or LES/BUS.

7.3.2 Classic IP Connectivity Problem Methodology

To diagnose problems that can occur when connecting to a logical IP subnet it is essential to investigate each stage of the Classical IP initialization process to determine where the problem occurs and hence locate the source of the problem. We used a similar approach to that described in 7.3.1, “Emulated LAN Connectivity Problem Methodology” on page 310.

- Stage 1 - Confirm that the essential Classical IP devices do not have a hardware or power problem.
- Stage 2 - Confirm that the ATM network is working correctly and all necessary devices are connected.
- Stage 3 - Confirm that you have not made a simple mistake with the configuration of your Classical IP environment and that the ATMARP server is working correctly.
- Stage 4 - Interrogate the ATMARP server (if used), the VCCs established and the ATMARP data flows.

7.3.2.1 Stage 1 - Hardware and Power Problems

The essential devices in a logical IP subnet are as follows:

- ATMARP clients, for example user workstations and servers
- ATMARP server (SVC environments only)

Check that all your components are powered on and if any have LEDs, check the LEDs to ensure they are working correctly. If possible, use internal commands to discover the status of the device providing the ATMARP server service to your logical IP subnet.

Refer to the troubleshooting section of your product-specific documentation to confirm products are not experiencing hardware problems and that the LEDs confirm power to the device providing the ATMARP server function.

7.3.2.2 Stage 2 - ATM Problems

Logical IP Subnets, such as emulated LANs, are dependent on the underlying ATM network to provide connectivity between all the devices attached. Internal hub commands or network traces may be necessary to confirm your ATM network is working correctly. The status of your ATM ports and modules, inter-hub links and the ESI addresses registered with your ATM switches will provide a useful guide in verifying the correct working of your ATM network.

In Classical IP environments using PVCs, PVCs must be defined between every client. Refer to your ATM switch documentation to ensure the PVCs have been defined and are working successfully. You can often check the status of these links by interrogating your ATM switch or by using a network management station.

Refer to Chapter 5, “Starting Problem Isolation in an ATM Network” on page 109 to make sure your ATM network is configured correctly, is working as designed, is providing the connectivity between all attached devices and that all the essential devices (listed in 7.3.2.1, “Stage 1 - Hardware and Power Problems” above) are connected to the network.

7.3.2.3 Stage 3 - Simple Configuration and ATMARP Server Status Problems

All the devices on a logical IP subnet must be configured for it to work correctly. This section acts as a guide to the main items to check, for both ATMARP servers and clients.

Searching for problems, check:

- IP address
All clients, including the ATMARP server, use a unique address within the range of IP addresses associated with the logical IP subnet.
- Subnet mask
All clients, including the ATMARP server, within a logical IP subnet use the same subnet mask.
- The status of the ATMARP server (if used)
In SVC environments the ATMARP server is essential for clients to communicate.
- ATM address of the ATMARP server specified for each client (if used)
Make sure clients have used the correct ATM address for their ATMARP server and it is in the correct format. IBM RISC System/6000s require the 20-byte ATM address to be input with each byte separated by periods.
- UNI version used by the clients and ATMARP server's interface
Make sure the UNI type specified in the client matches the one defined in the ATM switch.
- Service data unit (SDU) size
All clients, including the ATMARP server, within a logical IP subnet must use the same SDU size. We recommend using the default value (9188 bytes). The default maximum transmission unit value is 9180 bytes, the SDU becomes 9188 with the additional 8-byte LLC/SNAP header.
- ATMARP server's locally administered ATM address (if used)
ATMARP servers may be configured to use a locally administered address for the ATM ESI part of their address. If used, make sure this address is unique on your ATM hub.

7.3.2.4 Stage 4 - Interrogate ATMARP Server, VCCs and Data Flows

If possible, interrogate your ATMARP server to analyze statistics regarding its operation. This will only be possible if your ATMARP server supports internal commands, logs or MIB variables that allow you to query its operation.

In Classical IP environments using PVCs the connections between all devices are specifically defined in your ATM switch. Refer to 7.3.2.2, "Stage 2 - ATM Problems" on page 320 to ensure these connections are active.

In Classical IP environments using SVCs the ATM address of the ATMARP server is hardcoded in each client. Interrogate your hub or use a network management station to verify that all clients have established a direct VCC to the ATMARP server.

If you have no other way to discover which VCCs have been established, you may be able to use internal trace tools within your ATM switch or a network analyzer to examine the ATM SVC call setups as ATMARP clients set up their VCCs.

Trace the ATM network while you power on one of the ATMARP clients. Analyze the trace and search for the ATM address of the ATMARP client. You should be able to find all the SVC calls to the ATM switch.

To determine whether the InARP_REQUEST/REPLY and ARP_REQUEST/REPLY processes are working correctly you may be able to use the internal features of your ATMARP server. More commonly you will have to trace the data flowing from one of the clients to determine whether these process flows are working correctly.

Only trace the network as a last resort. Using traces can be difficult to set up and very time consuming.

7.3.2.5 Classical IP Connectivity Problems and Recommended Actions

You may discover the problem within any of the stages described above. The table below describes many of the common problems and their solutions, based on the information you gather above.

Symptom	Probable Cause	Recommended Actions
Abnormal LEDs on essential Classical IP or ATM device	Hardware problem	Refer to the product-specific documentation available for the device.
No LEDs lights on essential Classical IP or ATM device	No power to device	Check power cable and source and refer to the product-specific documentation available for the device.
Only Classical IP clients in one region of your network are affected	ATM connectivity problem	Refer to Chapter 5, "Starting Problem Isolation in an ATM Network" on page 109 to diagnose ATM network problems.
Clients fail to connect to their ATMARP server and the server's ATM address is not registered with the ATM switch	ATM connectivity problem (UNI version mismatch or other problem)	Refer to Chapter 5, "Starting Problem Isolation in an ATM Network" on page 109 to diagnose ATM network problems.
Clients fail to connect to their ATMARP server and the number of VCCs in your network exceeds the maximum for your switch	ATM maximum capacity exceeded	Refer to Chapter 5, "Starting Problem Isolation in an ATM Network" on page 109 to diagnose whether you have exceeded the maximum number of VCCs allowed in your network.

<i>Table 20 (Page 2 of 2). Diagnosing Problems Concerning Connecting to a Logical IP Subnet</i>		
Symptom	Probable Cause	Recommended Actions
Clients fail to connect to their ATMARP server but the maximum number of VCCs in your network has not been exceeded	ATMARP client configuration	Ensure ATMARP server ATM address is specified correctly, in the correct format.
	Adapter driver problem/microcode	Check to see if you are using the latest adapter drivers or microcode on your ATMARP clients and server. Check with your network adapter and ATMARP server manufacturer to discover if you are using the latest adapter drivers and whether they have any known problems.
	ATMARP server failure	<ol style="list-style-type: none"> 1. Restart your ATMARP server and check whether it is working correctly. Refer to documentation that came with your ATMARP server to perform self-tests to determine if it is working correctly. 2. Trace the data flows from a client to the ATMARP server to determine whether the client or server is not functioning correctly. Report the problem to the manufacturer of your ATMARP server or client's network adapter.

7.3.3 Data Transfer Problem Methodology

Once your LE clients or Classical IP clients are connected to their emulated LAN or logical IP subnet they may still experience numerous problems in communicating with other clients.

In LAN emulation, the LE clients may fail to establish each other's ATM address using the LE_ARP_REQUEST and LE_ARP_RESPONSE frame exchange with the LES. They may fail to establish a VCC between each other or they may fail to switch over their data path from using the BUS to the data direct VCC (using the FLUSH protocol).

In Classical IP, the ATMARP clients may fail to establish each other's ATM address via the ARP_REQUEST and ARP_RESPONSE frame exchange with the ATMARP server. They may also fail to establish a VCC.

Once clients have established a data direct VCC, problems may still occur with the MAC layer encapsulation of the higher layer LAN protocols or with the higher layer protocols themselves.

- Stage 1 - Check LE_ARP or ATMARP cache

If possible, check each client's LE_ARP or ATMARP cache to determine whether they have cached the ATM addresses of the clients with which they wish to communicate. This will establish whether the address resolution has completed successfully. If the problem affects a number of clients, investigate one pair of clients that are attempting to communicate to try and find the problem affecting all clients.

- Stage 2 - Check VCCs established

You should also check the VCCs established to and from the clients attempting to communicate. This should establish if the data direct VCC has been established. To check the VCCs established you may use a network management station (if available) or you may be able to discover them by interrogating your ATM switch. This may also identify problems in your ATM network, for example whether the maximum allowable number of VCCs by your ATM switch has been exceeded. Refer to Chapter 5, "Starting Problem Isolation in an ATM Network" on page 109 for more information on discovering whether you have exceeded the maximum number of VCCs available to your network.

- Stage 3 - Trace the LAN emulation data flows between clients

If no data direct VCC has been established or the LE_ARP or ATMARP cache does not include an entry for the partner client, examine the LE or ATMARP PDUs flowing between each of the clients and the LES/BUS or ATMARP server. This will determine exactly what is happening when clients attempt to resolve each other's addresses. It may be possible to examine these using the features of your LE service or ATMARP servers or more probably you may need to use a network analyzer.

If a data direct VCC has been established and the LE_ARP or ATMARP cache on both clients trying to communicate contains each other's ATM addresses, trace the PDUs flowing between the two clients. This will allow you to discover whether a problem exists with the higher layer communication protocols or the data encapsulation used. The diagnosis of these problems is outside the scope of this redbook. This will require the use of a network analyzer.

The network analyzer should be connected between one of the clients and its ATM port. Filters should be applied to discard all data except for the frames you are interested in.

7.3.3.1 Data Transfer Problems and Recommended Actions

The table below describes many of the common problems and their solutions, based on the information you gather above.

<i>Table 21. Diagnosing Problems Concerning Data Transfer</i>		
Symptom	Probable Cause	Recommended Actions
LE_ARP process fails on both LE clients or LE_FLUSH process fails on both clients	One of the LE clients is not a member of the emulated LAN	Refer to 7.3.1, "Emulated LAN Connectivity Problem Methodology" on page 310 to diagnose the problem.
	Adapter driver problem	Check to see if you are using the latest adapter drivers. Check with the network adapter card's manufacturer to discover if you are using the latest adapter drivers and whether they have any known problems.
	LES failure	<ol style="list-style-type: none"> 1. Check the microcode used in your LES to determine if it is using the latest drivers. Contact your network adapter manufacturer to check whether they know of any problems in their LES microcode. 2. Report the problem to the manufacturer of your LES.
ARP_REQUEST, ARP_RESPONSE process fails on one or both ATMARP clients	One of the ATMARP clients is not a member of the emulated LAN	Refer to 7.3.2, "Classic IP Connectivity Problem Methodology" on page 320 to diagnose the problem.
	Adapter driver problem	Check to see if you are using the latest adapter drivers. Check with the network adapter card's manufacturer to discover if you are using the latest adapter drivers and whether they have any known problems.
	ATMARP failure	<ol style="list-style-type: none"> 1. Check the microcode used in your ATMARP server to determine if it is using the latest drivers. Contact your network adapter manufacturer to check whether they know of any problems in their microcode. 2. Report the problem to the manufacturer of your ATMARP server.
Clients fail to set up data direct VCC and the number of VCCs in your network exceeds the maximum for your switch	ATM maximum capacity exceeded	Refer to Chapter 5, "Starting Problem Isolation in an ATM Network" on page 109 to diagnose whether you have exceeded the maximum number of VCCs allowed in your network.
Clients establish their data direct VCC but fail to communicate	MAC layer encapsulation or higher layer protocol problem	Use a network analyzer to decode the data flows from one of your failing clients. Refer to specific documentation on the protocol used to diagnose higher layer problems. This is outside the scope of this redbook.

7.3.4 Network Performance Problem Methodology

Many factors can affect the performance of your network. Effective network design can elevate the impact clients can have on the performance of other client's applications, by minimizing the resources shared between clients. This is discussed further in 7.11, "Case Studies Involving Network Performance Problems" on page 406.

The state of your LANE 1.0 and Classical IP environment, discussed in 7.1, "What Makes a Healthy LANE 1.0 or Classical IP Environment" on page 293, is determined by the condition of your ATM network. Two factors that can affect this are the address resolution process provided in LANE 1.0 and Classical IP, and the broadcast functions provided in LANE 1.0.

In order to diagnose problems with these two functions you need to monitor the LES, ATMARP server and BUS to ensure clients are not over-utilizing these functions and impacting the performance of other clients. Another useful guide to the excessive use of the BUS is to monitor the data frames it discards. If the BUS is over-utilized, it will not be able to transmit data frames before its maximum frame age timer. This will cause frames to be discarded.

Some implementations of these servers, such as IBM Nways Multiprotocol Switched Services (MSS) server, provide commands and features to monitor the operation of the LES and BUS. When these are not available you must trace the network and analyze the traffic flowing to your LES, BUS or ATMARP server. To do this you must use a network analyzer attached to the network between your LES, BUS or ATMARP server and its ATM switch.

If you discover a problem, it may be necessary to reduce the level of address resolution and broadcast traffic by splitting the emulated LAN and possibly using a second physical LES, BUS or ATMARP server.

7.3.5 LE Redundancy Failure Problem Methodology

The IBM Nways Multiprotocol Switched Services (MSS) server redundancy features provide essential extensions to the LANE 1.0 standard for implementing fault tolerance in your emulated LAN network. The redundancy feature must be configured correctly and is reliant on the redundancy VCC being established.

To solve problems with the redundancy feature we used a process similar to that described in 7.3.1, “Emulated LAN Connectivity Problem Methodology” on page 310.

Check in sequence:

- Stage 1 - Hardware and power problems
- Stage 2 - ATM problems
- Stage 3 - Simple configuration and status problems
- Stage 4 - Interrogate LECS, LES, BUS and VCCs

These stages are considered in more detail below.

7.3.5.1 Stage 1 - Hardware and Power Problems

The essential services used for redundancy are as follows:

- The LECS and backup LECS
- The LES/BUS and backup LES/BUS

Check the LEDs of both MSS servers used to provide these functions to ensure they are working correctly and are powered on. You should also use the MSS internal commands to confirm they are working correctly and their ATM interface is working.

Refer to the troubleshooting section of your MSS product guide to confirm the products are not experiencing hardware problems and that the LEDs confirm power to the device.

7.3.5.2 Stage 2 - ATM Problems

Confirm that the ATM network is working correctly and that the MSS servers used are connected to the ATM network and registered with their ATM switches. The status of the ATM ports, modules, inter-hub links and the ESIs registered with the ATM switches will again provide a useful guide to confirming that the ATM network is working correctly. Refer to 7.3.1.2, “Stage 2 - ATM Problems” on page 310 for more information.

7.3.5.3 Stage 3 - Simple Configuration Problems

This section considers the important configuration parameters and features to check when experiencing problems with LECS and/or LES/BUS redundancy.

- LECS redundancy

Check that your hubs are set up correctly to return both the primary LECS and the backup LECS address. Confirm that the status of both your primary and backup LECS are working correctly. Check the LECS parameters discussed in 7.3.1.3, “Stage 3 - Simple Configuration and LE Service Status Problems” on page 311 for each LECS, to ensure they are set up correctly and set up the same on both.

- LES/BUS redundancy

Check the configuration of your LECS, primary and backup LES/BUS to ensure the ATM addresses of the primary and backup LES/BUS are defined correctly.

7.3.5.4 Stage 4 - Interrogate LES/BUS and VCCs

Most problems with LECS redundancy will be due to the configuration of the hubs or LECS. For LES/BUS redundancy you need to interrogate the MSS servers used to determine whether the redundancy VCC is operational and if it is not determine why and resolve the problem. The statistics for the LES/BUS will show you the number of times the primary LES has tried to establish the redundancy VCC and failed to do so. This will give a good indication of whether the redundancy VCC is working or not.

To check the status of the redundancy VCC you can use the internal commands of the MSS to list the status of both the primary and backup LES/BUS. This will show the status of the redundancy VCC. Another way to check its status is to use a network management machine to determine whether a VCC exists between the MSS with the primary LES/BUS and the MSS containing the backup LES/BUS. If no redundancy VCC exists there could be a problem in your ATM network. Refer to Chapter 5, "Starting Problem Isolation in an ATM Network" on page 109 to establish whether a problem exists in your ATM network.

Finally you can use the MSS to display the LES LE PDUs. These will show you any events involving the establishment or failure of the redundant VCC.

7.4 Common Problems Specific to IBM-Compliant LAN Emulation

This section lists some of the most common problems that can occur with IBM-Compliant LAN Emulation.

<i>Table 22 (Page 1 of 2). Most Common Problems with LAN Emulation</i>		
Symptom	Probable Cause	Recommended Actions
LES Monitor Statistics: default VCCs counter oscillating, too few registered workstations.	The workstation knows its ATM address, but that address has been de-registered at the switch/control point level. This happens when the workstation is behind a concentrator (8282) that has been disconnected from the switch for a little while.	Wait for one or two minutes for the new registration to take place.
8260/8285 Clear table: a lot of SVCs were cleared with the cause 31 or cause 16.	A high bandwidth (100 Mbps or 155 Mbps) workstation or bridge has tried to call a low bandwidth (25 Mbps) workstation. The call was rejected by the low bandwidth workstation because the bandwidth specified in the Q2931 parameters (even for a UBR call) was too large. This is normal.	The source station or bridge will retry the destination station with a lower bandwidth bit rate, successfully. Nothing to do.
Some ATM stations cannot talk to LAN stations behind parallel bridges. The 8281 bridge has a limitation of 256 ATM connections. One could think that multiplying the number of 8281 bridges (in parallel) would multiply the number of available connections. But doing so will lead to the following problem: only around 256 stations can IMMEDIATELY establish connections with the bridges.	In a configuration with parallel 8281 bridges (those bridges register to the same LAN emulation server, and they connect to the same LAN), there may be collisions in terms of connections. Indeed each 8281 bridge will respond by establishing a connection to the originating ATM station. In a network where the number of ATM stations exceeds 256, which is the maximum number of SVCs per 8281, some stations will not be able to connect, until the bridges clear the SVCs that are unused (aging out process).	Wait 4 minutes (default aging time on the 8281 bridge) or avoid the parallel bridging.
LES monitor: after 3 minutes, the workstation is de-registered from the LES.	The workstation did not send the re-registration message during the 3-minute interval.	Ensure that the port for the workstation on the 8282 is green on ATMC. If not green, ensure that the cable between the 8260 and the 8282 is connected properly. Shut down, then power off the workstation and restart. If the problem persists, contact your workstation adapter retailer.
In a multi token-ring bridged configuration, a token-ring bridge cannot register to the LES.	Different ring numbers are assigned to the ATM ports of two bridges connected to the same LES.	Check the ring numbers of the ATM ports of all the bridges attached to the same LES; these numbers should not be equal. Change them if necessary.

<i>Table 22 (Page 2 of 2). Most Common Problems with LAN Emulation</i>		
Symptom	Probable Cause	Recommended Actions
LES monitor: bridge is on the general multicast tree, but not on the bridge multicast tree	The bridge did not send its route descriptors to the LES.	The bridge is faulty. Contact your IBM representative.
At workstation reboot: the ATM adapter initialization failed.	The switch (8260/8285) or concentrator (8282) port attached to the workstation is not enabled, or is not a UNI port.	From the terminal, or from the SNMP manager (ATMC), enable the corresponding port as a UNI port.
A workstation/bridge cannot connect to another workstation/bridge.	One of the workstations/bridges is not registered to the LES.	Using the LES monitor, check in the list of registered endstations that both workstations/bridges are there. If both addresses are registered then consider the next possible cause. If one workstation/bridge address is missing, use the call status history provided by the LES monitor to get the Q2931 cause of the failing call. The missing station/bridge probably has a wrong LES ATM address defined in its configuration. Check the missing station's configuration.
	Both workstations/bridges are registered to the LES, but one cannot call the other one, because the LES is not available any more (port disabled or NOT-IN SERVICE). Yet, the LES does not tell you that it lost its ATM address, because it only tells that after it gets the connection with the 8260/8285 back.	Ensure the LES cable is well plugged. Check that the LES port is enabled. If it stays enabled, and NOT-IN-SERVICE, then the LES is faulty. Contact your IBM representative for investigation, or re-boot the LES.
A station cannot register to a LES located behind a WAN (VP-tunnel). Some connections through the VP-tunnel work, but some do not, especially the ADD_PARTY to put the stations on the LES Multicast Tree. The 8260/8285 error-log is full of messages like invalid message length.	The WAN (public network providing the VP-tunnel) uses VCI=5 for its own purposes, and there is a conflict with the 8260/8285 which also uses VCI=5.	Ask your public network provider if they use VCI=5. If necessary, put ATM equipment between the WAN and the 8260/8285 to do the translation of signalling VCI to a value other than 5.
Note: These problems only relate to IBM-compliant LAN emulation and NOT to ATM Forum-compliant emulated LANs.		

7.5 A Guide for Using Commands with the IBM Nways Multiprotocol Switched Service (MSS) Server

In many cases when troubleshooting ATM networks containing an IBM Nways Multiprotocol Switched Services (MSS) server you need to issue commands at the MSS server command line interface. To get access to the command line you may directly attach a TTY console to the management port of the device or telnet to any of its internal IP addresses. The following are the default settings for the serial port:

- Speed 19.2 kbps
- Parity None
- DataBits 8
- StopBits 1

The MSS server command line interface is structured as a hierarchy of prompts. To use a specific command you must enter it at the correct prompt. This section acts as a guide to finding the correct prompt to enter commands.

The three key processes that are of importance when running commands are:

- talk 6 or t 6 which allows entry to the structure of commands to configure and query the saved version of the MSS server configuration
- talk 5 or t 5 which allows entry to the structure of commands to configure and query the active running version of the MSS server configuration
- talk 2 or t 2 which allows you to monitor events occurring on the MSS server.

To access these three processes enter:

t 6

t 5

or

t 2

at the OPCODE (*) prompt after you have logged on to your MSS server.

Note

For troubleshooting we recommend you usually use the talk 5 process, which is the monitoring process (GWCON), when checking the status of any of the IBM Nways Multiprotocol Switched Services (MSS) server features. This will show you the currently active status rather than the configured status of the MSS server. Configuration problems should be rectified using the talk 5 and talk 6 processes to ensure the problem does not reoccur on the next reIML of the MSS server

The menu structure of commands in the talk 5 and talk 6 processes is very similar. Always make sure you are in the correct process before issuing commands.

The menu structure for issuing the commands used in the following case studies is shown below. The format of this menu shows the last section of each prompt in () and the commands used to traverse the structure outside of the brackets. Variables, such as ELAN names, are shown in <>.

```
(*)
--> talk 2 ()
--> talk 5 (+)
--> event (ELS>)
--> net 0 (ATM+) where 0 is the physical ATM interface
--> le-services (LE-SERVICES+)
--> work <ELAN NAME> (EXISTING LES-BUS '<ELAN NAME>'+)
--> lecs (LECS console+)
--> elan (LECS ELANs+)
--> select <ELAN NAME> (ELAN '<ELAN NAME>' selected+)
--> policies (LECS POLICIES+)
--> protocol ip (IP>)
--> protocol arp (ARP>)
--> net 1 (LEC+) where 1 is the MSS LEC id
--> talk 6 (Config>)
--> net 0 (ATM Config>)
--> le-services (LE-SERVICES Config>)
--> les-bus <ELAN NAME> (LES-BUS config for ELAN '<ELAN NAME>')
--> lecs (LECS config>)
--> elan (LECS ELANs config>)
--> select <ELAN NAME> (ELAN '<ELAN NAME>' selected>)
--> policies (LECS POLICIES config>)
--> protocol ip (IP config>)
--> protocol arp (ARP config>)
```

Figure 219. MSS Command Menus

7.6 Gathering Information by Using IBM Nways Campus Manager

The IBM Nways Campus Manager is a very useful tool for gathering information about your network environment. It shows the physical and logical connectivity, and also ATM LAN emulation addresses and Classical IP addresses used in your network. This can significantly help in gathering the information you require to diagnose and locate your network problem.

In the following diagram you can see which ATM devices are connected to the IBM Nways 8260 Multiprotocol Switching Hub, and the status of each slot/port on the hub. You can get this screen by double-clicking on the following objects:

NetView for AIX Root Map -> ATM Campus -> Cluster Number -> Hub Symbol

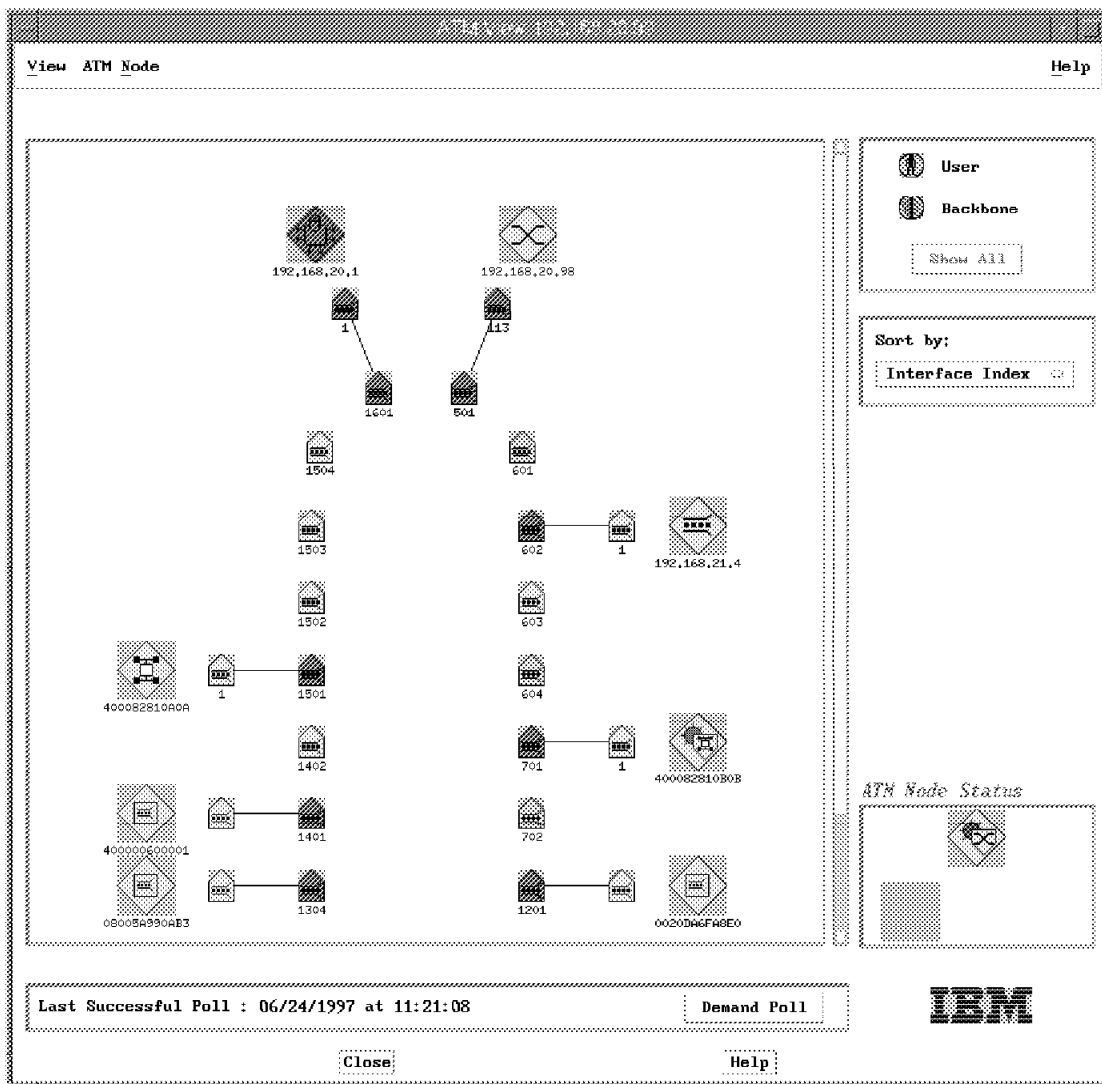


Figure 220. ATM Interface Submap in IBM Nways Campus Manager

You can see the status of LAN emulation components visually, and it's easy to get the configuration of these. You can get the following screen by double-clicking as follows: **NetView for AIX Root Map -> LAN Emulation -> Domain -> ELAN**

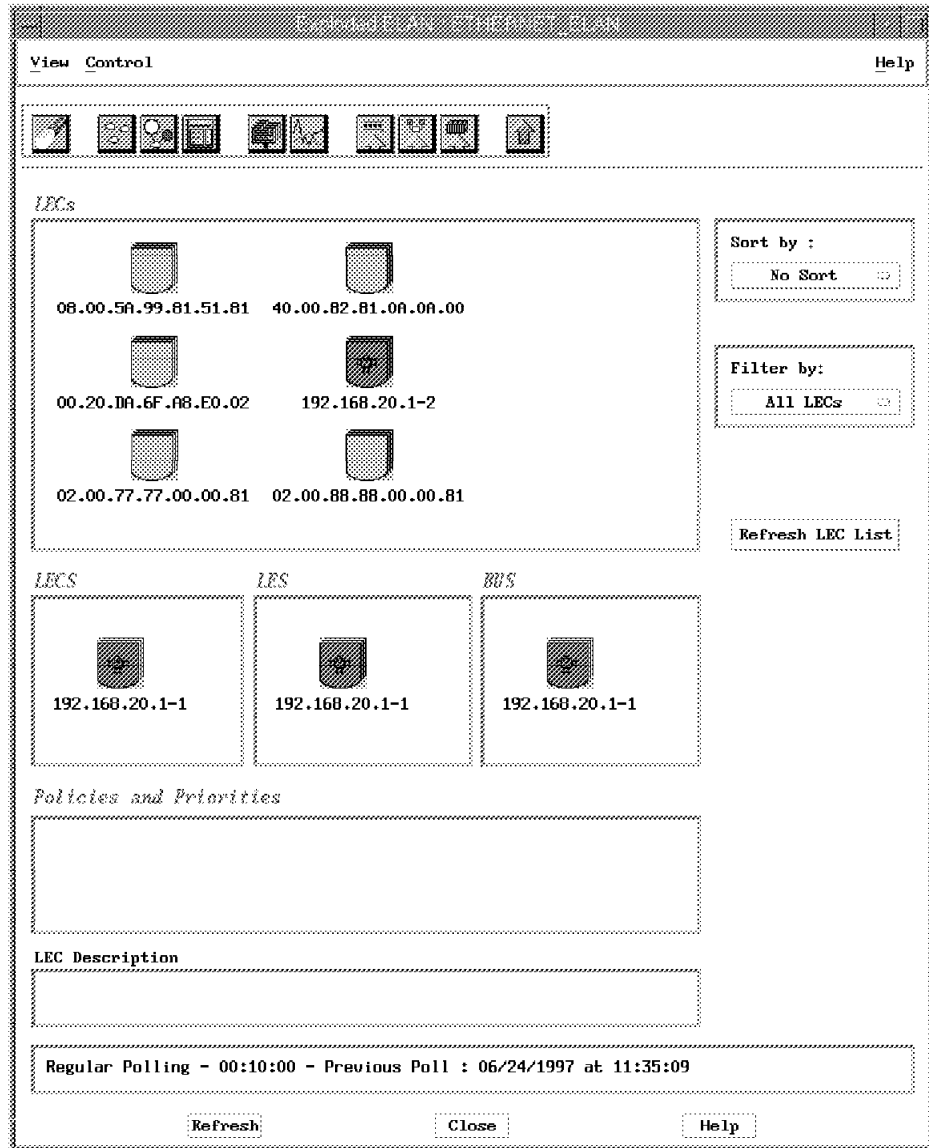


Figure 221. Exploded ELAN by IBM Nways Campus Manager

7.7 Hints and Tips with IBM Products

This section is intended to describe some useful commands, hints and tips for use with IBM LAN emulation and classical IP environments.

<i>Table 23 (Page 1 of 4). Hints and Tips with IBM LAN Emulation and Classic IP Networks</i>			
Topics	Techniques	Products	Examples and Descriptions
Problem scope	Commands	MSS	<ul style="list-style-type: none"> • MSS Existing LES-BUS ' < ELAN NAME>' + database list of all LECs : list LECs attached to an ELAN
LEDs on devices	Simple Test	All	: visual indication
Device status	Commands	MSS	<ul style="list-style-type: none"> • MSS + configuration : list the network interface information • MSS + int : display statistical information about the network interface • MSS + test 1 (1: interface #) : verify the state of an interface
	Commands	8274	<ul style="list-style-type: none"> • % vap 4/1 (4/1: slot #/port #) : view the configuration/status ATM port
The status of ATM Network (slot/port/ESI)	Commands	8260/8285	<ul style="list-style-type: none"> • ATM_SW> show module 16 verbose (16:slot #) : display the status/information of module/port • ATM_SW> show port 16.1 verbose (16.1: slot#.port#) : display configuration information for an ATM media port • ATM_SW> show atm_esi all (SHOW REACHABLE_ADDRESS ALL) : display ATM addresses registered in switch
Micro-code level	Commands	8260/8285	<ul style="list-style-type: none"> • ATM_SW> show device : display Boot/Flash EEPROM version
	Commands	MSS	<ul style="list-style-type: none"> • MSS + configuration : list software version and boot ROM version
	Screen view	Endstation	<ul style="list-style-type: none"> • See the screen on reboot : display version of ATM adapter device driver
Switch configuration	Commands	8260/8285	<ul style="list-style-type: none"> • ATM_SW> show lan_emul configuration_server : display the entries in the LECS address table

Table 23 (Page 2 of 4). Hints and Tips with IBM LAN Emulation and Classic IP Networks

Topics	Techniques	Products	Examples and Descriptions
LECS configuration	Commands	MSS	<ul style="list-style-type: none"> • MSS LECS console+ list : list the operating parameters of the LECS • MSS LECS POLICIES config> list : list LECS policy priorities • MSS LECS ELAN ' < ELAN NAME>' selected+ policy list name : list LECS ELAN name policies • MSS LECS ELAN ' < ELAN NAME>' selected+ les list : list LECS LES definitions
LES/BUS configuration	Commands	MSS	<ul style="list-style-type: none"> • MSS EXISTING LES-BUS ' < ELAN NAME>' + list :list the LES/BUS's status and current configuration parameters
LEC configuration	Commands	MSS	<ul style="list-style-type: none"> • MSS LEC+ list config : list the LEC configuration
	Configuration tool	8281	: Use SLIP/IP/Bridge connection to 8281
	Commands	8274	<ul style="list-style-type: none"> • % vas : display services defined on the 8274 • % mas 4/1 1 (4/1 1: slot#/port# service#) : modify a LAN emulation service
	Simple test	Endstation	: see the "protocol.ini" file
ATMARP client/server configuration	Commands	RS/6000	<ul style="list-style-type: none"> • smitty tcpip : display and set TCP/IP configuration for an interface • smitty atm : display and set parameters for an ATM interface
	Commands	MSS	<ul style="list-style-type: none"> • MSS IP> int : display IP addresses of MSS interfaces • MSS ARP config> list atm-arp-client-config : display ATMARP client/server configuration

<i>Table 23 (Page 3 of 4). Hints and Tips with IBM LAN Emulation and Classic IP Networks</i>			
Topics	Techniques	Products	Examples and Descriptions
LECS statistics and events	Commands	MSS	<ul style="list-style-type: none"> • MSS LECS CONSOLE+ statistics list : display many of LECS counters • MSS LECS policies+statistics list : display the counters associated with the LECS • MSS LECS ELANs+statistics list : display the ELAN counters • MSS ELS+ display subsystem lecs all : display the event logging messages of LECS
LES/BUS statistics and events	Commands	MSS	<ul style="list-style-type: none"> • MSS EXISTING LES-BUS '< ELAN NAME>' + database list specific LEC LEC ID 0003 : list detailed information about a user-specific LEC • MSS EXISTING LES-BUS '< ELAN NAME>' + database list specific LEC ATM : list detailed information about a user-specific LEC • MSS EXISTING LES-BUS '< ELAN NAME>' + statistics display LES-BUS LES : list many of the LES counters • MSS EXISTING LES-BUS '< ELAN NAME>' + statistics display LES-BUS BUS : list many of the BUS counters • MSS ELS+ display subsystem LES all : display the event logging messages of LES
LEC statistics	Commands	MSS	<ul style="list-style-type: none"> • MSS LEC+ MIB status : list MIB status of LEC • MSS LEC+ MIB server : display the LEC MIB Server VCC tables • MSS LEC+ list arp : display LE_ARP entries for an MSS LEC
	Commands	8274	<ul style="list-style-type: none"> • % vss 4/1 1 (4/1 1: slot#/port# service#) : view the LAN emulation statistics • % vlat 4/1 1 (4/1 1: slot#/port# service#) : view LE_ARP cache on 8274 service LEC

<i>Table 23 (Page 4 of 4). Hints and Tips with IBM LAN Emulation and Classic IP Networks</i>			
Topics	Techniques	Products	Examples and Descriptions
ATMARP statistics and events	Commands	MSS	<ul style="list-style-type: none"> • MSS ARP> dump 0 : display ATMARP cache entries • MSS ARP> display 0 : display ATMARP connection information • MSS ARP> statistics : display ATMARP statistics • MSS ELS+ display subsystem arp all : display ARP events on the MSS server

Some other general hints and tips are listed below:

- If using the MSS event logging system you may use Ctrl-S to pause scrolling and Ctrl-Q to resume scrolling. Another useful approach is to use a terminal application such as *Windows Terminal* which has a scroll bar so that you may scroll back to the events you are interested in. You may also be able to capture all activity from a telnet session to a text file and then review the text.
- When you have finished viewing events with the MSS event logging system (ELS) remember to disable all event displaying using the *nodisplay sub all* command from the (ELS command + prompt). Event displaying can consume considerable resource on the IBM Nways Multiprotocol Switched Services (MSS) server and should not be left running unless you are troubleshooting or monitoring a problem.
- A useful tip when designing your ATM emulated LAN environment is to include a separate logical IP subnet only used for hub and LE service device management. This is very useful as you can still manage your network even if the emulated LAN environment fails.

7.8 Case Studies Involving Problems with LAN Emulation Connectivity

This section contains various case studies about LAN emulation using IBM products to demonstrate our problem determination guidelines.

7.8.1 Network Environment

The test environment used for the LAN emulation connectivity case studies is described below:

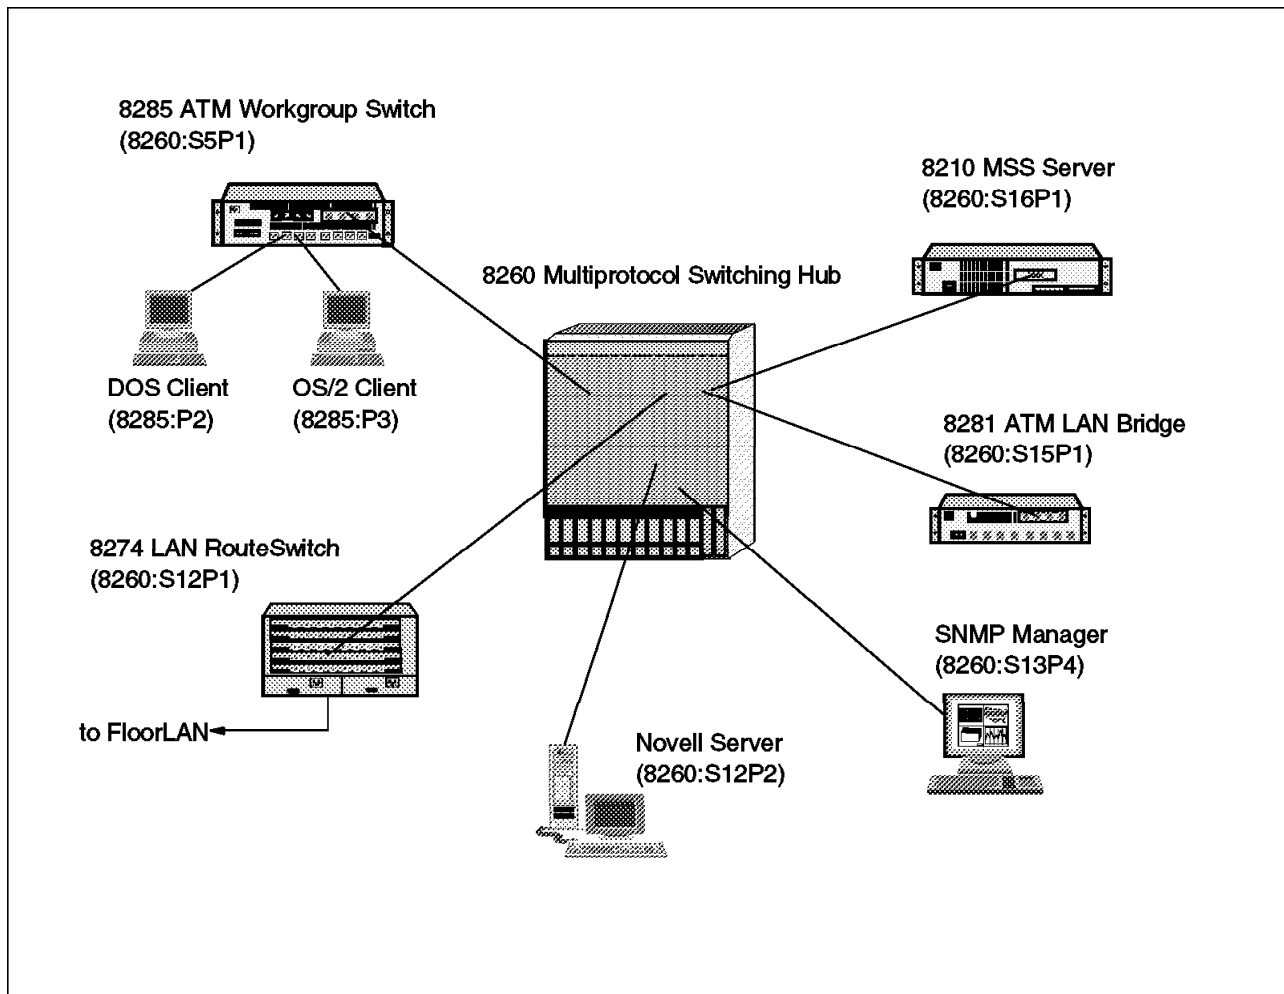


Figure 222. LAN Emulation Environment (Physical View)

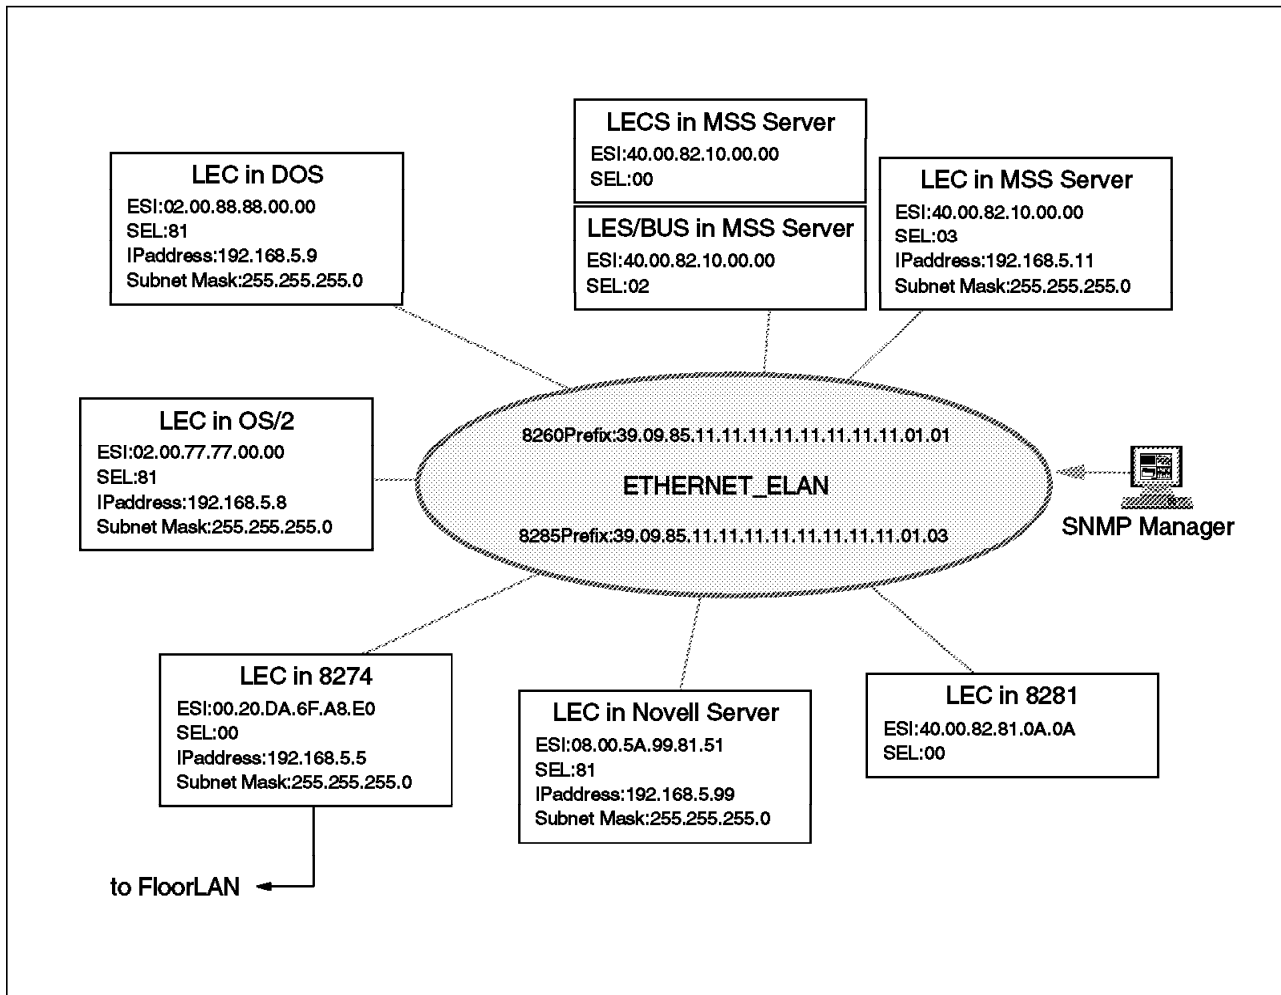


Figure 223. LAN Emulation Environment (Logical View)

We intend to use a simple ELAN environment for our case studies. In this example, the MSS server provides the LE services (LECS, LES and BUS functions). Any other information is as below:

- Single Ethernet ELAN (ELAN Name: ETHERNET_ELAN).
- ELAN polices are based on ELAN name and MAC address.
- SSI connection between IBM Nways 8260 Multiprotocol Switching Hub and IBM Nways 8285 ATM Workgroup Switch.
- Two proxy LE clients (IBM Nways 8274 LAN RouteSwitch, IBM Nways 8281 ATM LAN Bridge).

7.8.2 Symptom: LAN Emulation Clients Fail to Connect to Their ELANs

In LAN emulation environments most problems involve LE clients failing to connect to their ELAN. There are lots of possible causes, such as the ELAN name is incorrectly defined or a different maximum frame size is defined.

7.8.3 Troubleshooting Methodology in LAN Emulation Networks

The following diagram is intended to help you find the cause of problems involving LAN emulation connectivity.

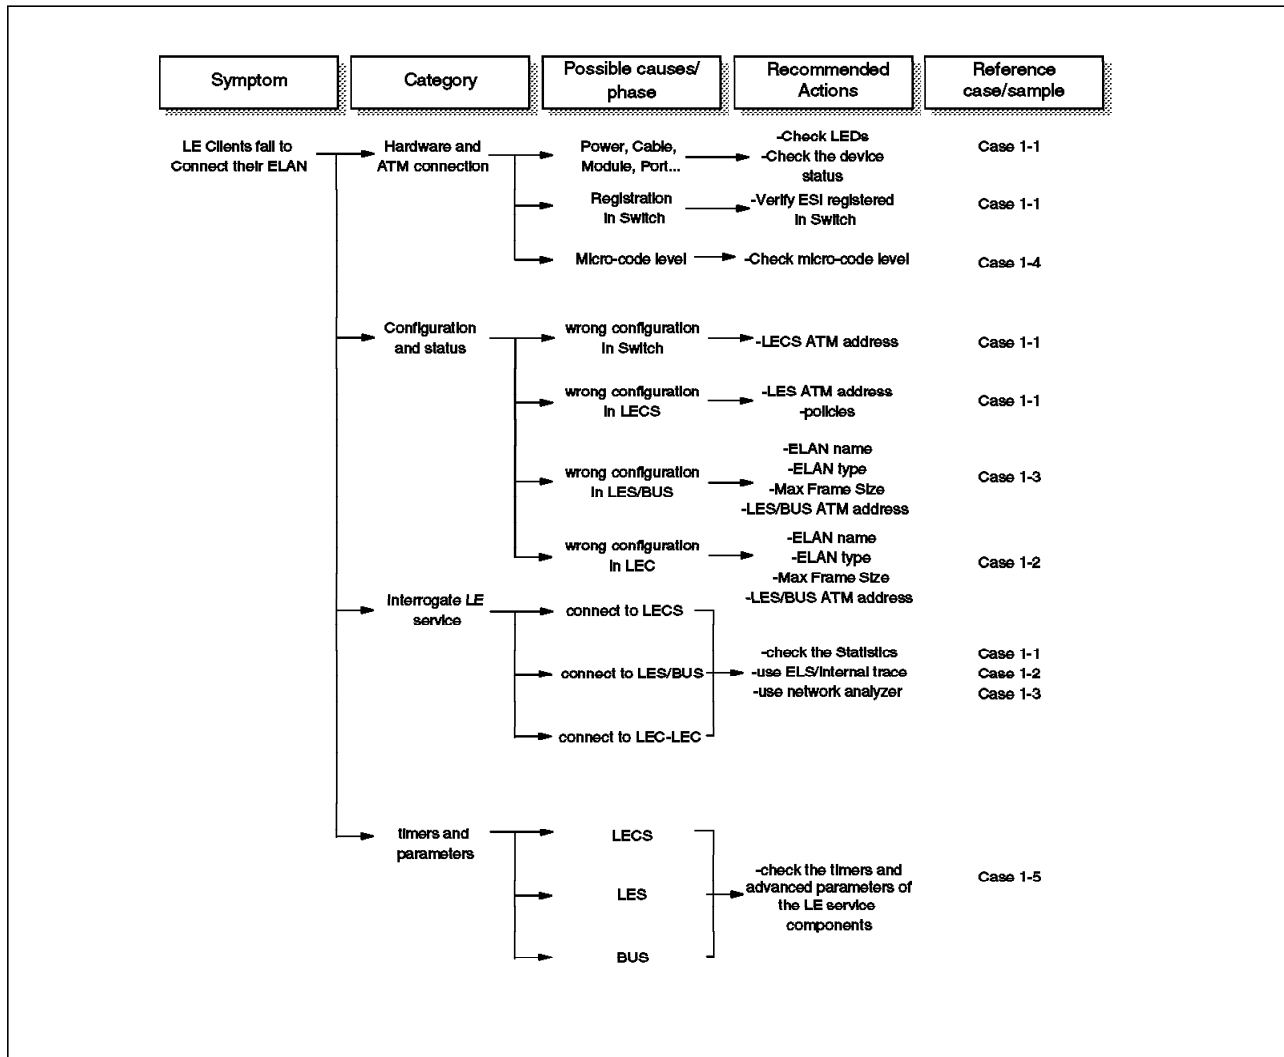


Figure 224. Troubleshooting LAN Emulation Connectivity Problems

7.8.4 Case (1-1): LECS Failure Problem

In the LAN emulation initialization phase, the LE clients establish a configuration direct VCC with the LECS. After registration, most LE clients will drop their configuration direct VCC, so ordinarily the LE clients will not have a connection with the LECS. Even if the LECS goes down, the LE clients aren't affected. If however the LE clients drop their connection to their ELAN, for example, by rebooting, they will fail to join their ELAN again.

To simulate this scenario, the LECS was deleted using the following command from the MSS server LECS console + prompt.

```
ITSO Lab MSS LECS console+delete
Delete LECS and all of its resources? [No]: yes
LECS deleted
ITSO Lab MSS LECS console+list
No LECS present
```

Figure 225. MSS Command to Delete the Active LECS

In this scenario five LECs (8210, 8274, 8281, DOS client, OS/2 client) fail to join their ELAN after the LECS goes down and they are rebooted. Only one LEC (Novell server) can join the ELAN. When you use the TURBOWAYS-155 MCA ATM Ethernet NetWare ODI driver in Novell NetWare, you may only specify the LES ATM address, you may not configure it to use a LECS.

The following command may be used on the MSS server to verify which LE clients have joined the ETHERNET_ELAN.

```
ITSO Lab MSS EXISTING LES-BUS 'ETHERNET_ELAN'+databases list all lec
Number of LEC's to display: 1

      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 LES ID  LES BUS Adrs MACs MACs
      ----- - ---- -- -- ---- -
39098511111111111111111010108005A99815181 N 0002 UP UP    1    1    0
```

Figure 226. MSS Command to Discover Clients Joined to an ELAN

7.8.4.1 Methodology

The following scenario uses the general methodology described in 7.3.1, “Emulated LAN Connectivity Problem Methodology” on page 310.

1. Stage 1 - Hardware and power problems

First check the LEDs and the status of the LE service devices. All LEDs look normal, so we concluded devices were working and powered on successfully. We checked the status of the MSS server and the IBM Nways 8274 LAN RouteSwitch using the following console commands. The MSS server (config, int and test) commands must be issued from the + prompt while the 8274 commands may be issued at any prompt.

```

ITSO Lab MSS +config

IBM 8210 Nways Multiprotocol Switching Server
Host name: ITSO Lab MSS
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
23 ASRT Adaptive Source Routing Transparent Enhanced Bridge
29 NHRP Next Hop Routing Protocol

Num Name Feature
2 MCF MAC Filtering
6 QOS Quality of Service

2 Networks:
Net Interface MAC/Data-Link Hardware State
0 ATM/0 ATM CHARM ATM Up 1
1 Eth/0 Ethernet/IEEE 802.3 CHARM ATM Up 2

ITSO Lab MSS +int

Nt Interface Slot-Port Self-Test Passed Self-Test Failed Maintenance Failed
0 ATM/0 Slot: 1 Port: 1 1 0 0 0
1 Eth/0 1 0 0 0

ITSO Lab MSS +test 0 3
Testing net 0 ATM/0...successful

ITSO Lab MSS +int

Nt Interface Slot-Port Self-Test Passed Self-Test Failed Maintenance Failed
0 ATM/0 Slot: 1 Port: 1 2 0 0 0
1 Eth/0 1 0 0 0
    
```

Figure 227. MSS Commands to Check Device and Interface Status

Both the physical ATM interface **1** and the LE client interface **2** show Up. We therefore concluded that the interfaces of the MSS server are working correctly. To make sure, we tested the physical ATM interface of the MSS **3** which tested correctly.

```

/ % vap 4/1

                          ATM Port Table

Slot Port      ATM Port Description      Conn Tran  Media UNI Max  VCI
=====
4    1    ATM PORT                        SVC  STS3c  Multi Pri 1023 10

Slot Port Loopback Cfg Tx Clk Source
=====
4    1    NoLoop      LocalTiming

Slot Port      ATM Network Prefix      End System  Sig Sig  ILMI  ILMI
=====
4    1    390985111111111111110101  0020da6fa8e0 3.0 5    True  16

                                          Status
Slot Port Tx Seg Sz Rx Seg Sz Tx Buff Sz Rx Buff Sz Oper SSCOP ILMI
=====
4    1      8192  8192  4600  4600 Enabled Up Up 1

```

Figure 228. 8274 Command to Check Device and Interface Status

The 8274 interface 4/1 (the ATM interface) **1** shows that SSCOP and ILMI are Up. We therefore also concluded that the interfaces of the 8274 are working correctly.

2. Stage 2 - ATM problems

We then checked the status of the ATM modules and ports on the IBM Nways 8260 Multiprotocol Switching Hub and confirmed the ATM connectivity of devices using the SHOW MODULE and SHOW PORT commands. The example below shows the output from the use of this command on module 16 which had the MSS server connected to port 1. Refer Chapter 5, “Starting Problem Isolation in an ATM Network” on page 109 for more detailed information.

```

8260ATM1> show module 16 verbose

Slot Install Connect Operation General Information
-----
16    Y    Y    Y    8260 ATM 2 Ports LAN 155 Mbps Module
      <-----<-----<----- 1

status: connected / hardware okay 2
      enable / Normal 3
P/N: 51H3635 EC level: E28056 Manufacture: VIME
Operational FPGA version : 81
Backup FPGA version : 81

Type Mode Status
-----
16.01: UNI enabled UP-OKAY 4
16.02: UNI failed ERROR 5

```

Figure 229. 8260/8285 Command to Check Module Status

We checked the following:

- **1** the ATM module is installed, connected and operating correctly.
- **2** / **3** the ATM module status has no problems and is normal.
- **4** slot16/port1 which is connected to MSS server is normal.

Note: Slot16/port2 **5** shows UNI failed ERROR since no device was attached to it.

```
8260ATM1> show port 16.1 verbose

Type Mode Status
-----
16.01:UNI enabled UP-OKAY

Signalling Version : with ILMI, forced 3.0 1
Flow Control      : Off
Frame format      : SONET STS-3c 2
Connector         : SC DUPLEX
Media             : Multimode fiber
Port speed        : 155000 Kbps
Remote device is active 3
IX status         : IX OK 4
Scrambling mode  : frame and cell
Clock mode        : internal
```

Figure 230. 8260/8285 Command to Check Port Status

We checked the following:

- **1** the signalling was the same as defined in the MSS server
- **2** the frame format (SONET STS-3c) was correct.
- **3** the remote device is active.
- **4** the status of IX is no problem.

We then made sure all clients were registered with the ATM switch. To check this we listed the ESI addresses registered with the 8260 and 8285 using the SHOW ATM_ESI ALL command. On the CPSW Version 3 software this command has been superseded by the SHOW REACHABLE_ADDRESS ALL command. The output of the command on the 8260 is shown below.

```
8260ATM1> show atm_esi all
Port  ATM_ESI          Type
-----
12.01 11.11.11.11.11    static (id= 1)
 6.02 50.00.00.82.82.A1 dynamic
 7.01 40.00.82.81.0B.0B dynamic
12.01 00.20.DA.6F.A8.E0 dynamic 1
12.02 08.00.5A.99.81.51 dynamic 2
13.04 08.00.5A.99.0A.B3 dynamic
14.01 40.00.00.60.00.01 dynamic
15.01 40.00.82.81.0A.0A dynamic 3
16.01 40.00.82.10.00.00 dynamic 4
```

Figure 231. 8260/8285 Command to Determine ATM Addresses Registered with the Switch

Note: **1** - 8274, **2** - Novell Server, **3** - 8281, **4** - MSS server

The other ESI addresses registered with the ATM switch are not used in this case study.

We concluded all devices had registered with the ATM switch correctly.

3. Stage 3 - Simple configuration problems

Incorrect configuration causes many common problems. Faster problem determination is possible by checking some simple configuration parameters for the main LAN emulation components.

- Switch configuration

We used the following commands to ensure that the ATM address of the LAN emulation configuration server (LECS) was defined in the 8260 and 8285 correctly. In LANE 1.0 three ways can be used by an LE client to determine the address of its LECS. These are described in 7.2.1.1, “LE Service Operation” on page 297. In our case study the actual LECS ATM address was defined as the well-known address in the 8260 and 8285 hubs.

```
8260ATM1> show lan_emul configuration_server
Index          ATM address
-----
1 WKA active 39.09.85.11.11.11.11.11.11.11.01.01.40.00.82.10.00.00.00
```

Figure 232. 8260 Command to Determine the Configured LECS ATM Address

```
8285> show lan_emul configuration_server
Index          ATM address
-----
1 WKA active 39.09.85.11.11.11.11.11.11.11.01.01.40.00.82.10.00.00.00
```

Figure 233. 8285 Command to Determine the Configured LECS ATM Address

- LECS configuration

We used the following command to check the status of the LECS, from the MSS LECS config> prompt.

```
ITSO Lab MSS LECS console+list
No LECS present
```

Figure 234. MSS Command to View the Status of the Running LECS

We discovered that the active LECS configuration had been deleted. So we checked the saved configuration using the following MSS command from the talk 6 process.


```
ITSO Lab MSS LECS config>list
LECS Detailed Configuration
  Lecs is                               Enabled      1
  ATM Device number:                     0
  ESI:                                    40.00.82.10.00.00 2
  Selector:                               0x00         3
  Validate Best Effort PCR:              No
  Configuration Direct Max Reserved BW (Kbps): 0
  Maximum number of simultaneous VCCs:    128
  Idle VCC Timeout (in seconds):         60
  .....
```

Figure 235. MSS Command to Check the Saved LECS Configuration

Note: The command screen has been shortened to only show the important section of the output.

We checked the following:

- 1 the LECS is enabled.
- 2 the end system identifier is the same as that defined in the 8260 and 8285.
- 3 the selector bytes are the same as those defined in the 8260 and 8285.

From the MSS ELANs config> prompt we checked the ELANs defined.

```
ITSO Lab MSS LECS ELANs config>list
ELAN Listing...

Name      Type      Packet Size  Enabled
=====  =====  =====
ETHERNET  Ethe      1516        Yes
  1        2        3        4
```

Figure 236. MSS Command to Check the LECS ELANs Defined

We checked the following:

- 1 the ELAN Name is correct.
- 2 the ELAN Type is correct.
- 3 Packet Size is as defined in each LE client.
- 4 the ELAN is enabled.

We confirmed the LECS policy priorities were correct by using the following command from the MSS LECS POLICIES config> prompt.

```
ITSO Lab MSS LECS POLICIES config>list
Policy Listing...

Enabled  Priority  Type
=====  =====  =====
Yes      10       byE1anNm
Yes      10       byMacAddr
```

Figure 237. MSS Command to Check the LECS Policy Priorities

In our case studies the LE clients were assigned to the emulated LAN by the LECS using the only ELAN name. To check the policy values defined in the LECS we used the following command from the MSS LECS ELANs + prompt.

```

ITSO Lab MSS LECS ELANs+select
  ( 1) ETHERNET_ELAN
Choice of ELAN [1]?
ELAN 'ETHERNET_ELAN' selected for detailed console
ITSO Lab MSS ELAN 'ETHERNET_ELAN' selected+policy list name

ELAN name => LES
=====
'ETHERNET_ELAN'
=> Local LES for: ETHERNET_ELAN
    
```

Figure 238. MSS Command to Determine LECS Policy Values

Note: If we had created a MAC address policy for the emulated LAN in the LECS and used the (policy list mac) command to view the policy values, we would see the list of MAC addresses defined for the MAC address policy.

We checked the LES to ensure that the LE clients assigned by the LECS were correct and that the ELAN name in the policy was correct.

Finally we checked the list of LESs defined for this ELAN using the command below from the MSS server ELAN 'ETHERNET_ELAN' selected + prompt.

```

ITSO Lab MSS ELAN 'ETHERNET_ELAN' selected+les list
LESs serving ELAN 'ETHERNET_ELAN'
=====

Primary ATM address: Local LES for: ETHERNET_ELAN
No backup LES provided
    
```

Figure 239. MSS Command to Determine the List of Defined LESs for an ELAN

We have discovered that the LECS was configured correctly but is not currently active. We restarted the MSS using the following command.

```

ITSO Lab MSS LECS console+restart
  ( 1) rebuild
  ( 2) retain
Retain current LECS databases or rebuild from SRAM [1]?
LECS restarted successfully
restarted all enabled LECS objects
    
```

Figure 240. MSS Command to Restart the LECS

If you find that your LECS has failed, it is possible to restart the MSS server LECS function using the (restart) command. The restart command stops an operating LECS, flushes its resources, and reinitializes/restarts the LECS. You are prompted to either keep the current database or to rebuild it from the static configuration data.

4. Stage 4 - Interrogate LE service, VCCs and data flows

We now interrogated the MSS server regarding the the LECS to ensure it was working correctly.

First we checked the status of the LECS, which should be Operating normally **1**, the size of the error log which should be empty **2** and ATM address of the LECS to make sure it was the same as defined in the 8260 and 8285 **3**.

```
<after restart the LECS>
.....
ITSO Lab MSS LECS console+list
Status of LECS:
ATM device number:          0
State:                      Operating normally(88) 1
  Time of last state change: 00.00.33.13
  Elapsed time since last change: 02.53.35.94
Error Log:                  no err (0) 2
Local ATM address:         390985111111111111111010140008210000000 3
Well-known address:       470079000000000000000000000000A03E00000100
UNI version:               UNI Version 3.0
Validate best effort PCR:  No
Maximum config direct VCC reserved bandwidth: 0 Kbps
Maximum number of config direct VCCs to LECS: 128
Seconds before VCC declared idle: 60
Trace ATM address value:0000000000000000000000000000000000000000
Trace ATM address mask: FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
```

Figure 241. MSS Command to Check the Active Status of the LECS

If the state of LECS is as below, the LECS must be restarted.

- Down due to error
- State unknown!!!

Note: To display the error log, you can use the MSS server firmware. You can access the firmware by stopping the boot process. To do this, you must have a TTY console directly attached to the serial port. When the MSS server starts its boot process, press and hold Ctrl +C at the terminal keyboard. For more information refer to the *Nways MSS Server Service*, GY37-0354.

You can also verify the configuration and status of the LECS by using IBM Nways Campus Manager as shown below:

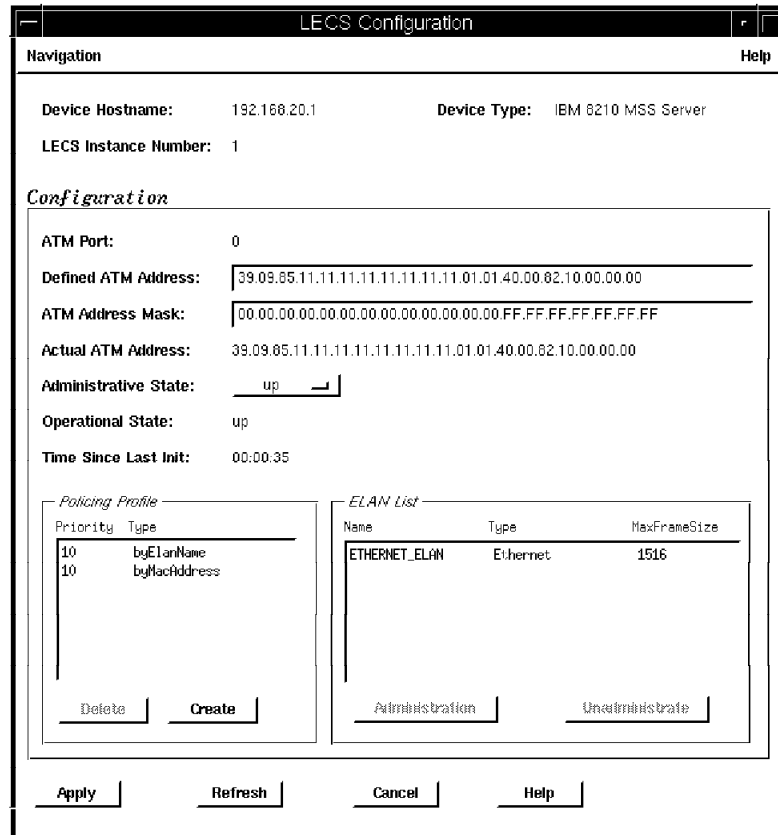


Figure 242. LECS Configuration in IBM Nways Campus Manager

Next verify the LECS is working correctly by viewing the LECS statistics. The command below illustrates how this can be done from the MSS LECS console prompt.

```

ITSO Lab MSS LECS console+statistics list
LECS has:
  1 ELAN(s)
  2 policy(ies) at 2 priority(ies)
LECS has 1 configure direct VCCs
  has accepted 1372 VCCs and rejected 0 VCCs
  0 VCCs have been dropped by LECS, 1371 dropped by caller
  has exceeded its maximum number of VCCs 0 times
LECS discarded frames: 0
LECS responses by status (zero responses if status not displayed)
  Success( 0) : 6
  No Configuration(20) : 1366
    
```

Figure 243. MSS Command to View LECS Statistics

Note: The numbers shown in brackets after the Success and No configuration LECS responses are the LANE status codes obtained from the LE data frames. For a full list of LANE status values refer to D.8, “Control Frame Status Values” on page 491.

You need to make sure that the number of successful LECS responses is more than the number of LECs in your LAN emulation network **1**. A number lower than the number of LECs would indicate that not all your LECs are connecting with your LECS. Also confirm that the number of No

Configuration LECS responses is relatively low **2**. A high value usually indicates that many LE clients are configured incorrectly. In our example above, a number of additional LE clients (not shown in our case study network diagram) were connected to the ATM network but not configured for our ELAN. These clients continually attempted to join the ELAN but failed because of their configuration.

Most LE clients will drop their configuration direct VCC to the LECS after they have completed their registration procedure. It is therefore difficult to verify whether LECs connect to their LECS or not by viewing the VCCs in the network. A better way to investigate the source of the connection problem is to use the MSS event logging system to display the LE configuration requests from clients to the LECS.

The following is the result of using the MSS server event logging system to display the LE configuration requests when the DOS client was rebooted after the LECS was restarted.

To use the event logging system do the following:

- a. Issue a `nodisplay subsystem all` command from the ELS + prompt
- b. Issue a `display sub <subsystem name> all`, for example: `display sub lecs all` from the same prompt.
- c. Press `CNTL + P` to return to the `OPCON *` prompt
- d. Enter `talk 2` or `t 2` to enter the event logging console

You may use `Ctrl+S` to pause scrolling and `Ctrl+Q` to resume scrolling. Another useful approach is to use a terminal application such as (Windows Terminal) which has a scroll bar so that you may scroll back to the events you are interested in. You may also be able to capture all activity from a telnet session to a text file and then review the text file.

Note

When you have finished viewing events remember to disable all event displaying using the `nodisplay sub all` command from the ELS + prompt. Event displaying can consume considerable resources on the IBM Nways Multiprotocol Switched Services (MSS) server and should not be left running for considerable periods of time.

```

LECS.106: LECS: incmng call: local addr
LECS.093: LECS: cfgtn drct frm
           x39098511111111111111111010302008888000081 estblshd
LECS.073: LECS: frm pssd vldtn chcks
LECS.074: LECS: LEC x39098511111111111111111010302008888000081 assgnd to
           LES xLocal LES for: ETHERNET_ELAN at 1 usng byElanNm
LECS.114: LECS: mem lkup success: 39098511111111111111111010302008888000081
           prim LES Local LES for: ETHERNET_ELAN:
           last LES 39098511111111111111111010140008210000002

LECS.121: LECS: lcl LES addr for ELAN 'ETHERNET_ELAN'
           mapped to LES: 39098511111111111111111010140008210000002

LECS.109: LECS: updtng LEC addr in mem:
           39098511111111111111111111111010302008888000081
           LES 39098511111111111111111010140008210000002 time 704
LECS.095: LECS: snt cfgrtn rspns for x39098511111111111111111010302008888000081
LECS.094: LECS: cfgtn drct frm
           x39098511111111111111111010302008888000081 dscnctd
    
```

Figure 244. MSS "talk 2" Events for the LECS Subsystem

The LEC establishes a configuration direct VCC to the LECS. The LECS assigns the suitable LES based, in our case, on ELAN name (ETHERNET_ELAN), then it gives the ATM address of LES to the LE client. Finally the configuration direct VCC is dropped.

7.8.4.2 Conclusion

We discovered the LECS had failed and restarted it using the configured LECS parameters. Most of the LE clients (IBM Nways 8274 LAN RouteSwitch, IBM Nways 8281 ATM LAN Bridge etc.) rejoin their ELANs after a few minutes; however, some of the LE clients (PC-DOS) need to be rebooted.

Finally all six LECs can join their ELAN as shown below.

```

ITSO Lab MSS EXISTING LES-BUS 'ETHERNET_ELAN'+databases list all lec
Number of LEC's to display: 6

      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
      -----
39098511111111111111111010140008210000003 N 0001 UP UP 1 1 0
39098511111111111111111010108005A99815181 N 0002 UP UP 1 1 0
3909851111111111111111101010020DA6FA8E000 Y 0003 UP UP 1 1 0
39098511111111111111111010302007777000081 N 0004 UP UP 1 1 0
39098511111111111111111010302008888000081 N 0005 UP UP 1 1 0
390985111111111111111110101400082810A0A00 Y 0006 UP UP 1 5 0
    
```

Figure 245. MSS Command to View LECs Assigned to an ELAN

7.8.5 Case (1-2): Wrong ELAN Parameters in LE Client

In this scenario four LE clients (8210, 8281, OS/2 client and Novell server) join their emulated LAN, but two LE clients (8274 and DOS client) fail to join. LE clients are assigned to their emulated LAN by the LECS using the ELAN name. If the wrong parameters are specified in the LE client's configuration, LE clients may fail to join the correct emulated LAN.

We used the database list all lec command from the MSS EXISTING LES-BUS 'ETHERNET_ELAN'+ prompt to illustrate which LE clients are joined to the emulated LAN at the start of the scenario.

```

ITSO Lab MSS EXISTING LES-BUS 'ETHERNET_ELAN'+databases list 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
-----
39098511111111111111111010140008210000003  N 0001  UP  UP    1    1    0
39098511111111111111111010108005A99815181  N 0002  UP  UP    1    1    0
39098511111111111111111010302007777000081  N 0004  UP  UP    1    1    0
390985111111111111111110101400082810A0A00  Y 0006  UP  UP    1    5    0
    
```

Figure 246. MSS Command to Determine Which Clients Are Connected to the ELAN

Note

Some LE client drivers don't support the use of the LECS, in which case the ATM address of the LES needs to be manually configured in the LE client.

7.8.5.1 Methodology

The following scenario uses the general methodology described in 7.3.1, "Emulated LAN Connectivity Problem Methodology" on page 310.

1. Stage 1 - Hardware and power problems

We verified the network devices were working normally, as described in 7.8.4, "Case (1-1): LECS Failure Problem" on page 342.

2. Stage 2 - ATM problems

We also verified all devices were registered with their ATM switch, as described in 7.8.4, "Case (1-1): LECS Failure Problem" on page 342.

3. Stage 3 - Simple configuration problems

- Switch and LECS configuration

We verified the configuration of the LECS address in each ATM switch, and the configuration of the LECS, as described in 7.8.4, "Case (1-1): LECS Failure Problem" on page 342.

- LES/BUS configuration

We confirmed the LES and BUS configuration was correct. This is described in more detail in 7.8.6, "Case (1-3): LES/BUS Failure Problem" on page 360.

- LEC and proxy LEC configuration

For the LECs that failed to connect to the emulated LAN, we interrogated their configuration. On the IBM Nways 8274 LAN RouteSwitch this can be done using the command shown below.

```

/ % mas 4/1 1

Slot 4 Port 1 Service 1 Configuration

1) Description (30 chars max)      : LAN Emulation Service 1
2) LAN Emulated Group              : 4
  21) LAN type { 802.3 (1),
                802.5 (2) }        : 802.3           1
  22) Change LANE Cfg { NO (1),
                YES (2) }         : NO
3) LECS Address (40-char-hex)      :
   470079000000000000000000000000A03E00000100 2
4) Admin Status { disable(1),
                enable(2) }      : Enable
6) Connection Type { PVC(1),
                SVC(2) }         : SVC
  60) SEL for the ATM address      : 00
Enter (option=value/save/cancel) : 22=2

Slot 4 Port 1 Service 1 LANE Configuration Parameters

1) Proxy { NO (1), YES (2) }       : YES
2) Max Frame Size { 1516 (1), 4544 (2),
                  9234 (3), 18190 (4) } : 1516           3
3) Use translation options{NO (1), YES (2) } : Yes (use Swch menu to set)
4) Use Fwd Delay time { NO (1), YES (2) } : NO
5) Use LE Cfg Server (LECS){ NO (1), YES (2)}: YES           4
6) Use Default LECS address { NO(1), YES (2)}: YES           5
7) Control Time-out (in seconds)      : 120
8) Max Unknown Frame Count            : 1
9) Max Unknown Frame Time (in seconds): 1
10) VCC Time-out Period (in minutes)  : 20
11) Max Retry Count                   : 1
12) Aging Time (in seconds)           : 300
13) Expectd LE_ARP Resp Time (in seconds) : 1
14) Flush Time-out (in seconds)       : 4
15) Path Switching Delay (in seconds) : 6
16) ELAN name (32 chars max)          : ETH_ELAN       6

```

Figure 247. 8274 Command to Check Its LE Client Configuration

Note: The ATM port in use on the IBM Nways 8274 LAN RouteSwitch was port 1 on module 4 (4/1) and the service number for the LE client assigned to our emulated LAN was number 1.

We checked the following:

- **1** The LAN type (802.3).
- **3** The Maximum frame size (1516).
- **4** The IBM Nways 8274 LAN RouteSwitch was set up to use a LECS (YES).
- **5** The IBM Nways 8274 LAN RouteSwitch was to use the default LECS WKA (YES).
- **6** The ELAN name used by the client (ETH_ELAN).

Note: The 8274 supports using the default LECS WKA or using a manually configured LECS address. The address shown in **2** above is only used if the parameter **4** is set to No.

On the DOS client its ATM configuration is stored in its *protocol.ini* file. The figure below shows the relevant section of the *protocol.ini* file used for ATM.

```
[AT25LED]
DRIVERNAME=AT25LEç
IO_ADDR=0X120
MAC_ADDR="400011110000"
LEC_AUTO_CFG="YES"
LAN_TYPE="802.3+V2"
ELAN_NAME="ETHERNET_ELAN"
MAX_FRAME_SIZE=4544
ENHANCED_MODE="YES"
UNI_VERSION="AUTO"
MAX_MULTIC_ADDR=8
BEST_E_PK_RATE=25600
CNTRL_P_PK_RATE=25600
AGING_TIMEOUT=300
CONN_COMP_TIMER=4
CONTROL_TIMEOUT=120
LE_ARP_RSP_TIME=4
FLUSH_TIMEOUT=4
FRWRD_DELAY_TME=15
LE_ARP_CACHE_SZ=16
MAX_CFG_RETRIES=1
MAX_RETRY_CNT=1
MAX_UNKN_FR_CNT=1
PATH_SWTCH_DELY=6
VCC_TIMEOUT=1200
ADAPTER_TYPE=2
```

1
2
3
4

We checked the following:

- **1** The DOS client was set to use a LECS (YES).
- **2** The LAN type (802.3+V2).
- **3** The ELAN name (ETHERNET_ELAN).
- **4** The Maximum frame size (4544).

We then compared the configurations of the two clients above with the configuration of a working client, the MSS LE client. The command shown below allows you to retrieve the MSS LE client configuration.

Note: The ETHERNET_ELAN LE client in our scenario was defined as interface 1. To confirm this we used the `int list` command from the `talk 5 (+)` prompt.

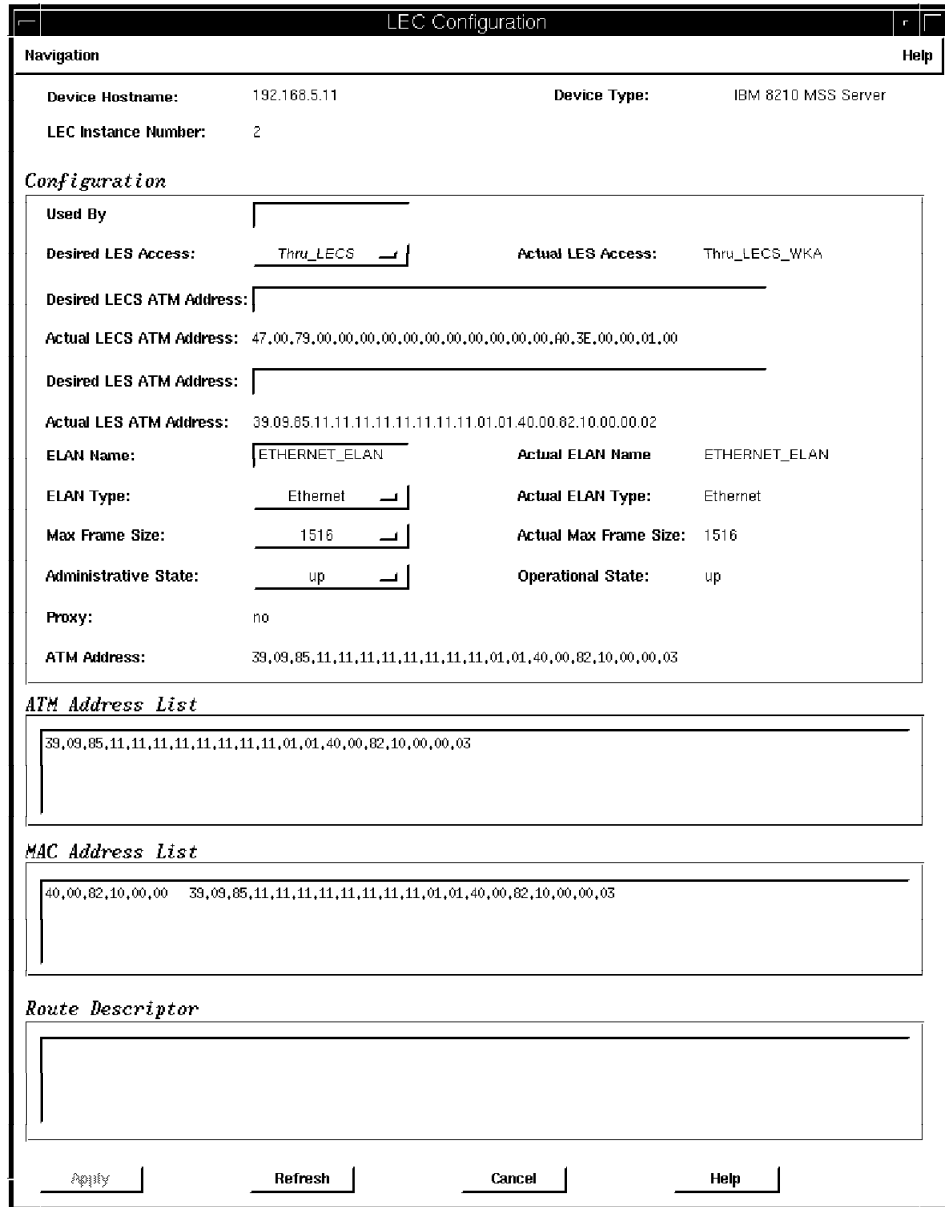


Figure 249. LEC Configuration in Campus Manager

We also checked the configuration of another working client, the 8281. To check the configuration of the 8281 set up a SLIP/IP connection to the 8281, and retrieve the current configuration using the IBM 8281 ATM/LAN Bridge Configuration Utility. In its LAN Emulation Settings screen, we verified the emulated LAN name and the maximum frame size.

From this we noticed that the ELAN name parameter used by the IBM Nways 8274 LAN RouteSwitch was incorrectly specified and the maximum frame size used by the DOS client was invalid for an Ethernet emulated LAN.

4. Stage 4 - Interrogate LE service, VCCs and data flows

To see what happens when the failing clients try to connect to the LECS, we looked at the statistics displayed by the LECS on the MSS. The command below demonstrates this and must be run from the LECS console + prompt.

First we cleared the statistics using the `stat clear` command and then rebooted the failing clients and looked at the statistics.

```
ITSO Lab MSS LECS console+stat list
LECS has:
  1 ELAN(s)
  2 policy(ies) at 2 priority(ies)
LECS has 0 configure direct VCCs
  has accepted 5 VCCs and rejected 0 VCCs
  0 VCCs have been dropped by LECS, 5 dropped by caller
  has exceeded its maximum number of VCCs 0 times
LECS discarded frames: 0
LECS responses by status (zero responses if status not displayed)
  Success( 0) : 3
  No Configuration(20) : 1 1
```

Figure 250. MSS Command to List the LECS Statistics

Note: The numbers shown in brackets after the Success and No configuration LECS responses are the LANE status codes obtained from the LE data frames. For a full list of LANE status values refer to D.8, “Control Frame Status Values” on page 491.

The No configuration responses line shows that the LECS has rejected one client. This is due to the IBM Nways 8274 LAN RouteSwitch client’s configuration containing an ELAN name that is not defined in the LECS. The LECS can’t therefore assign it to an ELAN.

The DOS client did not register in the statistics since its `MAX_FRAME_SIZE` parameter was out of bounds and hence it never tried to join the LECS. On the DOS client’s screen after boot up we saw:

```
Microsoft Protocol Manager version 2.1
IBM TURBOWAYS 25 ATM ISA Adapter Device Driver Version 2.20
(C)Copyright IBM Corp. 1994-1996. All rights reserved.
ATW0017 The value specified for the parameter MAX_FRAME_SIZE in PROTOCOL.INI is
        not in the allowable set of values.
ATW0002 An recoverable error occurred in the TURBOWAYS ATM device driver program.
ATW0025 Press any key continue . . .
```

To further examine what happens when the LECS can not assign a client to an emulated LAN we looked at the LECS configuration requests using the MSS event logging system. For a description of how to use the event logging system see 7.8.4, “Case (1-1): LECS Failure Problem” on page 342.

In this example we turned off all event displays using the `nodisplay sub all` from the `ELS +` prompt and then enabled events for the LECS by issuing a `display sub lecs all`. In talk 2 we saw the following events when the IBM Nways 8274 LAN RouteSwitch client tried to access the LECS.

```

LECS.106: LECS: incmng call: well-known addr 1
LECS.093: LECS: cfgtn drct frm 2
           x3909851111111111111111101010020DA6FA8E000 estblshd
LECS.073: LECS: frm pssd vldtn chcks 3
LECS.075: LECS: unbl to assign rqst frm 4
           x3909851111111111111111101010020DA6FA8E000
LECS.095: LECS: snt cfgrtn rspns for 5
           x3909851111111111111111101010020DA6FA8E000
LECS.094: LECS: cfgtn drct frm 6
           x3909851111111111111111101010020DA6FA8E000 dscnctd
    
```

Figure 251. MSS Talk 2 Events for the LECS Subsystem

The events above show that:

- a. **1** The IBM Nways 8274 LAN RouteSwitch tries to set up an ATM session with the LECS.
- b. **2** The configuration direct VCC is established.
- c. **3** The configuration request is processed by the LECS.
- d. **4** The LECS is unable to assign the LE client to an emulated LAN.
- e. **5** The LECS responds to the IBM Nways 8274 LAN RouteSwitch.
- f. **6** The configuration direct VCC is disconnected.

For a working example of the LECS events seen when clients successfully join their ELAN see 7.8.4, "Case (1-1): LECS Failure Problem" on page 342.

7.8.5.2 Conclusion

The DOS client fails at the initial LE stage since its maximum frame size is outside of the allowable bounds. The IBM Nways 8274 LAN RouteSwitch client connects to the LECS but fails to be assigned to an emulated LAN because its ELAN name is not defined in the LECS configuration. We corrected both problems and confirmed all LE clients could now join the emulated LAN. The command below shows the LE clients joined to the ETHERNET_ELAN at the end of the scenario.

```

ITSO Lab MSS EXISTING LES-BUS 'ETHERNET_ELAN'+databases list all lec
Number of LEC's to display: 6

LEC-LES and LEC-BUS State (UP=Up, ID=Idle, --. --.
**=Other; Show specific LEC to see actual)      v v
LEC Primary ATM Address      Proxy  ID   LES  State #ATM #Reg #Lrnd
-----
39098511111111111111111010140008210000003  N  0001  UP  UP    1    1    0
39098511111111111111111010108005A99815181  N  0002  UP  UP    1    1    0
3909851111111111111111101010020DA6FA8E000  Y  0003  UP  UP    1    1    0
3909851111111111111111101010302007777000081 N  0004  UP  UP    1    1    0
39098511111111111111111010302008888000081  N  0005  UP  UP    1    1    0
390985111111111111111110101400082810A0A00  Y  0006  UP  UP    1    5    0
    
```

Figure 252. MSS Command to Determine Which LECs Are Attached to an ELAN

7.8.6 Case (1-3): LES/BUS Failure Problem

The MSS implements a combined LES and BUS service. If the LES/BUS fails, clients will be unable to join the emulated LAN. For clients that have already joined their ELAN, before the LES/BUS failure and have started communicating with their servers via a data direct VCC, they will be unable to establish any new connections to other LE clients but their current data sessions will be unaffected.

In this scenario we stopped the MSS LES/BUS to simulate a LES/BUS failure.

We used the database `list all lec` command, from the MSS EXISTING LES-BUS 'ETHERNET_ELAN' + prompt, to illustrate which LE clients are joined to the emulated LAN at the start of the scenario.

```
ITSO Lab MSS EXISTING LES-BUS 'ETHERNET_ELAN'+data list all lec
Entry was not found
```

Figure 253. MSS Command to Determine Which LE Clients Are Attached to an ELAN

No clients were connected to the emulated LAN at the start of this scenario.

7.8.6.1 Methodology

The following scenario uses the general methodology described in 7.3.1, "Emulated LAN Connectivity Problem Methodology" on page 310.

1. Stage 1 - Hardware and power problems

We verified the network devices were working normally, as described in 7.8.4, "Case (1-1): LECS Failure Problem" on page 342.

2. Stage 2 - ATM problems

We also verified all devices were registered with their ATM switch, as described in 7.8.4, "Case (1-1): LECS Failure Problem" on page 342.

3. Stage 3 - Simple configuration problems

- Switch and LECS configuration

We verified the configuration of the LECS address in each ATM switch, and the configuration of the LECS, as described in 7.8.4, "Case (1-1): LECS Failure Problem" on page 342.

- LES/BUS configuration

We interrogated the MSS to find out the configuration of the LES/BUS using the command below from the MSS EXISTING LES-BUS 'ETHERNET_ELAN' + prompt.

```
ITSO Lab MSS EXISTING LES-BUS 'ETHERNET_ELAN'+list
ELAN Name:                ETHERNET_ELAN
ELAN Type:                Ethernet
ATM Device number:       0
# of Proxy LEC's:        0
# of Non-Proxy LEC's:    0
LES ATM Address:         (addr not valid in this LES-BUS State) 1

-Status-
LES-BUS State:           IDLE 2
Major Reason LES-BUS was last Down: none
Minor Reason LES-BUS was last Down: none
LES-BUS State last changed at: 00.02.17.75 (System Up Time)
LES-LEC Status Table changed at: 00.02.17.76 (System Up Time)
BUS-LEC Status Table changed at: 00.02.17.76 (System Up Time)
UNI Version:             unknown in this LES-BUS State
IP BCM:                  INACTIVE
IPX BCM:                 INACTIVE
NetBIOS BCM:            INACTIVE

-Current Configuration-
LES-BUS Enabled/Disabled: Enabled 3
ATM Device number:       0
End System Identifier (ESI): 40.00.82.10.00.00 4
Selector Byte:          0x02 5
ELAN Type:              (S2) Ethernet 6
Max Frame Size:         (S3) 1516 7
Control Timeout:        (S4) 120
Max Frame Age:          (S5) 1
LECID Range Minimum:    1
LECID Range Maximum:    65279
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
```

Figure 254. MSS Command to Verify the Status and Configuration of a LES/BUS

Note: The output from the command above, has been condensed to only show the important status and configuration information.

We checked the following:

- 1 The LES ATM address (not shown since LES/BUS has been stopped).
- 2 The LES/BUS status (IDLE).
- 3 The LES/BUS is enabled in the configuration.
- 4 / 5 The LES MAC and Selector byte were as defined in the LECS.
- 6 The LAN Type (ETHERNET).
- 7 The maximum frame size allowed (1516).

You can also verify the configuration and status of the LES by using IBM Nways Campus Manager as shown below:

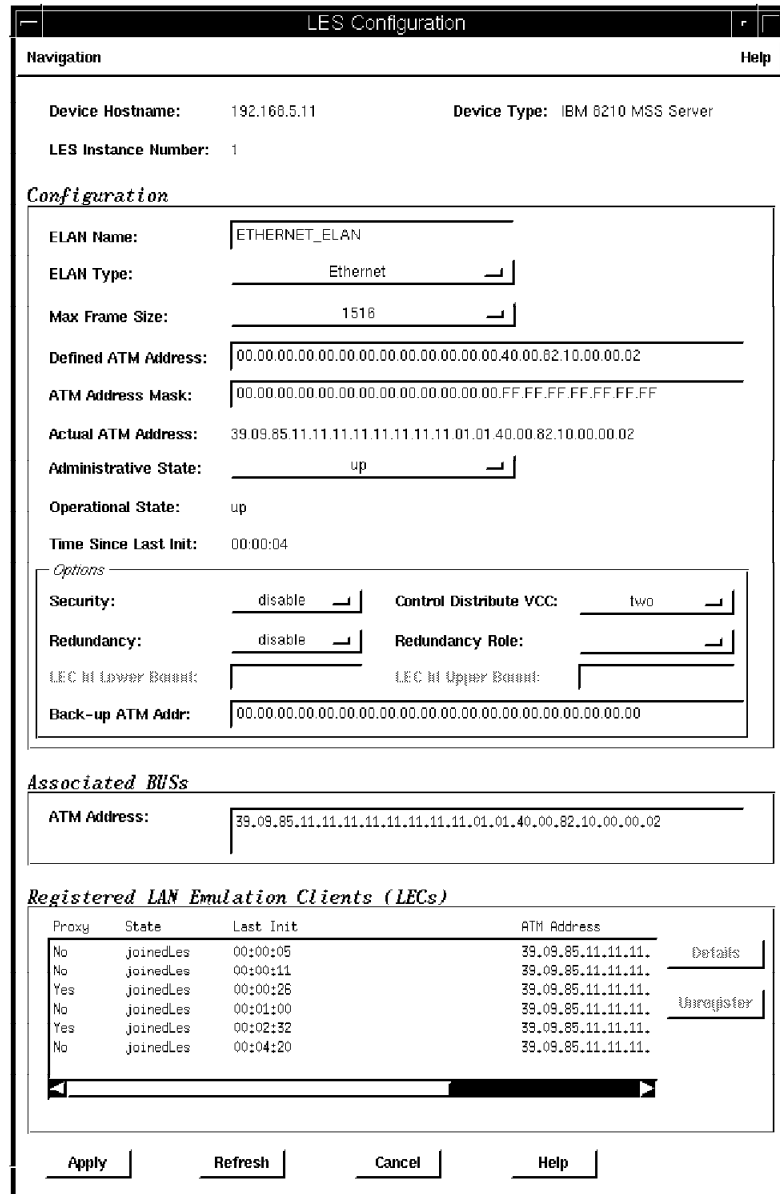


Figure 255. LES Configuration in Campus Manager

We concluded from the above that the LES/BUS looked to be set up and configured correctly but was in an incorrect state. IDLE is an indication that the LES/BUS is not operational. The valid status for a working LES/BUS is Operational.

4. Stage 4 - Interrogate LE service, VCCs and data flows

To further investigate what is happening when clients attempt to join the emulated LAN we first looked at the LECS statistics, using the command below from the MSS LECS console + prompt.

First we cleared the statistics using the stat clear command and then rebooted the failing clients and looked at the statistics.


```
ITSO Lab MSS LECS console+stat list
LECS has:
  1 ELAN(s)
  1 policy(ies) at 1 priority(ies)
LECS has 3 configure direct VCCs
  has accepted 5 VCCs and rejected 0 VCCs
  0 VCCs have been dropped by LECS, 5 dropped by caller
  has exceeded its maximum number of VCCs 0 times
LECS discarded frames: 0
LECS responses by status (zero responses if status not displayed)
  No Configuration(20) : 2
  LE_CONFIGURE Error(21) : 4 1
```

Figure 256. MSS Command to Check the LECS Statistics

Note: The numbers shown in brackets after the No configuration and LE_CONFIGURE error LECS responses are the LANE status codes obtained from the LE data frames. For a full list of LANE status values refer to D.8, “Control Frame Status Values” on page 491.

This shows that the LECS has failed to assign four LE clients **1** because of a problem with the LE configuration phase of the LE Initialization process. This is because it was configured to assign clients to a local LES but could not find a running local LES.

We then examined the LECS configuration requests and LES/BUS join requests using the MSS event logging system. For a description of how to use the event logging system see 7.8.4, “Case (1-1): LECS Failure Problem” on page 342.

In this example we turned off all event displays using the nodisplay sub all from the ELS + prompt and then enabled events for the LECS and LES/BUS by issuing a display sub lecs all followed by a display sub les all. In talk 2 we saw the following events when the OS/2 client tried to access the emulated LAN.

```
LECS.106: LECS: incmng call: local addr
LECS.093: LECS: cfgtn drct frm
           x39098511111111111111111010302007777000081 estblshd

LECS.073: LECS: frm pssd vldtn chcks
LECS.074: LECS: LEC x39098511111111111111111010302007777000081 assgnd to
           LES xLocal LES for: ETHERNET_ELAN at 10 usng byElanNm
LECS.115: LECS: mem lkup fld: 39098511111111111111111010302007777000081
LECS.122: LECS: unbl to find local LES for ELAN 'ETHERNET_ELAN'
           for LEC: 39098511111111111111111010302007777000081 1

LECS.095: LECS: snt cfgrtn rspns for x39098511111111111111111010302007777000081
LECS.094: LECS: cfgtn drct frm
           x39098511111111111111111010302007777000081 dscnntcd
```

Figure 257. MSS talk 2 Command to Examine Events for the LECS Subsystem

This demonstrates that the LECS is unable to locate a local LES for the ELAN ETHERNET_ELAN **1** so it informs the LEC and the LEC never attempts to join the LES.

As a final check we examined the VCCs established by the DOS client. To do this we use IBM Nways Campus Manager. You will only be able to use IBM Nways Campus Manager or any other SNMP manager if it still has an IP

connection to your ATM switches. In our case study we set up a Classical IP subnet for the 8260, 8285, MSS (ATMARP server) and the network management machine.

Note

A useful tip when designing your ATM emulated LAN environment is to include a separate logical IP subnet only used for hub and LE service device management. This is very useful as you can still manage your network even if the emulated LAN environment fails.

To view the current SVCs established by a client, double-click the following to get to the hub to which a device is attached: **NetView for AIX Root Map -> ATM Campus --> Cluster number --> Hub**

Then:

- a. Select the hub port that has the device of interest connected using the right-hand mouse button.
- b. A menu appears; select **Configuration**.
- c. The ATM Switch Configuration window appears, where you select **list** from the SVC menu.
- d. The list of active SVCs will appear in ATM SVC List window.

Figure 258 on page 365 shows that the DOS client currently has no SVCs established. It has therefore failed to establish any sessions with the LES and BUS servers.

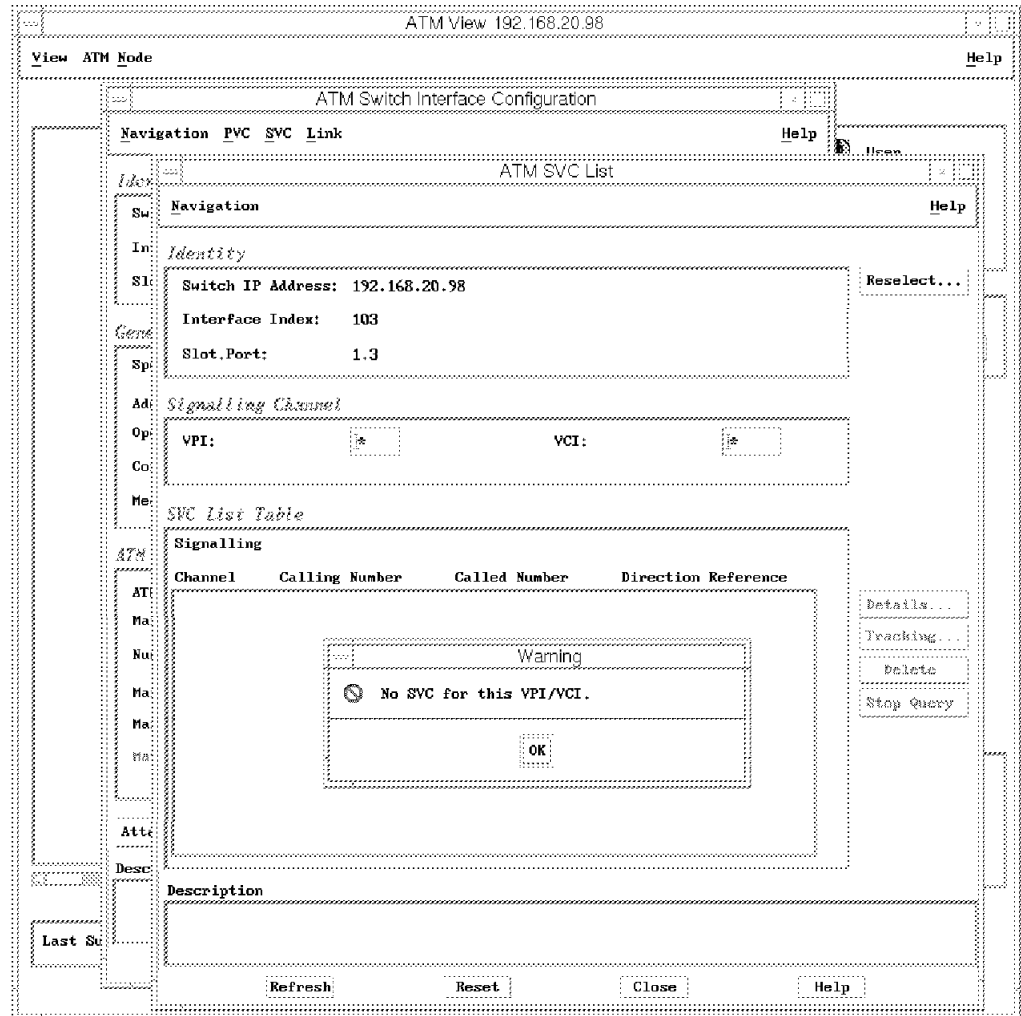


Figure 258. SVCs Established by DOS Client during LES Failure

7.8.6.2 Conclusion

We restarted the MSS LES/BUS and rechecked the *database list all lec* command, from the MSS EXISTING LES-BUS 'ETHERNET ELAN'+ prompt, to verify that all the LE clients rejoined the emulated LAN. After a few minutes all clients had rejoined.

To restart the LES/BUS we used the following command, also from the MSS EXISTING LES-BUS 'ETHERNET_ELAN'+ prompt.

```

ITSO Lab MSS EXISTING LES-BUS 'ETHERNET_ELAN'+restart t6
LES/BUS: 'ETHERNET_ELAN': RESTARTING
    
```

Figure 259. MSS Command to Restart a LES/BUS

Note: The LES/BUS can be restarted using the parameters configured under the running configuration t5 or the saved configuration t6.

We checked the VCCs now established by the DOS client, using the procedure described above.



Figure 260. SVCs Established by DOS Client after LES Restart

Using the MSS you can also list the SVCs set up by a specific LE client using the following command, from the EXISTING LES-BUS 'ETHERNET_ELAN' + prompt. In the example we have viewed the SVCs set up by the OS/2 client.

```

ITSO Lab MSS EXISTING LES-BUS 'ETHERNET_ELAN'+data list specific lec atm
ATM Address []? 39098511111111111111111111111111010302007777000081
LEC ID:                                0x0003
LEC ATM Address:                       39098511111111111111111111111111010302007777000081
Proxy:                                 No
LEC State at LES:                      OPERATIONAL
Entered LES State at:                  00.01.09.29 (System Up Time)
LEC State at BUS:                     OPERATIONAL
Entered BUS State at:                  00.01.09.37 (System Up Time)
Control Direct Vcc:                   OPERATIONAL 0/763
Control Distribute Vcc:               OPERATIONAL 0/738
Multicast Send Vcc:                   OPERATIONAL 0/764
Multicast Forward Vcc:               OPERATIONAL 0/740
MAC Address in Join Req:              02.00.77.77.00.00
Packet Tracing Eligible:             Yes
# ATM Address Mappings:               1
# MAC Address Mappings:               1
No TLVs registered with LEC
    
```

Figure 261. MSS Command to List VCCs for a Specific LEC

Some LE clients also allow you to view their status and the SVCs they set up via internal commands. In our scenario we interrogated the MSS client and 8274 client. The MSS commands must be used from the LEC+ prompt and these commands are shown below.

```

ITSO Lab MSS LEC+mib status

lecStatusTable:
  lecPrimaryAtmAddress                =
                                     39.09.85.11.11.11.11.11.11.11.01.01.40.00.82.10.00.00.03

  lecId                               = 1
  lecInterfaceState                   = Operational 1
  lecLastFailureRespCode              = None
  lecLastFailureState                 = Initial State
  lecProtocol                          = 1
  LecVersion                           = 1
  lecTopologyChange                   = False
  lecConfigServerAtmAddress           =
                                     47.00.79.00.00.00.00.00.00.00.00.00.00.00.A0.3E.00.00.01.00
  lecConfigSource                     = Used well known address
  lecActualLanType                    = 802.3 - Ethernet
  lecActualMaxDataFrameSize           = 1516
  lecActualLanName                    = ETHERNET_ELAN
  lecActualLesAtmAddress               =
                                     39.09.85.11.11.11.11.11.11.11.01.01.40.00.82.10.00.00.02

  lecProxyClient                      = False
    
```

Figure 262. MSS Command to Display the Status and Configuration of an MSS LEC

The command above indicates the status of the LE client **1** and many of its configuration parameters.

```

ITSO Lab MSS LEC+mib server

lecServerVccTable:
  lecConfigDirectInterface      = 0
  lecConfigDirectVpi           = 0
  lecConfigDirectVci           = 0
  lecControlDirectInterface     = 1
  lecControlDirectVpi          = 0
  lecControlDirectVci          = 308    1
  lecControlDistributeInterface = 1
  lecControlDistributeVpi      = 0
  lecControlDistributeVci      = 311    2
  lecMulticastSendInterface     = 1
  lecMulticastSendVpi          = 0
  lecMulticastSendVci          = 312    3
  lecMulticastForwardInterface  = 1
  lecMulticastForwardVpi       = 0
  lecMulticastForwardVci       = 315    4

```

Figure 263. MSS Command to View VCCs for an MSS LEC

The command above shows that the LE client has been connected to the four VCCs (**1** - **4**) it requires for correct operation of the emulated LAN. It also shows the VPI/VCI values used for each connection.

The 8274 command below illustrates the same information for its LE client.

```

/ % vss 4/1 1

  Status/Statistic for slot 4 interface 1 Service 1

Service      : LAN Emulation Service 1

LEC status   : Operational
ELAN Name    : ETHERNET_ELAN
ELAN Type    : 802.3
LEC ID       : 3
LES address  : 390985111111111111111010140008210000002 (learned)
BUS address  : 390985111111111111111010140008210000002
LECS address : 4700790000000000000000000000a03e0000100 (use WellKnown LECS addr)

BUS
MC Forward VPC/VCC : 0/ 821      MC Send VPC/VCC      : 0/ 820
Echo suppress      :          0

LES
Control Direct VPC/VCC: 0/ 818      Cntl Distribute VPC/VCC: 0/ 819
Control Frames Sent  : 14986      Control Frames Rcvd   : 14995
LE arps Sent        :          7      LE arps Received      :          19

LECS
Configuration VPC/VCC : 0/ 0
Packets Sent          :          0      Packet Received       :          0

```

Figure 264. 8274 Command to View a LEC Status and VCCs Established

Finally we examined the working events seen when displaying the LES subsystem in the event logging system. The following events were seen in talk 2 when the DOS client joins with the LES/BUS.

```
LES.087: LES/BUS: 'ETHERNET_ELAN': Ctrl Dir estblshd,  
        Calling ATM addr = x390985111111111111111010302008888000081 1  
LES.256: Trace LAN Emulation Control frame.  
LES.172: LES/BUS: 'ETHERNET_ELAN': adding Non-Proxy Ctrl Dist leaf,  
        LEC ATM addr = x39098511111111111111111010302008888000081 2  
LES.104: LES/BUS: 'ETHERNET_ELAN': Non-Proxy Ctrl Dist leaf estblshd,  
        LEC ATM addr = x39098511111111111111111010302008888000081 3  
LES.256: Trace LAN Emulation Control frame.  
LES.256: Trace LAN Emulation Control frame. 4  
LES.256: Trace LAN Emulation Control frame.  
LES.097: LES/BUS: 'ETHERNET_ELAN': Mcast Send estblshd,  
        LEC ATM addr = x39098511111111111111111010302008888000081 5  
LES.172: LES/BUS: 'ETHERNET_ELAN': adding Non-Proxy Mcast Fwd leaf,  
        LEC ATM addr = x39098511111111111111111010302008888000081 6  
LES.104: LES/BUS: 'ETHERNET_ELAN': Non-Proxy Mcast Fwd leaf estblshd,  
        LEC ATM addr = x39098511111111111111111010302008888000081 7
```

Figure 265. MSS talk 2 Events for the LES Subsystem

From the events above we can see:

- 1 The client establishes a control direct VCC to the LES.
- 2 The LES then attempts to add the non-proxy to its control distribute VCC.
- 3 The LES succeeds in adding the LE client to its VCC.
- 4 The client requests the ATM address of the BUS from the LES.
- 5 The client establishes the multicast send VCC to the BUS.
- 6 The BUS then attempts to add the client to its multicast forward VCC.
- 7 The BUS succeeds in adding the LE client to its VCC.

As an example the events below are seen on the MSS event logging system when the LES is stopped and client's connections fail.

```
LES.122: LES/BUS: 'ETHERNET_ELAN': Non-Proxy Mcast Fwd leaf rlsd: cause 47,  
        LEC ATM addr = x39098511111111111111111010302008888000081  
LES.119: LES/BUS: 'ETHERNET_ELAN': trmtnng LEC:  
        Non-Proxy Ctrl Dist leaf rlsd: cause 47,  
        LEC ATM addr = x39098511111111111111111010302008888000081
```

Figure 266. MSS talk 2 Events for the LES Subsystem

7.8.7 Case (1-4): Microcode/Driver Problem

The microcode used in your network devices and the drivers used on your clients can cause numerous problems when using emulated LAN networks. It is important to check that you are using the latest level to eliminate bugs that may have been fixed from older levels.

The following commands show how you can check the microcode and driver levels you are using.

- 8260 Switch

```
8260ATM1> show device
8260 ATM Control Point and Switch Module
.....
Manufacture id: VIME
Part Number: 58G9605 EC Level: C38846
Boot EEPROM version: v.1.2.0
Flash EEPROM version: v.2.0.4
Flash EEPROM backup version: v.2.0.4
Last Restart : 14:01:53 Mon 9 Jun 97 (Restart Count: 118)
.....
```

Figure 267. 8260/8285 Command to View Microcode Levels

- MSS Server

```
TSO Lab MSS +config

IBM 8210 Nways Multiprotocol Switching Server
Host name: ITSO Lab MSS
Version: 8210-MSS Feature 8706 V1 R1.1 PTF 0 RPQ 0 cc1_17b test-load
.....
```

Figure 268. MSS Command to View Microcode Levels

- Endstations (The following example is the DOS client.)

In endstations you can often see the driver level when the endstation is rebooted.

```
Microsoft Protocol Manager version 2.1
IBM TURBOWAYS 25 ATM ISA Adapter Device Driver Version 2.20
(C)Copyright IBM Corp. 1994-1996. All rights reserved.
ATW0004 TURBOWAYS ATM adapter is using a local address of 020088880000.
ATW0013 ATM Forum-Compliant emulated Ethernet LAN initialization complete.
.....
```

If the device driver isn't loaded correctly, the following message may appear.

```
Microsoft Protocol Manager version 2.1
IBM TURBOWAYS 25 ATM ISA Adapter Device Driver Version 2.20
(C)Copyright IBM Corp. 1994-1996. All rights reserved.
ATW0004 TURBOWAYS ATM adapter is using a local address of 020088880000.
ATW0002 An recoverable error occurred in the TURBOWAYS ATM device driver program.
ATW0025 Press any key continue . . .
```


In our case study the DOS client failed to connect with its emulated LAN once in every ten times which we believe was due to the driver level used. No events were seen on the MSS and so we can conclude the client never attempted to connect and must have failed to initialize its adapter.

7.8.8 Case (1-5): LECS Parameters

In large emulated LAN networks where you have a single LECS supporting a very large number of LE clients you may experience problems with the LECS dropping sessions to clients. This case study investigates how you determine whether you have a problem and will identify some of the parameters you may be able to change to fix the problem.

7.8.8.1 Methodology

The following scenario uses the general methodology described in 7.3.1, “Emulated LAN Connectivity Problem Methodology” on page 310.

First, as described in the previous case studies, look for more common problems by following:

- Stage 1 - Hardware and power problems
- Stage 2 - ATM problems
- Stage 3 - Simple configuration problems
- Stage 4 - Interrogate LE service, VCCs and data flows

Most problems will be solved by following this procedure. Once you have eliminated all the more common possibilities continue by investigating timers and parameters.

- Stage 5 - Investigate timers and parameters

To understand whether you are exceeding the maximum capacity of the LECS you need to look at the LECS statistics. To do this use the statistics list command from the MSS LECS console + command. An example is shown below.

```

ITSO Lab MSS LECS console+statistics list
LECS has:
  1 ELAN(s)
  2 policy(ies) at 2 priority(ies)
LECS has 1 configure direct VCCs
  has accepted 1372 VCCs and rejected 0 VCCs 1
  0 VCCs have been dropped by LECS, 1371 dropped by caller
  has exceeded its maximum number of VCCs 0 times 2
LECS discarded frames: 0
LECS responses by status (zero responses if status not displayed)
  Success( 0) : 6
  No Configuration(20) : 1366

```

Figure 269. MSS Command to View the LECS Statistics

If either the number of rejected VCCs **1** or the number of times the maximum number of VCCs was exceeded **2** are non-zero, you have exceeded the maximum capacity of the LECS in your network.

To determine why VCCs are being rejected, you should look at the following key parameters. These are discussed further in 7.3.1.5, “Stage 5 - Investigate LE Timers and Advanced Parameters” on page 314.

- Maximum number of configuration direct VCCs to a LECS
- LECS VCC idle time
- Validate best effort PCR

- Maximum Configuration Direct VCC reserved bandwidth

To check these parameters you can use the MSS list command from the talk 5 LECS console + prompt.

```
ITSO Lab MSS LECS console+list
Status of LECS:
  ATM device number:          0
  State:                      Operating normally(88)
    Time of last state change: 00.00.33.13
    Elapsed time since last change: 02.53.35.94
  Error Log:                  no err (0)
  Local ATM address:         390985111111111111111010140008210000000
  Well-known address:       4700790000000000000000000000A03E00000100
  UNI version:              UNI Version 3.0
  Validate best effort PCR:  No 1
  Maximum config direct VCC reserved bandwidth: 0 Kbps
  Maximum number of config direct VCCs to LECS: 128 2
  Seconds before VCC declared idle: 60 3
  Trace ATM address value:0000000000000000000000000000000000000000
  Trace ATM address mask: FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
```

Figure 270. MSS Command to List LECS Status and Configuration

We can see that the list command displays the validate PCR value **1**, the maximum VCCs it supports **2** and the time before a VCC is declared idle **3**.

7.8.8.2 Conclusion

In most cases the defaults for these values will be sufficient; however if LECS VCCs are being dropped, try disabling the validate PCR value, disabling the maximum reserved bandwidth, increasing the maximum number of configuration direct VCCs supported and decreasing the VCC idle time. Longer term you may need to split your LAN emulation environment into two separate domains controlled by separate LECS servers.

7.8.9 Case (1-6): LES/BUS Parameters

In a large LAN emulation network, when an LES or BUS supports a large number of LE clients at the same time, the maximum capacity of the LES or BUS or their ATM connections may be exceeded. This case study investigates how you can tell if the capacity of your LES/BUS is being exceeded and will identify some of the key parameters to check that may solve the problem.

7.8.9.1 Methodology

The following scenario uses the general methodology described in 7.3.1, "Emulated LAN Connectivity Problem Methodology" on page 310.

First, as described in the previous case studies, look for more common problems by following:

- Stage 1 - Hardware and power problems
- Stage 2 - ATM problems
- Stage 3 - Simple configuration problems
- Stage 4 - Interrogate LE service, VCCs and data flows

Most problems will be solved by following this procedure. Once you have eliminated all the more common possibilities continue by investigating timers and parameters.

- Stage 5 - Investigate timers and parameters

Use the following command to display statistics associated with LAN emulation MIBs of the LES. This command must be issued from the talk 5 MSS EXISTING LES-BUS 'ETHERNET_ELAN'+ prompt.

```

ITSO Lab MSS EXISTING LES-BUS 'ETHERNET_ELAN'+statistics display les-bus les
ATM Forum LES MIB Statistics:
  joinOK:                                6
  verNotSup:                              0
  invalidReqParam:                        0
  dupLanDest:                             0
  dupAtmAddr:                             0
  insRes:                                  0
  accDenied:                              0
  invalidReqId:                           0
  invalidLanDest:                         0
  invalidAtmAddr:                         0
  badPkts:                                 0
  outRegFails:                            0
  leArpIn:                                 9
  leArpFwd:                               0
Other Statistics:
  leArpAnswers:                           9
  leArpRspFwd:                            0
  topologyFwd:                             63
  narpFwd:                                  0
  flushRspFwd:                             2
  outJoinFails:                            0
  regOK:                                    5
  unRegOK:                                  0
  outUnRegFails:                           0
  proxyLeCs:                               2
  nonProxyLeCs:                            4
  macAddrMappings:                         10
  rdMappings:                              0
  atmAddrMappings:                         6
  joinRetransmits:                         0
  joinParmChanges:                         0
  joinTimeouts:                            0 1
  reRegs:                                   0
  ctDirRefused:                            0 2
  ctDirReleased_err:                       0
  ctDistFailure:                           0
  ctDistReleased_err:                      0
  ctDistPartyReleased_err:                 0
  redundancyVccRefused:                    0
  redundancyVccReleased:                   0
  redundancyVccFailure:                    0
  oam_droppedFrames:                       0
  invalidSize_droppedFrames:               0
  invalidMarker_droppedFrames:             0
  invalidProtocol_droppedFrames:          0
  verNotSup_droppedFrames:                 0
  invalidLecid_droppedFrames:              0
  unknownLecid_droppedFrames:              0
  invalidOpcode_droppedFrames:             0
  dupJoin_droppedFrames:                   0
  incompleteSourceJoin_droppedFrames:      0
  incompleteTargetJoin_droppedFrames:      0
  noProxy_droppedFrames:                   0

```

Figure 271. MSS Command to Display LES Statistics

The most important statistics to check are below:

- joinTimeouts **1**

If the LES is heavily utilized, the joins may not complete in the Control Timer value.

If not, the join fails.

Increase the control timer on the LES to see if this stops join failures.

– **ctldirRefused** **2**

If the LES is refusing control direct VCCs, it could be because it is validating best effort PCR or has a maximum reserved bandwidth set but clients are exceeding the maximums. Disable these parameters to see if the problem is corrected.

The important parameters to check to see if they are causing problems are discussed further in 7.3.1.5, “Stage 5 - Investigate LE Timers and Advanced Parameters” on page 314.

Use the following command to display statistics associated with LAN Emulation MIBs of the BUS. This command must be issued from the talk 5 MSS EXISTING LES-BUS 'ETHERNET_ELAN'+ prompt.

```

ITSO Lab MSS EXISTING LES-BUS 'ETHERNET_ELAN'+statistics display les-bus bus
ATM Forum BUS MIB Statistics:
  inDiscards:                0
  inOctets:                  13688
  inUcastFrms:              3
  inMcastFrms:             192
  frmTimeouts:              0
  mcastSendRefused:         0 1
  mcastFwdFailure:         0
Other Statistics:
  inExplorer:               0
  inFlushReq:               2
  outFlushReq_mcastFwd:    0
  outFlushReq_mcastSend:   2
  outUcastFrms_mcastFwd:   0
  outUcastFrms_mcastSend:  3
  outMcastFrms:            192
  outOctets:                13518
  mcastSendReleased:       0
  mcastFwdReleased:        0
  mcastFwdPartyReleased:   0
  invalidProtocol_droppedFrames: 0
  verNotSup_droppedFrames:  0
  invalidOpcode_droppedFrames: 0
  invalidLecid_droppedFrames: 0
  invalidSize_droppedFrames: 0
  flushToBus_droppedFrames: 0
  incompleteSourceConnect_droppedFrames: 1
  incompleteTargetConnect_droppedFrames: 0
  noProxy_droppedFrames:   0

```

Figure 272. MSS Command to Display BUS Statistics

The most important statistic to check is the **mcastSendRefused** **1**. If the BUS is refusing multicast send VCCs, it could be because it is validating best effort PCR or has a maximum reserved bandwidth set but clients are exceeding the maximums. Disable these parameters to see if the problem is corrected.

7.8.9.2 Conclusion

The cause of a client being dropped from an ELAN when the ELAN is heavily used may be due to some of the timers or parameters in the LES or BUS. Changing these parameters may solve the problem in the short term but in the longer term the emulated LAN will probably have to be split in two.

7.9 Case Studies Involving Problems with Classical IP Network Connectivity

This section investigates how to determine connection problems that can arise when using Classical IP logical IP subnets. We use IBM products to demonstrate our problem determination guidelines.

7.9.1 Network Environment

The test environment used for the Classic IP connectivity case study is described below:




Figure 273. Classical IP Environment (Physical View)

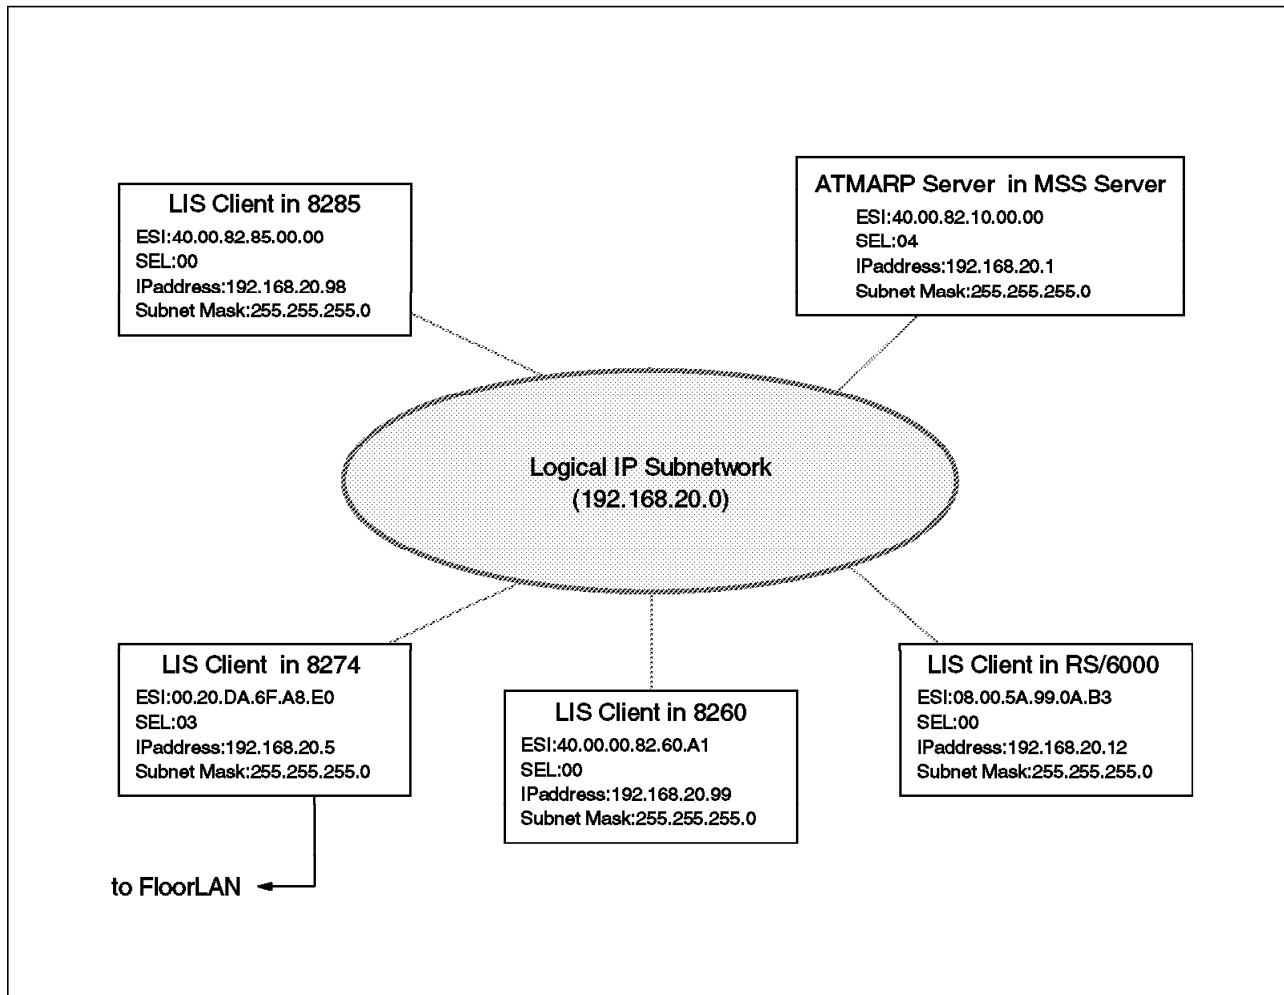


Figure 274. Classical IP Environment (Logical View)

We intend to use a simple Classical IP environment as a case study. In this example, the 8210 MSS Server provides the ATMARP server function. Any other information is as below:

- Single logical IP subnetwork (network address: 192.168.20.0).
- IP addresses are assigned in 8260, 8285, MSS, 8274 and our network management RISC System/6000.
- Every LIS client uses dynamic SVCs and an ATMARP server.
- An SSI connection links the 8260 and 8285.

Note: On the MSS server an ATMARP server performs the role of both ATMARP server and LIS client.

7.9.2 Symptom: Failure with a Logical IP Subnet Client-to-Client Connection

In Classic IP environments most problems occur when one LIS client can't communicate with other clients. To determine the cause of the problem, you must verify which method you are using to establish ATM VCCs between your clients. The available methods are discussed in 7.2.5, "Classical IP (RFC 1577) Rules" on page 303. Using SVCs minimizes the setup required on each client but requires an ATMARP server. Using PVCs requires all the inter-linking PVCs to be defined in the ATM switches in your network. PVCs are typically only used when the number of LIS clients is very small.

In our case studies we used SVCs and an ATMARP server. We intend to describe how to fix problems in this more common environment.

Note

The use of SVCs and an ATMARP server is recommended as it simplifies the installation, operation and management of your logical IP subnet subnetwork.

7.9.3 Troubleshooting Methodology in Classical IP

The following diagram is intended to help you diagnose the cause of your problems with Classic IP connectivity.

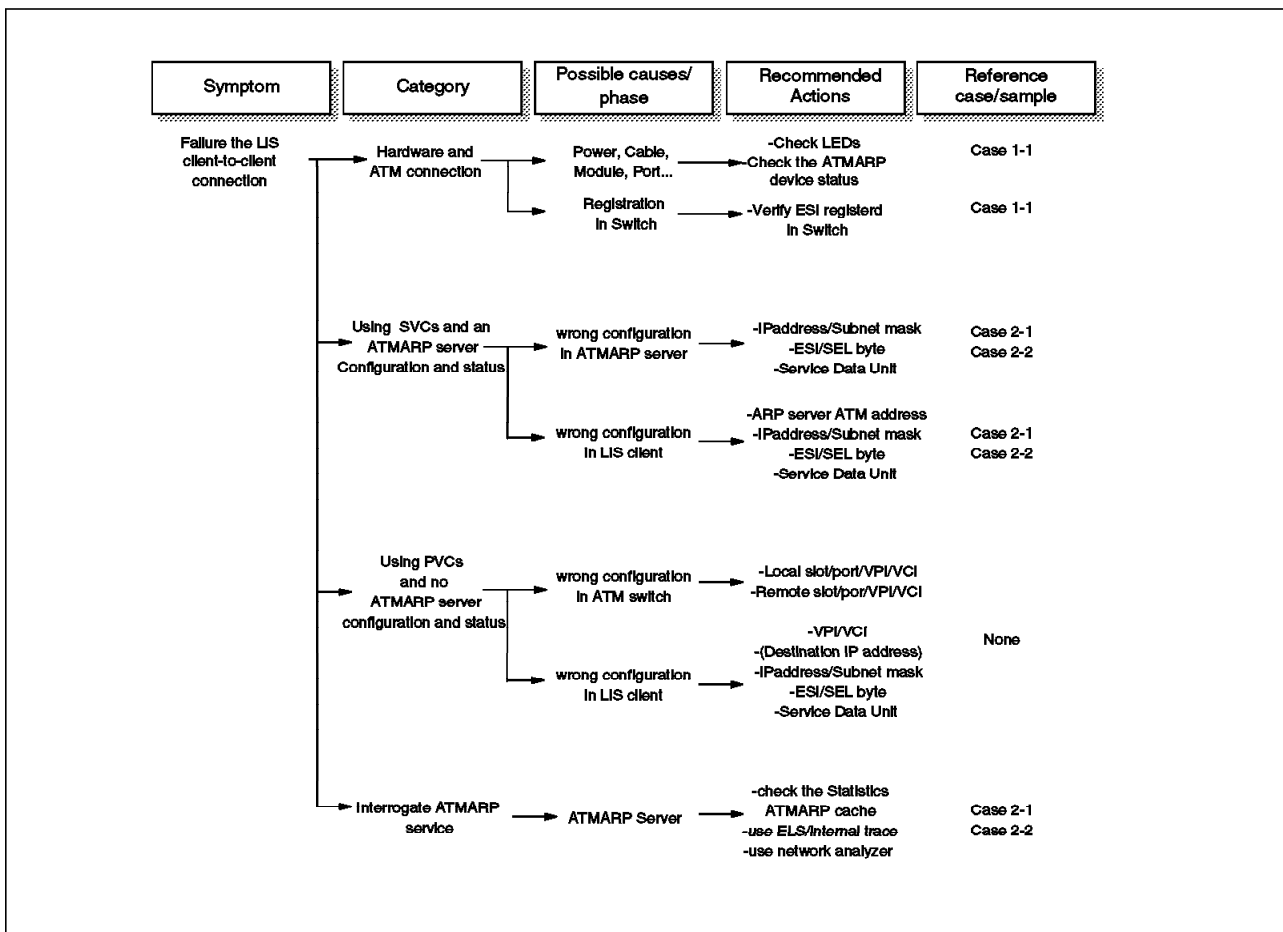


Figure 275. Troubleshooting Methodology for Classical IP Connectivity

7.9.4 Case (2-1): ATMARP Server Failure Problem

In our scenario the ATMARP server in the MSS server couldn't communicate with any other LIS clients. We tried to ping from the RS/6000 to ATMARP server, but couldn't communicate.

We checked the VCCs set up by clients, using IBM Nways Campus Manager as described in 7.8.6, "Case (1-3): LES/BUS Failure Problem" on page 360 and found that all clients still had SVC connections to the ARP server.

While pinging the ATMARP server we enabled the MSS event logging system to display ARP events. To use the event logging system we issued a `nodisplay sub all` command followed by the `display sub arp all`. For more information on using the event logging system see 7.8.4, "Case (1-1): LECS Failure Problem" on page 342. In talk 2 we saw the following events.

```
ARP.016: unkn dst prot ad nt 0 int ATM/0
ARP.051: ATM CIP atmArpRcvFrame: (prot = 800) nt 0 int ATM/0 1
ARP.051: ATM CIP atmArpRcvFrame: (prot = 800) nt 0 int ATM/0 2
ARP.051: ATM CIP atmArpRcvFrame: (prot = 806) nt 0 int ATM/0 3
ARP.002: Pkt in 1 13 800 nt 0 int ATM/0
ARP.099: Trace ARP/ATMARP frame
.....
```

Figure 276. MSS talk 2 Events for the ARP Subsystem

The RS/6000 still has a connection to the ATMARP server and sends an ARP_REQUEST frame to resolve the ATM address of the ATMARP server's IP address **1**. No ARP_REPLY or ARP_NAK frame is sent out by the MSS ATMARP server so it sends subsequent ARP_REQUESTs to resolve the address **2** and **3**.

7.9.4.1 Methodology

The following scenario uses the general methodology described in 7.3.2, "Classic IP Connectivity Problem Methodology" on page 320.

1. Stage 1 - Hardware and power problems

All LEDs looked normal, so we concluded devices were working and powered on successfully. We checked the status of the IBM Nways Multiprotocol Switched Services (MSS) server, as described in 7.8.4, "Case (1-1): LECS Failure Problem" on page 342 and found it was working correctly.

2. Stage 2 - ATM problems

We then checked the status of the ATM modules and ports on the IBM Nways 8260 Multiprotocol Switching Hub and confirmed the ATM connectivity of devices using the `SHOW MODULE` and `SHOW PORT` commands. See 7.8.4, "Case (1-1): LECS Failure Problem" on page 342 for more information. We also checked that all clients were registered with their ATM switch using the `SHOW ATM_ESI ALL` command on the 8260 and 8285. The output from using this command on the 8260 is shown below:

```

8260ATM1> show atm_esi all
Port    ATM_ESI                Type
-----
12.01   11.11.11.11.11.11     static (id= 1)
6.02    50.00.00.82.82.A1     dynamic
7.01    40.00.82.81.0B.0B     dynamic
12.01   00.20.DA.6F.A8.E0     dynamic 1
12.02   08.00.5A.99.81.51     dynamic
13.04   08.00.5A.99.0A.B3     dynamic 2
14.01   40.00.00.60.00.01     dynamic
15.01   40.00.82.81.0A.0A     dynamic
16.01   40.00.82.10.00.00     dynamic 3

```

Figure 277. 8260 Command to Verify ATM Addresses Registered with the Switch

- 1 - 8274
- 2 - RS/6000
- 3 - MSS ATMARP client/server

Note: None of the other addresses shown are used in this case study.

We concluded all devices were registered with their ATM switch.

3. Simple configuration

We checked the configuration and status of the ATMARP clients and server.

- RS/6000 LIS client configuration

To check the configuration of the RISC System/6000 enter the following from the command line:

- Enter smitty tcpip
- Select Minimum Configuration & Startup
- Select at0 ATM Network Interface

The configuration can then be seen, as shown in the screen below.

```

Minimum Configuration & Startup

To Delete existing configuration data, please use Further Configuration menus

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

* HOSTNAME [Entry Fields]
[marco]
* Internet ADDRESS (dotted decimal) [192.168.20.12] 1
Network MASK (dotted decimal) [255.255.255.0] 2
* Network INTERFACE
at0
NAMESERVER
Internet ADDRESS (dotted decimal) [9.24.104.108]
DOMAIN Name [itso.ral.ibm.com]
Default GATEWAY Address [9.24.104.1]
(dotted decimal or symbolic name)
Connection Type svc_c 3
ATM Server Address [39.09.85.11.11.11.11.11.11.11.01.01.40.00.82.>] 4
Alternate Device []
Idle Timer [60]
Best Effort Bit Rate (UBR) in Kbits/sec [0]
START TCP/IP daemons Now no

```

Note: When you specify `svc_c` as the virtual connection type to be configured for this network interface (`at0`), you must specify the ATM 20-byte address of the client's designated ATMARP server in the ATM Server Address field. This must be the full 20-byte ATM address with each byte separated by periods.

We checked the following:

- **1** IP address
- **2** Subnet mask
- **3** Connection type which should be `svc_c` for an AIX 4.x client

Note: An RS/6000 can also be set up as an ATMARP server, in which case the connection type would be `svc_s`

- **4** ATMARP server's ATM address

All values looked correct. We also verified the maximum PDU size used by the client. To check the configuration of the RISC System/6000 ATM interface use the following:

- a. Enter `smitty atm`
- b. Select Adapter
- c. Select Change / Show Characteristics of an ATM Adapter
- d. Select the ATM adapter `atm0 Available 00-03 100 Mbps ATM Fiber Adapter (8f7f)`

```
Change / Show Characteristics of an ATM Adapter

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

[TOP]                                     [Entry Fields]
Device name                               atm0
Description                               100 Mbps ATM Fiber Adapter
Status                                    Available
Location                                  00-03
Best Effort Peak Rate                     [1500]
Enable ALTERNATE ATM MAC address          no
ALTERNATE ATM MAC address (12 hex digits) [0x]
ATM Adapter PDU Size (bytes)              [9188] 1
ATM Adapter DMA bus width (bytes)         [0x1000000]
Maximum Small ATM mbufs                   [50]
Maximum Medium ATM mbufs                 [100]
Maximum Large ATM mbufs                   [300]
Maximum Huge ATM mbufs                    [50]
Maximum MTB ATM mbufs                     [4]
MTB ATM mbufs size                        [64]
Minimum Small ATM mbufs                   [20]
Minimum Medium ATM mbufs                  [30]
Minimum Large ATM mbufs                   [70]
Minimum Huge ATM mbufs                    [5]
Minimum MTB ATM mbufs                     [4]
Minimum ATM interface buffer size         [2048]
Minimum Guaranteed VCs Supported          [8]
Maximum number of VCs Needed              [32]
.....
```

The maximum PDU was the default 9188 bytes **1**.

- 8210 ATMARP server configuration

The ATMARP server defined on the MSS server is both a logical IP subnet client and ATMARP server. To verify the configuration of ATMARP server you can use the CONFIG process (talk 6) on the MSS. The GWCON monitoring process (talk 5) provides statistics for the ATMARP server but little information about its configuration.

We first checked the IP address and subnet mask used by the MSS. The command below shows how to do this from the (talk 5) IP> prompt.

```
ITSO Lab MSS IP>int
Interface IP Address(es) Mask(s)
ATM/0 192.168.20.1 255.255.255.0 1
Eth/0 192.168.5.11 255.255.255.0
```

The IP address and subnet mask were correct **1**.

We next checked the detailed configuration of the ATMARP server from the MSS ARP config> prompt.

```
ITSO Lab MSS ARP config>list atm-arp-client-configuration

ATM Arp Clients:
-----
If: 0 Prot: 0 Addr: 192.168.20.1 ESI: 40.00.82.10.00.00 Sel: 04 1
Server: no 2 Refresh T/O: 5 AutoRefr: no By InArp: yes Validate PCR: no
Use Best Effort: yes/yes (Control/Data) Max B/W(kbps): 0
Cell Rate(kbps): Peak: 155000/155000 Sustained: 0/ 0
Max SDU(bytes): 9188 3
```

We checked the following:

- **1** Whether the ESI and selector byte were correct.
- **2** Whether the ATMARP server was enabled.
- **3** The Max SDU size value.

We noticed that the ESI, selector and Max SDU were correct but the ATMARP server was not enabled.

4. Stage 4 - Interrogate ATMARP service, VCCs and data flows

We re-enabled the MSS ATMARP server function using the following MSS command from the (talk 6) ARP Config> prompt and reloaded the MSS.

```
ITSO Lab MSS ARP config>change ATM-ARP-CLIENT-CONFIG
Interface Number [0]?
Protocol [IP]?
Client IP Address [0.0.0.0]? 192.168.20.1
This client is also a server? [Yes]: yes
Refresh timeout (in minutes) [5]?
Enable auto-refresh? [No]:
Refresh by InAtmArp? [Yes]:
  ( 1) Use burned in ESI
  ( 2) 400082100000
Select ESI [2]?
Use internally assigned selector? [No]:
Selector Only, Range 00..FF [04]?
Validate PCR for best effort VCCs? [No]:
Maximum Reserved Bandwidth for incoming VCCs (Kbps) [0]?
Use Best Effort Service for Control VCCs? [Yes]:
Peak Cell Rate of outbound control VCCs (Kbps) [155000]?
Sustained Cell Rate of outbound control VCCs (Kbps) [0]?
Use Best Effort Service for Data VCCs? [Yes]:
Peak Cell Rate of outbound Data VCCs (Kbps) [155000]?
Sustained Cell Rate of outbound Data VCCs (Kbps) [0]?
Max SDU size (bytes) [9188]?
```

After the MSS has been reIMLed and the ATMARP server function is working, you can use the following command, from the MSS (talk 5) ARP> prompt, to access the ATM channel and ATMARP table entries.

```
ITSO Lab MSS ARP>dump 0
Hardware Address      IP Address           Refresh
0/734                 192.168.20.5        4
0/717                 192.168.20.99       5
0/731                 192.168.20.98       3
0/732                 192.168.20.12       4
```

Note: The Refresh time is the number of minutes until the ATMARP entry is aged out. If autorefresh is turned on, then an InARP request will be sent out 30 seconds before the expiration. If a reply is received before the expiration, the Refresh time is reset and the ARP entry remains. If no reply is received or autorefresh is turned off, the ARP entry will be deleted when it expires.

You can also use the following command to display all the active VCCs to the ATMARP server's ATM address and the parameters each client uses. The command must also be issued from the MSS (talk 5) ARP > prompt.

```

ITSO Lab MSS ARP>display 0
Active Channel List : Net 0
  P/S FLAGS LIST VPI/VCI FwdPcr FwdScr MaxSDUsz Control P2P
1) S 00 01 0/734 149760192 149760192 9188 F T
   Tgt Addr: 39.09.85.11.11.11.11.11.11.11.01.01.00.20.DA.6F.A8.E0.03
   Client Address (owner): 192.168.20.1
   Target Protocol Addresses: 192.168.20.5
2) S 00 01 0/717 101760000 101760000 9188 F T
   Tgt Addr: 39.09.85.11.11.11.11.11.11.11.01.01.40.00.00.82.60.A1.00
   Client Address (owner): 192.168.20.1
   Target Protocol Addresses: 192.168.20.99
3) S 00 01 0/731 101760000 101760000 9188 F T
   Tgt Addr: 39.09.85.11.11.11.11.11.11.11.01.03.40.00.82.85.00.00.00
   Client Address (owner): 192.168.20.1
   Target Protocol Addresses: 192.168.20.98
4) S 00 01 0/732 42400000 42400000 9188 F T
   Tgt Addr: 39.09.85.11.11.11.11.11.11.11.01.01.08.00.5A.99.0A.B3.00
   Client Address (owner): 192.168.20.1
   Target Protocol Addresses: 192.168.20.12
New Channel List : Net 0
PVC Channel List : Net 0

```

We can see that each client has registered with the ATMARP server, that the ATMARP server has learned their IP and ATM addresses and the SDU used by each client.

Finally we check the ATMARP cache for one of the clients (RS/6000).

```

> arp -t atm -a

SVC - at0 on device atm0 -
=====
at0(192.168.20.12) 39.9.85.11.11.11.11.11.11.11.1.1.8.0.5a.99.a.b3.0
IP Addr VPI:VCI Handle ATM Address
?(192.168.20.1) 0:732 5 39.9.85.11.11.11.11.11.11.11.1.1.40.0.82.10.0.0.4

```

Figure 278. ATMARP ATM Connections on the RS/6000

The RS/6000 displays that it has successfully resolved the ATM address and IP address for its ATMARP server.

7.9.4.2 Conclusion

We found that when the MSS ATMARP server was restarted all clients re-joined the logical IP subnet within a few minutes, but the RS/6000 required rebooting before it would join.

To check what is happening at the ATMARP server once it is working we once again enabled the MSS ARP event logging and observed the events below using (talk 2) on the MSS. To enable event logging for the ARP subsystem we issued a `nodisplay sub all` command followed by the `display sub arp all` from the ELS > prompt. For more information on using the event logging system see 7.8.4, “Case (1-1): LECS Failure Problem” on page 342.


```
ARP.051: ATM CIP atmArpRcvFrame: (prot = 800) nt 0 int ATM/0
ARP.051: ATM CIP atmArpRcvFrame: (prot = 800) nt 0 int ATM/0
ARP.051: ATM CIP atmArpRcvFrame: (prot = 800) nt 0 int ATM/0
ARP.095: ATM CIP: InArp Req sent 0/532 800 nt 0 int ATM/0 1
ARP.099: Trace ARP/ATMARP frame
ARP.051: ATM CIP atmArpRcvFrame: (prot = 806) nt 0 int ATM/0 2
ARP.002: Pkt in 9 13 800 nt 0 int ATM/0
ARP.099: Trace ARP/ATMARP frame
ARP.092: ATM CIP: Mk ent 800/192.168.20.99 nt 0 int ATM/0 3
.....
```

Figure 279. MSS talk 2 Events for the ARP Subsystem

We observed the following:

- **1** The ATMARP server now sends an InARP request to a Classical IP client over VCC VPI/VC1=0/532.
- **2** The client responds with an InARP reply. Protocol 806 in Ethernet Type format is an ARP frame.
- **3** The ATMARP server registers the source IP address (192.168.20.99) and its ATM address into its ATMARP table.

The ATMARP registration is complete. The client can now make ARP_REQUESTS from the ATMARP server and once it receives a reply, it can set up a direct VCC directly to the client it wishes to communicate with. Since all subsequent data transfer uses the direct VCC there is no load on the ATMARP server.

7.9.5 Case (2-2): Wrong Subnet Mask in LIS Client

In this scenario we changed the IP address of the IBM Nways 8274 LAN RouteSwitch from 192.168.20.5 to 192.168.20.133. The ATMARP server in the MSS then couldn't communicate with the IBM Nways 8274 LAN RouteSwitch. In the ATMARP cache on the MSS we see an entry for the new IP address for the IBM Nways 8274 LAN RouteSwitch. This can be seen using the *dump* command from the MSS talk 5 ARP > prompt.

```
ITSO Lab MSS ARP>dump 0
Hardware Address      IP Address           Refresh
0/913                 192.168.20.98       3
0/928                 192.168.20.133     4
0/1010                192.168.20.12      5
```

However if we ping the IBM Nways 8274 LAN RouteSwitch address from the ATMARP server we get no response. The *ping* command can also be issued from the ARP > prompt.

```
ITSO Lab MSS ARP>ping 192.168.20.133
PING 192.168.20.1 -> 192.168.20.133: 56 data bytes, ttl=64, every 1 sec.

----192.168.20.133 PING Statistics----
17 packets transmitted, 0 packets received, 100% packet loss
```

We also tried to ping the MSS ATMARP server from the IBM Nways 8274 LAN RouteSwitch and received an error message of network unreachable.

To see what was happening we enabled the MSS event logging system for logging ARP events and restarted the IBM Nways 8274 LAN RouteSwitch. To enable event logging for ARP events, issue a *nodisplay sub all* followed by a *display sub arp all* at the MSS talk 5 ELS > prompt. Then view the events in the talk 2 process.

```
ARP.040: ATM CIP ReceiveCall: CInt prot/addr 800/192 nt 0 int ATM/0
ARP.095: ATM CIP: InArp Req sent 0/928 800 nt 0 int ATM/0
ARP.099: Trace ARP/ATMARP frame
ARP.051: ATM CIP atmArpRcvFrame: (prot = 806) nt 0 int ATM/0
ARP.002: Pkt in 9 13 800 nt 0 int ATM/0
ARP.099: Trace ARP/ATMARP frame
ARP.092: ATM CIP: Mk ent 800/192.168.20.133 nt 0 int ATM/0
```

Comparing this set of events with the working example in the conclusion of 7.9.4, "Case (2-1): ATMARP Server Failure Problem" on page 381 we see that the client is registering with the ATMARP server correctly.

7.9.5.1 Methodology

The following scenario uses the general methodology described in 7.3.2, "Classic IP Connectivity Problem Methodology" on page 320.

1. Stage 1 and 2 - Hardware, power and ATM problems

We verified the status of all devices and their ATM connections, as in the previous case studies and found no problems.

2. Simple configuration

- IBM Nways 8274 LAN RouteSwitch LIS client configuration

First check the configuration of the IBM Nways 8274 LAN RouteSwitch using the commands below.

```

/ % vas

                                ATM Services

Slot Port Serv      Service      Service
  Num  Num  Num      Description  Type
====  ==  ==  =====
4    1    1    lec for gp 2    LANE
4    1    2    LAN Emulation Service 2    LANE
4    1    3    Classical IP Service 3    CIP 1
4    2    1    PTOP Bridging Service 1    PTOP Priv
    
```

We can see that the Classical IP service **1** is defined as service 3 on slot4 / port1. We checked the configuration of this service.

```

/ % mas 4/1 3

Slot 4 Port 1 Service 3 Configuration

1) Description (30 chars max)      : Classical IP Service 3
2) Classical IP Group              : 3
3) ARP Server address (40-char-hex) : 390985111111111111111010140008210000004 1
4) Admin Status { disable(1),
                  enable(2) }      : Enable
6) Connection Type { PVC(1),
                    SVC(2) }      : SVC 2
60) SEL for the ATM address        : 03 3

Enter (option=value/save/cancel) : save
Modifying service, please wait...

Resetting service, please wait...
    
```

We checked the following:

- **1** The ATM address of the ATMARP server
- **2** The connection type to ensure it was SVC
- **3** The selector byte of the IBM Nways 8274 LAN RouteSwitch LIS client

All items looked correct so we checked the IP configuration of the IBM Nways 8274 LAN RouteSwitch using the following command.

```
/ % modv1 3
Current values associated with GROUP 3.1 are as follows:

1) GROUP Number      - 3:1
2) Description       - for CIP
3) ATM CIP enabled   - Y
4) IP Network Address - 192.168.20.133    1
5) IP Subnet Mask    - 255.255.255.128   2
6) IP Broadcast Address - 192.168.20.255
7) Router Description -
8) RIP Mode          - Silent
   {Active(a), Inactive(i), Deaf(d), Silent(s)}

(save/quit/cancel)
: save
```

Note: The group ID for the Classic IP service was group 3.

We checked the following:

- **1** The IP address used by the IBM Nways 8274 LAN RouteSwitch
- **2** The subnet mask used

We found that the IP subnet mask was incorrect. It should have been 255.255.255.0. This would explain the network unreachable message we received when pinging the MSS from the IBM Nways 8274 LAN RouteSwitch.

We compared the client's configuration with the configuration of a working client, the 8260. To do this we used the show device command on the 8260.

```
8260ATM1>show device

A-CPSW
-----
ATM address: 39.09.85.11.11.11.11.11.11.11.01.01.40.00.00.82.60.A1.00

> Subnet atm: Up
IP address: 192.168.20.99    1
Subnet mask: FF.FF.FF.00   2
Default Gateway : OK
-----
IP address: 192.168.20.1

ARP Server:                3
-----
ATM address: 39.09.85.11.11.11.11.11.11.11.01.01.40.00.82.10.00.00.04

Diagnostics: ENABLED
```

The IP address **1**, subnet mask **2** and ATMARP server address **3** are all shown.

3. Stage 4 - Interrogate ATMARP service, VCCs and data flows

A useful method to investigate ATMARP problems on the 8274 is to check its Classical IP statistics. These can be viewed using the vss command.

```
/ % vss 4/1 3

Status/Statistic for slot 4 interface 1 Service 3

Service      : Classical IP Service 3

From IP:

  Packets Received = 170      Broadcast Packets Received = 0
  Packet Discarded = 0

To IP:

  Packets Sent      = 149

From net:

  Packets Received = 240      Packets Discarded      = 0
  ARP Response     = 4        ARP Request             = 0
  Inv ARP Response = 0        Inv ARP Request         = 87
  Negative ARP Reply = 0

To net:

  Packets Sent      = 253      Packet Discarded        = 4
  ARP Response     = 0        ARP Request             = 4
  Inv ARP Response = 87        Inv ARP Request         = 0
  ARP Acknowledge  = 0
```

These allow you to see whether any data is flowing between the ATMARP server and the 8274 and display the number of ARP Request/Response, InARP Request/Response and Negative ARP Reply (ARP_NAK) messages. From the statistics above we can see data is being sent to and from the IBM Nways 8274 LAN RouteSwitch. It is receiving and responding to InARP requests and has made some ARP requests and received some responses. We can conclude therefore that the problem is not with the client's connection to the ATMARP server but with some other feature, in this case its subnet mask.

On the MSS we can also check the ARP statistics using the statistics command from the talk 5 ARP > prompt.

```
ITSO Lab MSS ARP>statistics

ARP input packet overflows
  Net  Count
  ATM/0 0
  Eth/0 0
  IPPN/0 0
  BDG/0 0

ARP cache meters
Net Prot  Max Cur Cnt  Alloc Refresh: Tot  Failure TMOs: Refresh  Pending
  0  0    1  1  2    4    0    0    2    0
```

From these you can see the count of entries deleted from the ATMARP table due to a timeout of the refresh timer (Refresh), check for overflows (TMOs) and failures (Failure) due to unavailable internal resources.

7.9.5.2 Conclusion

Even if you specify the wrong IP parameters in your LIS client, the LIS client is still registered with the ATMARP server if you specify the correct ATM address of ATMARP server. A useful check would be to see what VCCs have been established. If VCCs to the ARP server have been established, the ATMARP server is working on the ATM network (but may not be functioning as an ATMARP server) and the clients are configured with the ATMARP server's correct ATM address.

7.10 Case Studies Involving Client-to-Client Data Transfer

Once you have established that clients can connect to their emulated LAN or logical IP subnet, problems may still occur when they attempt to transfer data. These problems are often with the higher layer protocols used to communicate between devices. Analysis of higher layer protocols is outside the scope of this troubleshooting redbook; however, we intend to demonstrate how console commands, IBM Nways Campus Manager and a network analyzer can be used to help you gather the necessary information to diagnose data transfer problems.

All clients are initially connected to the emulated LAN.

7.10.1 Case (3): LEC-LEC Connection

Once LE clients have connected with their emulated LAN and wish to communicate with another LE client they will request their partner's ATM address from the LES (by passing the LES the partner's MAC address). They then cache their partner's address in their LE_ARP table and establish a data direct VCC to the partner. They use the flush protocol to switch over unicast data from the BUS to the data direct VCC and then send data via the direct connection.

For this scenario we used the same network environment as described for the LAN emulation case studies. For more information refer to 7.8.1, "Network Environment" on page 339. To determine whether all clients are connected to the emulated LAN you may use the following MSS command.

```

ITSO Lab MSS_EXISTING LES-BUS 'ETHERNET_ELAN'+databases list all lec
Number of LEC's to display: 6

      LEC-LES and LEC-BUS State (UP=Up, ID=Idle, --. --.
      **=Other; Show specific LEC to see actual)          v  v
LEC Primary ATM Address          Proxy ID  LEC State #ATM #Reg #Lrnd
----- -- -- -- -- -- -- -- -- -- -- -- -- -- -- --
39098511111111111111111010140008210000003  N 0001  UP  UP    1    1    0
39098511111111111111111010108005A99815181  N 0002  UP  UP    1    1    0
3909851111111111111111101010020DA6FA8E000  Y 0003  UP  UP    1    1    0
39098511111111111111111010302007777000081  N 0004  UP  UP    1    1    0
39098511111111111111111010302008888000081  N 0005  UP  UP    1    1    0
390985111111111111111110101400082810A0A00  Y 0006  UP  UP    1    5    0
    
```

Figure 280. MSS Command to Determine Which Clients Are Connected to the Emulated LAN

The MAC addresses of the 8274 and MSS are shown below:

- MSS server MAC address: 40.00.82.10.00.00
- 8274 MAC address: 00.20.DA.75.8E.D0

7.10.1.1 Methodology

The methodology used is consistent with the general methodology described in 7.3.3, "Data Transfer Problem Methodology" on page 324.

- Stage 1 - Check LE_ARP tables

Entries in the LE_ARP table are aged and must be periodically refreshed. Each entry is refreshed when a data frame is received from its LAN

destination. The LE client also attempts to refresh entries in the absence of data frames.

To verify entries in the LE_ARP table console commands may sometimes be used. The following commands are examples of the commands that may be used for various clients and the LES. A ping request was issued from the MSS LE client to the 8274 LE client prior to checking the LE_ARP tables.

- MSS LE client

Use the list arp command from the talk 5 LEC + prompt.

```

ITSO Lab MSS LEC+list arp

          LEC Address Resolution Table

Max Table Size      = 10  1
Free Table Entries  = 8   2
Current Mac Entries = 2   3
Current RD Entries  = 0
Arp Aging Time     = 300 4
Verify Sweep Interval = 60

          Xmit  BUS  Arp
          Conn  Queue Frame Retry Aging
MAC Address  Remote Handle Depth Count Count Timer
          Destination ATM Address
-----
00.20.DA.75.8E.D5 True   968   0    1    0   240 5
          39.09.85.11.11.11.11.11.11.11.01.01.00.20.DA.6F.A8.E0.00
02.00.77.77.00.00 False  246   0    1    0    0
          39.09.85.11.11.11.11.11.11.11.01.03.02.00.77.77.00.00.81
    
```

Figure 281. MSS Command to Display the LE_ARP Cache for the MSS LEC

The display shows the LE_ARP table size **1**, the number of free **2** and used entries **3**, the aging timer **4** and the entries in the table including the aging timer value for the amount of time they have been inactive **5**.

- 8274 console

The vlat command may be used on the IBM Nways 8274 LAN RouteSwitch to view its LE_ARP entries.

```

/ % vlat 4/1 1

          ATM LANE LE_ARP Table

MAC Address      ATM Network Prefix      ESI      SEL VPI/VCI  Age Remote
=====
020077770000 390985111111111111111111111111110103020077770000081  0/ 830 10 False
400082100000 39098511111111111111111111111111010140008210000003  0/ 827 11 False
020088880000 390985111111111111111111111111110103020088880000081  0/ 825 140 False
    
```

Figure 282. 8274 Command to Verify the LE_ARP Cache for the 8274 LEC

Note: Our LE client was defined on service 1 of interface 4/1.

On the 8274 the table entries and their aging timer values are also shown.

If clients do not contain an entry for the client they are attempting to communicate with, they have not successfully completed the LE_ARP process. Numerous problems can cause this but the most common are due to the higher layer protocol used not being bound to one of the client's network adapters or the clients are using different encapsulation types.

- Stage 2 - Check data direct VCCs

On all clients the utilization of data direct VCCs is also monitored and the VCCs are released if there is no traffic for the VCC timeout period, which is 20 minutes in MSS LEC. Data direct VCCs are released in a least-recently used manner when establishment of a new data direct VCC fails due to insufficient resources being available.

Make sure the data direct VCC between communicating clients has been set up correctly.

On the MSS LEC you can verify the entries of data direct VCCs, using the following command from the MSS talk 5 LEC + prompt.

```

ITSO Lab MSS LEC+list data

      LEC Data Direct VCCs

Max Table Size      = 100  1
Current Size        = 2    2
Max Conn Handles    = 1024
Inactivity Timeout  = 1200 3
Sweep Interval      = 60

Conn      Inactive  User
Handle    VPI   VCI   Timer   Count   Destination ATM Address
-----
222      0    524   0        1
          39.09.85.11.11.11.11.11.11.11.11.01.03.02.00.77.77.00.00.81

251      0    553   0        1
          39.09.85.11.11.11.11.11.11.11.11.01.03.02.00.88.88.00.00.81
    
```

Figure 283. MSS Command to Display the VCCs for the MSS LEC

The display shows the maximum table size **1**, its current size **2**, the inactivity timer setting **3** and the table entries with the time they have been inactive.

For other clients, IBM Nways Campus Manager allows you to display the connections between ATM devices and hence verify if two clients have successfully set up a data direct VCC. To do this double-click on the following: **NetView for AIX Root Map -> ATM Campus -> Cluster Number -> Hub Symbol**

Then:

1. Select the port, with the device you are interested in, with the right mouse button.
2. Choose **Configuration** from the menu that appears and the ATM Connection Configuration window will appear.

3. Choose **List** from the SVC menu and the current list of SVCs is displayed in the ATM SVC List window.

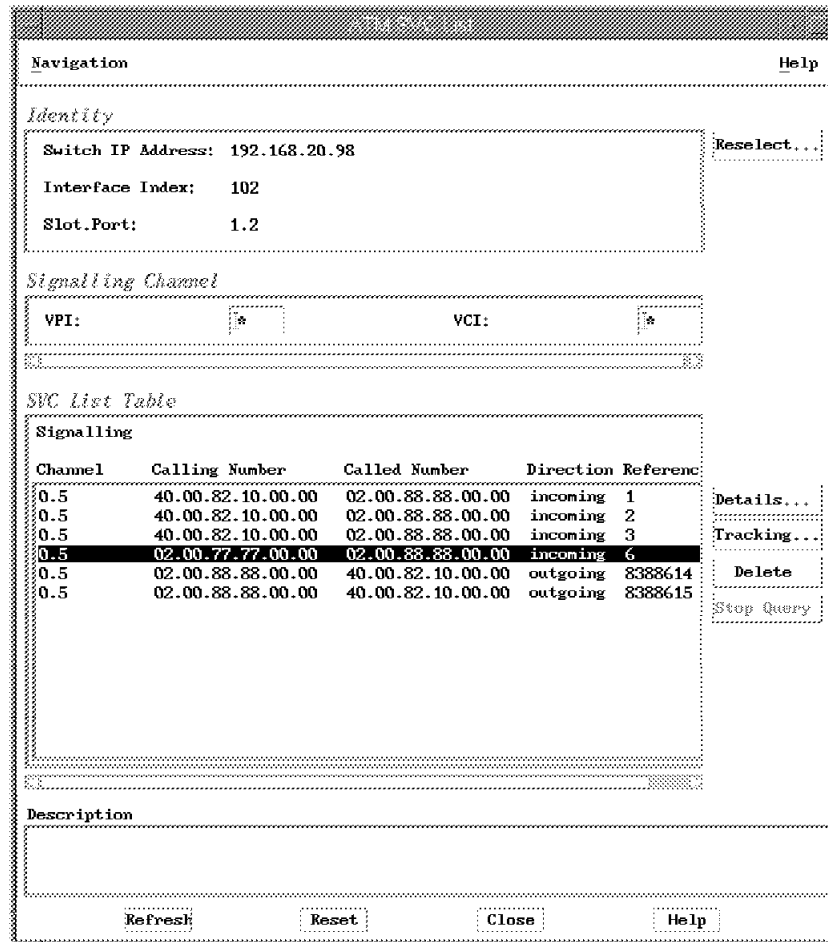


Figure 284. ATM SVC List in Campus Manager

You can then display the detailed information for a particular SVC by doing the following:

1. Select the SVC you are interested in.
2. Select the **Details...** button.

ATM SVC Details			
<u>Navigation</u>		<u>Help</u>	
<i>Identity</i>			
Switch IP Address:	192.168.20.99		
Interface Index:	1201		
Slot.Port:	12.1		
<i>Selection</i>			
Signalling Channel:	0.5	Call Reference:	27
VPI:	0	VCI:	741
<i>Direction</i>			
SVC Direction:	incoming		
<i>Calling Number</i>			
Network Prefix Part:	DCC/DFI/AA=0985/11/111111 RD=1111 AREA=01.01		
End System Part:	ESI=40.00.82.10.00.00 SELECTOR=02		
<u>Called Numbers</u>		<u>/Creation</u>	
DCC/DFI/AA=0985/11/111111 RD=1111 AREA=01.01 ESI=00.20.da.6f.a8.e0 SELECTOR=02			
<i>Parameters</i>			
Forward Traffic		Backward Traffic	
Type:	Best-Effort	Type:	Best-Effort
Service:	unspecified	Service:	unspecified
Parameters		Parameters	
no parameter		no parameter	
<i>Description</i>			
<input type="button" value="Refresh"/>		<input type="button" value="Close"/>	
<input type="button" value="Help"/>			

Figure 285. ATM SVC Show in Campus Manager

You can also view an end-to-end view of any SVC including the inter-hub links it uses by doing the following:

1. Select the SVC you are interested in from the ATM SVC List window.
2. Select the **Tracking...** push button.

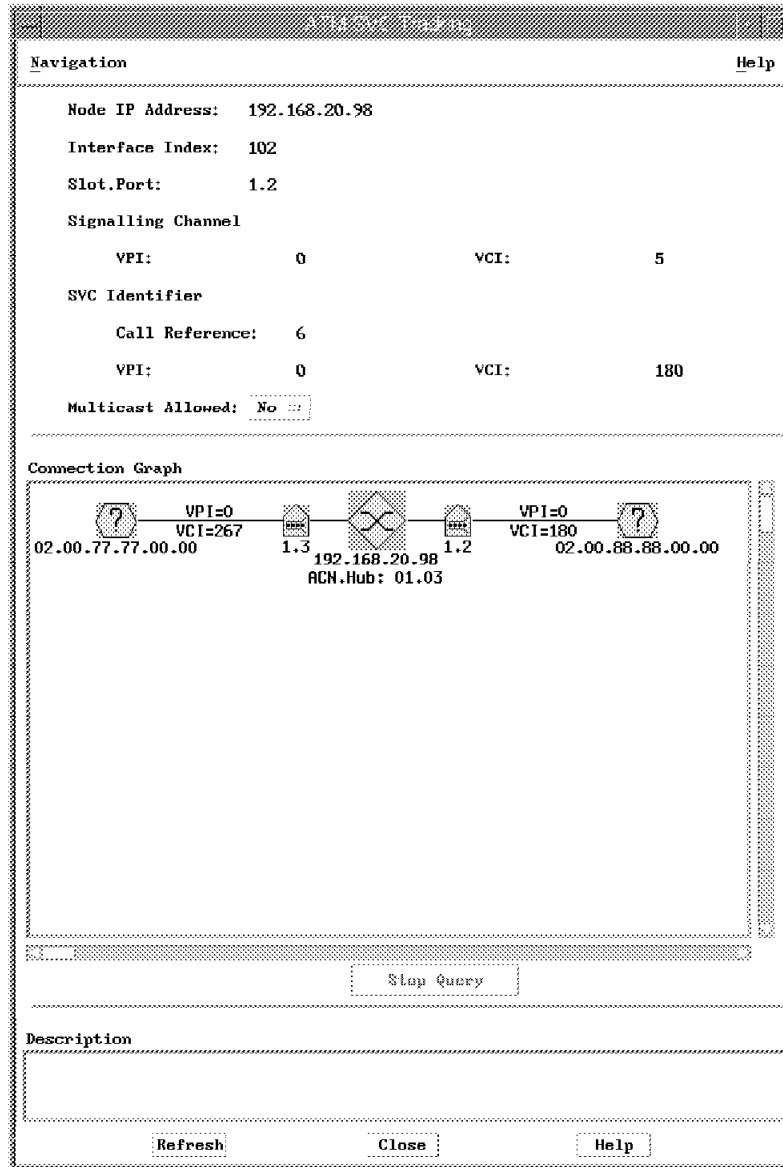


Figure 286. ATM SVC Tracking in Campus Manager

Numerous reasons can cause the data direct VCC to fail. The most common are again due to the higher layer protocol used not being bound to one of the client's network adapters or the clients are using different encapsulation types.

If you have a large ATM network, it may also be that you have exceeded the maximum number of VCCs in your network. See Chapter 5, "Starting Problem Isolation in an ATM Network" on page 109 for more information on whether you have exceeded the maximum number of allowable VCCs in your ATM network.

- Stage 3 - Trace the LAN emulation data flows between clients

Once you have established whether clients have learned each other's ATM addresses and set up an ATM connection you must use a network analyzer to analyze the LE data flows between one of the failing clients and its LE service servers and its partner client to diagnose exactly what the problem is related to.

We used a network analyzer, the InterWatch 95000, to capture traffic flowing from and to the 8274. The following screen shows a summary of the output. The VCC used for each frame has been added for information. The VCC information can be found in the detail for each frame.

```

***** Frame No 000000020 *****
000000020 CNTLDIR 7/10/97 10:38:32.549273 AAL5 LANE1.0
Control LE_ARP_REQUEST Success 0x00640443
0x0004
***** Frame No 000000021 *****
000000021 CNTRDIR 7/10/97 10:38:32.549504 AAL5 LANE1.0
Control LE_ARP_RESPONSE Success 0x00640443
0x0004
***** Frame No 000000022 *****
000000022 MLTISND 7/10/97 10:38:32.549766 AAL5 LANE1.0
Ethernet 0x0004 IP 192.168.5.9
192.168.5.22 ICMP ICMP Echo Request
***** Frame No 000000023 *****
000000023 MUTIFOR 7/10/97 10:38:32.550457 AAL5 LANE1.0
Ethernet 0x0001 IP 192.168.5.22
192.168.5.9 ICMP ICMP Echo Reply
***** Frame No 000000034 *****
000000034 MUTISND 7/10/97 10:38:32.610201 AAL5 LANE1.0
Control LE_FLUSH_REQUEST Success 0x00640444
0x0004
***** Frame No 000000038 *****
000000038 CNTLDIR 7/10/97 10:38:32.614294 AAL5 LANE1.0
Control LE_FLUSH_RESPONSE Success 0x00640444
0x0004
***** Frame No 000000041 *****
000000041 DATAD 7/10/97 10:38:32.630373 AAL5 LANE1.0
Ethernet 0x0004 IP 192.168.5.9
192.168.5.22 ICMP ICMP Echo Request
***** Frame No 000000042 *****
000000042 DATAD 7/10/97 10:38:32.630713 AAL5 LANE1.0
Ethernet 0x0001 IP 192.168.5.22
192.168.5.9 ICMP ICMP Echo Reply

```

The Interwatch 95000 analyzer was connected between the 8274 and the ATM network and we then started an IP ping between the 8274 and the MSS.

Refer to your network analyzer documentation to understand how to set up the network analyzer.

The data flow took place just after the IP ping was started between the 8274 and the MSS. It shows the following:

- Frame 20 - The 8274 sends an LE_ARP_REQUEST to the LES via the control direct VCC to determine the address of the MSS LE client.
- Frame 21 - The 8274 receives an LE_ARP_RESPONSE from the MSS via the control direct VCC.
- Frame 22 - In parallel the 8274 is sending an ICMP echo request (ping) to the MSS LE client via the BUS on the multicast send VCC.
- Frame 23 - The 8274 receives an ICMP echo reply (ping) from the BUS on the multicast forward VCC.
- Frame 34 - The 8274 sends an LE_FLUSH_REQUEST via the BUS on the multicast send VCC to the MSS LE client.

- Frame 38 - The 8274 receives an LE_FLUSH_RESPONSE from the LES on the control direct VCC.
- Frame 41 - The 8274 starts to use the data direct VCC for ICMP echo requests.
- Frame 42 - The 8274 starts receiving ICMP echo responses via the data direct VCC.

Problems can occur at any stage in the data flow shown above. If you find a request frame has been sent and no response received or a request frame should have been sent but hasn't been, you must look at the frame detail to examine the contents of the frame to understand why. Some problems may be caused by faulty microcode or driver levels but many others will be caused by the incorrect setup of the higher layer protocols. The diagnosis of higher layer protocols is outside the scope of this redbook.

Note: In tracing data flows it is important to verify the transaction ID. The transaction ID should be the same in the request frame and its subsequent response frame. The transaction ID is shown on the far right-hand side of the second line of each frame summary. For example, the transaction ID for frame 20 is 0x00640443.

The following screen shows the detail of frame 20 using the InterWatch 95000 Network Analyzer.

```

***** Frame No 000000020 *****

----- DETAIL -----

PDU Length 144 bytes sliced to 144 bytes
TE Side
AAL-5 CPCS      VCC = (0, 525)
AAL-5 CPCS Overhead:
  Number of padding octets    28
AAL-5 CPCS Trailer:
  User-to-User Indicator      0x00
  Common Part Indicator       0x00
  CPCS Payload Length        108
  Cyclic Redundancy Check     0x29faa41b

LANE1.0 - LAN Emulation 1.0 Frame
Marker:          0xff00    LANE 1.0 Control Frame
Protocol:        0x01
Version:         0x01
OP Code:         0x0006    LE_ARP_REQUEST
Status:          0x0000    Success
Transaction ID:  0x00640443
Requester LECID: 0x0004
Flags:           0x0000
Source LAN Destination:
  Tag:           0x0000    Not present
  Reserved:      0x000000000000
Target LAN Destination:
  Tag:           0x0001    MAC address
  MAC Address:   40:00:82:10:00:00
Source ATM Address:
  AFI:          39          DCC ATM Format
  DCC:          0985
  HO-DSP:       0x11111111111111110101
  ESI:          0x0020da6fa8e0
  SEL:          0x02
  Reserved:     0x00000000
Target ATM Address:
  AFI:          00
  Remaining Octets: 000000000000000000000000000000000000000000000000
  Reserved:
----- 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
----- 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
End of LANE 1.0 frame

```

Figure 287. LE_ARP_Request Frame Capture

Note: VCC 0/525 was the control direct VCC.

As another example we repeated the ping test but this time monitored the traffic using the WG DA-30C Network Analyzer.

```
47 BPDU BPDU_Type=Configuration
48 AAL 5 Frame Fragmented RX2 UU=0x00 CPI=Legal Len=62
49 BPDU BPDU_Type=Configuration
50 AAL 5 Frame Fragmented RX1 UU=0x00 CPI=Legal Len=62
51 LANE (ATM) RX2 OP-Code=LE_TOPOLOGY_REQUEST Status=Success
  Marker=Control Requester_LE_Client_ID=0x03
52 AAL 5 RX2 UU=0x00 CPI=Legal Len=0
53 AAL 5 Frame Fragmented RX2 UU=0x00 CPI=Legal Len=108
54 LANE (ATM) RX1 OP-Code=LE_TOPOLOGY_REQUEST Status=Success
  Marker=Control Requester_LE_Client_ID=0x03
55 AAL 5 RX1 UU=0x00 CPI=Legal Len=0
56 AAL 5 Frame Fragmented RX1 UU=0x00 CPI=Legal Len=108
57 ICMP Echo
58 AAL 5 Frame Fragmented RX2 UU=0x34 CPI=Illegal/Reserved Len=13879
59 AAL 5 Frame Fragmented RX2 UU=0x00 CPI=Legal Len=108
60 ICMP Echo Reply
61 AAL 5 Frame Fragmented RX1 UU=0x34 CPI=Illegal/Reserved Len=13879
62 AAL 5 Frame Fragmented RX1 UU=0x00 CPI=Legal Len=108
```

Figure 288. Summary Report Produced from Network Analyzer

Note: We used the Wandel and Goltermann DA-30C Inter-network-analyzer. The analyzer was connected between the 8274 and the ATM network and we then started an IP ping between the 8274 and the MSS.

We collected all the ATM data sent and received by the 8274 and then used the Examine application to decode the results. Refer to your network analyzer documentation to understand how to set up the network analyzer.

Each number in the list refers to a frame sent or received. We can see that the 8274 is sending Bridge Protocol Data Units (BPDUs) as it actively participates in the spanning tree. It is also sending some LE_TOPOLOGY_REQUEST frames to the LES. Frame 57 shows the ICMP ping request to the MSS and frame 60 shows the MSS ICMP ping reply.

You can then look at the detail of any frames in the data flow shown above.

The ICMP echo frame is shown below.

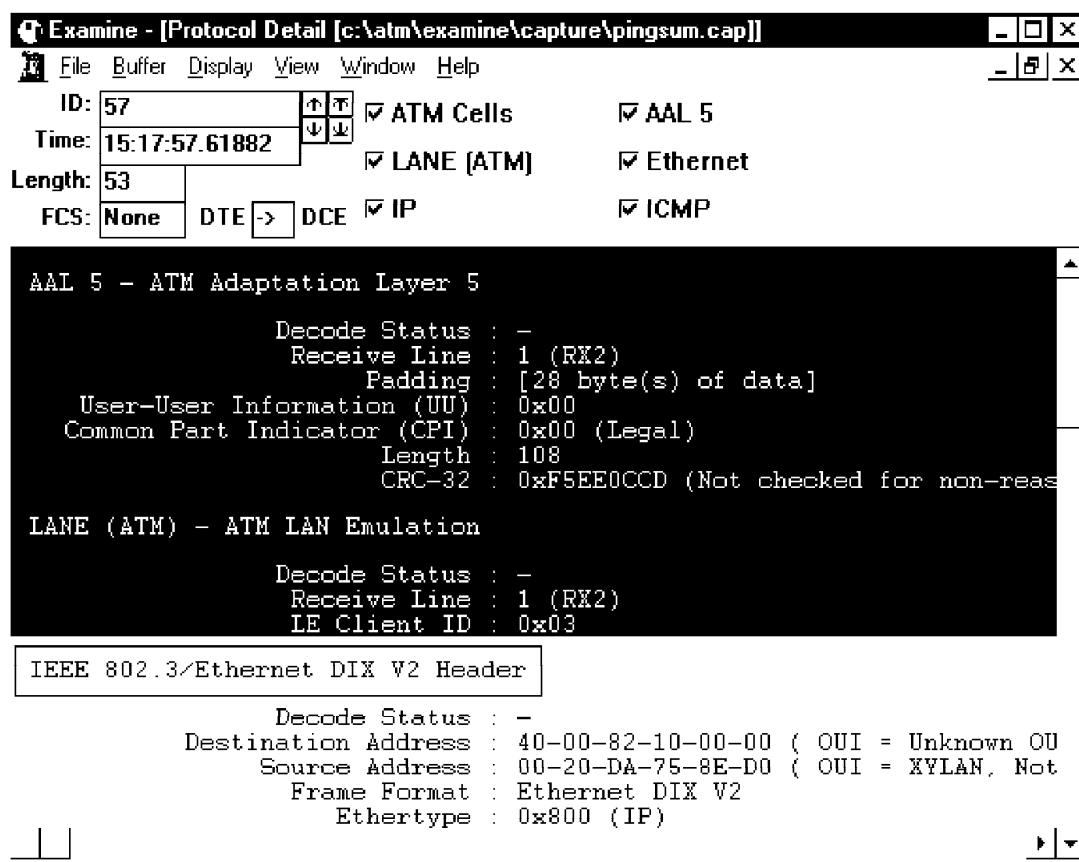


Figure 289. ICMP Echo Frame Screen 1 of 2

```

Examine - [Protocol Detail [c:\atm\examine\capture\pingsum.cap]]
File Buffer Display View Window Help

IP - Internet Protocol

    Decode Status : -
      Version      : 4.  Header length : 20
    Type of Service : 0x00
      000. .... = Routine Precedence
      ...0 .... = Normal Delay
      .... 0... = Normal Throughput
      .... 0... = Normal Reliability
      .... .0.. = Normal Reliability
      .... ..00 = Reserved

    Total length  : 92 bytes
    Identification : 4262
    Fragment Control : 0x00
      0... .. = Reserved
      .0... .. = May Fragment
      ..0... .. = Last Fragment
      ...0 0000 0000 0000 = Fragment Offset

    Time to Live  : 32 seconds/hops
    Protocol      : 1 (ICMP)
    Checksum      : 0xFE9A (Checksum Good)
    Source Address : 192.168.5.5
    Destination Address : 192.168.5.11
    No IP options

ICMP - Internet Control Message Protocol

    Decode Status : -
      Type        : 8 (Echo)
      Code        : 0
    ICMP Checksum : 0xAEFC (Checksum Good)
    Identifier    : 61256
  
```

Figure 290. ICMP Echo Frame Screen 2 of 2

The full AAL 5 PDU is decoded. The frame detail shows the AAL 5, LANE, Ethernet, IP and ICMP sections of the ICMP ECHO frame.

Both of the network analyzers used can be used for capturing traffic, but can also be used to monitor the network. Some of the extra features they include are the following:

- Network statistics
- Protocol distribution
- Frame-size and utilization charts
- Error statistics
- Traffic generation

7.10.1.2 Conclusion

In the case of the MSS server and 8274, you can use console commands to monitor the status of the LE client LE_ARP table. This can be used to make sure the LE clients have successfully learned the ATM addresses of the partner clients with which they wish to communicate.

IBM Nways Campus Manager is a very useful tool for discovering and tracking the status of a client's VCCs. This can help to determine whether clients have successfully established a data direct VCC.

Finally a network analyzer is essential for problems involving data transfer between two clients, once the clients have established a direct connection. All the frames through an ATM port are captured so the effective use of filters is important to locate the specific information in which you are interested.

7.11 Case Studies Involving Network Performance Problems

This section contains a case study explaining how to monitor the BUS to diagnose performance problems.

For this scenario we used the same network environment as described for the LAN emulation case studies. For more information refer to 7.8.1, "Network Environment" on page 339.

7.11.1 Case (4): BUS Performance

The state of your LANE 1.0 environment, discussed in 7.1, "What Makes a Healthy LANE 1.0 or Classical IP Environment" on page 293, is determined by the condition of your ATM network. When a network contains a large number of clients or where the level of broadcast traffic in an emulated LAN is excessive the BUS may not be capable of sustaining the traffic throughput required. This may result in it becoming a bottleneck in the performance of the network resulting in slow performance to users.

7.11.1.1 Methodology

The methodology used is consistent with the general methodology described in 7.3.4, "Network Performance Problem Methodology" on page 326.

It is therefore important to monitor the level of broadcast traffic in the network and to identify the top users of the BUS. This can be done using the BUS monitor function in the MSS. The following screen is an example of the output produced by the BUS monitor.

```

ITSO Lab MSS EXISTING LES-BUS 'ETHERNET_ELAN'+statistics display bus monitor
-BUS Monitor Status-
Currently in a sample interval ?           no
Next sample interval scheduled in:        0 minute(s), 36 second(s)

-Results of Last Complete Sample-
BUS Monitor sample interval started at:    00.12.04.85 (System Up Time)
Duration of sample interval:              10 second(s)
# Top Hosts Actually Recorded:            1
# Frames Received in sample interval:     17
# Frames Sampled in sample interval:      2
Frame sampling rate:                      1 out of 10

Rank  Source MAC Addr.   Associated LEC ATM Address          # frames
-----
  1  02.00.77.77.00.00  390985111111111111111010302007777000081  2

```

Figure 291. MSS Command to Display the BUS Monitor Statistics

The BUS monitor shows the top users of the BUS in the sample interval and the number of frames they have sent during the sample interval. This is very valuable information and could help to determine if one user was flooding the network. For information on how to set up the BUS monitor refer to the *Nways MSS Server Command Line Interface User*, SC30-3818.

Once you have identified an LE client that is potentially flooding the BUS you can display the detailed statistics for a particular user. This is done using the following MSS command.

```
ITSO Lab MSS
EXISTING LES-BUS 'ETHERNET_ELAN'+statistics display les-bus lec-bus lecid 0003

ATM Forum LEC-BUS MIB Statistics:
  recvs:                4687
  discards:             0
```

Figure 292. MSS Command to Display the Statistics for a Specific LEC

A final useful check to see if the BUS is being over utilized is to check the BUS statistics using the following command.

```
ITSO Lab MSS EXISTING LES-BUS 'ETHERNET_ELAN'+statistics display les-bus bus
ATM Forum BUS MIB Statistics:
  inDiscards:           0
  inOctets:             13688
  inUcastFrms:         3
  inMcastFrms:         192
  frmTimeouts:         0 1
  mcastSendRefused:    0
  mcastFwdFailure:     0
Other Statistics:
  inExplorer:          0
  inFlushReq:          2
  outFlushReq_mcastFwd: 0
  outFlushReq_mcastSend: 2
  outUcastFrms_mcastFwd: 0
  outUcastFrms_mcastSend: 3
  outMcastFrms:        192
  outOctets:           13518
  mcastSendReleased:  0
  mcastFwdReleased:   0
  mcastFwdPartyReleased: 0
  invalidProtocol_droppedFrames: 0
  verNotSup_droppedFrames: 0
  invalidOpcode_droppedFrames: 0
  invalidLecid_droppedFrames: 0
  invalidSize_droppedFrames: 0
  flushToBus_droppedFrames: 0
  incompleteSourceConnect_droppedFrames: 1
  incompleteTargetConnect_droppedFrames: 0
  noProxy_droppedFrames: 0
```

Figure 293. MSS Command to Display the Statistics for the BUS

The statistics list the frmTimeouts statistic **1** which is the number of frames discarded due to the maximum frame age timer timing out. The Max Frame Age (default: 1s) is the maximum time the BUS must take to transmit a data frame. If the timer expires, the data frame is discarded. A non-zero value would indicate the BUS performance has been lowered to the point that frame transmission is timing out.

7.11.1.2 Conclusion

Many factors can affect the low performance of an emulated LAN network. The BUS is just one factor. The MSS provides a number of features that allow the BUS performance and utilization to be monitored to allow problems with the BUS to be identified quickly. To solve the problem you may need to replace faulty network adapter cards or more likely redesign your network.

7.12 Case Studies Involving Redundancy

This section contains a case study explaining how to determine redundancy problems. We will use IBM campus network products to demonstrate our problem determination guidelines.

7.12.1 Case (5): LES/BUS Redundancy Problem

The IBM Nways Multiprotocol Switched Services (MSS) server provides a proprietary redundancy feature for the LES/BUS. This test environment investigates how to diagnose problems that can occur when using this redundancy feature.

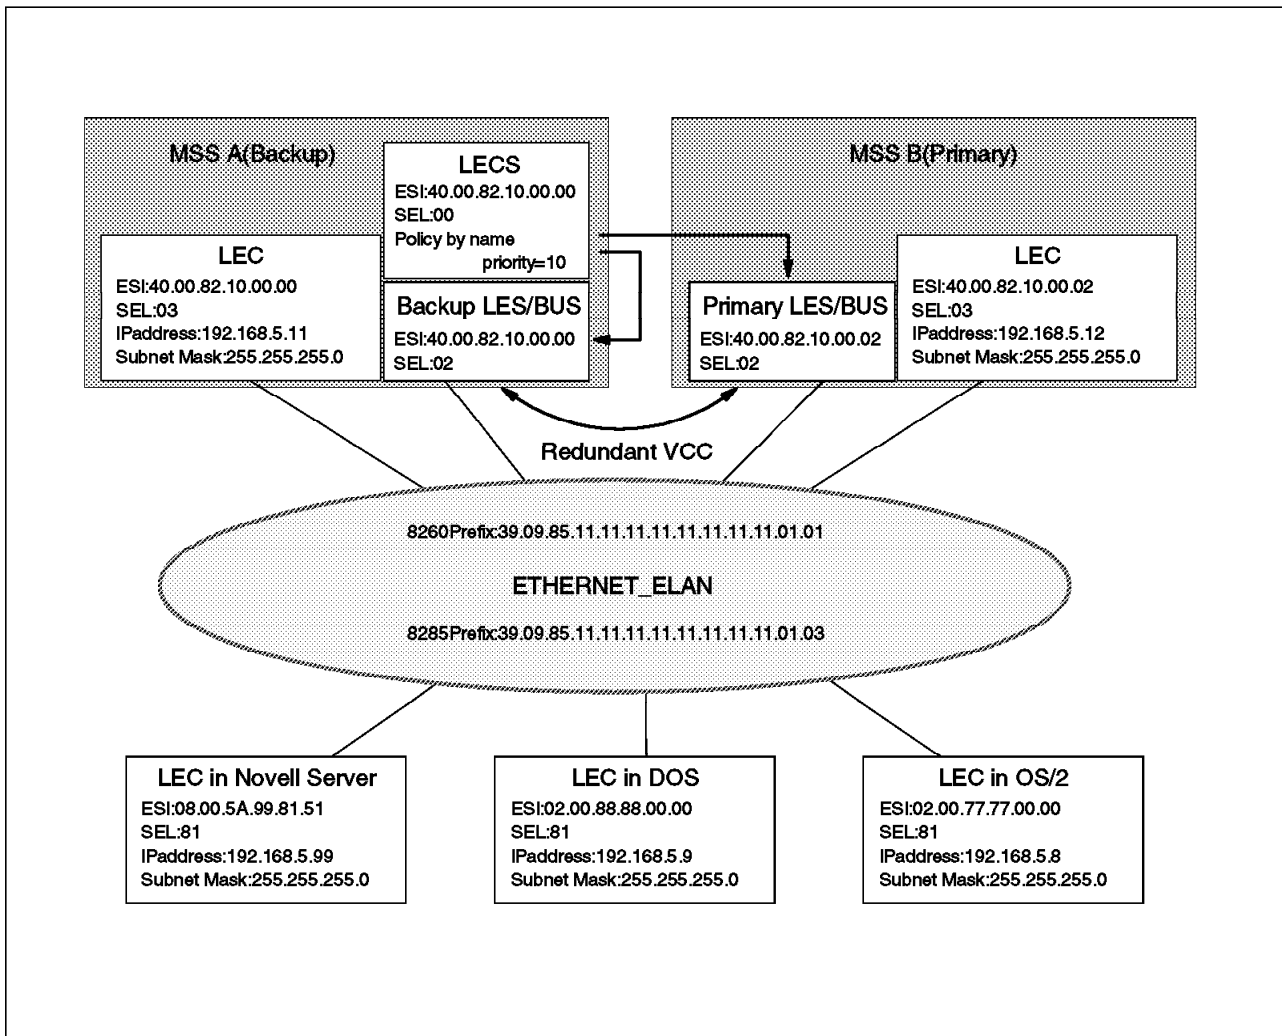


Figure 294. LES/BUS Redundancy Test Environment

The above figure shows a single Ethernet ELAN (ELAN name: ETHERNET_ELAN) with two IBM Nways Multiprotocol Switched Services (MSS) servers. On MSS-A the LECS function has been configured. Also, MSS-A is configured as the backup LES/BUS for this ELAN, while MSS-B is configured as the primary LES/BUS. The Novell server doesn't support using the LECS, so we configured it to use the primary LES ATM address (MSS-B).

At the start of this scenario only one LE client (the Novell Server) was registered with the primary LES/BUS. This can be seen using the following command from MSS-B.

```

:MSS-B EXISTING LES-BUS 'ETHERNET_ELAN'+data list all lec
Number of LEC's to display: 1

      LEC-LES and LEC-BUS State (UP=Up, ID=Idle, --. --.
      **=Other; Show specific LEC to see actual)      v  v
      LEC State #ATM #Reg #Lrnd
LEC Primary ATM Address      Proxy ID LES BUS Adrs MACs MACs
-----
390985111111111111111111010108005A99815181 N 0001 UP UP 1 1 0
    
```

Figure 295. MSS Command to Display LECs Attached to an ELAN

On the backup MSS server four LE clients were registered (DOS client, OS/2 client, MSS-A LEC and MSS-B LEC). This can be seen using the following command from MSS-A.

```

MSS-A EXISTING LES-BUS 'ETHERNET_ELAN'+data list 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
LEC Primary ATM Address      Proxy ID LES BUS Adrs MACs MACs
-----
390985111111111111111111010302008888000081 N 0001 UP UP 1 1 0
390985111111111111111111010302007777000081 N 0002 UP UP 1 1 0
390985111111111111111111010140008210000003 Y 0003 UP UP 1 1 0
390985111111111111111111010140008210000203 N 0004 UP UP 1 1 0
    
```

Figure 296. MSS Command to Display LECs Attached to an ELAN

7.12.1.1 Methodology

We used the methodology described in 7.3.5, "LE Redundancy Failure Problem Methodology" on page 327 to diagnose and fix the problem.

1. Stage 1 - Hardware and power problems

We verified the network device LEDs and status which indicated that the devices were operating normally and were powered on, as described in 7.8.4, "Case (1-1): LECS Failure Problem" on page 342.

2. Stage 2 - ATM problems

We also verified all devices were registered with their ATM switch and that the SSI links were active and available, as described in 7.8.4, "Case (1-1): LECS Failure Problem" on page 342.

3. Stage 3 - Simple configuration problems

- LECS configuration

When the primary LES/BUS fails, the control direct VCCs to the primary LES/BUS will fail and the LE clients have to reinitialize. During the re-initialization the LE clients obtain the LECS ATM address from the ATM switch, and connect to the LECS to learn the LES ATM address.

The LECS learns which LES is active, and returns the correct ATM address to the LE clients. The LECS must be configured with both the primary and backup LES/BUS ATM addresses for it to assign LE clients to the correct LES/BUS.

We used the following command to check the configuration of LECS in the MSS.

```
MSS-A ELAN 'ETHERNET_ELAN' selected+les list
LESS serving ELAN 'ETHERNET_ELAN'
=====
Primary ATM address:
          39.09.85.11.11.11.11.11.11.11.01.01.40.00.82.10.00.02.02  1
Backup LES with IBM LES redundancy:  2
Local LES for: ETHERNET_ELAN
```

Figure 297. MSS Command to Verify the Configuration of the LES

We checked that the primary LES ATM address was correct **1**, and the local LES was defined as the backup LES **2**. We compared this configuration with the configuration for the primary and backup LES/BUS.

- Primary LES/BUS configuration (MSS-B)

Using the MSS server talk 6 process we can view the defined configuration of the LES/BUS.


```
MSS-B LES-BUS config for ELAN 'ETHERNET_ELAN'>list
LES-BUS Detailed Configuration
Name: ETHERNET_ELAN
LES-BUS Enabled/Disabled: Enabled
ATM Device number: 0
End System Identifier (ESI): 40.00.82.10.00.02
Selector Byte: 0x02
ELAN Type: (S2) Ethernet
Max Frame Size: (S3) 1516
Control Timeout: (S4) 120
Max Frame Age: (S5) 1
LECID Range Minimum: 1
LECID Range Maximum: 65279
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): Disabled
Partition LE_ARP_REQUEST Forwarding Domain: Yes
LE_ARP RESPONSE Destination: One client
Partition Unicast Frame Domain: Yes
Redundancy: Enabled 1
Redundancy Role: Primary LES-BUS 2
ATM address of Backup LES-BUS: 390985111111111111111111010140008210000003 3
ATM address trace filter value: 0000000000000000000000000000000000000000000000000000000000000000
ATM address trace filter mask: FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
.....
```

Figure 298. MSS Command to Verify the Configuration of the LES/BUS

We checked the following:

- **1** The redundancy feature was enabled (Redundancy: Enable).
 - **2** The role of the LES/BUS (Redundancy Role: Primary).
 - **3** ATM address of backup LES-BUS (390985111111111111111111010140008210000003).
- Backup LES/BUS configuration (MSS-A)

We also checked the configuration on the backup LES/BUS MSS-A using the same command from the MSS talk 6 process.


```
MSS-B EXISTING LES-BUS 'ETHERNET_ELAN'+list
ELAN Name:                ETHERNET_ELAN
ELAN Type:                Ethernet
ATM Device number:       0
# of Proxy LEC's:        0
# of Non-Proxy LEC's:    0
LES ATM Address:         390985111111111111111111010140008210000202

-Status-
LES-BUS State:            OPERATIONAL      1
Redundancy VCC State:     IDLE             2
Major Reason LES-BUS was last Down: none
Minor Reason LES-BUS was last Down: none
LES-BUS State last changed at: 00.00.09.19 (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
.....
```

Figure 300. MSS Command to Verify the Configuration of the LES/BUS

We checked that MSS-B was operational **1**. The redundancy VCC was idle **2**, indicating that the backup LES/BUS was inactive. MSS-A was also in the same status as MSS-B; LES/BUS status was operational, and the redundancy VCC was idle.

When an error is encountered during the redundancy VCC establishment or the VCC is released, the primary LES/BUS will retry the call every 5 seconds. We checked the MIB statistics of the LES/BUS on MSS-B to determine if this was the case.

```

MSS-B EXISTING LES-BUS 'ETHERNET_ELAN'+statistics display les-bus les
ATM Forum LES MIB Statistics:
  joinOK:                                0
  verNotSup:                              0
  invalidReqParam:                        0
  dupLanDest:                             0
  dupAtmAddr:                             0
  insRes:                                  0
  accDenied:                              0
  invalidReqId:                           0
  invalidLanDest:                         0
  invalidAtmAddr:                         0
  badPkts:                                0
  outRegFails:                            0
  leArpIn:                                 0
  leArpFwd:                               0
Other Statistics:
  leArpAnswers:                           0
  leArpRspFwd:                            0
  topologyFwd:                             0
  narpFwd:                                 0
  flushRspFwd:                            0
  outJoinFails:                           0
  regOK:                                   0
  unRegOK:                                 0
  outUnRegFails:                          0
  proxyLeCs:                              0
  nonProxyLeCs:                           0
  macAddrMappings:                        0
  rdMappings:                             0
  atmAddrMappings:                        0
  joinRetransmits:                        0
  joinParmChanges:                        0
  joinTimeouts:                           0
  reRegs:                                  0
  ctDirRefused:                            0
  ctDirReleased_err:                      0
  ctDistFailure:                          0
  ctDistReleased_err:                     0
  ctDistPartyReleased_err:                0
  redundancyVccRefused:                    0
  redundancyVccReleased:                   0
  redundancyVccFailure:                    1350  1
  oam_droppedFrames:                      0
  invalidSize_droppedFrames:               0
  invalidMarker_droppedFrames:             0
  invalidProtocol_droppedFrames:           0
  invalidLecid_droppedFrames:              0
  unknownLecid_droppedFrames:              0
  invalidOpcode_droppedFrames:             0
  dupJoin_droppedFrames:                   0
  incompleteSourceJoin_droppedFrames:      0
  incompleteTargetJoin_droppedFrames:      0
  noProxy_droppedFrames:                   0

```

Figure 301. MSS Command to Display the BUS Statistics

You can see that the number of redundancyVccFailure is very high **1**. This indicates that the redundancy VCC is continually being retried.

Finally we looked at the MSS LES/BUS events using the event logging system on MSS-B. To do this we issued a `nodisplay sub all` followed by a `display sub les all` from the ELS + prompt. The events were observed in

talk 2. For more information on using the event logging system refer to 7.8.4, "Case (1-1): LECS Failure Problem" on page 342.

```

LES.261: LES/BUS:'ETHERNET_ELAN':plcng Rndncy call
        Called ATM addr = x390985111111111111111010140008210000003
LES.056: LES/BUS:'ETHERNET_ELAN':Rndncy call fld:cause 88
        Called ATM addr = x39098511111111111111111010140008210000003
LES.261: LES/BUS:'ETHERNET_ELAN':plcng Rndncy call
        Called ATM addr = x39098511111111111111111010140008210000003
LES.056: LES/BUS:'ETHERNET_ELAN':Rndncy call fld:cause 88
        Called ATM addr = x39098511111111111111111010140008210000003
.....
    
```

Figure 302. MSS talk 2 Display of Events for the LES Subsystem

Note: The UNI 3.0/3.1 cause codes for a connection failure are shown on each redundancy call. Refer to Appendix C, "UNI 3.0-3.1 Cause Maintenance Error Codes" on page 471 for a list of the status codes.

You can see that the LES/BUS is constantly trying to establish the redundancy VCC but fails to do so. Notice that the ATM address that is being called is incorrect. So the configuration of the LES/BUS on MSS-B must have an incorrect address for the backup LES/BUS.

7.12.1.2 Conclusion

The redundant VCC is not being established between the primary LES/BUS and the backup LES/BUS. The LECS on MSS-A interpreted that the primary LES/BUS on MSS-B has failed, so it assigned the backup LES/BUS (MSS-A) as the active LES/BUS for this ELAN. The Novell server LE client does not use the LECS function, so it was assigned to the primary LES/BUS correctly.

We changed the backup LES/BUS ATM address in the primary LES/BUS configuration on MSS-B and the redundancy VCC was then established.

```

MSS-B EXISTING LES-BUS 'ETHERNET_ELAN'+list
ELAN Name:          ETHERNET_ELAN
ELAN Type:          Ethernet
ATM Device number:  0
# of Proxy LEC's:   0
# of Non-Proxy LEC's: 1
LES ATM Address:    39098511111111111111111010140008210000202

-Status-
LES-BUS State:      OPERATIONAL
Redundancy VCC State: ESTABLISHED 1
Major Reason LES-BUS was last Down: none
Minor Reason LES-BUS was last Down: none
LES-BUS State last changed at: 00.00.09.19 (System Up Time)
LES-LEC Status Table changed at: 02.46.30.64 (System Up Time)
BUS-LEC Status Table changed at: 02.46.30.74 (System Up Time)
UNI Version:        3.0
IP BCM:             INACTIVE
IPX BCM:            INACTIVE
NetBIOS BCM:        INACTIVE
.....
    
```

Figure 303. MSS Command to Verify the Configuration of the LES/BUS

You can now see that the redundancy VCC has been established **1**.

Finally all of the five LECs can register in the primary LES.

```

MSS-B EXISTING LES-BUS 'ETHERNET_ELAN'+data list all lec
Number of LEC's to display: 5

    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
    Primary ATM Address      Proxy ID  LES BUS Adrs MACs  MACs
-----
39098511111111111111111010302007777000081 N 0001  UP  UP    1    1    0
39098511111111111111111010302008888000081 N 0002  UP  UP    1    1    0
39098511111111111111111010140008210000003 Y 0003  UP  UP    1    1    0
39098511111111111111111010140008210000203 N 0004  UP  UP    1    1    0
39098511111111111111111010108005A99815181 N 0005  UP  UP    1    1    0
    
```

Figure 304. MSS Command to Display LECs Attached to an ELAN

We then tested the redundancy feature by stopping the LES/BUS on MSS-B (primary) and after a few minutes used the following command on MSS-A (backup) to determine which clients were now registered with it.

```

MSS-A EXISTING LES-BUS 'ETHERNET_ELAN'+data list 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 ID      State #ATM #Reg #Lrnd
    Primary ATM Address      Proxy ID  LES BUS Adrs MACs  MACs
-----
39098511111111111111111010302008888000081 N 0001  UP  UP    1    1    0
39098511111111111111111010302007777000081 N 0002  UP  UP    1    1    0
39098511111111111111111010140008210000203 N 0003  UP  UP    1    1    0
39098511111111111111111010140008210000003 Y 0004  UP  UP    1    1    0
    
```

Figure 305. MSS Command to Display LECs Attached to an ELAN

Important

The LE clients will only connect to the backup LES/BUS if they learn the LES address from the LECS. For the LE clients that use a hard-coded LES address such as the Novell server, the automatic backup mechanism described will not work.

When the primary LES/BUS becomes active again, it will restore the redundancy VCC, the backup LES/BUS will drop all VCCs to the LE clients and will no longer accept calls to avoid the LE clients being connected to the two different LES/BUSs in the same ELAN.

Chapter 8. ATM Bridging and Routing

This chapter presents a methodology for troubleshooting problems commonly found in ATM networks containing bridges and routers. A classification of the symptoms experienced and their possible causes is also worth mentioning. In addition some case studies are used to illustrate how the methodology can be applied to resolve problems efficiently and effectively.

8.1 What Is a Healthy ATM Network with Bridging and Routing

As with networks based on other technologies, ATM networks containing bridges and routers can be considered as healthy when the network is stable and properly responds to the demands (traffic) of its users, helping them to do their jobs. In other words the network must:

- Be stable, that is without significant (or perceptible) interrupts; this can be achieved using redundant components in various ways and places in the network, and in general with a careful topology design.
- Perform according to business needs and user demands; that is the right (powerful) components should be at the right place and in sufficient number to satisfy load sharing and redundancy requirements.
- Adapt easily to evolving needs; this might be one of the most challenging targets to meet. In general we can say that this can be achieved when choosing the right technologies (no dead-end streets), paying special attention to network design and network conventions, and having a good networking team (and/or a good service company) to accompany and conduct the network evolution in well-planned steps, apart from additional teams for help desk and day-to-day troubleshooting matters. Looking at this last item more closely, we identify this not to be a technical issue, but actually more a question of implementing the right organization with sufficient and qualified resources.

ATM can help to achieve the above goals thanks to its inherent capabilities:

- It offers a lot of reliability, with redundant links not visible from the layers atop (LAN emulation, Classical IP, and other network models), adding to network stability.
- It offers higher speeds than those used in legacy LAN technologies, allowing added network performance and availability where it is needed most,
- Being a highly scalable technology, ATM adapts well to changing and increasing needs of network users and uses.
- Since the same technology is used in ATM for LAN and WAN, it blurs the former borders between the two, making it easier to expand the LAN over the WAN in response to business needs.
- Finally, it makes developing new networking models (NHRP, IP switching, MPOA, etc.) possible for enhanced performance and reliability of networks.

8.2 Rules of ATM Bridging and Routing

Most of the concepts needed for tackling bridging and routing problems are well covered in many publications available from IBM and other sources on the market. Nevertheless, some of the most relevant ones have already been briefly reviewed in Chapter 2, “Main Concepts of Legacy LANs” on page 7 and Chapter 3, “Main Concepts of ATM (Asynchronous Transfer Mode)” on page 43.

8.3 PD Guidelines for Bridging and Routing over ATM

This chapter outlines a specific approach to troubleshooting a network that involves bridges and routers.

8.3.1 Categories of ATM Problems with Bridging and Routing

First we want to think of the kinds of problems we may find in an ATM network with bridges and routers. This helps us to find the right methodology for troubleshooting them.

Most problems found in ATM networks with bridges and routers are the same as those found in conventional networks, which means that they will frequently occur in the layers above ATM. Some of them are even common to bridges and routers, since routers implement bridging techniques, too. We can say that problems possibly encountered in both ATM and above layers are of the following general nature:

- HW-related problems

These are actual HW defects, product bugs, as well as old or incompatible microcode versions used on one or several of the bridges and routers in the network. Insufficient CPU power and memory can be counted here too, although this will more accurately translate into performance and data transfer problems.

- ATM layer, emulated LAN, and Classical IP connectivity problems

These have been covered in the previous chapters. Please refer to Chapter 5, “Starting Problem Isolation in an ATM Network” on page 109 for ATM layer problems, to Chapter 7, “ATM Emulated LANs and Logical IP Subnets (LANE 1.0, RFC 1577)” on page 293 for emulated LAN problems, and to Chapter 6, “LAN Switches in Campus Networks” on page 237 for LAN switch problems.

- Data transfer problems

These are the kinds of problems we will focus on in this chapter. Data transfer problems include all those cases where all or some devices on a local segment (emulated LAN or logical IP segment) cannot communicate with devices on another segment, be it using some or all of the protocols installed on those devices. Most of these problems are due to configuration errors in the intermediate bridges or routers along the communication path connecting those devices willing to communicate with each other.

The symptoms perceived by the users of the network usually fall into the following categories:

- A bridge blocks and does not transfer any traffic, although it should do so.

- A router does not forward data to some parts of the network, while it does to others.
- Some stations can establish a session with other stations on a remote ring, while others cannot.
- Some servers can be reached from a certain segment, while others (using a different protocol) cannot.
- Part of the network cannot be reached after moving some equipment from one place on the network to another.
- Performance is perceived as very slow, although there is not much work being done by the users. And the bridges in the network display constant activity (LEDs invariably steady on).

And the reasons or causes for those problems may be as varied as the ones listed below:

- Duplicate ring numbers (token-ring)
 - Duplicate MAC addresses
 - Duplicate network addresses
 - Wrong static entries in MAC filtering databases (transparent bridges)
 - Insufficient cache allocation for filtering databases (transparent bridges)
 - Inconsistent or insufficient timer values for spanning tree (both transparent and source route bridges)
 - Wrong or insufficient network address filters
 - Wrong or insufficient network protocol filters
 - Wrong frame translation between unlike topologies (for example Ethernet to token-ring, and vice versa)
 - Wrong subnet masks
 - Wrong static routes (including missing ones)
 - Wrong import/export of routes (for example, between OSPF and RIP, etc.)
 - Inconsistent or insufficient timer values (for example, within OSPF area)
- Performance problems

These include problems where bridges and routers in the network communicate, but performance is lower than expected, and some devices may even quit service (and possibly reboot). These kinds of problems often translate into data transfer problems, and are first recognized as such. They may be alleviated with some tuning of components, but the real cure for the problem is often a partial redesign of the network, which might be as simple as adding components or splitting a segment, or might go as far as a partial or full redesign of the network, or even a migration to a more powerful technology. Sometimes performance problems may arise because of wrong configurations, which is the case, for example, with spanning tree problems. Many of the mentioned problem areas are most often observed in fast growing networks where not enough attention could be paid to both pro-active planning and consolidation of achieved results (technical and user documentation, help desk and user training, etc.).

Some of the problems we may find here are due to:

- Insufficient CPU power (filtering may consume much of this).
- Insufficient device memory for data buffers with respect to the (possibly huge) size of the network.
- Speed accommodation: bridges and routers buffer frames sent from a faster to a slower segment if the slower one cannot take them over quickly enough. (The slower one might be at the opposite end of a WAN link, or just running at a different speed.)

- Too many users and/or devices on a segment, leading to session losses and other problems due to overcommitment of address or route caches.
- Topology or architecture-related problems

These are usually perceived as broadcast storms in the case of spanning tree problems, or slow convergence of routing tables, for example, with protocols such as RIP (whether or not you are using the split horizon scheme), and will therefore translate into data transfer problems. This means again that the network users will usually perceive this situation as a data transfer problem. The cure here is a (partial) redesign of the network with respect to the underlying architecture, which might mean repositioning the root bridge or splitting collision domains (spanning tree), or splitting big router network areas into smaller ones (which on the other hand might affect performance adversely at some other place).
- Wrong protocol definitions on server or user workstations

These also translate into data transfer problems which might be attributed to bridges and routers, although the source of the problem is the server or the workstation. Therefore we include the troubleshooting of this kind of problem into our overall methodology for data transfer-related problems. Server and workstation originated problems include:

 - Wrong network addresses, subnet masks, frame encapsulation types, etc.
 - Incorrect or insufficient definitions for properly joining the desired ELAN.

Figure 306 on page 421 offers a classification of the bridging problems by perceived symptoms and their possible causes, while Figure 307 on page 422 does the same for the routing problems.

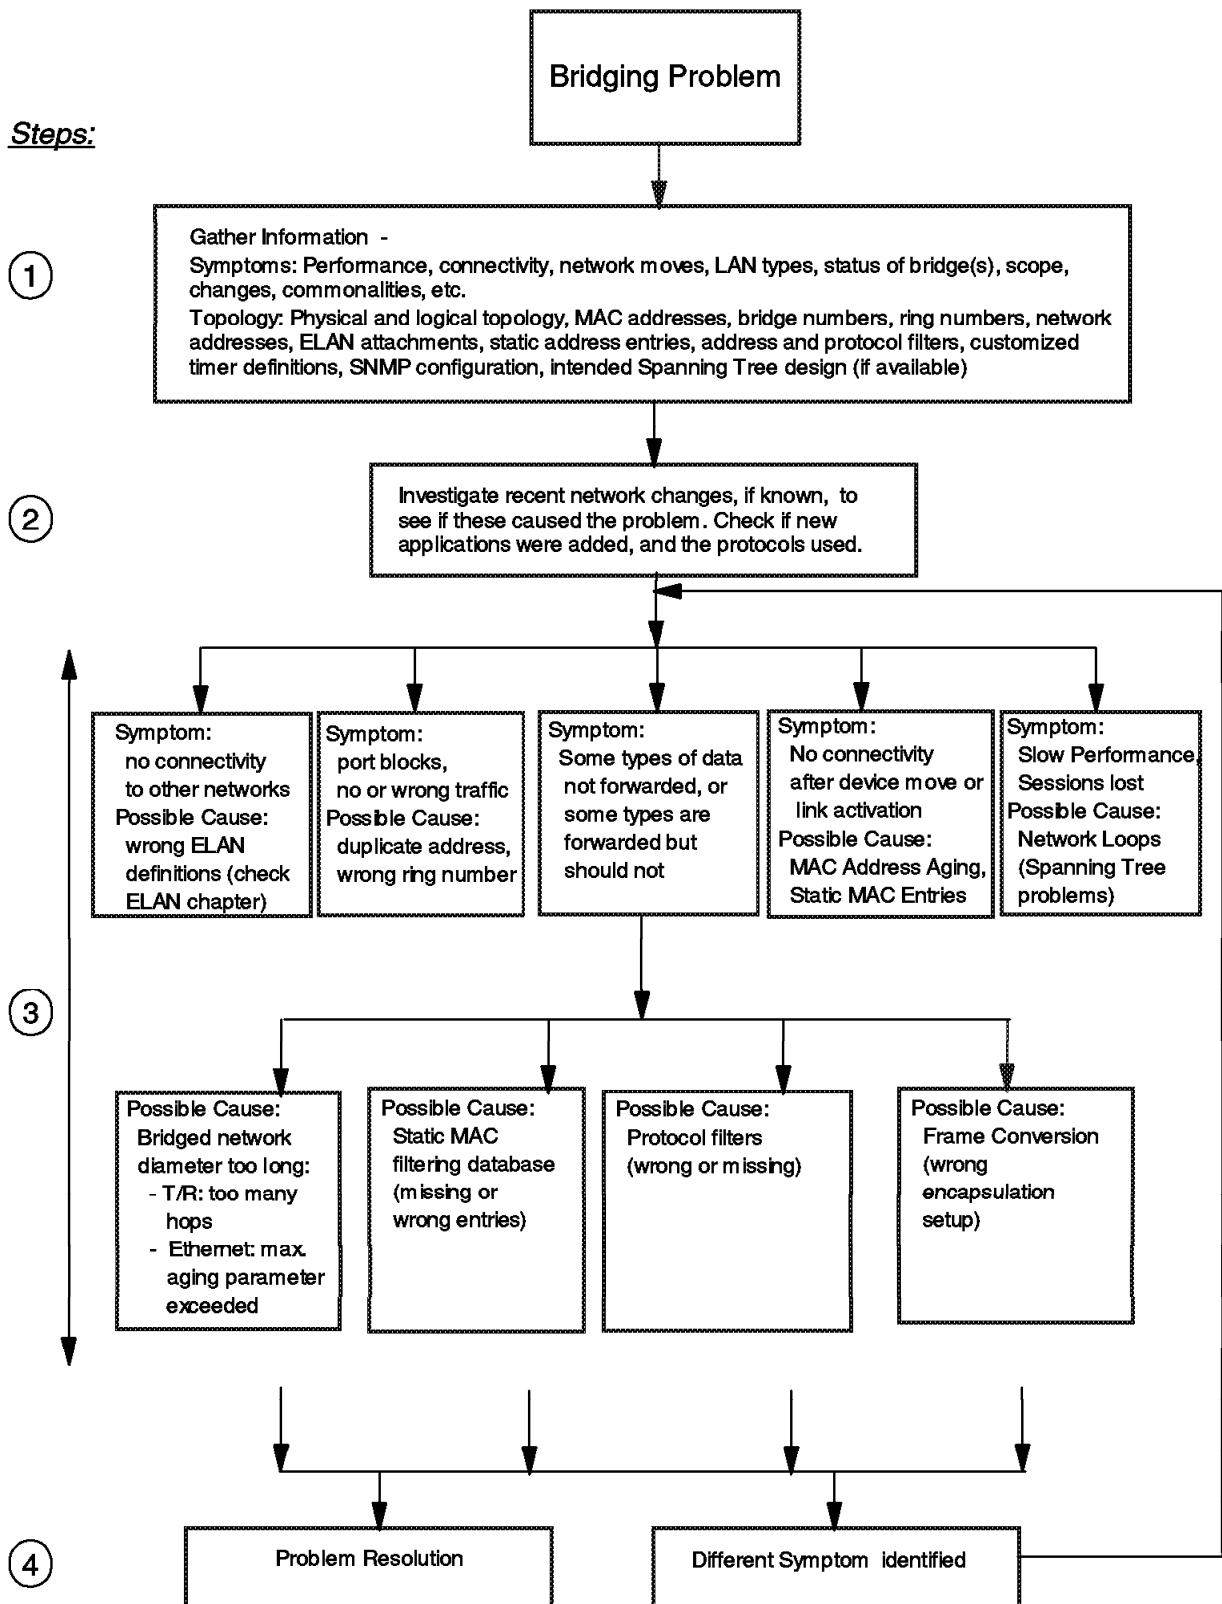


Figure 306. Symptoms and Causes for Bridging Problems

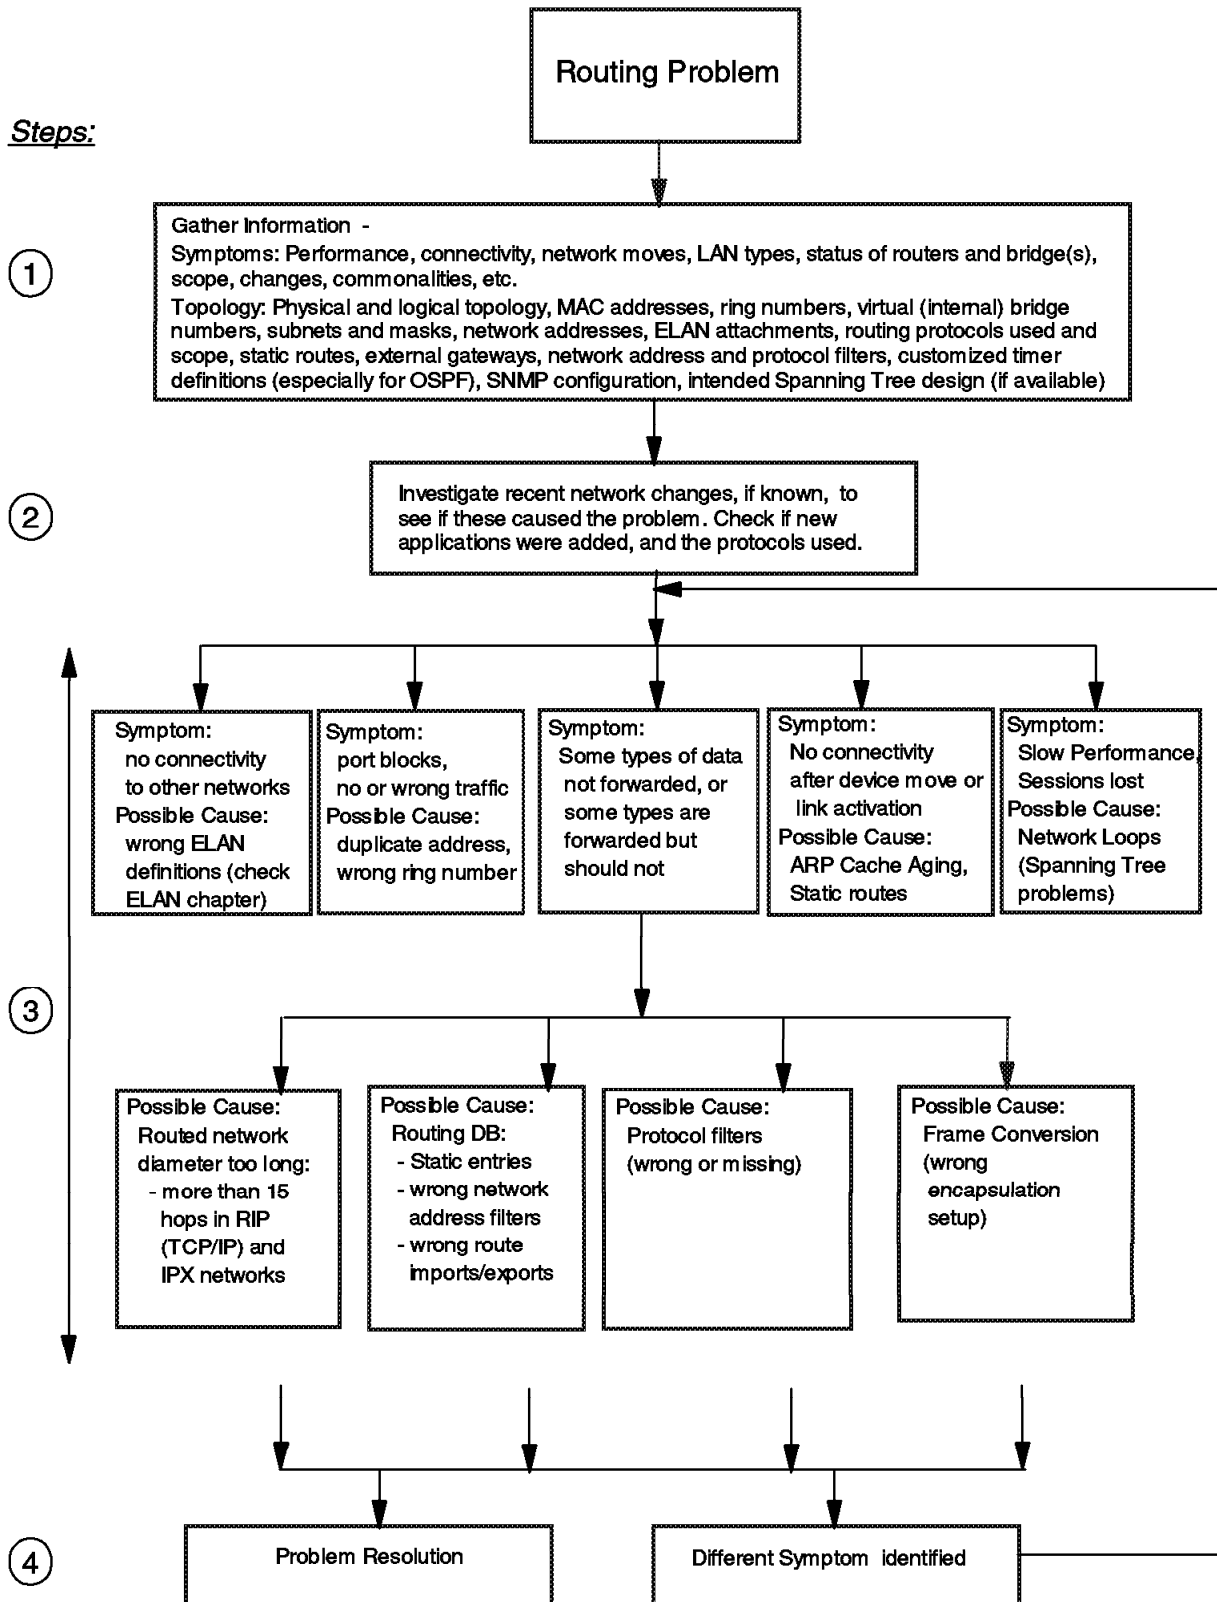


Figure 307. Symptoms and Causes for Routing Problems

8.3.2 PD Methodology for ATM with Bridging and Routing

Having looked at the types of problems that might be found, we developed the following troubleshooting methodology, specifically targeted at data transfer problems, but useful also for the other problem areas.

Phase 1: Gather information

Gathering information at an early stage of the troubleshooting process will effectively help you find adequate clues to the nature of the underlying reasons for the network problems currently being experienced.

First of all, ask network users, help-desk teams, network supporters, network management responsables, and IT management about overall uses of the network and the symptoms currently experienced. This will provide you with a more sound picture of the overall situation: what is the network used for (types of users, their their needs, and priorities), types of workstations, applications, protocols, etc., and finally find as much information about symptoms experienced. Different users might have different views, and sometimes only some users might be aware of certain symptoms.

- What devices are affected by the problem? Are all of the devices attached to the same emulated LAN or logical IP segment? Are they all affected or only a subset of them? Where are the devices connected to the ATM network? What are their ATM, MAC and network addresses (IP/SNA/NetBIOS/IPX)?
- What do the affected devices have in common? Are they on the same hub or connected to a certain portion of the network? Are they all using a common device, for example, a common LECS, LES, BUS or ATMARP server, or a common file, print or application server?
- What happens when the problem occurs? Do the affected devices always fail to get their required connectivity or do they fail only sometimes? Do they always connect to their servers but then at some point lose their connectivity? Can they connect but receive lower performance than usual?
- What has changed in the network since the problem started happening? Have devices been added or reconfigured?
- When does the problem occur? Does it only happen at a particular time of the day or when users are doing something specific, such as using a particular application or doing a large file transfer?

Secondly, gather as much information as possible regarding the physical and logical topology of the network. Some of the key questions to answer are:

- How is your ATM network physically connected together? How are the hubs and switches interconnected? What other LAN segments/networks are connected to the ATM? Are there redundant components and paths, where, why and what for? Remember, for proper problem isolation you need a good overview of the complete network. Have you heard of LAN segments nobody knew existed? How about that bridge that remained unnoticed under the floor for many years, and failed the very minute it was needed most (and nobody knew where it was physically)?
- How is your network logically configured? What emulated LANs or logical IP segments have been defined and how are these logically interconnected between each other and to other LAN segments/networks? If redundancy was designed, which components provide it, and how were they configured?

- Does the network extend over WAN connections?
- Which technologies and protocols are used there and what for?
- Find out about application and database/file servers, gateways to other networks, etc., to make your picture of the network more complete.
- Ask about addressing range and other conventions. Maybe you will need them later on to identify a possible problem source.
- Become more specific towards the network devices used, their type and make, components, versions, links between them, protocols, their management platforms, etc., before digging into the details of function and performance of components in the overall network.
- What are the ATM, MAC and network addresses (IP/SNA/NetBIOS/IPX) of the key network components? These includes hubs, bridges, routers, switches, LECS, LES, BUS and ATMARP servers.
- Which protocols are in use on the network, and for which applications? Do all workstations use all protocols, or only some workstations in certain departments, maybe even in only some segments?
- Where are the key LAN emulation or Classical IP network devices physically connected to your network? Where are the bridges and routers, and which ones are built-in components within hubs? Where are redundant paths (and possible loops)? Which routing protocols are in use, and how was the router network topology defined (OSPF areas, external routes, firewalls)? Have protocol or filter barriers of any kind been implemented, where and what for?
- Are the key components all managed by the same department, or by different ones? If more than one, is the information you gathered from them consistent? Otherwise ask again.

Good quality network documentation should provide the answers to these questions. If your documentation is not up-to-date or not maintained, a network management station, configured for your networking environment should provide many of the answers.

Phase 2: Investigate recent network changes

Most problems occur when you have changed something in your network. Effective change control and management can make sure one change does not impact the rest of your network. It also allows you to easily identify changes that may have introduced problems and allows you to back out these changes quickly and hence resolve the problem. You can then examine why the changes caused the problem.

Phase 3: Interrogate the network based on the problem symptoms

If the change causing the problem cannot be identified or is unknown, you must interrogate the network to discover the problem. The key information you need to discover depends on the problem you are experiencing.

In general we can say that we will need to perform some or all of the following tests depending on the nature of the problem. Usually tests further below become necessary if preceding ones were not successful, while some test steps may need reiteration. We will exit the sequence of tests listed below, as soon as we have discovered the reason for the problem, and will continue then with phase 4.

1. Check the correct configuration of your client, and if it is properly connected to the local segment from where you start your tests.
2. Ping a bridge, a router, a server, or another TCP/IP workstation to confirm layer 3 TCP/IP connectivity within the segment and especially to other segments. Does your test client have correct network address and subnet mask definitions? How about protocol encapsulation? Can your target be reached with increasingly larger IP packets?
3. Log on to any file server (for example LAN Server, Novell, or NT) to confirm connectivity for other protocols (NetBIOS and/or IPX). Is your client using the right protocol encapsulation? Can it exchange big data chunks (for example, files) with its server?
4. Start host emulation and connect to the host to confirm SNA connectivity. If this is not the case, and you have some tool to test the reachability of your gateway (37xx, Communications Server, etc.) with XID test frames, try using it to confirm layer 2 connectivity, and if the tools support different packet lengths, try increasing them up to the maximum needed for your applications. If there is no connectivity at all, check your SNA configuration before continuing.
5. Interrogate hubs to check for proper LECS address definitions.
6. Check status of ATM hub port (UNI) where your bridge(s) and/or router(s) is/are connected.
7. Interrogate the status of your bridge(s):
 - Are MAC addresses, bridge and ring numbers correct?
 - Has participation in spanning tree been enabled? Are the definitions correct regarding your overall spanning tree design? Most important in the design is to put the root bridge at a central position in the network, and to define priorities of other bridges accordingly.

Note: We need to make certain here that we do not have a spanning tree loop in the network, since continuing from here without checking that would be equivalent to trying to stabilize a tall building on moving sands. The most indicative symptoms for such a problem are:

- Traffic LEDs of ATM modules steady lit, indicating very heavy traffic. There will not be any indication on this when using only the command interface from a remote location. You can notice this type of problem only being physically in front of the IBM Nways 8260 Multiprotocol Switching Hub. On the other hand, checking the status of other devices such as IBM Nways MSS Server will give sufficient information to find out (IBM Nways MSS Server's most relevant counter here is the number of topology changes).
- Clients cannot communicate with their servers or sessions go lost, due to very high use of the bandwidth, although there are not many users currently active.
- Novell servers beep continuously displaying a message similar to "loop discovered, received own frames", indicating that they have received, for example, their own SAP frames.
- Counters in the bridges (if supported) indicate excessive topology changes of the spanning tree. We suggest you review the concepts presented in 2.5.1, "Transparent Bridging (TB)" on page 27 and 2.5.1.1, "Spanning Tree Protocol (STP)" on page 28 in case you are not familiar with this.

- Use of some network analyzer tool indicates that there are continuous topology changes of the spanning tree.

The only cure is to carefully review the spanning tree parameters of all bridge configurations, and to set them according to the Spanning Tree scheme you need. In general, the root bridge should be in the middle of your network, or where the biggest backbone bandwidth is available.

Following are some important rules for the spanning tree:

$$2 * (\text{forward delay} - 1) \geq \text{max age} \geq 2 * (\text{hello time} + 1)$$

The values for the IBM Nways MSS Server product are (in seconds):

- 2 secs hello time (range: 1 - 10)
- 15 secs forward delay (range: 4 - 30)
- 20 secs max age (range: 6 - 40)
- Are the ports in forwarding state when expected to be so?

Note that bridge and router ports on IBM Nways MSS Server are the IBM Nways MSS Server LECs attached to the corresponding emulated LANs or logical IP segments.

- Has a bridge port on the MSS server inadvertently been set for routing?

Note that assigning an IP address to a port makes it automatically able to route IP data traffic.

- Is the hop count correct (token-ring)?
- Is the maximum age count correct (Ethernet)? The farthest bridges might not receive Hello BPDUs from the root bridge in time if the value is set too low.
- Is the address cache full (and configured not to add new entries)?
- Are there any unexpected static entries defined (acting as filters)?
- Are MAC and protocol filters being used correctly?

Note that in spite of a possible NetBIOS name length of up to 32 characters, most bridging devices (whether bridges or routers) will recognize only up to 15 or 16 leading characters as unique differentiators. If you are using NetBIOS name filters check on this.

- When using source route translational bridging (SR-TB), have group addresses (Ethernet) been properly mapped to functional addresses (token-ring) and vice versa?
- Were the right encapsulation schemes selected for frame conversion?
- Are there long transmit queues (buffered frames waiting to be sent)?
- Is a high percentage of the receiving buffers filled up?
- How about rapidly increasing error counters and frame discards?
- If the bridge provides for error logs, check them for more clues.

Note: The IBM Nways 8281 ATM LAN Bridge configuration tool offers you the possibility of getting an overall bridge status report containing code version, up time, and port information such as MAC address, ATM address, half or full-duplex operation, transmit queue length, number of receive buffers, and port-related spanning tree status (blocked, learning, forwarding). In addition, much information about the IBM Nways 8281 ATM LAN Bridge can be retrieved using Nways Campus Manager.

8. Log on to your router and ping other stations or routers from there (TCP/IP).
9. IPXping other routers from your router (IPX).

10. Interrogate the status of your router(s):

- If router does bridging, are bridging items all correct? Make certain you check MAC addresses anyway.
- Are the ports of the router up (check forwarding and receiving counters)?
- Are the correct network addresses and subnet masks in use?
- Does the routing table contain a route to the desired target network? Check both network addresses and subnet masks. Remember that RIP (V1) does not use subnet masks. Could this be the problem?
- Does your router see all of its router partners ok as expected? Are the overall routing protocol definitions correct with regards to the overall router design (router areas, protocol borders for exports/imports, etc.)?

Note that all interfaces of OSPF routers attached to the same segment must have been configured in the same way for the hello interval and router dead interval parameters.

- Is the line utilization of some interfaces driving constantly towards 100% (usage of the available bandwidth of those interfaces)?
- Are there long transmit queues (data awaiting to exit the router)?
- Is a high percentage of the receiving buffers filled up?
- How about rapidly increasing error counters and packet discards?
- Check router error logs for more clues.

Note: The status of IBM routers can be interrogated via various ways such as a VT100 session through asynchronous port or over Telnet, some via a WWW browser (SLIP or over network), and of course via SNMP management.

11. Interrogate the LES/BUS to determine clients connected to the emulated LAN.

12. Interrogate the ATMARP server to determine clients connected to the logical IP segment.

Note: Only some LES/BUS/ATMARP products will support commands to determine what devices are linked to them. This function is supported by the IBM Nways MSS Server.

These tests should help to confirm whether the devices are attached to their emulated LANs or logical IP segments, whether bridges are configured to forward frames correctly or not, and if routers contain the appropriate definitions to do their job. Where this is not the case, this methodology will help you to quickly find the possible problem sources.

Phase 4: Resolve problem or investigate a different symptom

By following the problem methodology described in the above steps you will usually identify and resolve the problem. This may mean correcting wrong configurations, checking that the applied correction works OK, and reiterating the above test steps where necessary to verify that the problem has been solved. In some cases you will find that the problem is actually related to a different problem area than expected originally, in which case you may need to repeat the entire problem determination process to investigate this new area to resolve the problem. Other times you will have to take a different path in your investigation than the one you originally had in mind.

8.4 Hints and Tips with IBM Products

There are bridge and router command reference manuals and user's guides, each of them often covering many volumes. We would like therefore to refer the reader to them, and to only summarize below some useful commands applied when using certain IBM products used in our case studies.

Many useful commands especially for the underlying layers up to the LAN Emulation layer have already been covered in 7.5, "A Guide for Using Commands with the IBM Nways Multiprotocol Switched Service (MSS) Server" on page 331. We repeat here only the most relevant ones, and add those which are felt most usable for specific bridging and routing problems.

Table 24 (Page 1 of 4). Hints and Tips with Bridged and Routed ATM Networks

Topics	Techniques	Products	Examples: Descriptions
Problem scope	Commands	MSS	<ul style="list-style-type: none"> MSS Existing LES-BUS '<code>< ELAN NAME></code>' + database list all lec (list LECs attached to an ELAN)
LEDs on devices	Simple Test	all	: visual indication
Device status	Commands	MSS	<ul style="list-style-type: none"> MSS * uptime : display time elapsed since last reboot MSS + configuration : list the network interface information MSS + interface : display statistical information about the network interface MSS + test 1 (1: interface #) : verify the state of an interface
	Commands	8274	<ul style="list-style-type: none"> % vap 4/1 (4/1: slot #/port #) : view the configuration/status ATM port
Status of ATM network (slot/port/ESI)	Commands	8260/8285	<ul style="list-style-type: none"> ATM_SW> show module 16 verbose (16:slot #) : display the status/information of module/port ATM_SW> show port 16.1 verbose (16.1: slot#.port#) : display configuration information for an ATM media port ATM_SW> show atm_esi all (SHOW REACHABLE_ADDRESS ALL) : display ATM addresses registered in switch
Switch configuration	Commands	8260/8285	<ul style="list-style-type: none"> ATM_SW> show lan_emul configuration_server : display the entries in the LECS address table
LES/BUS configuration	Commands	MSS	<ul style="list-style-type: none"> MSS EXISTING LES-BUS '<code>< ELAN NAME></code>' + list :list the LES/BUS's status and current configuration parameters

<i>Table 24 (Page 2 of 4). Hints and Tips with Bridged and Routed ATM Networks</i>			
Topics	Techniques	Products	Examples: Descriptions
LEC configuration	Commands	MSS	<ul style="list-style-type: none"> • MSS LEC+ list config : list the LEC configuration
	Configuration tool	8281	<ul style="list-style-type: none"> : Use SLIP/IP/Bridge connection to 8281 : Use Get Bridge Status Report menu option to retrieve overall status report from bridge
	Commands	8274	<ul style="list-style-type: none"> • % vas : display services defined on the 8274 • % mas 4/1 1 (4/1 1: slot#/port# service#) : modify a LAN Emulation service
	Simple Test	Endstation	: see the (protocol.ini) file
ATMARP client/server configuration	Commands	RS/6000	<ul style="list-style-type: none"> • smitty tcpip : Display and set TCP/IP configuration for an interface • smitty atm : Display and set parameters for an ATM interface
	Commands	MSS	<ul style="list-style-type: none"> • MSS IP> int : display IP addresses of MSS interfaces • MSS ARP config> list atm-arp-client-config : display ATMARP client/server configuration

Table 24 (Page 3 of 4). Hints and Tips with Bridged and Routed ATM Networks

Topics	Techniques	Products	Examples: Descriptions
Bridging and spanning tree protocol	Commands	MSS	<ul style="list-style-type: none"> • MSS + protocol asrt : enter submenu for (adaptive) source-route transparent bridging • MSS ASRT> list bridge : list bridge ID, priority, number of ports, their state, MAC addresses, and maximum SDU sizes, as well as spanning tree type; additionally for token-ring: bridge number and virtual segment • MSS ASRT> list port : list detailed port information, including status, type of bridging supported, etc. • MSS ASRT> list spanning-tree-protocol counters : list number of topology changes and time since last one : list BPDUs sent and received
	Commands	8274	<ul style="list-style-type: none"> • % sts : display spanning tree parameters for the selected group (implies bridging function), such as priority, bridge ID, designated root bridge, etc. • % stc : configure spanning tree parameters for the selected group • % stps : display spanning tree port parameters, like port status, designated bridge (in segment), etc. • % stpc : configure spanning tree port parameters

<i>Table 24 (Page 4 of 4). Hints and Tips with Bridged and Routed ATM Networks</i>			
Topics	Techniques	Products	Examples: Descriptions
Basic IP routing information	Commands	MSS	<ul style="list-style-type: none"> • MSS + protocol ip : enter submenu for IP • MSS IP> interface : list addresses and masks of interfaces • MSS IP> dump <i>network_number</i> : list known IP routes on MSS interface number <i>network_number</i> • MSS IP> ping 192.168.5.99 : ping a certain IP host • MSS IP config> list all : list interface addresses and masks, routing protocols used, and other info (in the static configuration, that is when using talk 6).
	Commands	8274	<ul style="list-style-type: none"> • % ipr : view IP routes • % aizr : add static routes (also default gateway) • % ping 192.168.5.99 : ping a certain IP host • % traceroute 192.168.5.99 : trace an IP route to a certain host

8.5 Case Studies for Bridging and Routing over ATM

The following scenarios provide you cases where the troubleshooting procedure described in the previous sections is applied.

8.5.1 Case Study 1: Bridging and Routing IPX in ATM

For a better understanding of the case study, we recommend that you briefly review the concepts presented in 2.6.3, "IPX/SPX (Novell's NetWare)" on page 37 before continuing.

8.5.1.1 Network Topology

This case study is based on the network represented in Figure 308 and Figure 309 on page 433.

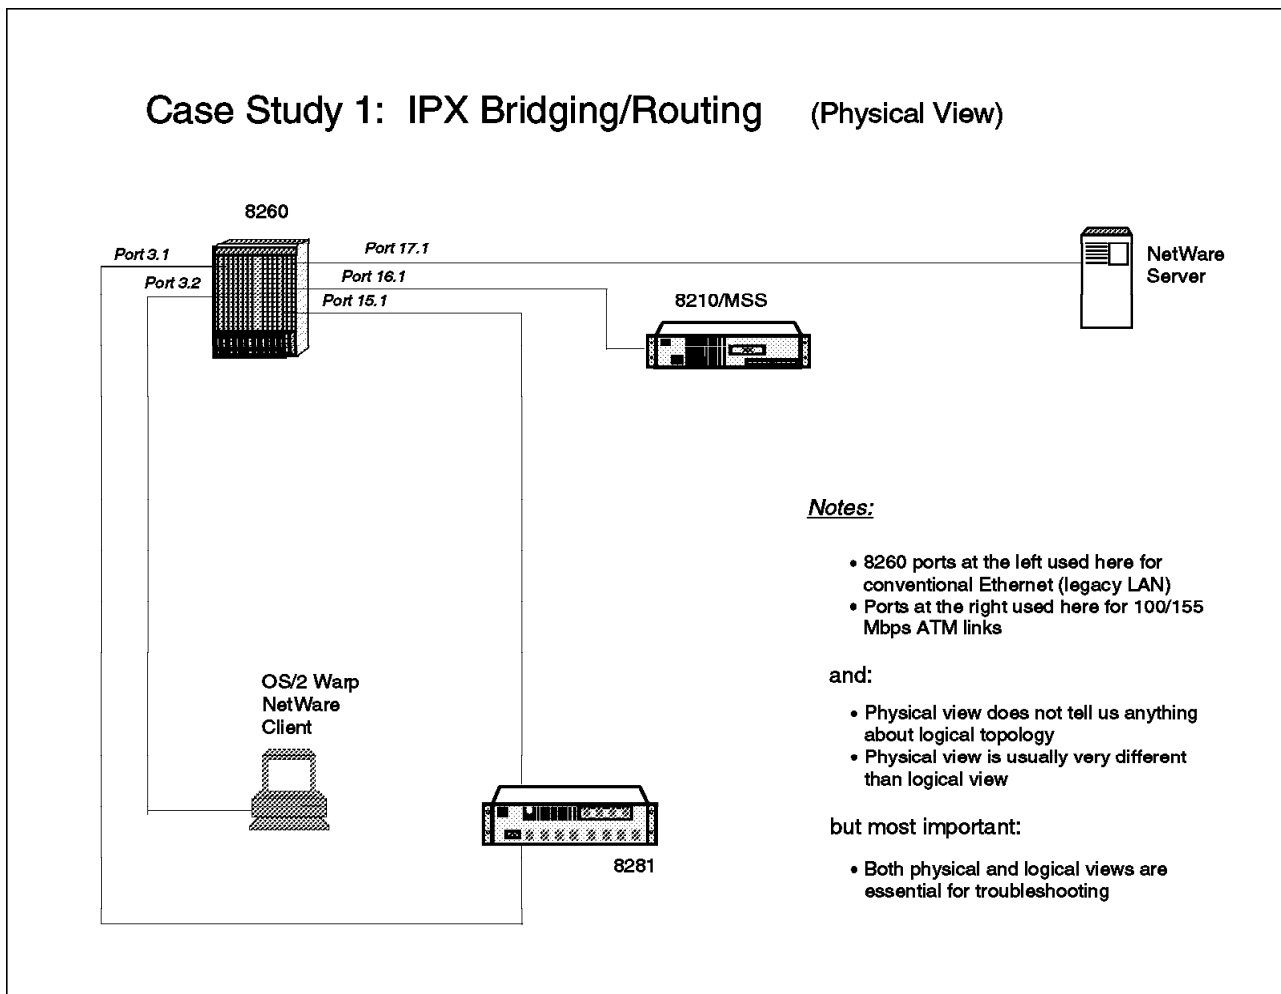


Figure 308. IPX Bridging/Routing (Physical View)

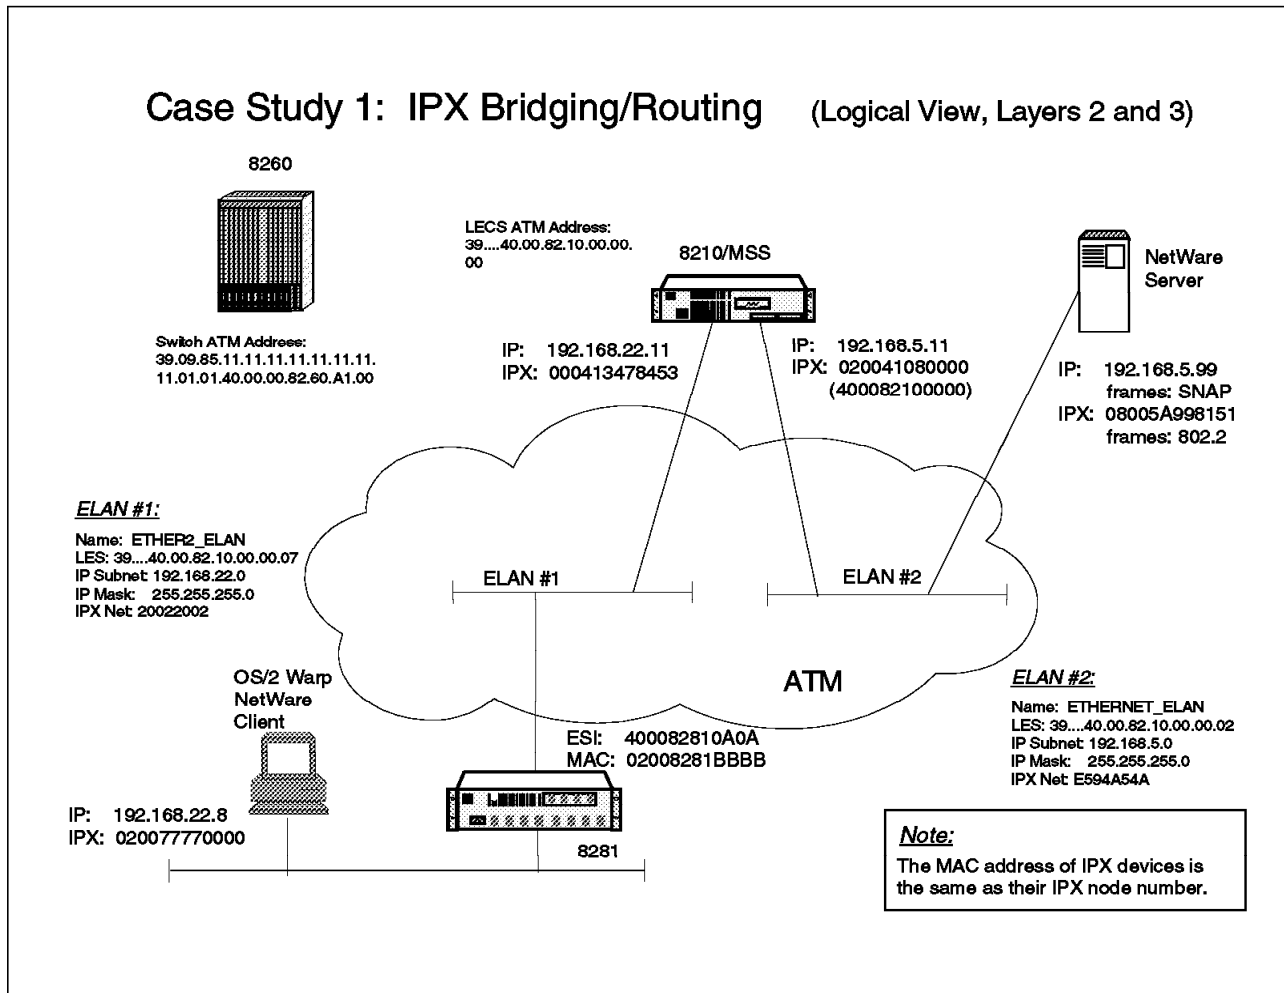


Figure 309. IPX Bridging/Routing (Logical View)

We start our case study by assuming that the network had been functioning well for quite a while, and that recently a new department (represented by the NetWare client at the bottom) has been connected to the ATM environment. The users of the new department complain that they never have been able to communicate with their NetWare file server nor with other TCP/IP users outside their segment, ever since the IBM Nways 8281 ATM LAN Bridge was installed (last week).

8.5.1.2 Methodology

In troubleshooting this problem situation, we follow the methodology we developed in 8.3.2, "PD Methodology for ATM with Bridging and Routing" on page 423.

Phase 1: Gather as much information as possible about environment, symptoms experienced, applications affected, etc. This phase was summarized in the mentioned pictures and in the introduction above.

Phase 2: Investigate recent network changes. This phase, too, has been summarized above. (We were told that the IBM Nways 8281 ATM LAN Bridge was installed last week.)

Phase 3: Interrogate the network based on problem symptoms (1st iteration).

This is the most substantial phase in our case study. In real life, it could perfectly be the other way round (phases 1 and 2 may take a substantial part of the time due to lack of documentation and other appropriate info sources).

We now follow the test steps defined for phase 3 in our methodology:

- **Step 3-1** - Check configuration of test client, and if it is properly connected to the local segment.

We choose a certain user station (our OS/2-based NetWare client in the picture) to make some preliminary end-to-end connectivity tests. First, we need to check if the NetWare client is reasonably set up. So we start by entering the following command at an OS/2 window on the NetWare client to check if the client connects well to the Ethernet LAN:

```
[C:MAGRONLABSIPX]type ibmcomlantran.log
...
...
IBM LANDD is accessing IBM 802.3 LAN Interface.
Adapter 0 was initialized and opened successfully.
Adapter 0 is using node address 020077770000. The
Token-Ring format is 4000EEEE0000.
IBM LANDD was successfully bound to MAC: IBMENI_nif->VECTOR.
ODI2NDI.OS2 was successfully bound to MAC VECTOR.
The current node address for ODI2NDI.OS2 Adapter 0
is 020077770000. The Universal node address is 020077770000.
```

1 shows that the network adapter of the NetWare client can access the Ethernet LAN.

2 and **3** point respectively to the canonical (Ethernet) and noncanonical representations of the MAC address being used by the NetWare client.

4 tells us that the MAC address being used by the NetWare client as part of its IPX network address equals the MAC address we have already seen, which is correct.

The following command shows us that the TCP/IP definitions for the test machine are correct, too (correct address, correct subnet mask, correct frame format, since SNAP means here IEEE 802.3, and correct default gateway).

```
[C:MAGRONLABSIPX]ifconfig lan0
lan0: flags=b863<UP,BROADCAST,NOTRAILERS,RUNNING,BRIDGE,SNAP>
      inet 192.168.22.8 netmask fffffff0 broadcast 192.168.22.255

[C:\MAGRON\LABS\IPX]netstat -r

      destination          router      refcnt      use flags      snmp  intrf
                               metric
      default      192.168.22.11      0           0  U      -1  lan0
      192.168.22.0  192.168.22.8      0          1207  U      -1  lan0
      192.168.22.0  192.168.22.11     0           0  U      -1  lan0
```

- **Step 3-2** - Ping others in local and remote segments to confirm layer 3 connectivity.

We try to reach the NetWare server using IP ping, but notice that it cannot be reached from our NetWare client station:


```
[C:MAGRONLABSIPX]ping 192.168.5.99
PING 192.168.5.99: 56 data bytes

----192.168.5.99 PING Statistics----
5 packets transmitted, 0 packets received, 100% packet loss
```

- **Step 3-3** - Log on to a file server:

We checked if our client workstation was able to do a (default IPX) logon to some server in the network when it came up, but we found that this did not work.

```
[C:MAGRONLABSIPX]nlist server /b

NLIST-4.19-260: You are not attached to a server.
```

- **Step 3-4** - Start host emulation and log on to host to confirm SNA connectivity.

We skip this step, since we are not using SNA in this example.

- **Step 3-5** - Interrogate hub for proper LECS address.

By now we know that we need to do a thorough check through the OSI layers, bottom-up, including of course ATM as our physical and data link layers. For this we use the techniques learned in Chapter 5, "Starting Problem Isolation in an ATM Network" on page 109 and Chapter 7, "ATM Emulated LANs and Logical IP Subnets (LANE 1.0, RFC 1577)" on page 293 and proceed to check the LANE definitions, expecting that the underlying ATM links are okay, since we heard that the rest of the network users are working okay, and only this new department is affected. If the ATM links were not okay, we would have to come back to that and resolve them first.

We continue with the mentioned next step in our methodology which is to check if there is an appropriate pointer (ATM address) to the LECS at the ATM switch where our device (bridge or router) is connected. So we check at the ATM console of the IBM Nways 8260 Multiprotocol Switching Hub if the ATM address of the switch corresponds to the network map we received, and if the switch contains a pointer with the correct address of the LECS:

```

8260ATM1> show device
8260 ATM Control Point and Switch Module
...
...
...
Last Restart : 21:27:43 Sun 6 Jul 97 (Restart Count: 127)

A-CPSW
-----
ATM address: 39.09.85.11.11.11.11.11.11.01.01.
              40.00.00.82.60.A1.00

8260ATM1> show lan_emul configuration_server
Index          ATM address
-----
1 WKA active  39.09.85.11.11.11.11.11.11.01.01.
              40.00.82.10.00.00.00

```

This obviously seems to be the case. This means that LEC will be able to find the LECS to ask for the address of its LES.

- **Step 3-6** - Check status of UNI hub ports where your devices are connected.

We are using an external IBM Nways 8281 ATM LAN Bridge and see here the following status when asking for the port status of the UNI port where the bridge is connected to our IBM Nways 8260 Multiprotocol Switching Hub:

```

8260ATM1> show port 15.1 verbose

      Type  Mode      Status
-----
15.01:UNI enabled  UP-OKAY 1

Signalling Version : with ILMI
Flow Control       : Off
Connector          : MIC
Media              : fiber
Port speed         : 100000 kbps
Remote device is active
IX status          : IX OK 2

```

1 and **2** tell us that the port status is OK. That is the IBM Nways 8281 ATM LAN Bridge was able to register its ATM address with the hub, and it's communicating at least at the lower levels of the ATM hierarchy.

- **Step 3-7** - Interrogate the status of the bridge(s).

We proceed to the next major test step which is to interrogate the status of our bridge (blocking ports, etc.) looking for symptoms which could possibly explain why our NetWare client workstation cannot communicate with the rest of the network. To accomplish this, we utilize the configuration tool for the IBM Nways 8281 ATM LAN Bridge and request a bridge status. The configuration tool shows us the information in Figure 310 on page 437. It shows that the bridge is obviously communicating properly. That is, it forwards the frames as we would expect it to do. In addition, the MAC addresses correspond to those on the network map.

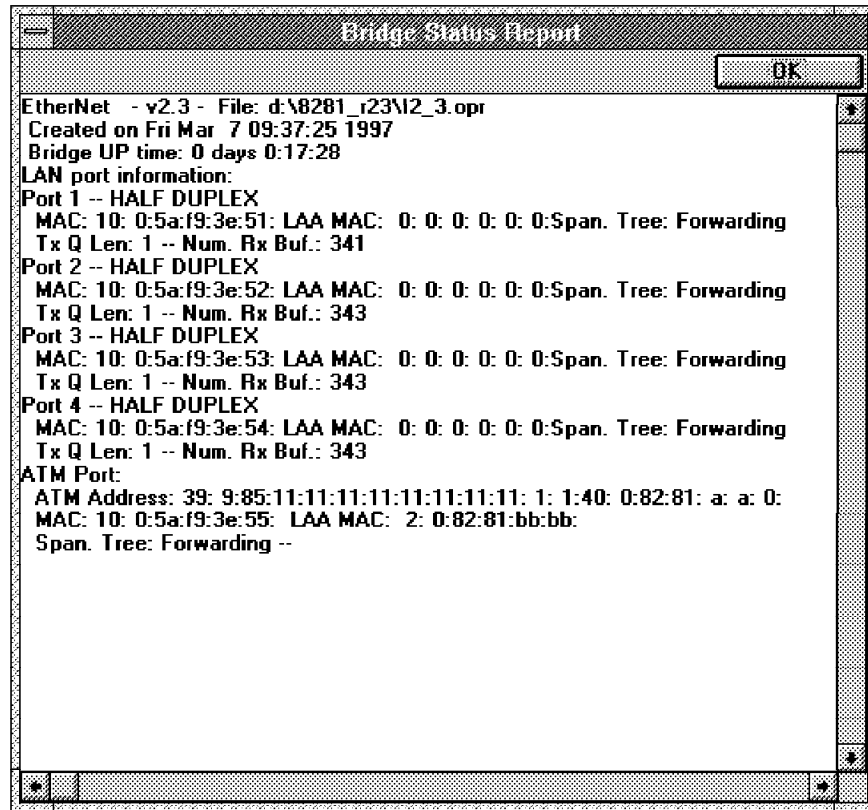


Figure 310. Bridge Status Report (IBM Nways 8281 ATM LAN Bridge)

Everything seems to be working, and nevertheless we cannot reach our file server.

- **Step 3-8** - Log on to your router, and ping the other stations.

Since everything seems to be okay at layer 2 (bridging), we take a closer look at the routing function (layer 3) in the MSS server, which leads us to the mentioned step in our methodology. This will help us to make sure that our router (the MSS server) routes properly between both emulated LANs, called ETHER2_ELAN and ETHERNET_ELAN, respectively. After all, we know for sure that user workstations in the same ELAN as the NetWare server work fine (ETHERNET_ELAN), but we have not heard from any workstations in the other ELAN yet (ETHER2_ELAN).

We log on to the IBM Nways MSS Server and are presented with the following view:

```

login: ibm8210
Password:

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

Cary_MSS * 1

Cary_MSS *status 2
  Pid  Name      Status TTY  Comments
  1    COpCon  IDL   TTY0 ibm8210
  2    Monitr  DET   --
  3    Tasker  IDL   --
  4    MOSDBG  DET   --
  5    CGWCon  IOW   --
  6    Config  DET   --
  7    ROpCon  IDL   TTY1
  8    ROpCon  IDL   TTY2
  9    WEBCon  IDL   --
3
4
5

```

1 is the prompt of the IBM Nways MSS Server.

2 shows the command we entered to display the tasks (programs) running in the IBM Nways MSS Server. The most important ones for troubleshooting are:

- **3** Monitr (task number 2), used to display the (running) event log.
- **4** GWCon (task number 5), used to display and modify the running (dynamic) configuration of the IBM Nways MSS Server.
- **5** Config (task number 6), used to display and modify the static configuration of the MSS server, which is used at the next restart of the system.

When interacting with these tasks we refer to them by their number, so we enter talk 5 or just t 5 to contact the GWCon task (to display data of the running configuration), and talk 6 or just t 6 to contact the Config task (to display data of the static configuration).

We proceed to ping the test client from our router:

```

Cary MSS IP>ping 192.168.22.8
PING 192.168.22.11 -> 192.168.22.8: 56 data bytes, ttl=64, every 1 sec.

----192.168.22.8 PING Statistics----
5 packets transmitted, 0 packets received, 100% packet loss

```

We notice that it cannot be reached.

- **Step 3-9** - IPXping other routers:

We do not have any IPX routers other than the IBM Nways MSS Server in the sample network, so we skip this step.

- **Step 3-10** - Interrogate the status of your router(s).

Here we want to check our IP definitions and connectivity:

```
Cary_MSS *talk 6 1
Gateway user configuration

Cary MSS Config>protocol ip 2
Internet protocol user configuration

Cary MSS IP config>list all 3
Interface addresses
IP addresses for each interface: 4
  intf 0                ... IP disabled on this interf ...
  intf 1                ... IP disabled on this interf ...
  intf 2  192.168.5.11  255.255.255.0  ...
  intf 3  192.168.22.11 255.255.255.0  ...

Routing

Protocols
BOOTP forwarding: disabled
IP Time-to-live: 64
Source Routing: disabled
Echo Reply: enabled
Directed broadcasts: enabled
ARP subnet routing: enabled
ARP network routing: enabled
Per-packet-multipath: disabled
OSPF: disabled 5
BGP: disabled 6
RIP: disabled 7

Cary MSS IP config>exit
Cary MSS Config><CTRL/P> 8
Cary_MSS *
```

1 This command allows us to talk to the task for displaying and modifying the static configuration of the MSS server (Config task). We do not want to make any changes, but to only display some parameters which we know are common to both the static configuration and the active (running) configuration. We look at them here, since this is the better interface to display them with the least number of commands.

2 With this command we enter the IP submenu.

3 This command is used to display the known IP configuration of the running system that:

4 Shows the IP addresses of the interfaces.

5 Shows that OSPF is not being used as the IP routing protocol.

6 Shows that RIP is not being used as the IP routing protocol.

7 Shows that BGP is not being used as the IP routing protocol. Since neither OSPF, nor RIP, nor BGP are being used, we know that the workstations in the network can only communicate with stations in other subnets if they properly define the IBM Nways MSS Server as their default gateway.

8 The <CTRL/P> command allows us to exit the Config task, so we can later communicate with some other MSS server task. Please note

that the <CTRL/P> was added for clarity, and that the MSS server will not display it.

```

Cary_MSS *talk 5 9
CGW Operator Console

Cary_MSS +protocol ip 10

Cary_MSS IP>dump 0 11
Type  Dest net      Mask      Cost    Age      Next hop(s)
Dir*  192.168.5.0    FFFFFFF0  1       35689   Eth/0
Dir*  192.168.22.0   FFFFFFF0  1       35689   Eth/1

Routing table size: 768 nets (49152 bytes), 2 nets known

Cary_MSS IP>interface 12
Interface  IP Address(es)  Mask(s)
Eth/0     192.168.5.11   255.255.255.0
Eth/1     192.168.22.11  255.255.255.0

Cary_MSS IP>ping 192.168.5.99 13
PING 192.168.5.11 -> 192.168.5.99: 56 data bytes, ttl=64, every 1 sec.
56 data bytes from 192.168.5.99: icmp_seq=0. ttl=128. time=0. ms
56 data bytes from 192.168.5.99: icmp_seq=1. ttl=128. time=0. ms

----192.168.5.99 PING Statistics----
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max = 0/0/0 ms

Cary_MSS IP>ping 192.168.22.8 14
PING 192.168.22.11 -> 192.168.22.8: 56 data bytes, ttl=64, every 1 sec.

----192.168.22.8 PING Statistics----
9 packets transmitted, 0 packets received, 100% packet loss

Cary_MSS IP>exit 15
Cary_MSS Config><CTRL/P> 16

Cary_MSS *

```

9 This command is used to contact the GWCon task to look at the active (running) MSS server configuration.

10 We use this command to enter the TCP/IP submenu.

11 With this command, we display the known routes.

12 With this command, we display the interface addresses.

13 Here we attempt to ping the NetWare server (and reach it).

14 And our NetWare client test workstation (and miss it).

15 Finally, we leave the IP submenu.

16 Leave the Config task, too.

- **Step 3-11** - Interrogate LES/BUS to determine clients connected to the ELAN:

Test steps on layer 2 (bridging) and layer 3 (routing) functions have shown the components work, so something below those layers (connections at physical and up to the MAC layer) could be the cause of the error. This is usually some kind of logical error, made inadvertently.

We proceed with the mentioned new step, to check if the LAN emulation clients (LECs) have joined the emulated LAN properly. We are especially interested to see if the IBM Nways 8281 ATM LAN Bridge has properly joined the ELAN called ETHER2_ELAN according to our network diagram (see Figure 309 on page 433).

Back at the MSS server, we enter the submenus and commands appropriate to check the desired ELAN definitions (hidden below network 0 which is the ATM interface of the MSS server holding the definitions for LAN emulation).

```

Cary_MSS +network ?
0 : CHARM ATM
1 : SK-NET FDDI
2 : ATM Ethernet LAN Emulation: ETHERNET_ELAN
3 : ATM Ethernet LAN Emulation: ETHER2_ELAN
4 : IP Protocol Network
5 : Bridge Application
Network number [0]? 0
ATM Console

Cary_MSS ATM+le-services
LE-Services Console

Cary_MSS LE-SERVICES+list
ELAN Type (E=Ethernet/802.3, T=Token Ring/802.5)
] Interface #
] ] LES-BUS State (UP=Up, ID=Idle, ND=Net Down, ...
] ] ] **=Other; Work with specific LES-BUS ...
] ] ]
] ] ] ELAN Name          LES ATM Addr
-----
E 0 UP ETHER2_ELAN      3909851111111111111111111101014000821000007
E 0 UP ETHERNET_ELAN   3909851111111111111111111101014000821000002

```

So far, everything seems correct. The LES for both ELANS (ETHER2_ELAN and ETHERNET_ELAN) are up and running. We continue by checking if the IBM Nways 8281 ATM LAN Bridge has joined ETHER2_ELAN properly, and display the LAN emulation clients (LECs) which have joined that ELAN.

```

Cary_MSS LE-SERVICES+work ETHER2_ELAN
LE-Services Console for an existing LES-BUS Pair
Cary_MSS EXISTING LES-BUS 'ETHER2_ELAN'+data list all lec
Number of LEC's to display: 1

    LEC-LES and LEC-BUS State (UP=Up, ID=Idle, --. --.
    **=Other; Show specific LEC to see actual)      v  v
                                                    LEC   State   ...
LEC Primary ATM Address          Proxy ID   LES BUS ...
-----
39098511111111111111111111010100041347845308  N 0001  UP  UP   ...

Cary_MSS EXISTING LES-BUS 'ETHER2_ELAN'+exit

```

To our big surprise, the 8281 ATM LAN bridge is missing. The only LEC present is the (router) LEC of the IBM Nways MSS Server in that ELAN (check our network map with the addresses). But, where is the 8281 then? We saw earlier that the 8281 had its ATM port in forwarding state, so it had joined an ELAN. Let us check the other ELAN:

```

Cary_MSS LE-SERVICES+work ETHERNET_ELAN
LE-Services Console for an existing LES-BUS Pair

Cary_MSS EXISTING LES-BUS 'ETHERNET_ELAN'+data list 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  ID   LES  State  ...
----- - ---- -- --
39098511111111111111111010140008210000003  N 0001  UP  UP    ...
390985111111111111111110101400082810A0A00  Y 0002  UP  UP    ...
39098511111111111111111010108005A99815181  N 0003  UP  UP    ...

Cary_MSS EXISTING LES-BUS 'ETHERNET_ELAN'+exit

Cary_MSS LE-SERVICES+exit
Cary_MSS ATM+ <CTRL/P>

Cary_MSS *

```

The IBM Nways 8281 ATM LAN Bridge joined the wrong ELAN, and therefore our NetWare client was in the wrong subnet, and could not be reached nor reach others with a ping. The other two LECs are the LEC from the MSS server in that ELAN (first line displayed) and the NetWare file server (third line displayed). This does not explain though, why the client could not communicate with its NetWare server, but we will check this later. It is obvious that something must be wrong with the setup of the 8281, so the first thing we need to do is get it connected to the right ELAN. Since we found the source of the problem, we exit this iteration of phase 3 tests and continue with phase 4.

Phase 4: Resolve problem (make and test corrections), or investigate a different symptom.

This is the phase where we correct the wrong set up items found during our diagnostic test steps. We go back to the configuration tool of the IBM Nways 8281 ATM LAN Bridge, and check the ELAN definitions.

Control time-out (10-300 seconds) :	10
VCC time-out (1 - 65,535 seconds) :	20
Flush time-out (1-4 seconds) :	4
Path Switching Delay (1-8 seconds) :	6
Connection Completion timeout (1-10 seconds) :	4
Max LE_ARP Response time (1-30 seconds) :	1
Max LE_ARP Retry count (0-2 retries) :	1
Max Unknown frame count (1-10 frames) :	1
Data Direct VCC Up after control failure (0-3600) :	0
Emulated LAN name (0-32 chars) :	ETHERNET_ELAN
LE_ARP cache aging timeout	
Aging Time (also used for TB filtering database):	300
Forward delay time (also used for spanning tree):	15

OK Cancel Help

Control time-out period.

Figure 311. IBM Nways 8281 ATM LAN Bridge: Wrong LAN Emulation Settings

We see in Figure 311 that the ELAN name parameter is wrong (it reads ETHERNET_ELAN instead of ETHER2_ELAN), and correct it as shown in Figure 312 on page 444.

Parameter	Value
Control time-out (10-300 seconds) :	120
VCC time-out (1 - 65,535 seconds) :	20
Flush time-out (1-4 seconds) :	4
Path Switching Delay (1-8 seconds) :	6
Connection Completion timeout (1-10 seconds) :	4
Max LE_ARP Response time (1-30 seconds) :	1
Max LE_ARP Retry count (0-2 retries) :	1
Max Unknown frame count (1-10 frames) :	1
Data Direct VCC Up after control failure (0-3600) :	0
Emulated LAN name (0-32 chars) :	ETHER2_ELAN
LE_ARP cache aging timeout	
Aging Time (also used for TB filtering database):	300
Forward delay time (also used for spanning tree):	15

Buttons: OK, Cancel, Help

Footer: Emulated LAN name

Figure 312. IBM Nways 8281 ATM LAN Bridge: Corrected LAN Emulation Settings

We then download the new configuration to the IBM Nways 8281 ATM LAN Bridge which makes it restart using the new parameter.

We then need to verify if the correction has given the correct results, so we enter some commands at the IBM Nways MSS Server console to check if the 8281 has now joined ETHER2_ELAN:

```

Cary_MSS *talk 5
CGW Operator Console

Cary_MSS +network 0
ATM Console

Cary_MSS ATM+le-services
LE-Services Console

Cary_MSS LE-SERVICES+work ETHER2_ELAN
LE-Services Console for an existing LES-BUS Pair
Cary_MSS EXISTING LES-BUS 'ETHER2_ELAN'+data list 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  State    ...
LEC Primary ATM Address          Proxy ID    LES BUS    ...
-----
39098511111111111111111111111111010100041347845308  N 0001  UP  UP    ...
390985111111111111111111111111110101400082810A0A00  Y 0002  UP  UP    ...

Cary_MSS EXISTING LES-BUS 'ETHER2_ELAN'+exit

```

We see that this is the case. The 8281 has now been accepted by the corresponding LES and joined the ELAN.

Just for the sake of completeness, we check the other ELAN, too.

```

Cary_MSS LE-SERVICES+work ETHERNET_ELAN
LE-Services Console for an existing LES-BUS Pair

Cary_MSS EXISTING LES-BUS 'ETHERNET_ELAN'+data list 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  State    ...
LEC Primary ATM Address          Proxy ID    LES BUS    ...
-----
39098511111111111111111111111111010140008210000003  N 0001  UP  UP    ...
39098511111111111111111111111111010108005A99815181  N 0003  UP  UP    ...

Cary_MSS EXISTING LES-BUS 'ETHERNET_ELAN'+exit
Cary_MSS LE-SERVICES+exit
Cary_MSS ATM+ <CTRL/P>

Cary_MSS *

```

We see that now only the IBM Nways MSS Server and the Novell server have joined the ELAN, as expected. Well, now everything seems OK, since the IBM Nways 8281 ATM LAN Bridge forwards frames, and as we know, now to the correct ELAN.

According to our test methodology, after phase 4 where we corrected the error and verified the correction, we should go to phase 3 for the next test iteration. We proceed accordingly.

Phase 3 (2nd iteration) - Next test steps

- **Step 3-1** - Check configuration and proper LAN connection of test client.

Here again we re-check the definitions of our NetWare client. (Is it connected to the network?)

```
[C:\MAGRONLABSIPX]ifconfig lan0
lan0: flags=b863<UP,BROADCAST,NOTRAILERS,RUNNING,BRIDGE,SNAP>
      inet 192.168.22.8 netmask ffffffff broadcast 192.168.22.255

[C:\MAGRON\LABS\IPX]netstat -r

      destination          router      refcnt      use flags      snmp intrf
                        metric
      default  192.168.22.11      0           0  U           -1  lan0
192.168.22.0  192.168.22.8        0          1253  U           -1  lan0
192.168.22.0  192.168.22.11      0           0  U           -1  lan0
```

- **Step 3-2** - Ping other network participants.

We verify if our test client can reach its Novell server over IP, and find that this is now the case.

```
[C:\MAGRONLABSIPX]ping 192.168.5.99
PING 192.168.5.99: 56 data bytes
64 bytes from 192.168.5.99: icmp_seq=1. time=0. ms
64 bytes from 192.168.5.99: icmp_seq=2. time=0. ms

----192.168.5.99 PING Statistics----
2 packets transmitted, 2 packets received, 0% packet loss
round-trip (ms)  min/avg/max = 0/0/0
```

- **Step 3-3** - Log on to some LAN server.

We check again if our test client station can reach its server (make a default logon).

```
[C:\MAGRONLABSIPX]nlist server /b

NLIST-4.19-260: You are not attached to a server.
```

It is still not connected. Well, we expected this to occur, since we could not reach the server via IPX in our first test iteration. If everything was okay, our client should then have reached its server, since both were sitting on the same ELAN, due to the IBM Nways 8281 ATM LAN Bridge which was incorrectly connected to that ELAN, too.

To illustrate the value of proceeding along the steps of our methodology, we assumed so far that we had forgotten to check the IPX configuration on our test client (actually the first step in phase 3). We did this only for IP, right? Even more, we do not know which IPX frame encapsulation should be in use on ETHER2_ELAN. We only know this for ETHERNET_ELAN, since the Novell server was shown to use Ethernet_802.2 (mentioned on Figure 309 on page 433). Since the IBM Nways MSS Server is the IPX router for ETHER2_ELAN, all clients need to use the same frame encapsulation as the IBM Nways MSS Server; this could be the same as the one used on ETHERNET_ELAN (best option for maintainability

of the network), or a different one (not recommended, unless dictated by some application). In order to review the frames used in NetWare networks we refer back to 2.6.3, "IPX/SPX (Novell's NetWare)" on page 37.

It seems that we need to gather more data (IPX encapsulation used by IBM Nways MSS Server), and this leads us back to phase 1 of our methodology.

Phase 1 Gathering (more) data

At the IBM Nways MSS Server console we check to see which of the MSS server interfaces (LECs) is connected to the same ELAN (ETHER2_ELAN) as our NetWare client station, so we can afterwards look at how it had been configured:

```
Cary_MSS +network ?
0 : CHARM ATM
1 : SK-NET FDDI
2 : ATM Ethernet LAN Emulation: ETHERNET_ELAN
3 : ATM Ethernet LAN Emulation: ETHER2_ELAN
4 : IP Protocol Network
5 : Bridge Application
Network number [0]? 3
ATM Emulated LAN Console
```

1
2

1 tells us that interface number 2 of the MSS server is being used by the MSS server LEC used as the router leg into ETHERNET_ELAN.

2 tells us that interface number 3 is used by the LEC for ETHER2_ELAN.

We therefore review the configuration of interface number 3:

```
Cary_MSS LEC+list config
```

```
ATM LEC Configuration
```

```
Physical ATM interface number = 0
LEC interface number         = 3
LECS auto configuration      = Yes
Default LECS ATM address     = 00.00.00.00.00.00.00.00.00.00.
                             00.00.00.00.00.00.00.00.00
```

```
C1: Primary ATM address
```

```
    ESI address               = Use burned in addr
    Selector byte             = 0x8
C2: Emulated LAN type        = Ethernet
C3: Maximum frame size       = 1516
C5: Emulated LAN name        = ETHER2_ELAN
C6: LE Client MAC address    = Use burned in addr
C7: Control timeout          = 30
C10: Maximum unknown count   = 10
C11: Maximum unknown time    = 1
C12: VCC timeout period      = 1200
C13: Maximum retry count     = 1
C17: Aging time              = 300
C18: Forward delay time      = 15
C20: LE ARP response time    = 1
C21: Flush timeout           = 4
C22: Path switch delay       = 6
C24: Multicast send VCC type = Best-Effort
C25: Multicast send VCC avg rate = 155000
C26: Multicast send VCC peak rate = 155000
C28: Connection completion timer = 4
```

```
LE ARP queue depth           = 5
LE ARP cache size            = 10
Best effort peak rate        = 155000
Maximum config retries       = 3
Packet trace                  = No
NetWare IPX encapsulation    = ETHERNET_802.2
IP Encapsulation              = IEEE-802.3
```

```
1
2
```

```
Cary_MSS LEC+exit
Cary_MSS +
```

1 shows that the MSS server indeed uses the expected ETHERNET_802.2 frame encapsulation for IPX, and

2 shows that it uses IEEE 802.3 frame encapsulation for IP. IEEE 802.3 in the language of IBM Nways MSS Server is the same as ETHERNET_SNAP in the language of NetWare networks. You might want to take a glance at Figure 17 on page 39 for a review of NetWare's frame types.

Well, now that we have the information, we can proceed with the applicable phase of the methodology which is phase 3 with the tests. (We jump over phase 2, investigation of recent changes, assuming that there have not been more changes.)

Phase 3 (3rd iteration) - More test steps

- **Step 3-1** - Check configuration and proper LAN connection of test client.

At an OS/2 window of our client's screen we check the PROTOCOL.INI file in the IBMCOM directory, and see the following:

```
[ODI2NDI_nif]
DriverName = odi2ndi$
Bindings = IBMENI_nif
NETADDRESS = "020077770000"
TOKEN-RING = "no"
TOKEN-RING_SNAP = "no"
ETHERNET_802.3 = "yes"
ETHERNET_802.2 = "no"
ETHERNET_II = "no"
ETHERNET_SNAP = "no"
TRACE = 0x0
```

1
2

Well, it could not be more wrong:

- 1 shows that the client uses ETHERNET_802.3 encapsulation for IPX.
- 2 shows that it does not use ETHERNET_802.2, as we actually want it to do.

Being very strict on the methodology, we now exit phase 3, proceed to phase 4 where we correct and verify the configuration, and then we re-enter the tests of phase 3 at this very step again. For a better overview, and since we are at the very first test step of this phase, we just correct these simple errors, and continue with the next test steps.

We therefore correct the client definitions as follows:

```
[ODI2NDI_nif]
DriverName = odi2ndi$
Bindings = IBMENI_nif
NETADDRESS = "020077770000"
TOKEN-RING = "no"
TOKEN-RING_SNAP = "no"
ETHERNET_802.3 = "no"
ETHERNET_802.2 = "yes"
ETHERNET_II = "no"
ETHERNET_SNAP = "no"
TRACE = 0x0
```

1

1 now shows that the client uses the correct encapsulation, so we restart the machine.

- **Step 3-2** - Ping others.

From the previous iteration we know that IP ping will work, so we continue directly with the next applicable step.

- **Step 3-3** - Log on to a LAN server.

Here we check if a default IPX logon has taken place (the client connects by default to the nearest server in NetWare networks), and enter the following command at the terminal of our test client workstation:

```
[C:MAGRONLABSIPX]nlist server /b
Object Class: Server
Known to Server: PCSRV320
Active NetWare Server= The NetWare Server that is currently running
Address                = The network address
Node                   = The network node
Status                 = The status of your connection

Active NetWare Server      Address      Node      Status
-----
PCSRV320                  [A76ACA43][          1]Default
One server object was found.
```

It seems that we have mastered the problem now.

8.5.1.3 Conclusion

This case study illustrates all phases and most of the test steps of our troubleshooting methodology.

It also shows that these kinds of problems are not always due to a single configuration error, and may affect several devices and go through all possible OSI layers:

- Our IBM Nways 8281 ATM LAN Bridge was attached to the wrong segment, due to wrong ELAN definitions, just above the ATM layer (which would equate to a wrong cable connection at the physical layer on legacy LANs).
- The client station used the wrong IPX frame type, which was bad both for layer 2 and layer 3, since the LLC layer (in layer 2) could not properly pass the frame to the IPX subsystem (layer 3) due to the wrong encapsulation method. This could have been a wrong set up of the router (IBM Nways MSS Server), but was not the case here; this time the frame conversion error was at the client.
- Due to the IBM Nways 8281 ATM LAN Bridge joining the wrong ELAN, the client station was originally attached to the wrong IP subnet. A routing problem from the perspective of the client, since it could not reach its default IP gateway. This was therefore a layer 3 problem.

We could also observe that a good understanding of the LAN concepts and frame structures can greatly help to quickly pinpoint the reasons for network problems, and bring them to a satisfactory resolution.

8.5.2 Case Study 2: Spanning Tree Loop

For a better understanding of the case study, we recommend that you briefly review the concepts presented in 2.5.1, “Transparent Bridging (TB)” on page 27 and 2.5.1.1, “Spanning Tree Protocol (STP)” on page 28, before continuing.

8.5.2.1 Network Topology

This case study is based on the network represented in Figure 313 on page 451 and Figure 314 on page 452.

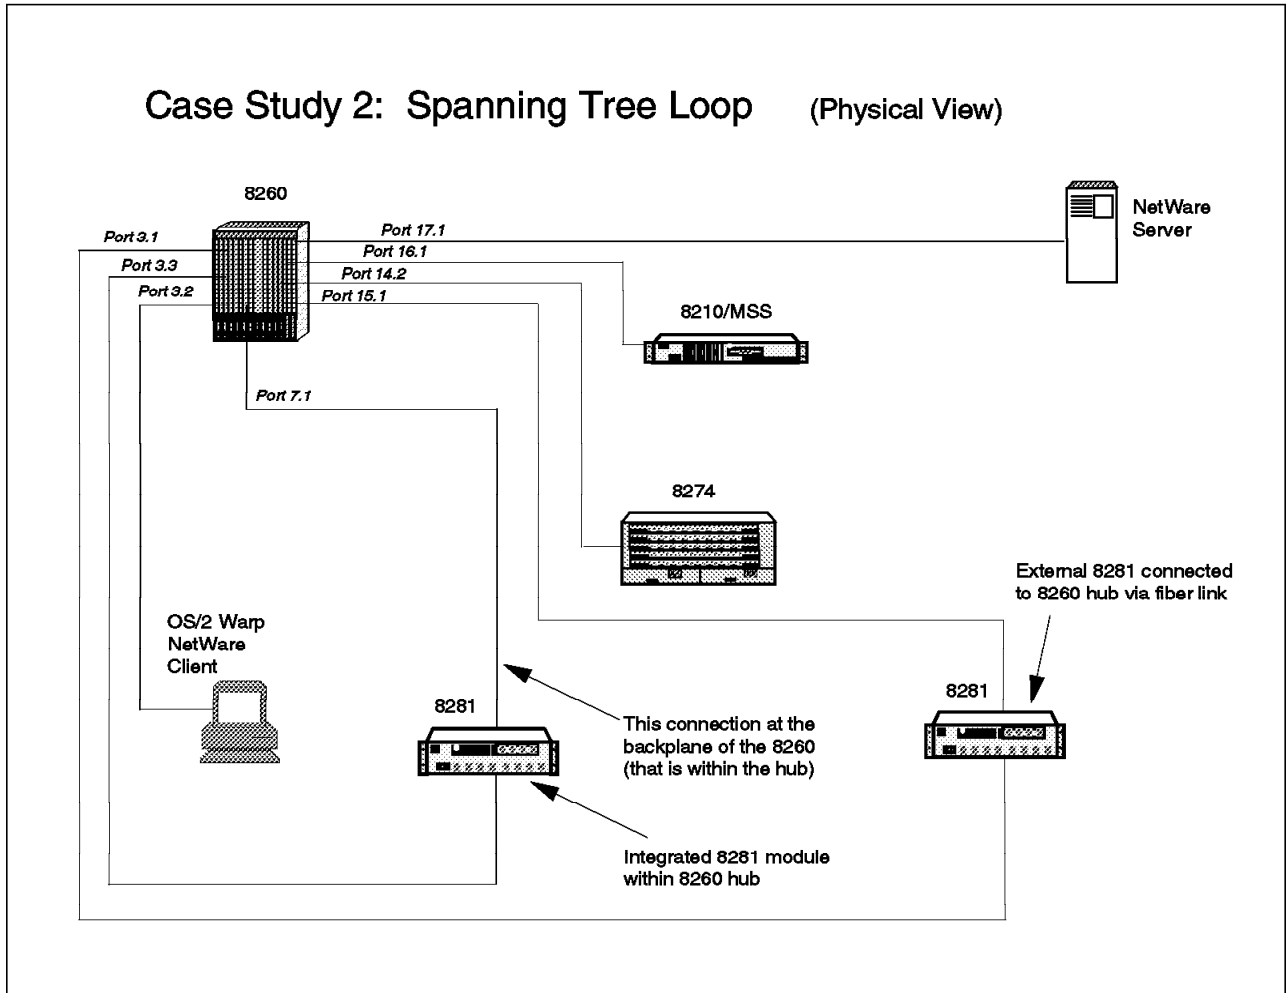


Figure 313. Spanning Tree Loop (Physical View)

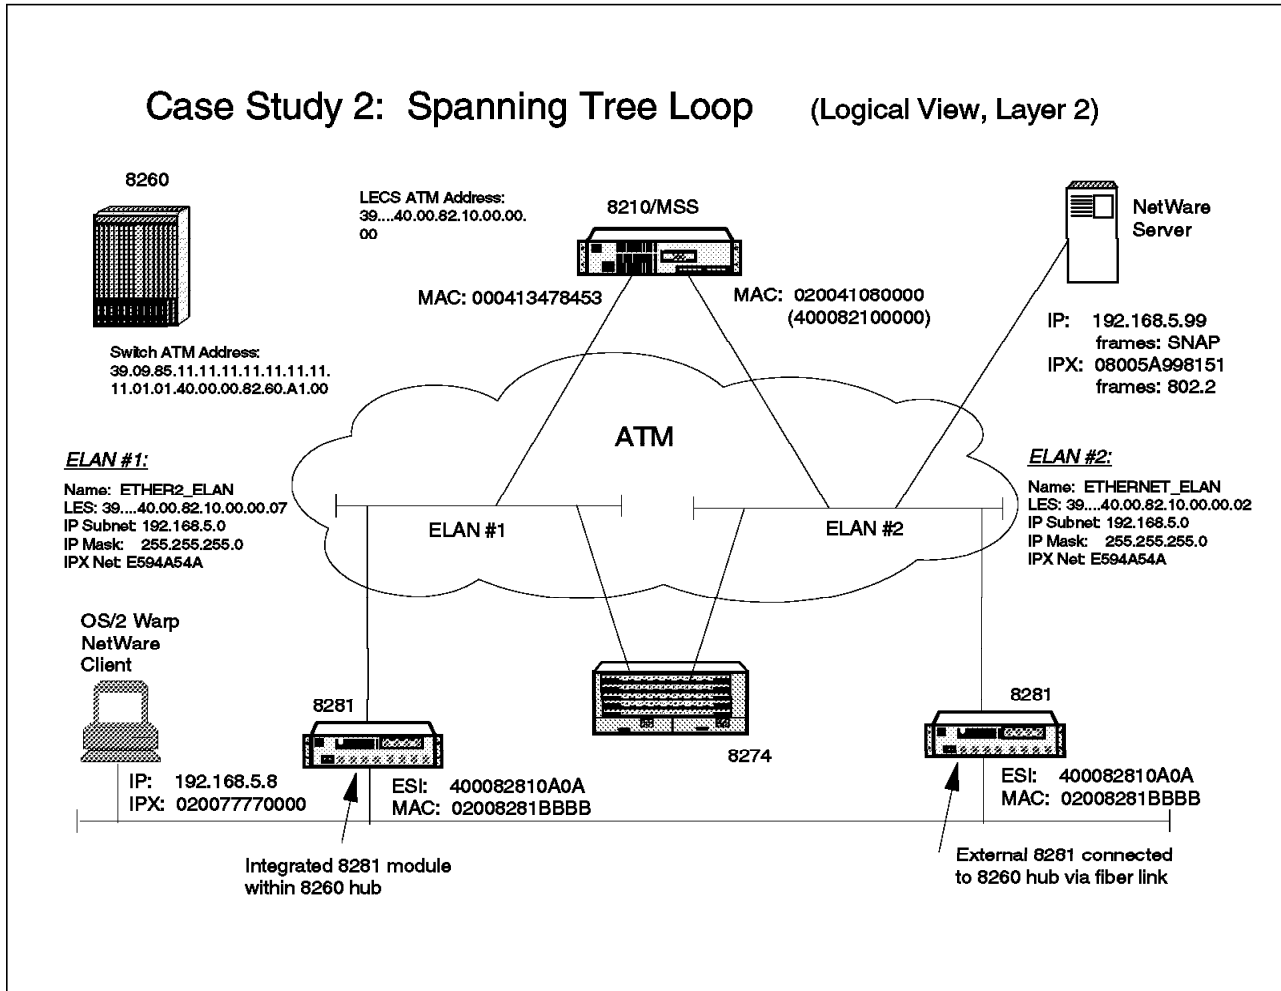


Figure 314. Spanning Tree Loop (Logical View)

We start our case study assuming that the network had been functioning well for quite a while. We have added an IBM Nways 8274 RouteSwitch to the network schemes for clarity of the following explanations, although our original information did not contain any reference to this box being included in the network. We were called to the place to resolve a severe performance problem being experienced the whole day, making use of part of the network impossible (the part displayed in our schemes).

8.5.2.2 Methodology

In troubleshooting this problem situation, we follow the methodology we developed in 8.3.2, "PD Methodology for ATM with Bridging and Routing" on page 423.

Phase 1: Gather as much information as possible about environment, symptoms experienced, applications affected, etc. This phase was summarized in the mentioned pictures and in the introduction above.

Phase 2: Investigate recent network changes. This phase, too, has been summarized above, although we do not yet know about the IBM Nways 8274 RouteSwitch being attached to the network.

Phase 3: Interrogate the network based on problem symptoms (1st iteration).

We again choose a test station (OS/2 Warp NetWare client) to check the symptoms experienced in the network (bad performance), in order to get a better clue as to the reasons for this network behavior. We jump then directly into the first step of our methodology.

- **Step 3-1** - Check configuration and proper LAN connection of test client.

We quickly confirm that our test client is indeed properly connected to the LAN, since the following command says that the interface is operational.

```
[C:\MAGRONLABS\STP]ifconfig lan0
lan0: flags=b863<UP,BROADCAST,NOTRAILERS,RUNNING,BRIDGE,SNAP>
      inet 192.168.5.8 netmask fffffff0 broadcast 192.168.5.255

[C:\MAGRON\LABS\STP]netstat -r

      destination          router      refcnt      use flags      snmp  intrf
                        metric
      default      192.168.5.11      0           0  U      -1  lan0
      192.168.5.0   192.168.5.8      0          5801  U      -1  lan0
      192.168.5.0   192.168.5.11     0           0  U      -1  lan0
```

- **Step 3-2** - Ping other network participants.

We then try to reach the NetWare server using IP ping, but notice that it cannot be reached from our NetWare client station.

```
[C:\MAGRONLABS\STP]ping 192.168.5.99
PING 192.168.5.99: 56 data bytes

----192.168.5.99 PING Statistics----
7 packets transmitted, 0 packets received, 100% packet loss
```

- **Step 3-3** - Log on to a LAN server

Checking the (default IPX) logon to some server in the network yields a similar result, even rebooting the machine produced analogous results.

```
[C:\MAGRONLABS\STP]nlist server /b

NLIST-4.19-260: You are not attached to a server.
```

- **Step 3-4** - Start host emulation and log on to host to confirm SNA connectivity.

We skip this step, since it is not applicable to our test environment.

- **Step 3-5** - Interrogate hub for proper LECS address.

Here, too, we check to see if the LECS address has been added to the ATM switch where our bridges are connected, and find everything okay:

```

8260ATM1> show device
8260 ATM Control Point and Switch Module
...
...
...
Last Restart : 21:27:43 Sun 6 Jul 97 (Restart Count: 127)

A-CPSW
-----
ATM address: 39.09.85.11.11.11.11.11.11.11.01.01.
              40.00.00.82.60.A1.00

8260ATM1> show lan_emul configuration_server
Index          ATM address
-----
1 WKA active 39.09.85.11.11.11.11.11.11.11.01.01.
              40.00.82.10.00.00.00

```

- **Step 3-6** - Check status of UNI hub ports where your devices are connected.

We find that the UNI ports for the external and the internal IBM Nways 8281 ATM LAN Bridge look okay. For illustration purposes, we display here the one for the external model, knowing that the other would look somewhat similar for our purposes. (Actually the internal IBM Nways 8281 ATM LAN Bridge port would show more information since the device is integrated into the IBM Nways 8260 Multiprotocol Switching Hub.)

```

8260ATM1> show port 15.1 verbose

      Type  Mode      Status
-----
15.01:UNI enabled UP-OKAY 1

Signalling Version : with ILMI
Flow Control       : Off
Connector          : MIC
Media              : fiber
Port speed         : 100000 kbps
Remote device is active
IX status          : IX OK 2

```

1 and **2** mean that the port status is okay. That is, the IBM Nways 8281 ATM LAN Bridge was able to register its ATM address with the hub, and is communicating at least at the lower levels of the ATM hierarchy.

But when we look at the traffic LEDs on the corresponding IBM Nways 8281 ATM LAN Bridge modules, we begin to understand. They are steadily lit, meaning that there is heavy traffic on the network. Since we know that the users cannot work, this might be a spanning tree problem. We will have to check the status of our bridges in the corresponding steps.

When checking those ATM modules where the traffic LEDs are steadily lit, we discover that the following ports are affected:

- Port 7.1 with the integrated IBM Nways 8281 ATM LAN Bridge
- Port 14.2 with an unknown device (fiber connection leads to a test lab according to the site network manager)
- Port 15.1 with the external IBM Nways 8281 ATM LAN Bridge

- Port 17.1 with the NetWare server

We definitely need to know what is connected to port 14.2, in order to have a better picture of the situation. Remember that gathering all possible kinds of information is part of phase 1 (and phase 2) in our methodology. But since that device is in another room on a different floor, we decide to first check the status of the bridges installed just within or next to the IBM Nways 8260 Multiprotocol Switching Hub. After all, it only has one fiber connection to the IBM Nways 8260 Multiprotocol Switching Hub and this should not do any harm, should it? We therefore jump over to the next step, keeping the device at port 14.2 in mind.

- **Step 3-7** - Interrogate the status of the bridge(s).

Here we quickly check the status of our IBM Nways 8281 ATM LAN Bridges with the corresponding configuration tool. We request a bridge status report from both bridges, as shown in Figure 315 and Figure 316 on page 456, which tell us about port addresses and especially their status.

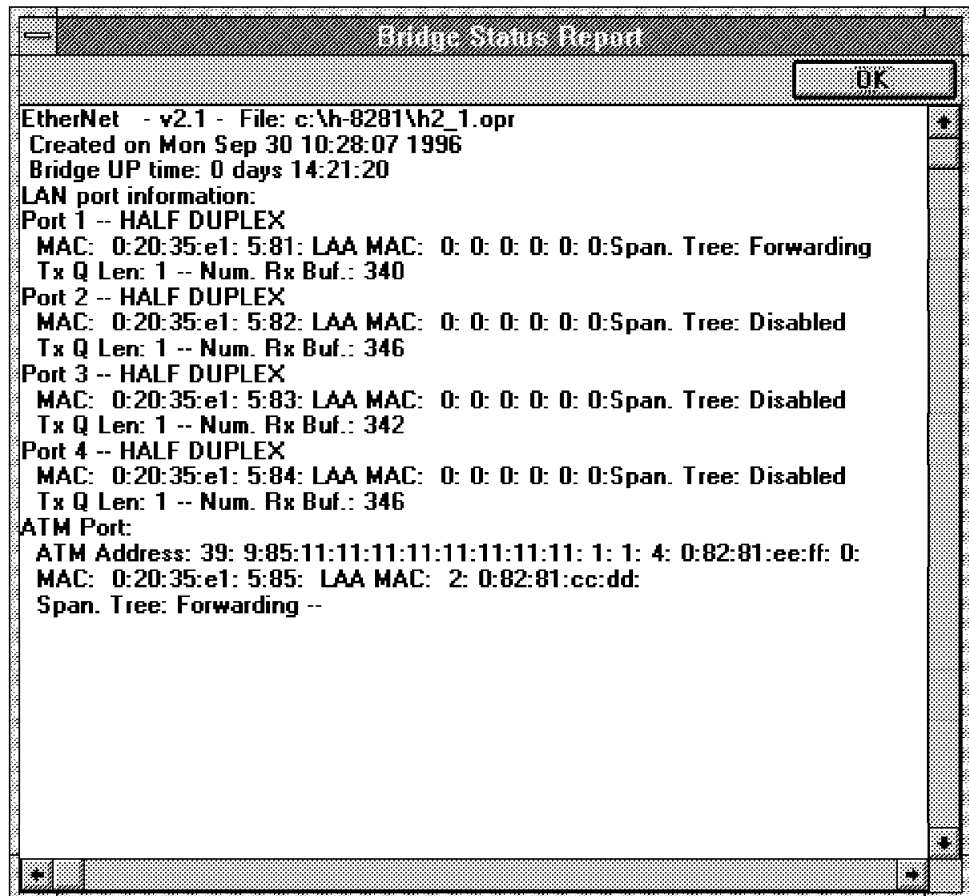


Figure 315. Internal 8281: Bridge Status Report




Figure 316. External 8281: Bridge Status Report

Both 8281 bridges have their respective ports in forwarding state. We know from the IBM Nways 8281 ATM LAN Bridge configuration, that there is no way to disable participation in the spanning tree for this device without explicitly disabling a port. Therefore, the reason for our problems should not be with either of the two 8281 bridge devices, at least in a first approximation.

We then check the configuration of the next bridging device, which is the IBM Nways MSS Server:

```

Cary_MSS *t 5

Cary_MSS +
Cary_MSS +prot asrt

Cary_MSS ASRT>list bridge
Bridge ID (prio/add):      32768/40-00-82-10-00-00      1
Bridge state:              Enabled                    2
UB-Encapsulation:         Disabled
Bridge type:               STB                        3
Number of ports:          2                          4
STP Participation:        IEEE802.1d                  5

Port  Interface  State  MAC Address          Modes  Maximum  Segment  Flags
   1  Eth/0       Up    40-00-82-10-00-00   T      1520    RD
   2  Eth/1       Up    00-04-13-47-84-53   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

```

- 1** tells the bridge priority and its address.
- 2** says that bridging is enabled.
- 3** says that spanning tree transparent bridging is in use (denoted here by STB).
- 4** Tells the number of ports used (as we will see both over the same ATM link).
- 5** says that spanning tree participation is using the scheme for Ethernet (group addresses for 802.1d bridges) and not for token-ring. If you want to review the relevant concepts, please refer to 2.5.1.1, "Spanning Tree Protocol (STP)" on page 28 and 2.5.2.1, "Spanning Tree Protocol (in Token-Ring)" on page 32.

We proceed to check the status of the ports, which confirms that they are both configured for transparent bridging. This tells us that one is in blocking state (port 1) and the other in forwarding state (port 2):

```

Cary_MSS ASRT>list port
Port Id (dec)      : 128: 1, (hex): 80-01
Port State        : Blocking
STP Participation: Enabled
Port Supports     : Transparent Bridging Only
Duplicates Frames Allowed:  STE: Yes, TSF: Yes
Assoc Interface #/name : 2/Eth/0
Super ELAN bridging: No      Super ELAN ID: 0
+++++
Port Id (dec)      : 128: 2, (hex): 80-02
Port State        : Forwarding
STP Participation: Enabled
Port Supports     : Transparent Bridging Only
Duplicates Frames Allowed:  STE: Yes, TSF: Yes
Assoc Interface #/name : 3/Eth/1
Super ELAN bridging: No      Super ELAN ID: 0
+++++

```

The following command provides additional information, part of which could be key for detecting spanning tree problems in the case of constant topology changes:

```
Cary_MSS ASRT>list spanning-tree-protocol counters
802.Id Spanning Tree Counters:

Time since topology change (seconds)      2132      1
Topology changes:                          16        2
BPDUs received:                           1966
BPDUs sent:                                45

Port  Interface  BPDUs received  BDPUs input overflow  Forward transitions
-----
1  Eth/0      985             0                      0
2  Eth/1      981             0                      1
```

1 tells that the last topology change (from the perspective of the IBM Nways MSS Server) took place 2132 seconds ago. This is way too long ago (approximately 36 minutes) to be indicative for constant topology changes in the spanning tree. (Imagine a network stabilizing every half an hour after an actual topology change.)

2 the number of topology changes, on the other hand, seems relatively high, but this depends on the number of bridges in the network, and on the time the network is up. A rapidly increasing counter would be something to worry about, but our counter stays constant for the next few minutes.

It seems that if there is a loop in the network, the resultant topology in our example does not result in constant topology changes but in a steady loop. On the other hand, let us check how long the MSS server has been operational, just for our own information:

```
Cary_MSS ASRT><CTRL/P>      1

Cary_MSS +uptime            2
Last Reload: 40 minutes ago  3

Cary_MSS + t 5              4
```

1 <CTRL/P> brings us back to the main IBM Nways MSS Server command level.

2 we interrogate the IBM Nways MSS Server for the length of time it has been already running.

3 and it tells us, this has been only for 40 minutes.

4 with t 5 we return to the place we left the GWcon submenu (at the bridging submenu).

From this we conclude that there must have been a reboot of the IBM Nways MSS Server just when we arrived. The network manager confirms that it did restart the IBM Nways MSS Server. So far we have not found many clues, other than the ones with the steady LEDs at the IBM Nways 8260 Multiprotocol Switching Hub, which let us assume there is a loop in the network. Let us continue.

Since we have a NetWare server in the network, we check its status, especially for being in the room next to our IBM Nways 8260 Multiprotocol Switching Hub and IBM Nways 8281 ATM LAN Bridge equipment. And indeed, it beeps and beeps telling that it has received its own frames due to a likely loop in the network. This confirms our suspicions.

Now is the time to go and check the device at port 14.1 of our IBM Nways 8260 Multiprotocol Switching Hub. With the network manager we go to the test room on the 3rd floor of the building, and see that at the end of the fiber there is indeed an IBM Nways 8274 RouteSwitch connected, which is the one displayed in Figure 313 on page 451. Somebody has been testing this recently purchased device.

We check the bridging configuration of the IBM Nways 8274 RouteSwitch, and notice that it was configured for transparent bridging, but for a very, very transparent one. Let us check its configuration.

The following command displays the services being used. LAN emulation services in IBM Nways 8274 RouteSwitch terminology are for LAN emulation the same as LEC interfaces in IBM Nways MSS Server terminology:

```

/ % vas
                                ATM Services
Slot Port Num      Serv      Service      Service
      Num          Num      Description  Type
=====
4     1     1     PTOP Bridging Service 1    PTOP Priv
4     1     2     LAN Emulation Service 2    LANE
4     1     3     LAN Emulation Service 3    LANE
4     2     1     PTOP Bridging Service 1    PTOP Priv
    
```

From **1** and **2** above, we now know that on slot 4, port 1, we have two LECs with service numbers 2 and 3. And from the physical connections in the test lab, we know that port 1 in slot 4 is the ATM port connected to the IBM Nways 8260 Multiprotocol Switching Hub, that is to our emulated LAN environment, too.

```

/ % stpc
Spanning Tree Port Configuration for Group 1 (Default GROUP (#1))
Index Slot/Intf/Service/Inst  Port  Path  Enable
                             Priority Cost  Spanning Tree
                             (a)   (b)   (c)
-----
1     3/ 1/ Brg/ 1           128   100   y
....  ....  ....  ....  ....
14    4/ 1/ Lne/ 1           128    8    n
15    4/ 1/ Lne/ 2           128    8    n
16    4/ 2/ Brg/ 1           128    7    y
save]cancel]next]prev : 14c=y
save]cancel]next]prev : 15c=y
save]cancel]next]prev : save
All items saved
save]cancel]next]prev : cancel
    
```

- 1** shows that this service (LEC) is not participating in spanning tree.
- 2** shows that this service (LEC) is not participating either.

3 here we change the configuration of this service to force its participation in the spanning tree protocol.

4 here we do the same.

5 then we save the new configuration.

6 we exit the configuration step.

Phase 4: Resolve problem (make and test corrections), or investigate a different symptom.

We did this part already during the last step of our testing, since the IBM Nways 8274 RouteSwitch interface permitted us to display and correct definitions in one go. As can be seen in Figure 314 on page 452 the 8274 was acting like a short-circuit between the emulated LANs called ETHERNET_ELAN and ETHER2_ELAN.

The next thing to do is to enter the next iteration of phase 3 for further tests.

Phase 3 (2nd iteration) - Next test steps

- **Step 3-1** - Check configuration and proper LAN connection of test client.

We repeat this step for completeness:

```
[C:\MAGRONLABS\STP]ifconfig lan0
lan0: flags=b863<UP,BROADCAST,NOTRAILERS,RUNNING,BRIDGE,SNAP>
      inet 192.168.5.8 netmask ffffffff broadcast 192.168.5.255

[C:\MAGRON\LABS\STP]netstat -r
```

destination	router	refcnt	use	flags	snmp metric	intrf
default	192.168.5.11	0	0	U	-1	lan0
192.168.5.0	192.168.5.8	0	5801	U	-1	lan0
192.168.5.0	192.168.5.11	0	0	U	-1	lan0

- **Step 3-2** - IP ping other network participants.

We try to reach the NetWare server using IP ping, and succeed.

```
[C:\MAGRONLABS\STP]ping 192.168.5.99
PING 192.168.5.99: 56 data bytes
64 bytes from 192.168.5.99: icmp_seq=1. time=0. ms
64 bytes from 192.168.5.99: icmp_seq=2. time=0. ms

----192.168.5.99 PING Statistics----
2 packets transmitted, 2 packets received, 0% packet loss
round-trip (ms)  min/avg/max = 0/0/0
```

- **Step 3-3** - Log on to a LAN server.

Here we check again the (default IPX) logon to some server in the network, and see that it works well this time:

```
[C:MAGRONLABSSTP]nlist server /b
Object Class: Server
Known to Server: PCSRV320
Active NetWare Server= The NetWare Server that is currently running
Address                = The network address
Node                   = The network node
Status                 = The status of your connection

Active NetWare Server      Address      Node      Status
-----
PCSRV320                   [A76ACA43][          1]Default
One server object was found.
```

With the above correction of the spanning tree parameters and the tests confirming connectivity, we conclude our case study.

8.5.2.3 Conclusion

With the above corrections, we resolved the problem, and did conclude our case study. But in real life, the work would have only been half done. Let us see why:

- In our case study we resolved the spanning tree loop, but we did not collect the spanning tree configuration data of all bridges involved after correcting the configuration of the IBM Nways 8274 RouteSwitch, so we still do not know which bridge of those three below in our logical network view has gone into the blocking state (internal 8281 ATM LAN bridge, external 8281 ATM LAN bridge, or 8274 RouteSwitch).
- Note also that one of the ports of the MSS server was in blocking state, although it is being used to connect high-speed segments (the emulated LANs). For the data traffic between our test client and its server it did not matter, but in an operational network this would be a bad spanning tree design. All the traffic from ETHER2_ELAN to the file server would be transferred over the slow 10 Mbps Ethernet segment connecting the two 8281 devices.

Therefore the next obvious step would be to revise the spanning tree design (in this environment there is obviously none), document it, and reconfigure the network components accordingly.

If a spanning tree problem arises, a temporary measure during operational peak hours might be the shutting down of redundant path bridges (eliminating loops). This would bring the desired relief, but at the same time opens the door for further network problems (outages) where manual recovery would become needed at the most unexpected and inconvenient time of day. Therefore, spanning tree problems should be resolved quickly, to guarantee a solid foundation for the rest of the network.

Spanning tree problems are not always easy to detect. The best thing to do is to plan ahead, and to do a careful spanning tree design, setting the relevant priority variables according to the design at the time of implementation. Since routers may also implement bridging, they should be included in the design and implementation of the spanning tree as well. This will help in finding the source of any unexpected problem in the future much more quickly, but above all careful spanning tree design will ensure that no big problems will arise at all.

Appendix A. Ports and Cable Pinouts

This appendix gives information on ports and cable pin assignments.

A.1 Pinouts for ATM 25 Mbps versus Common Network Connectors

Most networking standards have developed specifications for using shielded or unshielded twisted-pair cabling with RJ-45 modular plugs to connect devices together. Table 25 illustrates the differences between the following cabling specifications:

- ATM25.6 (ATM Forum standard)
IBM adapters for this standard have an orange dot with a white line across it to easily distinguish from the next two types.
- ATM25.6 (Pre-standard used by some early ATM devices)
This adapter has a green dot on it indicating that it uses standard token-ring pinouts.
- Token-ring
This adapter has a green dot on it indicating that it uses standard token-ring pinouts.
- Ethernet (10Base-T)

Table 25. RJ-45 Pin Assignments by Network Type

Pin Number	ATM25 (Forum- Compliant)	ATM25 (Pre-Standard)	Token Ring	Ethernet (10Base-T)
1	RD+			TD+
2	RD-			TD-
3		TD+	TD+	RD+
4		RD+	RD+	
5		RD-	RD-	
6		TD-	TD-	RD-
7	TD+			
8	TD-			

A.2 Other Cabling Considerations

Special cables are required in two specific instances:

- When connecting to pre-standard devices
- When connecting between two ATM switches

Both of these instances are discussed below.

A.2.1 Converter Cables

Some early ATM 25 Mbps adapters, such as the IBM TURBOWAYS 25 ATM adapter (P/N 04H7370), use a pre-standard pin assignment scheme based on the token-ring network cabling standard. To make these adapters compatible with the other ATM-compliant product ports, it is necessary to use a token-ring-to-ATM converter cable, available from IBM as P/N 10H3904. The pinouts for this cable are shown in Table 26.

Table 26. Pin Assignments for Converter Cable (P/N 10H3904)

Signal	Port Pin	Adapter Pin
RD+	1	4
RD-	2	5
TD+	7	3
TD-	8	6

A.2.2 Hubs Crossover Wiring

The hub ports are designed to connect user devices and require a switch-to-switch crossover cable to connect to other ATM switches, just as a 10Base-T hub does. The pinouts for this cable are shown in Table 27.

Table 27. Pin Assignments for Switch-to-Switch Crossover Cable

Signal	Port Pin	Adapter Pin
RD+	1	7
RD-	2	8
TD+	7	1
TD-	8	2

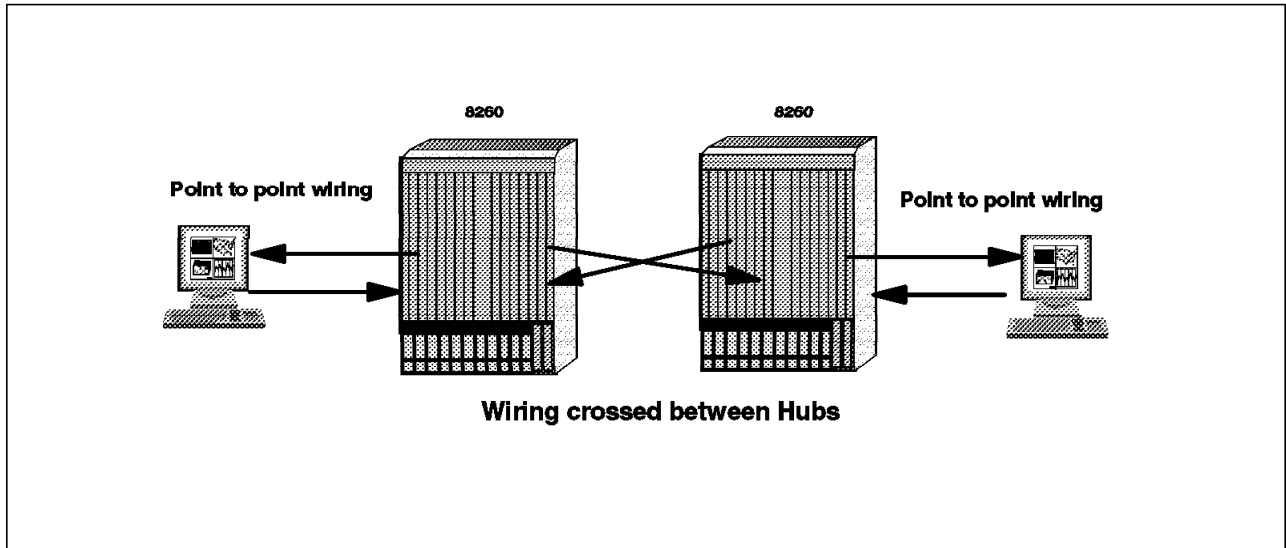


Figure 317. Wires Crossed between Hubs

Appendix B. Hub Code Level History

This appendix gives specific product information on code levels. A list of Web home pages is also provided to get information on IBM products from the Internet.

B.1 8260

Follow the 8260 code levels history.

Table 28. 8260 Code Level History

8260 code version	CPSW 5000 5100	A4 FB100 (MIC) 5004	A4 FB100 (SC) 5104	A2 MB155 5002	A3 MB155 5003	A12 TP25 5102	A CMU1 5102	A CMU2 5202	AMB BRG 5204	A MSS 5300	A2 WAN 5302
V1.1.5	4	4	4	na	na	na	na	na	na	na	na
V1.2.9	6	6	6	7	na	na	B2E4	B2E4	B2E4	na	na
V2.0.4	6/8	6/8	6/8	7/81	na	na	B2E4 B3E4	B2E4 B3E4	B2E4 B3E4	na	na
V2.0.8	6/8	6/8	6/8	7/81	na	na	B3E4 B3F3	B3E4 B3F3	B3E4 B3F3	na	na
V2.1.0	9	6/8	6/8	7/81	na	na	B3E4 B3F3	B3E4 B3F3	B3E4 B3F3	B3F3	B3F3
V2.2.0	9	6/8	6/8	7/81	na	1	B3E4 B3F3	B3E4 B3F3	B3E4 B3F3	B3F3	B3F3
V2.2.2	9	6/8	6/8	7/81	na	1	B3E4 B3F3	B3E4 B3F3	B3E4 B3F3	B3F3	B3F3
V2.3.0	9	6/8	6/8	7/81	1	1	B3E4 B3F3	B3E4 B3F3	B3E4 B3F3	B3F3	B3F3
V2.4.0	B40	B40	B40	B40	C10*	C10*	B40	B40	B40	B40	B40
V2.4.3	B40	B40	B40	B40	C21	C20	B40	B40	B40	B40	B40
V2.5.0	B40 B50	B40 B50	B40 B50	B40 B50	C21 C31	C20 C30	B40 B50	B40 B50	B40 B50	B40 B50	B40 B50
V2.5.1	B40 B50	B40 B50	B40 B50	B40 B50	C21 C31	C20 C30	B40 B50	B40 B50	B40 B50	B40 B50	B40 B50
V2.5.2	B40 B50	B40 B50	B40 B50	B40 B50	C21 C31	C20 C30	B40 B50	B40 B50	B40 B50	B40 B50	B40 B50
V3.0.0	B50 B51	B40 B50	B40 B50	B40 B50	C21 C31	C20 C30	B40 B50	B40 B50	B40 B50	B40 B50	B40 B50
V3.1.0	B51	B40 B50	B40 B50	B40 B50	C21 C31	C20 C30	B40 B50	B40 B50	B40 B50	B40 B50	B40 B50

Note: The CPSW must have 16 megabytes of memory to install and use code level Version 2.1.0 and above.

In the table the * indicates an FPGA problem.

B.2 8285

Follow the 8285 code level history.

8260 code version	CPSW 5000 5100	A4 FB100 (MIC) 5004	A4 FB100 (SC) 5104	A2 MB155 5002	A3 MB155 5003	A12 TP25 5102	A CMU1 5102	A CMU2 5202	AMB BRG 5204	A MSS 5300	A2 WAN 5302
V1.0.0	1	6/8	6/8	7/81	na	na	B3F3	B3F3	B3F3	B3F3	B3F3
V1.0.1	1	6/8	6/8	7/81	na	na	B3F3	B3F3	B3F3	B3F3	B3F3
V1.2.0	3	6/8	6/8	7/81	na	1	B3F3	B3F3	B3F3	B3F3	B3F3
V1.3.0	3	6/8	6/8	7/81	1	1	B3F3	B3F3	B3F3	B3F3	B3F3
V1.3.1	3	6/8	6/8	7/81	1	1	B3F3	B3F3	B3F3	B3F3	B3F3
V1.4.0	C10*	B40	B40	B40	C11*	C10*	B40	B40	B40	B40	B40
V1.4.3	C10*	C20	B40	B40	C21	C20	B40	B40	B40	B40	B40
V1.5.0	C20	B40	B40	B40	C21	C20	B40	B40	B40	B40	B40
	C30	B50	B50	B50	C31	C30	B50	B50	B50	B50	B50
V1.5.1	C20	B40	B40	B40	C21	C20	B40	B40	B40	B40	B40
	C30	B50	B50	B50	C31	C30	B50	B50	B50	B50	B50
V1.5.2	C20	B40	B40	B40	C21	C20	B40	B40	B40	B40	B40
	C30	B50	B50	B50	C31	C30	B50	B50	B50	B50	B50
V3.0.0	C30	B40	B40	B40	C21	C20	B40	B40	B40	B40	B40
	C31	B50	B50	B50	C31	C30	B50	B50	B50	B50	B50
V3.1.0	C31	B40	B40	B40	C21	C20	B40	B40	B40	B40	B40
		B50	B50	B50	C31	C30	B50	B50	B50	B50	B50

Note: A memory upgrade on the 8285 base is required to be able to use code Version 3.0.0 and above. 12 megabytes of memory is required.

In the table the * indicates an FPGA problem.

B.3 IBM Networking Information Home Pages

<i>Table 30. Useful Web Home Pages</i>	
Subject	Web Home Pages
Networking Hardware Products	http://www.raleigh.ibm.com/netprod.html
IBM Networking	http://www.raleigh.ibm.com/nethome.html
IBM Networking Server	http://www.raleigh.ibm.com/netinfo.html The IBM Networking Server home page is the best access entry point for obtaining NHD hardware and software information.
Routers	
2210 Nways Multiprotocol Router	http://www.raleigh.ibm.com/220/220prod.html
Network Services	
2217 Nways Multiprotocol Concentrator	http://www.raleigh.ibm.com/227/227prod.html
2220 Nways Broadband Switch	http://www.raleigh.ibm.com/222/222prod.html
3172 Nways Interconnect Ctlr	http://www.raleigh.ibm.com/312/312prod.html
3174 Establishment Ctlr	http://www.raleigh.ibm.com/314/314prod.html
3745 Communications Ctlr and	http://www.raleigh.ibm.com/375/375prod.html
3746 Nways Multinetwork Ctlr	http://www.raleigh.ibm.com/376/376prod.html
5394/5494 Remote Ctrl Units	http://www.raleigh.ibm.com/549/549prod.html
6611 Network Processor and	http://www.raleigh.ibm.com/611/611prod.html
Multiprotocol Network Program	
8222 Ethernet Workgroup Hub	http://www.raleigh.ibm.com/82n/82nprod.html
8224 Ethernet Stackable Hub	http://www.raleigh.ibm.com/84n/84nprod.html
8226 T-R RJ45 Connection MAU	http://www.raleigh.ibm.com/861/861prod.html
8229 Bridge	http://www.raleigh.ibm.com/829/829prod.html
8230 T-R Net Cont Acc Unit	http://www.raleigh.ibm.com/823/823prod.html
8235 Dials Server	http://www.raleigh.ibm.com/825/825prod.html
8238 T-R Stackable Hub	http://www.raleigh.ibm.com/82t/82tprod.html
8244 FDDI Workgroup Concentrator	http://www.raleigh.ibm.com/824/824prod.html
8250 Multiprotocol Intel Hub	http://www.raleigh.ibm.com/825/825prod.html
8260 Multiprotocol Intel Sw Hub	http://www.raleigh.ibm.com/826/826prod.html
8271 Nways Ethernet LAN Switch	http://www.raleigh.ibm.com/821/821prod.html
8272 Nways T-R LAN Switch	http://www.raleigh.ibm.com/822/822prod.html
8281 Nways ATM LAN Bridge	http://www.raleigh.ibm.com/828/828prod.html
8282 Nways ATM Workgroup Concentrator	http://www.raleigh.ibm.com/882/882prod.html
8285 Nways ATM Workgroup Switch	http://www.raleigh.ibm.com/82a/82aprod.html

<i>Table 31. Microcode and Software Upgrades</i>	
Subject	Web Home Pages
6611	http://www.raleigh.ibm.com/611/611fix.html
6611	http://www.raleigh.ibm.com/611/611ptf.html
8230	http://www.raleigh.ibm.com/823/823fix.html
8260	http://www.raleigh.ibm.com/826/826fix.html

<i>Table 32. Software Management Products</i>	
References	Web Home Pages
Management Products	
Nways Campus Mgr LAN for xxxxx	http://www.raleigh.ibm.com/abm/abmprod.html
AIX, HP-UX Environment	
Nways Campus Mgr ATM for xxxxx	http://www.raleigh.ibm.com/cma/cmprod.html
AIX, HP-UX Environment	
Nways LAN Remote Monitor for xxxxxx	http://www.raleigh.ibm.com/rmn/rmnprod.html
Windows, AIX, and HP-UX Environments	
Nways Manager for Windows	http://www.raleigh.ibm.com/nmw/nmwprod.html
IBM Intelligent Hub Manager for AIX	http://www.raleigh.ibm.com/hma/hmaprod.html
AIX Ver 2 and Entry Environments	
Intelligent Hub Management Prog/	http://www.raleigh.ibm.com/hmd/hmdprod.html
DOS Entry Version 2 Environments	
LAN Network Manager for AIX	http://www.raleigh.ibm.com/nma/nmaprod.html
Nways Broadband Switch Manager	http://www.raleigh.ibm.com/sar/sarprod.html
for AIX	
AIX LAN Management Utilities/6000	http://www.raleigh.ibm.com/lm6/lm6prod.html
IBM LAN NetView Management Utilities	http://www.raleigh.ibm.com/lmu/lmuover.html
for OS/2	
LNM Network Manager	http://www.raleigh.ibm.com/lnm/lnmnews.html

Appendix C. UNI 3.0-3.1 Cause Maintenance Error Codes

This appendix lists error or maintenance codes available on LAN/ATM campus networks.

C.1 ATM Forum UNI Cause Codes

Follow codes defined in the ATM Forum for events.

<i>Table 33 (Page 1 of 5). Cause Codes</i>		
HEX	DEC	Definitions
01	1	<p>Unallocated (unassigned) number. ATM address unassigned.</p> <p>This cause indicates that the called party cannot be reached because, although the number is in a valid format, it is not currently assigned (allocated). Check the destination ATM address.</p>
02	2	<p>No route to specified transit network.</p> <p>This cause indicates that the equipment sending this cause has received a request to route this call through a particular network which it does not recognize. The equipment sending this cause does not recognize the transit network either because the transit network does not exist or because that particular transit network, while it does exist, does not serve the equipment which is sending this cause. The diagnostic field contains a copy of the contents of the transit selection information identifying the unreachable network. This cause is supported on a network-dependent basis.</p>
03	3	<p>No route to destination.</p> <p>This cause indicates that the called party cannot be reached because the network through which the call has been routed does not serve the destination desired. Probable cause is that no station has registered this ATM address with the switch. Check configuration at both ends. This cause is supported on a network-dependent basis.</p>
0A	10	<p>VPI/VCI unacceptable.</p> <p>This cause indicates that the virtual channel most recently identified is not acceptable to the sending entity for use in this call.</p>
10	16	<p>Normal call clearing.</p> <p>This cause indicates that the call is being cleared because one of the users involved in the call has requested that the call be cleared. Under normal situations, the source of this cause is not the network.</p>
11	17	<p>Called party busy.</p> <p>This cause indicates that the called party is unable to accept another call because the user busy condition has been encountered. This cause value may be generated by the called user or the network.</p>
12	18	<p>No user responding.</p> <p>This cause is used when a called party does not respond to a call establishment message with a connect indication within the prescribed period of time allocated. For example, called party does not respond to SETUP message by the time timer T303 expired. The call is cleared.</p>
15	21	<p>Call rejected.</p> <p>This cause indicates that the equipment sending this cause does not wish to accept this call, although it could have accepted the call because the equipment sending this cause is neither busy nor incompatible</p>

Table 33 (Page 2 of 5). Cause Codes		
HEX	DEC	Definitions
16	22	<p>Number changed, ATM address changed.</p> <p>This cause is returned to a calling party when the called party number indicated by the calling user is no longer assigned. The new called party number may optionally be included in the diagnostic field. If the network does not support this capability, cause #1 will be used instead.</p>
17	23	<p>User rejects all calls with calling line identification restriction (CLR).</p> <p>Caller address is required by the called party. This cause is returned by the called party when the call is offered without calling party number information and the called party requires this information.</p>
1B	27	<p>Destination out of order.</p> <p>This cause indicates that the destination indicated by the user cannot be reached because the interface to the destination is not functioning correctly. The term <i>not functioning correctly</i> indicates that a signalling message was unable to be delivered to the remote user, for example, a physical layer or SAAL failure at the remote user, or user equipment is off line. Cause posted when T309 expires before AAL signalling can be re-established with the destination. At the remote endstation, layer 2 (QSAAL) is down and/or the ATM address is not registered with the switch.</p>
1C	28	<p>Invalid number format or invalid ATM address.</p> <p>This cause indicates that the called user cannot be reached because the called party number is not in a valid format or is not complete.</p>
1E	30	<p>Response to STATUS ENQUIRY.</p> <p>This cause is included in the STATUS message when the reason for generating the STATUS message was the prior receipt of a STATUS ENQUIRY message. Reports the current call state. It does not directly affect the call state of the sender or the receiver.</p>
1F	31	<p>Normal, unspecified.</p> <p>This cause is used to report a normal event only when no other cause in the normal class applies. This is sending in call clearing (RELEASE COMPLETE) and results from an ADD PARTY reject when there are no other active parties.</p>
23	35	<p>Requested VPI/VCI not available.</p> <p>The network allocates a VPCI/VCI value and includes this value in the SETUP message. The user receiving the SETUP message accepts the indicated VPCI/VCI for the call. This cause indicates that the VCI is not available within the indicated VPCI. The user sends a RELEASE COMPLETE and this cause. Layer 3 (SVC) sends this code if the switch tries to assign to a call a VPI/VCI that is already in use. The AIX SVC device driver (L4 +) sends this code if the switch assigns a VCI that is reserved for a PVC.</p>
24	36	<p>VPI/VCI assignment failure.</p> <p>The network shall allocate a VPCI/VCI value and include this value in the SETUP message. The user receiving the SETUP message accepts the indicated VPCI/VCI for the call. If the VPCI/VCI values the user sends in its first response are not the values offered by the network, the network sends a RELEASE message to the user with this cause.</p>
25	37	<p>User cell rate unavailable.</p> <p>The network cannot route the call due to insufficient bandwidth. The network initiates call clearing.</p>
26	38	<p>Network out of order.</p> <p>This cause indicates that the network is not functioning correctly and the condition is likely to last a relatively long period of time, for example, immediately re-attempting the call is not likely to be successful.</p>

<i>Table 33 (Page 3 of 5). Cause Codes</i>		
HEX	DEC	Definitions
29	41	<p>Temporary failure.</p> <p>This cause indicates that the network is not functioning correctly and that the condition is not likely to last long. The user may wish to try another call attempt immediately. An in-process or established call was cleared due to a layer 2 (QSAAL) disconnection and re-establishment. For example, this cause can be posted if there is no STATUS response received for the STATUS ENQUIRY message (T322 expires) and the STATUS ENQUIRY has been retransmitted the maximum number of times. The "maximum number of times" is implementation-specific. The call is cleared to the local interface and the network may also clear the connection.</p>
2B	43	<p>Access information discarded.</p> <p>This cause indicates that the network could not deliver access information to the remote user as requested, that is, ATM adaptation layer parameters, broadband low layer information, broadband high layer information, or sub-address as indicated in the diagnostic. This can be content errors in non-mandatory fields. See also cause 100.</p>
2D	45	<p>No VPCI/VCI available.</p> <p>The network allocates a VPCI/VCI value and includes this value in the SETUP message. The network selects any available VPCI and VCI. The user receiving the SETUP message accepts the indicated VPCI/VCI for the call. This cause indicates that the network is not able to allocate VCI in any VPCI. The network sends RELEASE COMPLETE.</p>
2F	47	<p>Resource unavailable, unspecified.</p> <p>This cause is used to report a resource unavailable event only when no other cause in the resource unavailable class applies. (Some layer above layer 3 at the called party rejected the call due to lack of bandwidth or some other resource. Layer 3 never originates this cause code.) For example, with an ADD PARTY message, the QOS and bandwidth must be the same as the connection and are not explicitly indicated in the ADD PARTY message. If the user is not able to support the requested ATM traffic descriptor, the user will reject the call with this cause.</p>
31	49	<p>Quality of Service unavailable.</p> <p>This cause is used to report that the requested Quality of Service cannot be provided.</p>
33	51	<p>User cell rate not available.</p> <p>This cause is used to report that the requested ATM Traffic Descriptor is unobtainable. The network rejects the call. The diagnostics field of the cause information element should indicate those parameters that exceed the capacity of the network.</p>
39	57	<p>Bearer capability not authorized.</p> <p>This cause indicates that the user has requested a bearer capability which is implemented by the equipment that generated this cause, but the user is not authorized to use.</p>
3A	58	<p>Bearer capability not presently available.</p> <p>This cause indicates that the user has requested a bearer capability which is implemented by the equipment which generated this cause, but which is not available at this time. Sent by the network back to the user.</p>
3F	63	<p>Service or option not available, unspecified.</p> <p>This cause is used to report a service or option not available event only when no other cause in the service or option not available class applies. For example, the parameters specified in SETUP message should be consistent. Table F-1 in Appendix F of the UNI 3.1 Specs shows allowable combinations of some parameters. This cause will be returned when illegal combinations are specified. The network clears the call.</p>
41	65	<p>Bearer capability not implemented.</p> <p>This cause indicates that the equipment sending this cause does not support the bearer capability requested.</p>

<i>Table 33 (Page 4 of 5). Cause Codes</i>		
HEX	DEC	Definitions
49	73	<p>Unsupported combination of traffic parameters.</p> <p>This cause indicates that the combination of traffic parameters contained in the ATM traffic descriptor information elements is not supported.</p>
4E	78	<p>AAL parameter cannot be supported.</p> <p>When the calling endpoint wishes to indicate to the called endpoint the AAL common part parameters and service part to be used during the call, the calling endpoint includes ATM adaptation layer parameter information in the SETUP message. This information element is conveyed by the network and delivered to the called user. If the called user does not include the ATM adaptation layer parameters in the CONNECT message, the calling user shall assume that the called user accepts the values of the Forward and Backward Maximum CPCS-SDU size indicated by the caller in the SETUP message. If the calling party cannot use the Forward and Backward CPCS-SDU size indicated in the CONNECT message (that is, because the value negotiated by the called party is unacceptably small), the call will be cleared with this cause. See Appendix F of UNI 3.1 specifications.</p>
51	81	<p>Invalid call reference value.</p> <p>This cause indicates that the equipment sending this cause has received a message with a call reference that is not currently in use on the user-network interface. Whenever any message except SETUP, RELEASE COMPLETE, STATUS ENQUIRY, or STATUS is received that specifies a call reference that is not recognized as belonging to an active call or a call in progress this cause is returned. Also sent if an ADD PARTY, ADD PARTY ACKNOWLEDGE, ADD PARTY REJECT, DROP PARTY or DROP PARTY ACKNOWLEDGE message is received while in the null link state. The diagnostic field specifies the call reference.</p>
52	82	<p>Identified channel does not exist, VPI/VCI does not exist.</p> <p>This cause indicates that the equipment sending this cause has received a request to use a channel not activated on the interface for a call.</p>
58	88	<p>Incompatible destination.</p> <p>This cause indicates that the equipment sending this cause (usually a user) has received a request to establish a call which has broadband low layer information, broadband high layer information, or other AAL attributes which cannot be accommodated. Refer to Annex C on UNI 3.1 specification. Some layer above layer 3 at the called party rejected the call parameters. Check configuration at both ends.</p>
59	89	<p>Invalid endpoint reference value.</p> <p>The purpose of the endpoint reference IE is to identify the individual endpoints of a point-to-multipoint connection. A value of 0 in the Endpoint Reference identifier always is used to identify the first party of the point-to-multipoint call. A non-zero value is always used to identify subsequent parties of the call. This cause indicates that the equipment sending this cause has received a message with an endpoint reference which is currently not in use on the user-network interface. Whenever any message except SETUP, ADD PARTY, or DROP PARTY ACKNOWLEDGE is received for a party in the null party state, dropping is initiated by sending a DROP PARTY ACKNOWLEDGE with this cause and the sender will remain in the null party state.</p>
5B	91	<p>Invalid transit network selection, transit net does not exist.</p> <p>This cause indicates that a transit network identification was received which is of an incorrect format as defined in Annex D of the UNI 3.1 specification. Some networks may provide screening to the transit network (for example, to ensure that a business relationship exists between the user and the transit network). Should the screening fail, this cause will be returned.</p>
5C	92	<p>Too many pending add party requests.</p> <p>This cause indicates a temporary condition when the calling party sends an add party message but the network is unable to accept another add party message because its queues are full.</p>
5D	93	<p>AAL parameters cannot be supported.</p> <p>This cause indicates that the equipment sending this cause has received a request to establish a call which has ATM adaptation layer parameters which cannot be accommodated.</p>

<i>Table 33 (Page 5 of 5). Cause Codes</i>		
HEX	DEC	Definitions
60	96	<p>Mandatory information element is missing.</p> <p>This cause indicates that the equipment sending this cause has received a message (SETUP, RELEASE, DROP PARTY, etc.) which is missing an information element which must be present in the message before the message can be processed. Could also be because a RELEASE message was received with the cause information element missing. The responding RELEASE COMPLETE will have this cause. Refer to the diagnostic field in the cause code element.</p>
61	97	<p>Message type non-existent or not implemented.</p> <p>This cause indicates that the equipment sending this cause has received a message with a message type it does not recognize either because this is a message not defined or one defined but not implemented by the equipment sending this cause. (If the offending message type were received by the sender of this cause when it was in the null state, it would not have sent this cause code, but ignored the message instead.)</p>
63	99	<p>Information element non-existent or not implemented.</p> <p>This cause indicates that the equipment sending this cause has received a message which includes information element(s) not recognized because the information element identifier(s) are not defined or are defined but not implemented by the equipment sending the cause. This cause indicates that the information element(s) were discarded. Action will be taken on the message and those information elements which are recognized and have valid content. However, the information element is not required to be present in the message in order for the equipment sending this cause to process the message. The diagnostic field, if present, will contain more information about the unrecognized element. A possible reason for the cause is if VP/VC information were included in the SETUP message from the user.</p>
64	100	<p>Invalid information elements contents.</p> <p>This cause indicates that the equipment sending this cause has received an information element which it has implemented; however, one or more of the fields in the information element are coded in a way that cannot be implemented by the equipment sending this cause.</p> <p>A common reason for this cause is incompatible UNI versions.</p>
65	101	<p>Message not compatible with call state.</p> <p>This cause indicates that a message has been received which is incompatible with the call state. For example, if a STATUS message indicating any call state except the null state is received in the null state, then the receiving entity shall send a RELEASE COMPLETE with this cause and remain in the null state.</p>
66	102	<p>Recovery on timer expiry.</p> <p>This cause indicates that a procedure has been initiated by the expiration of a timer in association with error handling procedures. Layer 3 (SVC) sent a message and no response was received when the defined timer expired. This can be a retry of a previous event, and the cause indicates that the reason for the retry is that the response timer had expired.</p>
68	104	<p>Bad message length.</p>
6F	111	<p>Protocol error, unspecified, SVC protocol error.</p> <p>This cause is used to report a protocol error event only when no other cause in the protocol error class applies.</p>

C.2 Maintenance Codes Valid on 8260/8285 ATM Hubs

The following is a list of prompts available in the maintenance mode, and their corresponding meanings.

<i>Table 34. 8260/8285 Maintenance Codes</i>	
Codes	Explanation
0020	The NVRAM diagnostics failed, the battery may be low .
0021	Bad checksum, the loading or the de-compression of the operational code failed.
0022	After three retries, the switch FPGAs did not initialize properly (8260).
0023	After three retries, the base FPGAs did not initialize properly (8285).
0030	The initialization or the diagnostics failed for the switch, the 8032, or the serial link .
0031	The ATM Wrap test from control point board to switch board failed.
0032	The initialization of the operational code stopped due to a lack of memory.
0033	The initialization of the operational code stopped due to a lack of memory.
0034	The initialization of the operational code stopped due to a lack of memory.
0035 (8285)	FAT + other signals tested failed with wrap expansion plug.
0036 (8285)	After three retries, the base FPGAs did not initialize properly.
0040 (8260)	Active to backup CPSW polling does not work. SPI serial link may fail.
0050 (8260)	No FPGA picocode level (active or backup) in the A-CPSW module matches the active microcode level, and the backup microcode of the A-CPSW module is either unavailable or identical to the active one.
0051	The SWAP of the ATM control point FPGA picocode terminated in error.
0052	A connected ATM media module has no FPGA picocode matching the A-CPSW microcode level. This is a normal condition for the first A-CPSW of a redundant 8260 during the automatic migration process to level B50. It makes the second A-CPSW active, allowing the upgrade of the rest of the 8260. Once the whole 8260 is upgraded, the A-CPSW displaying >>0052>> becomes either active or standby at the next reset.
00BA	Maintenance mode is running with the backup daemon.

C.3 IBM LAN Emulation Server (LES) 8260/8285 Error Reason Codes

The following are error reason codes used on the 8260/8285 in a LAN emulation environment.

- 1 Network cause
- 2 Internal cause
- 3 Memory exhausted
- 4 Network is down

Table 35. Cause Code for LAN-Emulation Server Monitor

Reason	Cause Code	Explanation
1	Any	Network cause (see Q.2931 Cause code list).
2	31	The LAN emulation server either received an adapter failure message or a deregister ILMI address message.
	41	The SAAL layer went down and was re-established within 90 seconds, but a status response message was not received.
	102	The LAN Emulation server received a call proceeding message, but did not receive an acknowledgment that the connection is ready.
4	27	The SAAL layer went down and was not re-established within 90 seconds.
	41	This message is sent immediately after the SAAL layer goes down and whenever the LAN emulation server receives a restart PDU message. All existing and pending calls are cleared.
	other	The LAN emulation server did not receive a response during call setup.

Appendix D. ATM Forum-Compliant LANE Frame Formats

The following sections provide you with the frame formats used in ATM Forum LAN emulation environments.

D.1 ATM Forum LAN Emulation Server Parameters

S1 LAN Emulation Server's ATM Address

The LAN emulation server (LES) must know the ATM addresses of its LAN emulation clients (LECs). The ATM address cannot be removed when any LEC is connected through it to the LES.

S2 LAN Type

This is the type of ATM emulated LAN. This can be either IEEE 802.3 (Ethernet) or IEEE 802.5 (token-ring).

S3 Maximum Data Frame Size

This is the maximum AAL-5 Service Data Unit (SDU) that the LAN emulation service can guarantee not to drop because it is too large. It is also the minimum AAL-5 SDU that every LEC must be able to receive. Valid values are 1516, 4544, 9234, or 18190 octets.

S4 Control Timeout

This parameter sets the period used for timing-out request/response control frame interactions. Once a LEC establishes a control direct VCC to the LES, the join phase must complete within the join timeout time. If this is not the case the LAN emulation service should release any control VCCs to that LEC thereby terminating the join phase.

S5 Maximum Frame Age

The broadcast and unknown server (BUS) must discard a frame if it has not sent the frame to all relevant multicast send VCCs or multicast forward VCCs within the maximum frame age following the BUS's receipt of the frame over a multicast send VCC.
Values: minimum=1 second, maximum=4 seconds, default=1 second

S6 Broadcast and Unknown Server's ATM Address

A BUS must know at least one of its own ATM addresses for LECs to be able to establish connections to it. A BUS can have several ATM addresses. The address can be added while the BUS is operational but cannot be removed while any LEC has a connection to the BUS through the address.

D.2 ATM Forum LAN Emulation Client Parameters

C1 LE Client's ATM Address

The primary ATM address used to connect to the LE server and the BUS. This must be known before the configuration and join phases can start and cannot change without restarting the configuration and join phases. The primary ATM address must be used to establish the LE client's control direct VCC and multicast VCC, and must be specified as the SOURCE-ATM-ADDRESS in the client's LE_JOIN_REQUESTs. An LE client may have additional ATM addresses for use with data direct VCCs. These addresses do not need to be known at join time and can be removed without restarting the join phase.

C2 LAN Type

The type of LAN that the client wishes to join or is a member of. This must be either IEEE 802.3, IEEE 802.5 or unspecified. This must not be unspecified after a successful join. This parameter must not be changed without terminating the client and returning to the Initial state.

C3 Maximum Data Frame Size

The maximum AAL-5 SDU size of a data frame that the LE client wishes to send on the multicast send VCC or to receive on the multicast forward VCC. This parameter also specifies the maximum AAL-5 SDU of all LE clients' data direct VCCs. This value must not be unspecified after a successful join and cannot be changed without terminating the LE client and returning to the initial state.

Value: 1516, 4544, 9234, 18190 or unspecified.

C4 Proxy

This indicates whether the LE client may have remote unicast MAC addresses in C27. This must be known before the join phase can start and cannot be changed without restarting the configuration phase.

C5 ELAN Name

The identity of the emulated LAN the LE client wishes to join, or to which the LE client last joined. This may be unspecified before a join, but can never be specified after a successful join.

C6 Local Unicast MAC Address(es)

Each LE client has zero or more local unicast MAC addresses. In an operational LE client every address in this variable must have been registered with the LE server. Two LE clients joined to the same emulated LAN cannot have the same local unicast MAC address. An LE client's MAC address may change during normal operations.

C7 Control Timeout

Timeout period used for timing out most request/response control frame interactions, as specified elsewhere.

Value: minimum=10 seconds, maximum=300 seconds, default=120 seconds.

C8 Route Descriptors

Route descriptors exist only in source routed IEEE 802.5 LE clients that are source route bridges. All route descriptors in any given emulated LAN must be unique. An LE client may have zero or more route descriptors.

Route descriptors can change during normal operation. If an LE client has route descriptors it must register all of them with the LE server.

C9 LE Server ATM Address

The ATM address of the LE server is used to establish the control direct VCC. This is obtained in the configuration phase from the LECS. It must be known before the join phase can start.

C10 Maximum Unknown Frame Count

Value: minimum=1, maximum=10 default=1.
(see parameter C11).

C11 Maximum Unknown Frame Time

Within the time period defined by the maximum unknown frame time an LE client will send no more than maximum unknown frame count frames to the BUS for a given unicast LAN destination, and it must also initiate the address resolution protocol to resolve that LAN destination.
Value: minimum=1second, maximum=60 seconds, default=1 second.

C12 VCC Timeout Period

The LE client should release any data direct VCC that it has not used to transmit or receive any data frames for the length of the VCC timeout. This parameter only applies to SVC data direct VCCs.
Value: minimum=none, maximum=unlimited, default=20 minutes.

C13 Maximum Retry Count

An LE client must not retry an LE_ARP_REQUEST for a given LAN destination more than the maximum retry count times, after the first LE_ARP_REQUEST for the same frame's LAN destination.
Value: minimum=0, maximum=2, default=1.

C14 LE Client Identifier

This is a unique identifier for the LE client assigned by the LE server. The LECID is placed in all control requests by the LE client and may be used for echo suppression on multicast data frames sent by that client. This value cannot change without first terminating the LE client and returning to the initial state.
Value: X'0001' to X'FEFF'.

C15 LE Client Multicast MAC Address

Each LE client may have a list of multicast MAC addresses that it wishes to receive and pass up to the higher layers. The broadcast address should be included in the list.

C16 LE_ARP Cache

A table of entries, each of which establishes a relationship between a LAN destination external to the LE client and the ATM address to which data frames for that destination will be sent.

C17 Aging Time

The maximum time that an LE client will maintain an entry in its LE_ARP cache in the absence of verification of that relationship.
Value: minimum=10 seconds, maximum=300 seconds, default=300 seconds.

C18 Forward Delay Time

The maximum time that an LE client will maintain an entry for a non-local MAC address in its LE_ARP cache in the absence of a verification of that relationship, as long as the topology flag (C19) is true.

Value: minimum=4 seconds, maximum=30 seconds, default=15 seconds.

C19 Topology Change

Boolean indication that the LE client is using the forward delay time (C18), instead of the aging time (C17) to age non-local entries in its LE_ARP cache.

C20 Expected LE_ARP Response Time

The maximum time the LE client expects an LE_ARP_REQUEST/LE_ARP_RESPONSE cycle to take. It is used for retries and verifies.

Value: minimum=1 second, maximum=30 seconds, default=1 second.

C21 Flush Timeout

Time limit to wait to receive an LE_FLUSH_RESPONSE after the LE_FLUSH_REQUEST has been sent before taking recovery action.

Value: minimum=1 second, maximum=4 seconds, default=4 seconds.

C22 Path Switching Delay

The time since sending a frame to the BUS after which the LE client may assume that the frame has been either discarded or delivered to the recipient.

Value: minimum=1 second, maximum=6 seconds, default=8 seconds.

C23 Local Segment ID

The segment ID of the emulated LAN. Only required in IEEE 802.5 LE clients that are source route bridges.

C24 Multicast Send VCC Type

Signalling parameter that should be used by the LE client when establishing the multicast send VCC. This is the method used by the LE client when specifying traffic parameters when it sets up the multicast send VCC for the emulated LAN.

C25 Multicast Send VCC AvgRate

Signalling parameter that should be used by the LE client when establishing the multicast send VCC. Forward and backward peak cell rate to be requested by the client when setting up the multicast send VCC if using variable bit encoding.

C26 Multicast Send VCC PeakRate

Signalling parameter that should be used by the LE client when establishing the multicast send VCC. Forward and backward peak cell rate to be requested when setting up the multicast send VCC when using either constant or variable bit rate coding.

C27 Remote Unicast MAC Address(es)

The MAC address for which this LE client will answer LE_ARP_REQUESTs that are not registered with the LE server. This list must be empty in any LE client that did not join the emulated LAN as a proxy agent.

C28 Connection Completion Timer

Optional. In connection establishment this is the time period in which data or READY_IND message is expected from a calling party.

Value:Minimum=1 second, Maximum=10 seconds, Default=4 seconds.

D.3 Configuration Frame Format

Table 36. Configuration Frame Format

Offset	Size(Bytes)	Name	Function
0	2	MARKER	Control frame = X'FF00'.
2	1	PROTOCOL	ATM LAN emulation protocol = X'01'.
3	1	VERSION	ATM LAN emulation protocol version = X'01'.
4	2	OP-CODE	Type of request: X'0001' LE_CONFIGURE_REQUEST X'0101' LE_CONFIGURE_RESPONSE
6	2	STATUS	Always X'0000' in requests. In responses refer to Table 43 on page 491 for a list of values.
8	4	TRANSACTION-ID	Arbitrary value supplied by the requester and returned by the responder.
12	2	REQUESTER-LEC-ID	Always X'0000' in requests, ignored on response.
14	2	FLAGS	Always X'0000' when sent, ignored on receipt.
16	8	SOURCE-LAN-DESTINATION	MAC address or route descriptor of prospective LE client. May be encoded as (not present).
24	8	TARGET-LAN-DESTINATION	Always X'0000' when sent, ignored on receipt.
32	20	SOURCE-ATM-ADDRESS	Primary ATM address of perspective LE client for which information is requested.
52	1	LAN-TYPE	X'00' Unspecified X'01' Ethernet/IEEE 802.3 X'02' IEEE 802.5
53	1	MAXIMUM-FRAME-SIZE	X'00' Unspecified X'01' 1516 X'02' 4544 X'03' 9234 X'04' 18190
54	1	NUMBER-TLVs	Number of Type/Length/Value elements encoded in request/response.
55	1	ELAN-NAME-SIZE	Number of octets in ELAN_NAME - may be 0.
56	20	TARGET-ATM-ADDRESS	ATM address of the LE server to be used for the LE client described in the request if Configure Response and STATUS='SUCCESS', else X'00'.
76	32	ELAN-NAME	Name of emulated LAN.
108	4	ITEM_1-TYPE	Three octets of OUI, one-octet identifier.
112	4	ITEM_1-LENGTH	Length in octets of VALUE field. Minimum=0
113	Variable	ITEM_1-VALUE	

D.4 Join Frame Format

Offset	Size(Bytes)	Name	Function
0	2	MARKER	Control frame = X'FF00'.
2	1	PROTOCOL	ATM LAN emulation protocol = X'01'.
3	1	VERSION	ATM LAN emulation protocol version = X'01'.
4	2	OP-CODE	Type of request: X'0002' LE_JOIN_REQUEST X'0102' LE_JOIN_RESPONSE
6	2	STATUS	Always X'0000' in requests. In responses refer to Table 43 on page 491 for a list of values.
8	4	TRANSACTION-ID	Arbitrary value supplied by the requester and returned by the responder.
12	2	REQUESTER-LECID	Assigned LECID of joining client if join response and STATUS = 'SUCCESS', else X'0000'.
14	2	FLAGS	Each bit of the FLAGS file has a separate meaning if set: X'0080' Proxy Flag: LE client server has not registered MAC addresses and therefore wishes to receive LE_ARP requests for non-registered LAN destinations.
16	8	SOURCE-LAN-DESTINATION	Optional MAC address to register as a pair with the SOURCE_ATM_ADDRESS.
24	8	TARGET-LAN-DESTINATION	Always X'00' when sent, ignored on receipt.
32	20	SOURCE-ATM-ADDRESS	Primary ATM address of LE client issuing join request.
52	1	LAN-TYPE	X'00' Unspecified X'01' Ethernet/IEEE 802.3 X'02' IEEE 802.5
53	1	MAXIMUM-FRAME-SIZE	X'00' Unspecified X'01' 1516 X'02' 4544 X'03' 9234 X'04' 18190
54	1	NUMBER-TLVs	Always X'00' when sent, ignored on receipt.
55	1	ELAN-NAME-SIZE	Number of octets in ELAN_NAME. X'00' indicates empty ELAN_NAME.
56	20	TARGET-ATM-ADDRESS	Always X'00' when sent, ignored on receipt.
76	32	ELAN-NAME	Name of emulated LAN.

D.5 Registration Frame Format

Offset	Size	Name	Function
0	2	MARKER	Control frame = X'FF00'.
2	1	PROTOCOL	ATM LAN emulation protocol = X'01'.
3	1	VERSION	ATM LAN emulation protocol version = X'01'.
4	2	OP-CODE	Type of request: X'0004' LE_REGISTER_REQUEST X'0104' LE_REGISTER_RESPONSE X'0005' LE_UNREGISTER_REQUEST X'0105' LE_UNREGISTER_RESPONSE
6	2	STATUS	Always X'0000' in requests. In responses refer to Table 43 on page 491 for a list of values.
8	4	TRANSACTION-ID	Arbitrary value supplied by the requester and returned by the responder.
12	2	REQUESTER-LECID	LECID of LE client issuing the register or unregister request and returned by the responder.
14	2	FLAGS	Always X'00' when sent, ignored on receipt.
16	8	SOURCE-LAN-DESTINATION	Unicast MAC address or route descriptor LE client is attempting to register.
24	8	TARGET-LAN-DESTINATION	Always X'00' when sent, ignored on receipt.
32	20	SOURCE-ATM-ADDRESS	An ATM address of LE client issuing register or unregister request.
52	56	RESERVED	Always X'00' when sent, ignored on receipt.

D.6 Address Resolution Frame Format

<i>Table 39. LE_ARP Frame Format</i>			
Offset	Size	Name	Function
0	2	MARKER	Control frame = X'FF00'.
2	1	PROTOCOL	ATM LAN emulation protocol = X'01'.
3	1	VERSION	ATM LAN emulation protocol version = X'01'.
4	2	OP-CODE	Type of request: X'0006' LE_ARP_REQUEST X'0106' LE_ARP_RESPONSE
6	2	STATUS	Always X'0000' in requests. In responses refer to Table 43 on page 491 for a list of values.
8	4	TRANSACTION-ID	Arbitrary value supplied by the requester.
12	2	REQUESTER-LECID	LECID of LE client issuing the LE_ARP request.
14	2	FLAGS	Each bit of the FLAGS file has a separate meaning if set: X'0001' Remote address. The TARGET_LAN_DESTINATION is not registered with the LE server.
16	8	SOURCE-LAN-DESTINATION	Source MAC address from data frame that triggered this LE_ARP sequence. May be encoded with (not present) LAN destination tag.
24	8	TARGET-LAN-DESTINATION	Destination unicast MAC address or next route descriptor for which an ATM address is being sought.
32	20	SOURCE-ATM-ADDRESS	ATM address of originator of LE_ARP request.
52	4	RESERVED	Always X'00' when sent, ignored on receipt.
56	20	TARGET-ATM-ADDRESS	X'00' in LE_ARP request. ATM address of LE_Client responsible for target LAN destination in LE_ARP response.
76	32	RESERVED	Always X'00' when sent, ignored on receipt.

<i>Table 40. LE_NARP Frame Format</i>			
Offset	Size	Name	Function
0	2	MARKER	Control frame = X'FF00'.
2	1	PROTOCOL	ATM LAN emulation protocol = X'01'.
3	1	VERSION	ATM LAN emulation protocol version = X'01'.
4	2	OP-CODE	Type of request: X'0008' LE_NARP_REQUEST
6	2	STATUS	Always X'0000'.
8	4	TRANSACTION-ID	Arbitrary value supplied by the requester.
12	2	REQUESTER-LECID	LECID of LE client issuing the LE_NARP request.
14	2	FLAGS	Always X'00'
16	8	SOURCE-LAN-DESTINATION	Not used. Encoded as X'00'.
24	8	TARGET-LAN-DESTINATION	Destination unicast MAC address or next route descriptor for which the target ATM address no longer applies.
32	20	SOURCE-ATM-ADDRESS	ATM address of originator of LE_NARP request.
52	4	RESERVED	Always X'00'when sent, ignored on receipt.
56	20	TARGET-ATM-ADDRESS	Target ATM address of LE_Client which was previously representing the target LAN destination.
76	32	RESERVED	Always X'00'when sent, ignored on receipt.

<i>Table 41. Topology Change Frame Format</i>			
Offset	Size	Name	Function
0	2	MARKER	Control frame = X'FF00'.
2	1	PROTOCOL	ATM LAN emulation protocol = X'01'.
3	1	VERSION	ATM LAN emulation protocol version = X'01'.
4	2	OP-CODE	Type of request: X'0009' LE_TOPOLOGY_REQUEST
6	2	STATUS	Always X'0000'.
8	4	TRANSACTION-ID	Arbitrary value supplied by the requester.
12	2	REQUESTER-LECID	LECID of LE client issuing the topology change request.
14	2	FLAGS	Each bit of the FLAGS field has a separate meaning if set: X'0100' Topology Change Flag. A network topology change is in progress.
16	92	RESERVED	Always X'00' when sent, ignored on receipt.

D.7 Flush Frame Format

Offset	Size	Name	Function
0	2	MARKER	Control frame = X'FF00'.
2	1	PROTOCOL	ATM LAN emulation protocol = X'01'.
3	1	VERSION	ATM LAN emulation protocol version = X'01'.
4	2	OP-CODE	Type of request: X'0007' LE_FLUSH_REQUEST X'0107' LE_FLUSH_RESPONSE
6	2	STATUS	Always X'0000' in requests. In responses refer to Table 43 on page 491 for a list of values.
8	4	TRANSACTION-ID	Arbitrary value supplied by the requester.
12	2	REQUESTER-LECID	LECID of LE client issuing the LE_ARP request.
14	2	FLAGS	Always 0 when sent, ignored on receipt.
16	8	SOURCE-LAN-DESTINATION	Always X'00' when sent, ignored on receipt.
24	8	TARGET-LAN-DESTINATION	Always X'00' when sent, ignored on receipt.
32	20	SOURCE-ATM-ADDRESS	ATM address of originator of the flush request.
52	4	RESERVED	Always X'00' when sent, ignored on receipt.
24	8	TARGET-LAN-DESTINATION	Destination unicast MAC address or next route descriptor for which an ATM address is being sought.
56	20	TARGET-ATM-ADDRESS	ATM address of LE_Client to which flush request is directed.
76	32	RESERVED	Always X'00' when sent, ignored on receipt.

D.8 Control Frame Status Values

Table 43. Control Frame Status Values

Code (dec)	Name	Meaning	Responses
0	Success	Successful response	All responses
1	Version not supported	Version field of request contains a value higher than that supported by the responder.	All responses
2	Invalid request parameters	The parameters given are incompatible with the ELAN.	All responses
4	Duplicate LAN destination registration	SOURCE-LAN-DESTINATION duplicates a previously registered LAN destination.	Join or Register
5	Duplicate ATM address	SOURCE-ATM-ADDRESS duplicates a previously registered ATM address	Join or Register
6	Insufficient resources to grant request	Responder is unable to grant request for reasons such as insufficient table space or ability to establish VCCs.	Configure, Join or Register
7	Access denied	Request denied for security reasons.	Configure or Join
8	Invalid REQUESTER-ID	LECID field is not zero (Configure or Join) or is not LE client's LECID.	Configure, Join, Register, Unregister, ARP
9	Invalid LAN destination	LAN destination is a multicast address or on an Ethernet/IEEE 802.3 emulated LAN, a route descriptor.	Configure, Join, Register, ARP, Flush
10	Invalid ATM address	ATM address is not in a recognizable format.	Configure, Join, Register, ARP, Flush
20	No configuration	LE client is not recognized.	Configure
21	LE_CONFIGURE error	Parameters supplied give conflicting answers. May also be used to refuse service without giving a specific reason.	Configure
22	Insufficient information	LE client has not provided sufficient information to allow the LECS to assign it to a specific emulated LAN.	Configure

Note: The responses field indicates the commands for which the control status code value is valid.

Appendix E. Traces and MIBs References

The following MIB references could be useful if you see them in formatted traces.

E.1 Trace Formatter Object Names

<i>Table 44 (Page 1 of 2). Object Identification</i>	
Object Identification	Symbolic Name in Formatted Trace
1.3.6.1.2.1.1	"MIB2.System"
1.3.6.1.2.1.1.1	"sysDescr"
1.3.6.1.2.1.1.2	"sysObjectID"
1.3.6.1.2.1.1.5	"sysName"
1.3.6.1.2.1.1.6	"sysLocation"
1.3.6.1.4.1.353.1.5.1	"atmfSrvcRegLecs"
1.3.6.1.4.1.353.2	"atmForumUni"
1.3.6.1.4.1.353.2.1.1.1.7	"atmfPortMyIfName"
1.3.6.1.4.1.353.2.1.2	"MyIpNmAddress"
1.3.6.1.4.1.353.2.1.3	"MyOsiNmNsapAddress"
1.3.6.1.4.1.353.2.2.1.1.6	"MaxVpiBits"
1.3.6.1.4.1.353.2.2.1.1.7	"MaxVciBits"
1.3.6.1.4.1.353.2.2.1.1.9	"UniVersion"
1.3.6.1.4.1.353.2.6.1.1	"atmfAddressEntry"
1.3.6.1.4.1.353.2.6.1.1.1	"atmfAddressPort"
1.3.6.1.4.1.353.2.6.1.1.2	"atmfAddressAtmAddress"
1.3.6.1.4.1.353.2.6.1.1.3	"AddressTable"
1.3.6.1.4.1.353.2.7.1.1	"atmfNetPrefixEntry"
1.3.6.1.4.1.353.2.7.1.1.1	"atmfNetPrefixPort"
1.3.6.1.4.1.353.2.7.1.1.2	"atmfNetPrefixPrefix"
1.3.6.1.4.1.353.2.7.1.1.3	"NetPrefixTable"
1.3.6.1.4.1.353.2.8.1.1.3	"atmfSrvcRegATMAddress"
1.3.6.1.4.1.2.6.33.1.5.1.1	"nbrEntry"
1.3.6.1.4.1.2.6.33.1.5.1.1.1	"nbrlocalIndex"
1.3.6.1.4.1.2.6.33.1.5.1.1.2	"nbrIpAddress1"
1.3.6.1.4.1.2.6.33.1.5.1.1.3	"nbrIpAddress2"
1.3.6.1.4.1.2.6.33.1.5.1.1.4	"nbrAtmAddress"
1.3.6.1.4.1.2.6.33.1.5.1.1.5	"nbrIndex"
1.3.6.1.4.1.2.6.33.1.5.1.1.6	"nbrDescriptor"
1.3.6.1.4.1.2.6.33.1.5.1.1.7	"nbrOid"
1.3.6.1.4.1.2.6.33.1.5.1.1.8	"nbrName"
1.3.6.1.4.1.2.6.33.1.5.1.1.9	"nbrLocation"

Object Identification	Symbolic Name in Formatted Trace
1.3.6.1.4.1.2.6.33.1.5.1.1.10	"nbrtrunkId"
1.3.6.1.4.1.2.6.33.1.1.1	"dateTime"
1.3.6.1.4.1.2.6.33.1.1.2	"lastChange"
1.3.6.1.4.1.2.6.33.1.6.2	"transferDate"
1.3.6.1.4.1.2.6.33.1.9.1.3.1.9	"atmSvcCreationTime"
1.3.6.1.4.1.2.6.33.1.9.1.6.1.5	"atmSvcLogCreationTime"
1.3.6.1.4.1.2.6.33.1.9.1.6.1.6	"atmSvcLogTime"
1.3.6.1.4.1.2.6.33.1.10.3.1.10	"atmPvcEpLastActive"

E.2 ATM Forum MIB Variable Listing

<i>Table 45 (Page 1 of 5). MIB Listing</i>	
Formatted Trace	Object Identification
sysDescr	1.3.6.1.2.1.1.1.
sysObjectID	1.3.6.1.2.1.1.2.
sysUpTime	1.3.6.1.2.1.1.3.
sysContact	1.3.6.1.2.1.1.4.
sysName	1.3.6.1.2.1.1.5.
sysLocation	1.3.6.1.2.1.1.6.
sysServices	1.3.6.1.2.1.1.7.
ifNumber	1.3.6.1.2.1.2.1.
ifIndex	1.3.6.1.2.1.2.2.1.1.
ifDescr	1.3.6.1.2.1.2.2.1.2.
ifType	1.3.6.1.2.1.2.2.1.3.
ifMtu	1.3.6.1.2.1.2.2.1.4.
ifSpeed	1.3.6.1.2.1.2.2.1.5.
ifPhysAddress	1.3.6.1.2.1.2.2.1.6.
ifAdminStatus	1.3.6.1.2.1.2.2.1.7.
ifOperStatus	1.3.6.1.2.1.2.2.1.8.
ifLastChange	1.3.6.1.2.1.2.2.1.9.
ifInOctets	1.3.6.1.2.1.2.2.1.10.
ifInUcastPkts	1.3.6.1.2.1.2.2.1.11.
ifInNUcastPkts	1.3.6.1.2.1.2.2.1.12.
ifInDiscards	1.3.6.1.2.1.2.2.1.13.
ifInErrors	1.3.6.1.2.1.2.2.1.14.
ifInUnknownProtos	1.3.6.1.2.1.2.2.1.15.
ifOutOctets	1.3.6.1.2.1.2.2.1.16.
ifOutUcastPkts	1.3.6.1.2.1.2.2.1.17.
ifOutNUcastPkts	1.3.6.1.2.1.2.2.1.18.
ifOutDiscards	1.3.6.1.2.1.2.2.1.19.
ifOutErrors	1.3.6.1.2.1.2.2.1.20.
ifOutQLen	1.3.6.1.2.1.2.2.1.21.
ifSpecific	1.3.6.1.2.1.2.2.1.22.
atIfIndex	1.3.6.1.2.1.3.1.1.1.
atPhysAddress	1.3.6.1.2.1.3.1.1.2.
atNetAddress	1.3.6.1.2.1.3.1.1.3.
ipForwarding	1.3.6.1.2.1.4.1.
ipDefaultTTL	1.3.6.1.2.1.4.2.
ipInReceives	1.3.6.1.2.1.4.3.
ipInHdrErrors	1.3.6.1.2.1.4.4.

<i>Table 45 (Page 2 of 5). MIB Listing</i>	
Formatted Trace	Object Identification
ipInAddrErrors	1.3.6.1.2.1.4.5.
ipForwDatagrams	1.3.6.1.2.1.4.6.
ipInUnknownProtos	1.3.6.1.2.1.4.7.
ipInDiscards	1.3.6.1.2.1.4.8.
ipInDelivers	1.3.6.1.2.1.4.9.
ipOutRequests	1.3.6.1.2.1.4.10.
ipOutDiscards	1.3.6.1.2.1.4.11.
ipOutNoRoutes	1.3.6.1.2.1.4.12.
ipReasmTimeout	1.3.6.1.2.1.4.13.
ipReasmReqds	1.3.6.1.2.1.4.14.
ipReasmOKs	1.3.6.1.2.1.4.15.
ipReasmFails	1.3.6.1.2.1.4.16.
ipFragOKs	1.3.6.1.2.1.4.17.
ipFragFails	1.3.6.1.2.1.4.18.
ipFragCreates	1.3.6.1.2.1.4.19.
ipAdEntAddr	1.3.6.1.2.1.4.20.1.1.
ipAdEntIfIndex	1.3.6.1.2.1.4.20.1.2.
ipAdEntNetMask	1.3.6.1.2.1.4.20.1.3.
ipAdEntBcastAddr	1.3.6.1.2.1.4.20.1.4.
ipAdEntReasmMaxSize	1.3.6.1.2.1.4.20.1.5.
ipRouteDest	1.3.6.1.2.1.4.21.1.1.
ipRouteIfIndex	1.3.6.1.2.1.4.21.1.2.
ipRouteMetric1	1.3.6.1.2.1.4.21.1.3.
ipRouteMetric2	1.3.6.1.2.1.4.21.1.4.
ipRouteMetric3	1.3.6.1.2.1.4.21.1.5.
ipRouteMetric4	1.3.6.1.2.1.4.21.1.6.
ipRouteNextHop	1.3.6.1.2.1.4.21.1.7.
ipRouteType	1.3.6.1.2.1.4.21.1.8.
ipRouteProto	1.3.6.1.2.1.4.21.1.9.
ipRouteAge	1.3.6.1.2.1.4.21.1.10.
ipRouteMask	1.3.6.1.2.1.4.21.1.11.
ipRouteMetric5	1.3.6.1.2.1.4.21.1.12.
ipRouteInfo	1.3.6.1.2.1.4.21.1.13.
ipNetToMediaIfIndex	1.3.6.1.2.1.4.22.1.1.
ipNetToMediaPhysAddress	1.3.6.1.2.1.4.22.1.2.
ipNetToMediaNetAddress	1.3.6.1.2.1.4.22.1.3.
ipNetToMediaType	1.3.6.1.2.1.4.22.1.4.
ipRoutingDiscards	1.3.6.1.2.1.4.23.
icmpInMsgs	1.3.6.1.2.1.5.1.
icmpInErrors	1.3.6.1.2.1.5.2.

<i>Table 45 (Page 3 of 5). MIB Listing</i>	
Formatted Trace	Object Identification
icmpInDestUnreachs	1.3.6.1.2.1.5.3.
icmpInTimeExcds	1.3.6.1.2.1.5.4.
icmpInParmProbs	1.3.6.1.2.1.5.5.
icmpInSrcQuenchs	1.3.6.1.2.1.5.6.
icmpInRedirects	1.3.6.1.2.1.5.7.
icmpInEchos	1.3.6.1.2.1.5.8.
icmpInEchoReps	1.3.6.1.2.1.5.9.
icmpInTimestamps	1.3.6.1.2.1.5.10.
icmpInTimestampReps	1.3.6.1.2.1.5.11.
icmpInAddrMasks	1.3.6.1.2.1.5.12.
icmpInAddrMaskReps	1.3.6.1.2.1.5.13.
icmpOutMsgs	1.3.6.1.2.1.5.14.
icmpOutErrors	1.3.6.1.2.1.5.15.
icmpOutDestUnreachs	1.3.6.1.2.1.5.16.
icmpOutTimeExcds	1.3.6.1.2.1.5.17.
icmpOutParmProbs	1.3.6.1.2.1.5.18.
icmpOutSrcQuenchs	1.3.6.1.2.1.5.19.
icmpOutRedirects	1.3.6.1.2.1.5.20.
icmpOutEchos	1.3.6.1.2.1.5.21.
icmpOutEchoReps	1.3.6.1.2.1.5.22.
icmpOutTimestamps	1.3.6.1.2.1.5.23.
icmpOutTimestampReps	1.3.6.1.2.1.5.24.
icmpOutAddrMasks	1.3.6.1.2.1.5.25.
icmpOutAddrMaskReps	1.3.6.1.2.1.5.26.
tcpRtoAlgorithm	1.3.6.1.2.1.6.1.
tcpRtoMin	1.3.6.1.2.1.6.2.
tcpRtoMax	1.3.6.1.2.1.6.3.
tcpMaxConn	1.3.6.1.2.1.6.4.
tcpActiveOpens	1.3.6.1.2.1.6.5.
tcpPassiveOpens	1.3.6.1.2.1.6.6.
tcpAttemptFails	1.3.6.1.2.1.6.7.
tcpEstabResets	1.3.6.1.2.1.6.8.
tcpCurrEstab	1.3.6.1.2.1.6.9.
tcpInSegs	1.3.6.1.2.1.6.10.
tcpOutSegs	1.3.6.1.2.1.6.11.
tcpRetransSegs	1.3.6.1.2.1.6.12.
tcpConnState	1.3.6.1.2.1.6.13.1.1.
tcpConnLocalAddress	1.3.6.1.2.1.6.13.1.2.
tcpConnLocalPort	1.3.6.1.2.1.6.13.1.3.
tcpConnRemAddress	1.3.6.1.2.1.6.13.1.4.

<i>Table 45 (Page 4 of 5). MIB Listing</i>	
Formatted Trace	Object Identification
tcpConnRemPort	1.3.6.1.2.1.6.13.1.5.
tcpInErrs	1.3.6.1.2.1.6.14.
tcpOutRsts	1.3.6.1.2.1.6.15.
udpInDatagrams	1.3.6.1.2.1.7.1.
udpNoPorts	1.3.6.1.2.1.7.2.
udpInErrors	1.3.6.1.2.1.7.3.
udpOutDatagrams	1.3.6.1.2.1.7.4.
udpLocalAddress	1.3.6.1.2.1.7.5.1.1.
udpLocalPort	1.3.6.1.2.1.7.5.1.2.
egpInMsgs	1.3.6.1.2.1.8.1.
egpInErrors	1.3.6.1.2.1.8.2.
egpOutMsgs	1.3.6.1.2.1.8.3.
egpOutErrors	1.3.6.1.2.1.8.4.
egpNeighState	1.3.6.1.2.1.8.5.1.1.
egpNeighAddr	1.3.6.1.2.1.8.5.1.2.
egpNeighAs	1.3.6.1.2.1.8.5.1.3.
egpNeighInMsgs	1.3.6.1.2.1.8.5.1.4.
egpNeighInErrs	1.3.6.1.2.1.8.5.1.5.
egpNeighOutMsgs	1.3.6.1.2.1.8.5.1.6.
egpNeighOutErrs	1.3.6.1.2.1.8.5.1.7.
egpNeighInErrMsgs	1.3.6.1.2.1.8.5.1.8.
egpNeighOutErrMsgs	1.3.6.1.2.1.8.5.1.9.
egpNeighStateUps	1.3.6.1.2.1.8.5.1.10.
egpNeighStateDowns	1.3.6.1.2.1.8.5.1.11.
egpNeighIntervalHello	1.3.6.1.2.1.8.5.1.12.
egpNeighIntervalPoll	1.3.6.1.2.1.8.5.1.13.
egpNeighMode	1.3.6.1.2.1.8.5.1.14.
egpNeighEventTrigger	1.3.6.1.2.1.8.5.1.15.
egpAs	1.3.6.1.2.1.8.6.
snmpInPkts	1.3.6.1.2.1.11.1.
snmpOutPkts	1.3.6.1.2.1.11.2.
snmpInBadVersions	1.3.6.1.2.1.11.3.
snmpInBadCommunityNames	1.3.6.1.2.1.11.4.
snmpInBadCommunityUses	1.3.6.1.2.1.11.5.
snmpInASNParseErrs	1.3.6.1.2.1.11.6.
snmpInTooBigs	1.3.6.1.2.1.11.8.
snmpInNoSuchNames	1.3.6.1.2.1.11.9.
snmpInBadValues	1.3.6.1.2.1.11.10.
snmpInReadOnlys	1.3.6.1.2.1.11.11.
snmpInGenErrs	1.3.6.1.2.1.11.12.

<i>Table 45 (Page 5 of 5). MIB Listing</i>	
Formatted Trace	Object Identification
snmpInTotalReqVars	1.3.6.1.2.1.11.13.
snmpInTotalSetVars	1.3.6.1.2.1.11.14.
snmpInGetRequests	1.3.6.1.2.1.11.15.
snmpInGetNexts	1.3.6.1.2.1.11.16.
snmpInSetRequests	1.3.6.1.2.1.11.17.
snmpInGetResponses	1.3.6.1.2.1.11.18.
snmpInTraps	1.3.6.1.2.1.11.19.
snmpOutTooBigs	1.3.6.1.2.1.11.20.
snmpOutNoSuchNames	1.3.6.1.2.1.11.21.
snmpOutBadValues	1.3.6.1.2.1.11.22.
snmpOutGenErrs	1.3.6.1.2.1.11.24.
snmpOutGetRequests	1.3.6.1.2.1.11.25.
snmpOutGetNexts	1.3.6.1.2.1.11.26.
snmpOutSetRequests	1.3.6.1.2.1.11.27.
snmpOutGetResponses	1.3.6.1.2.1.11.28.
snmpOutTraps	1.3.6.1.2.1.11.29.
snmpEnableAuthenTraps	1.3.6.1.2.1.11.30.

E.3 ILMI Objects

<i>Table 46 (Page 1 of 2). ILMI</i>	
List	References
atmfPortIndex	1.3.6.1.4.1.353.2.1.1.1.1.
atmfPortAddress	1.3.6.1.4.1.353.2.1.1.1.2.
atmfPortTransmissionType	1.3.6.1.4.1.353.2.1.1.1.3.
atmfPortMediaType	1.3.6.1.4.1.353.2.1.1.1.4.
atmfPortOperStatus	1.3.6.1.4.1.353.2.1.1.1.5.
atmfPortSpecific	1.3.6.1.4.1.353.2.1.1.1.6.
atmfPortMyIfName	1.3.6.1.4.1.353.2.1.1.1.7.
atmfMyIpNmAddress	1.3.6.1.4.1.353.2.1.2. (Internet)
atmfAtmLayerIndex	1.3.6.1.4.1.353.2.2.1.1.1.
atmfAtmLayerMaxVPCs	1.3.6.1.4.1.353.2.2.1.1.2.
atmfAtmLayerMaxVCCs	1.3.6.1.4.1.353.2.2.1.1.3.
atmfAtmLayerConfiguredVPCs	1.3.6.1.4.1.353.2.2.1.1.4.
atmfAtmLayerConfiguredVCCs	1.3.6.1.4.1.353.2.2.1.1.5.
atmfAtmLayerMaxVpiBits	1.3.6.1.4.1.353.2.2.1.1.6.
atmfAtmLayerMaxVciBits	1.3.6.1.4.1.353.2.2.1.1.7.
atmfAtmLayerUniType	1.3.6.1.4.1.353.2.2.1.1.8.
atmfAtmLayerUniVersion	1.3.6.1.4.1.353.2.2.1.1.9.
atmfAtmStatsIndex	1.3.6.1.4.1.353.2.3.1.1.1.
atmfAtmStatsReceivedCells	1.3.6.1.4.1.353.2.3.1.1.2.
atmfAtmStatsDroppedReceivedCells	1.3.6.1.4.1.353.2.3.1.1.3.
atmfAtmStatsTransmittedCells	1.3.6.1.4.1.353.2.3.1.1.4.
atmfVpcPortIndex	1.3.6.1.4.1.353.2.4.1.1.1.
atmfVpcVpi	1.3.6.1.4.1.353.2.4.1.1.2.
atmfVpcOperStatus	1.3.6.1.4.1.353.2.4.1.1.3.
atmfVpcTransmitTrafficDescriptorType	1.3.6.1.4.1.353.2.4.1.1.4.
atmfVpcTransmitTrafficDescriptorParam1	1.3.6.1.4.1.353.2.4.1.1.5.
atmfVpcTransmitTrafficDescriptorParam2	1.3.6.1.4.1.353.2.4.1.1.6.
atmfVpcTransmitTrafficDescriptorParam3	1.3.6.1.4.1.353.2.4.1.1.7.
atmfVpcTransmitTrafficDescriptorParam4	1.3.6.1.4.1.353.2.4.1.1.8.
atmfVpcTransmitTrafficDescriptorParam5	1.3.6.1.4.1.353.2.4.1.1.9.
atmfVpcReceiveTrafficDescriptorType	1.3.6.1.4.1.353.2.4.1.1.10.
atmfVpcReceiveTrafficDescriptorParam1	1.3.6.1.4.1.353.2.4.1.1.11.
atmfVpcReceiveTrafficDescriptorParam2	1.3.6.1.4.1.353.2.4.1.1.12.
fVpcReceiveTrafficDescriptorParam3	1.3.6.1.4.1.353.2.4.1.1.13.
atmfVpcReceiveTrafficDescriptorParam4	1.3.6.1.4.1.353.2.4.1.1.14.
atmfVpcReceiveTrafficDescriptorParam5	1.3.6.1.4.1.353.2.4.1.1.15.
atmfVpcQoSCategory	1.3.6.1.4.1.353.2.4.1.1.16.
atmfVpcTransmitQoSClass	1.3.6.1.4.1.353.2.4.1.1.17.

<i>Table 46 (Page 2 of 2). ILM1</i>	
List	References
atmfVpcReceiveQosClass	1.3.6.1.4.1.353.2.4.1.1.18.
atmfVccPortIndex	1.3.6.1.4.1.353.2.5.1.1.1.
atmfVccVpi	1.3.6.1.4.1.353.2.5.1.1.2.
atmfVccVci	1.3.6.1.4.1.353.2.5.1.1.3.
atmfVccOperStatus	1.3.6.1.4.1.353.2.5.1.1.4.
atmfVccTransmitTrafficDescriptorType	1.3.6.1.4.1.353.2.5.1.1.5.
atmfVccTransmitTrafficDescriptorParam1	1.3.6.1.4.1.353.2.5.1.1.6.
atmfVccTransmitTrafficDescriptorParam2	1.3.6.1.4.1.353.2.5.1.1.7.
atmfVccTransmitTrafficDescriptorParam3	1.3.6.1.4.1.353.2.5.1.1.8.
atmfVccTransmitTrafficDescriptorParam4	1.3.6.1.4.1.353.2.5.1.1.9.
atmfVccTransmitTrafficDescriptorParam5	1.3.6.1.4.1.353.2.5.1.1.10.
atmfVccReceiveTrafficDescriptorType	1.3.6.1.4.1.353.2.5.1.1.11.
atmfVccReceiveTrafficDescriptorParam1	1.3.6.1.4.1.353.2.5.1.1.12.
atmfVccReceiveTrafficDescriptorParam2	1.3.6.1.4.1.353.2.5.1.1.13.
atmfVccReceiveTrafficDescriptorParam3	1.3.6.1.4.1.353.2.5.1.1.14.
atmfVccReceiveTrafficDescriptorParam4	1.3.6.1.4.1.353.2.5.1.1.15.
atmfVccReceiveTrafficDescriptorParam5	1.3.6.1.4.1.353.2.5.1.1.16.
atmfVccQoSCategory	1.3.6.1.4.1.353.2.5.1.1.17.
atmfVccTransmitQoSClass	1.3.6.1.4.1.353.2.5.1.1.18.
atmfVccReceiveQoSClass	1.3.6.1.4.1.353.2.5.1.1.19.
atmfAddressPort	1.3.6.1.4.1.353.2.6.1.1.1.
atmfAddressAtmAddress	1.3.6.1.4.1.353.2.6.1.1.2.
atmfAddressStatus	1.3.6.1.4.1.353.2.6.1.1.3.
atmfNetPrefixPort	1.3.6.1.4.1.353.2.7.1.1.1.
atmfNetPrefixPrefix	1.3.6.1.4.1.353.2.7.1.1.2.
atmfNetPrefixStatus	1.3.6.1.4.1.353.2.7.1.1.3.
atmfSrvRegPort	1.3.6.1.4.1.353.2.8.1.1.1.
atmfSrvRegServiceID	1.3.6.1.4.1.353.2.8.1.1.2.
atmfSrvRegATMAddress	1.3.6.1.4.1.353.2.8.1.1.3.
atmfSrvRegAddressIndex	1.3.6.1.4.1.353.2.8.1.1.4.

Appendix F. ATM/LAN/WAN Analyzers Used

Several analyzers are now offered on the market; two of these on which screens were captured are listed here.

F.1 InterWATCH 95000

This analyzer was used during the ITSO residency.

F.1.1 Overview

Designed by the GN Nettetst Navtel division, the interWatch 95000 is a portable multiport LAN - WAN - ATM protocol analyzer designed especially for the high performance requirements of broadband ATM inter-network testing. This equipment can be controlled from a remote X-terminal.

A worldwide product pricing agreement reference "1075067032" was signed between GN Nettetst and IBM in 1996.

F.1.2 Features

Operator Guide

- X-Window/Motif guide with multiple window display
- Save and launch application setups from icon
- One-click setup: line monitor - traffic generator - testing

Monitoring

- High-speed direct capture (to OC3 rates) with time stamping
- Full-duplex operation
- Higher layer protocol decodes (ILMI, IP, RFC1577, RFC1483, etc.)
- Filters, statistics and events (trigger/action) defined at each layer

Inter-Network Correlation

- Correlation by time, or time and data, between all data streams captured during simultaneous multi-port monitoring

Traffic Generation

- Over 100,000 global and individual frame/cell combinations
- VBR (burst), CBR (constant) and normal rate distribution patterns
- Recipe mixing control - up to 100% bandwidth utilization (to OC3)
- Error insertion

ATM

- Monitor signalling emulation and decode options
- Monitor LAN emulation (LANE) option
- AAL and AAL5 header higher layer operations
- UNI3.0 and UNI3.1 (c/w Q.SAAL or Q.2100)
- GCRA verification (policing)

- Cell tests: delay, error, loss, missinsertion, QoS, single and dual leaky bucket

LAN

- Filtering - Ethernet, Novell, 802.2, 802.3, (including SNAP), octet string match

F.1.3 Network Interfaces Available

- OC3/STM1 single mode and multimode
- Taxi/100 Mbps multimode fiber
- DS3 and E3
- Dual port token-ring
- Dual port Ethernet
- E1 and DS1(T1)

F.1.4 Tests and Applications

1. QoS verification

Can perform cell tests in real-time at line rate for up to 72 hours. This provides the accuracy and large sample sizes needed to identify problems with switch architectures, physical media errors, propagation delays, queuing and routing. To ensure accuracy, QoS tests are usually done in out-of-service conditions.

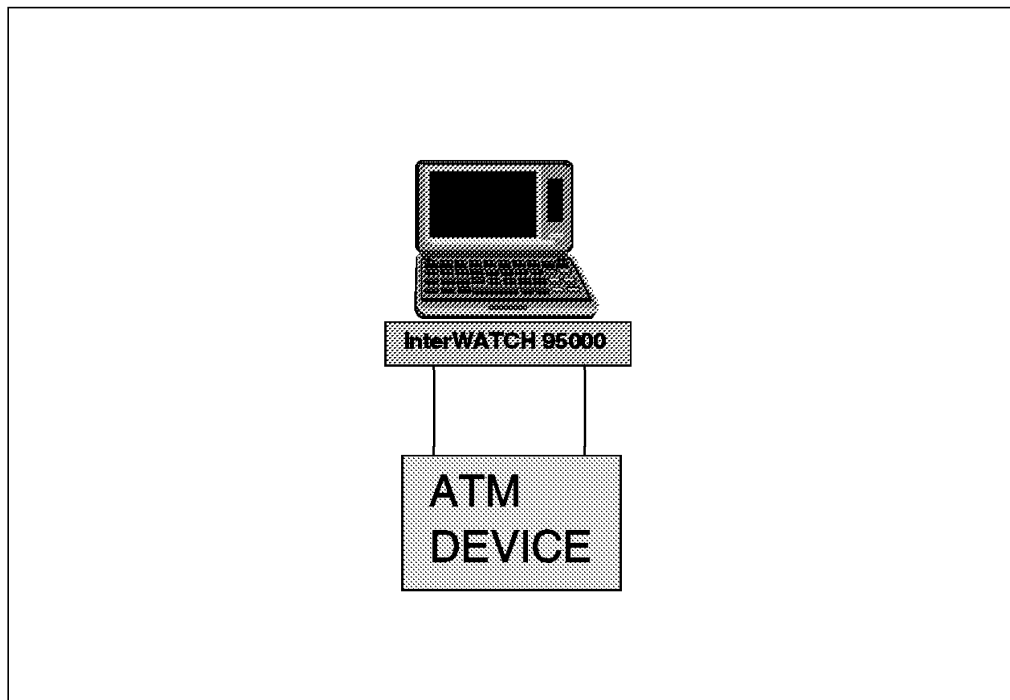


Figure 318. QoS Test

2. In-service test data insertion

The analyzer is directly attached to the ATM switch. Test cells are generated on specific, previously unassigned VCIs or VPIs and multiplexed with existing

user traffic. Measurements can be done on either the user traffic or the test traffic.




Figure 319. In-Service Test Data Insertion

3. Regenerative monitoring

If the ATM switches do not provide test ports, the link can be broken and the signal can be regenerated by the analyzer. In this case the analyzer acts like a repeater by regenerating the received signal. The data received is forwarded and regenerated by the analyzer without any modification. To minimize impact on delay, clock jitter and any other parameters, this generation is done at the line level.

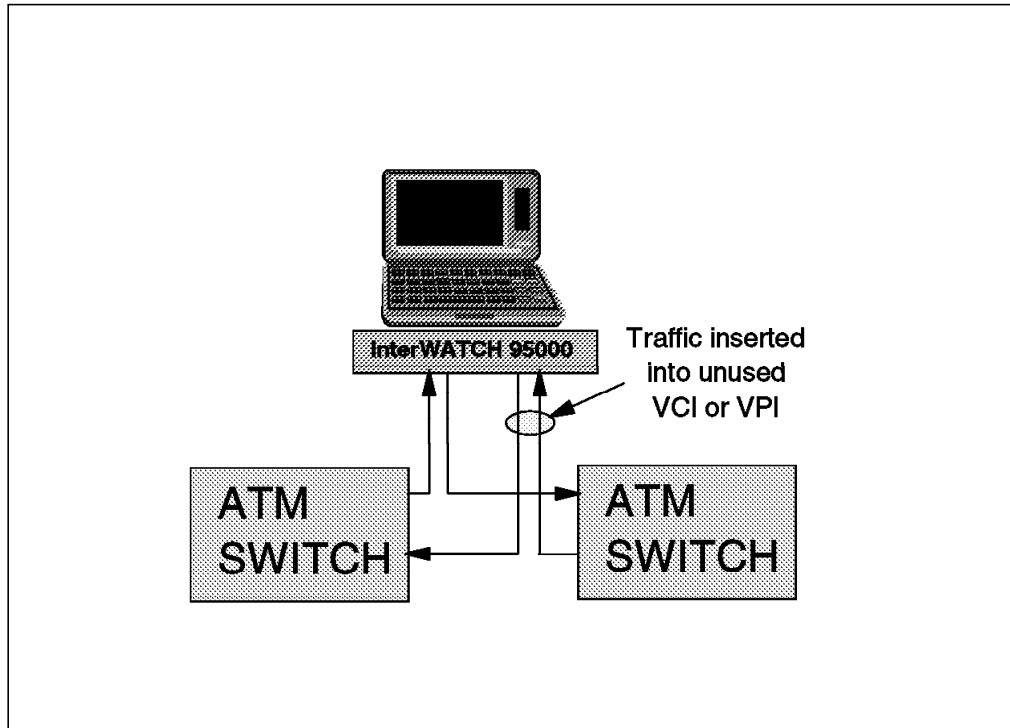


Figure 320. Regenerative Monitoring

4. Out-of-service analysis

This test is performed over dedicated test links with no user data present. This permits rigorous stress testing and analysis of the system under test. The analyzer can generate test cells and background traffic on a mix of VPIs and VCIs. Test traffic is generated by the analyzer itself.




Figure 321. Out-of-Service Analysis

5. Multiport LAN monitoring and analysis

The analyzer can be simultaneously connected to several segments. This setup is useful to determine router throughput performance.

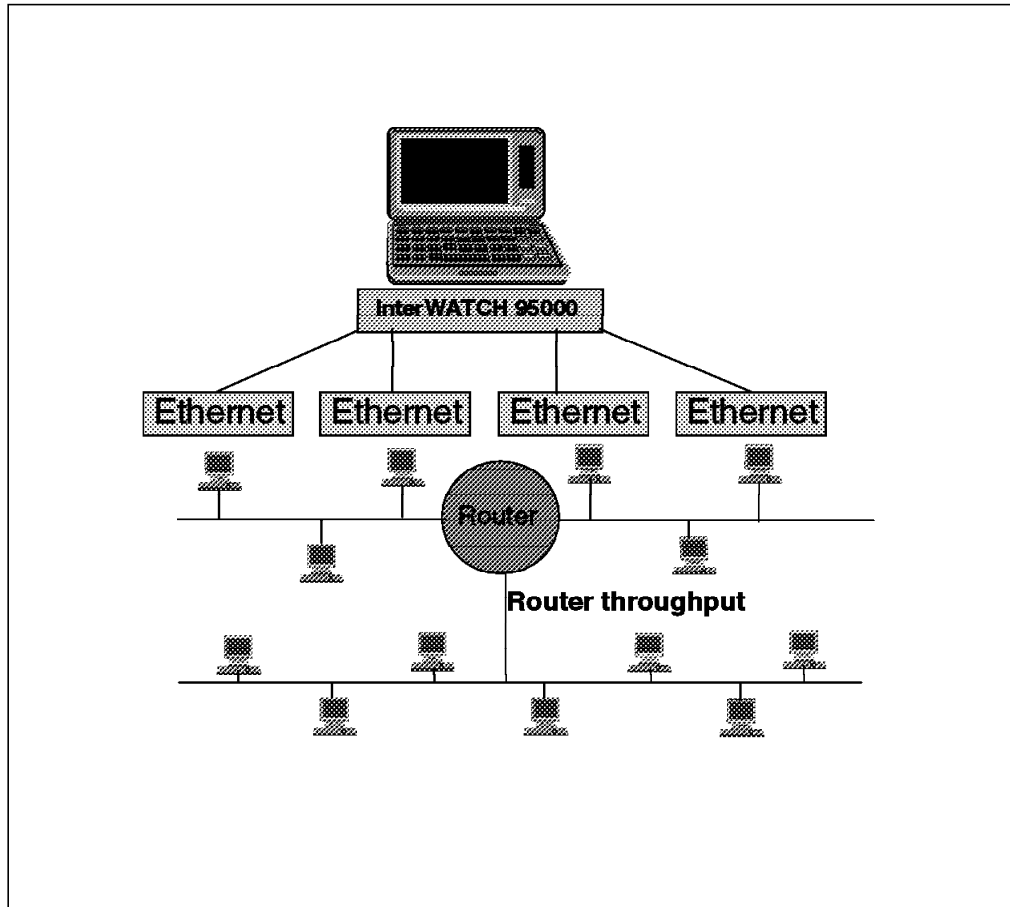


Figure 322. Multiport LAN Monitoring and Analysis

6. ATM switch verification

In this application, legacy LAN traffic and ATM traffic are carried on an ATM switch. The analyzer can connect all ports simultaneously to verify switch operation.

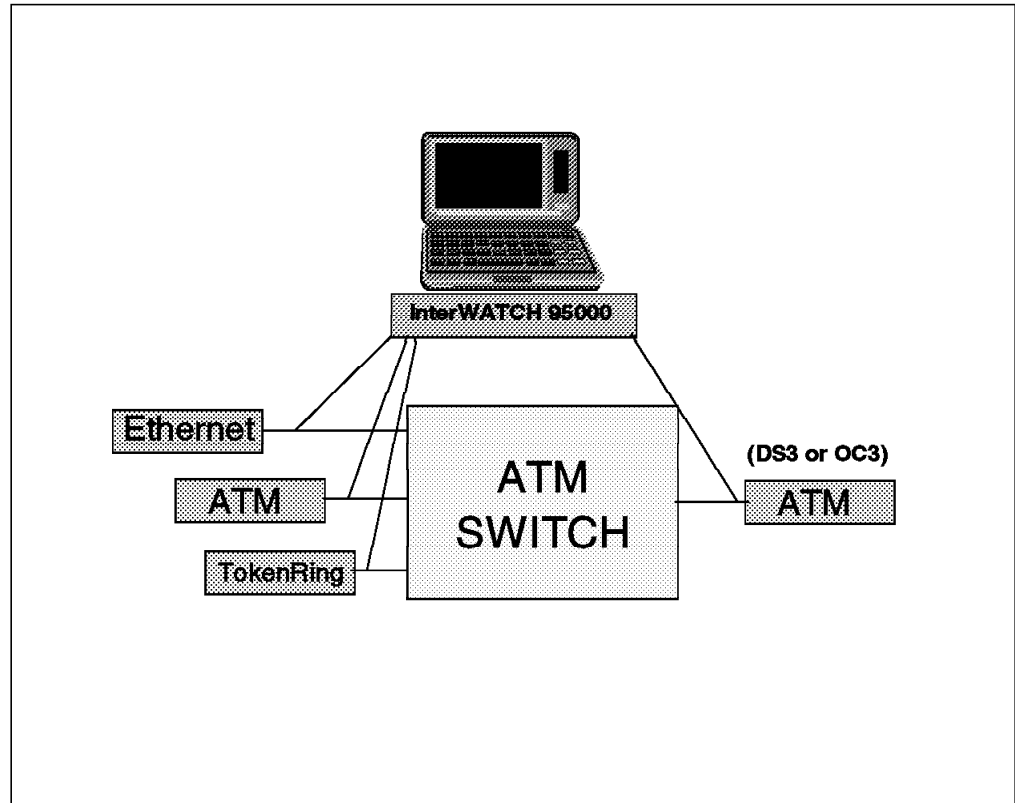


Figure 323. ATM Switch Verification

7. Multiport, multitopology monitoring

The analyzer performs simultaneously non-intrusive monitoring and analysis across in local and wide area ports. It can display and correlate data between these ports to within 1 microsecond to identify a range of potential problems at various layers in the protocol stack. The 1 microsecond resolution accuracy is important for local area ATM switching applications where latency is about 30 microseconds. This capability is especially useful when the network contains a blend of legacy WAN circuits and new ATM circuits as the tester can be used to verify cut-over activities as the circuits are moved from the legacy systems to the new facilities.

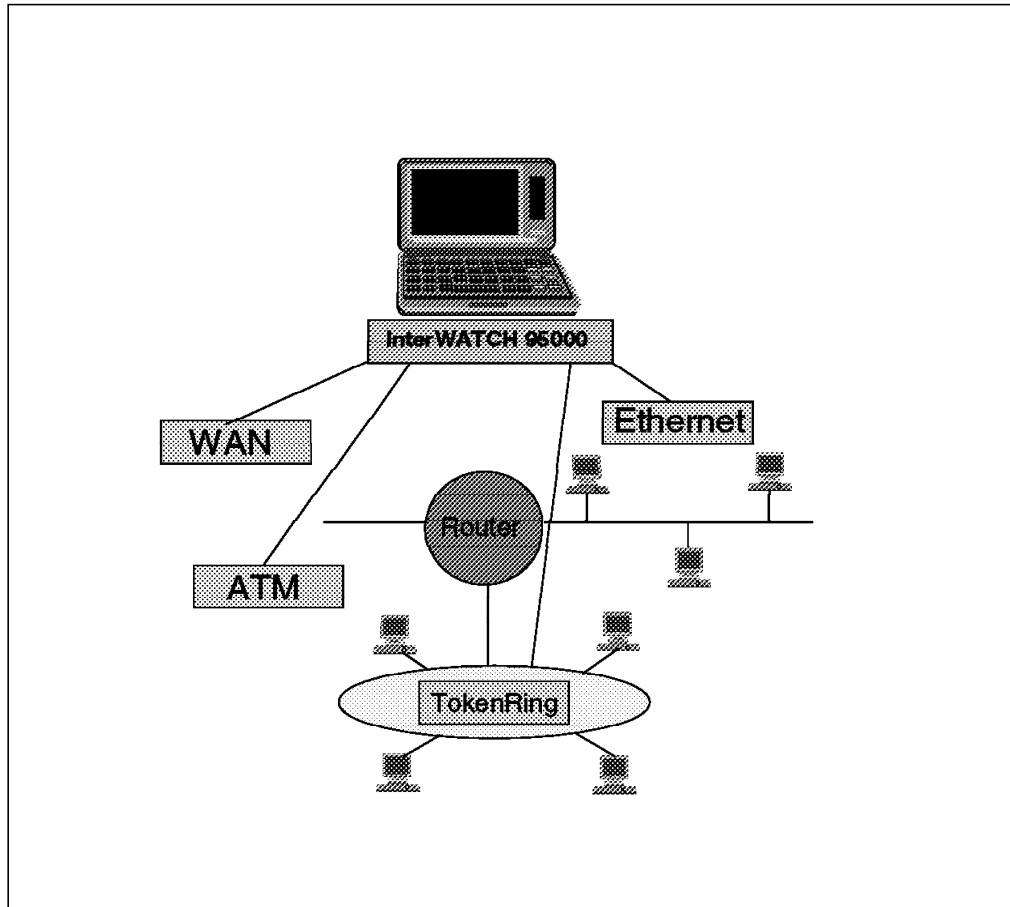


Figure 324. Multiport, Multitopology Monitoring

8. Connectivity testing

The analyzer can perform a number of functions in this application. It can test whether devices with native ATM can communicate with each other. Next, it can test if devices without native ATM can communicate. Then it can test if devices with native ATM can communicate with devices requiring a conversion to ATM.

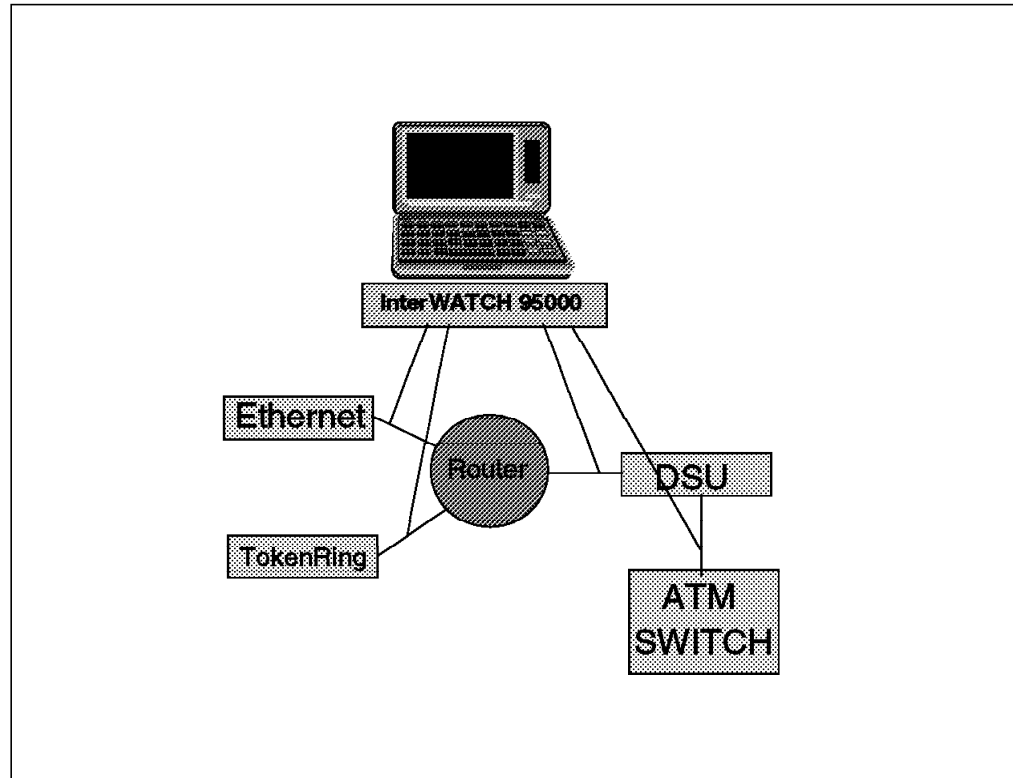


Figure 325. Connectivity Testing

9. LAN emulation operation

A LAN emulation option must be added to provide LANE specific decode, statistics, filters and events to facilitate LAN system troubleshooting over an ATM backbone network. When installed, the LAN emulation function is integrated into the existing monitor.

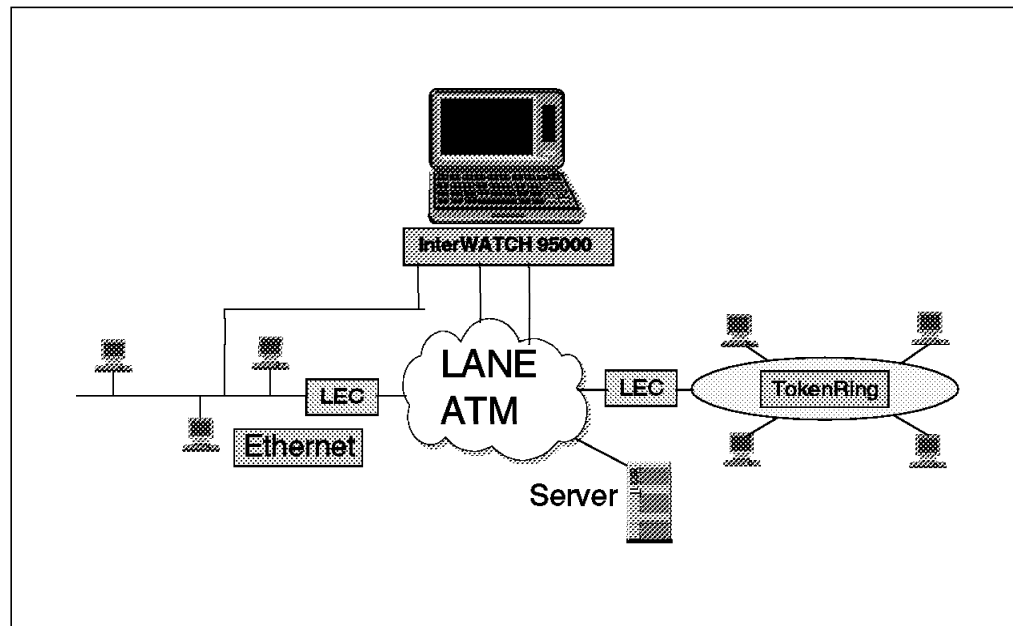


Figure 326. LAN Emulation Option

F.2 DA-30C

This tester was used during the tests.

F.2.1 Overview

This Wandel and Goltermann Internetwork analyzer is known for its portability. It employs a Microsoft Windows-based interface.

A product pricing agreement reference: "C12764" was signed between IBM and Wandel and Goltermann.

F.2.2 Features

Capture

- Filter/trigger/capture
- Post trigger capture

Monitor

- E3/D3/OC3 interfaces
- ATM cell statistics
- ATM traffic statistics
- Filtering and capture of all traffic matching the filter characteristics
- AAL5 statistics.

Each reassembled channel shows total reassembled frames, bandwidth utilization percent, as well as the number of frame, payload, overhead, and total bits per second.

Graphical Statistic Displays

- Histogram graphic displays available for all network statistics.
- Up to eight different statistics can be overlaid for purpose of comparison.
- Scaling can be automatic or user selected.

Real-Time Decode

- Expert system can analyze online frames versus standards.
- Cells or AAL PDUs that are filtered into the capture RAM can be decoded and displayed as they happen in real time or completely examined offline after capturing.

Transmit

- Can transmit user-defined sequences of cells or segmented AAL PDUs at up to full line rate.
- Binary or text files can be imported for use as AAL5 payload sequences.
- Alarm/error insert.
- VC database.
- Analyzer network connection available in emulate and monitor modes.

F.2.3 DA-30C Analyzer and E3/DS3/OC3 Analysis Applications

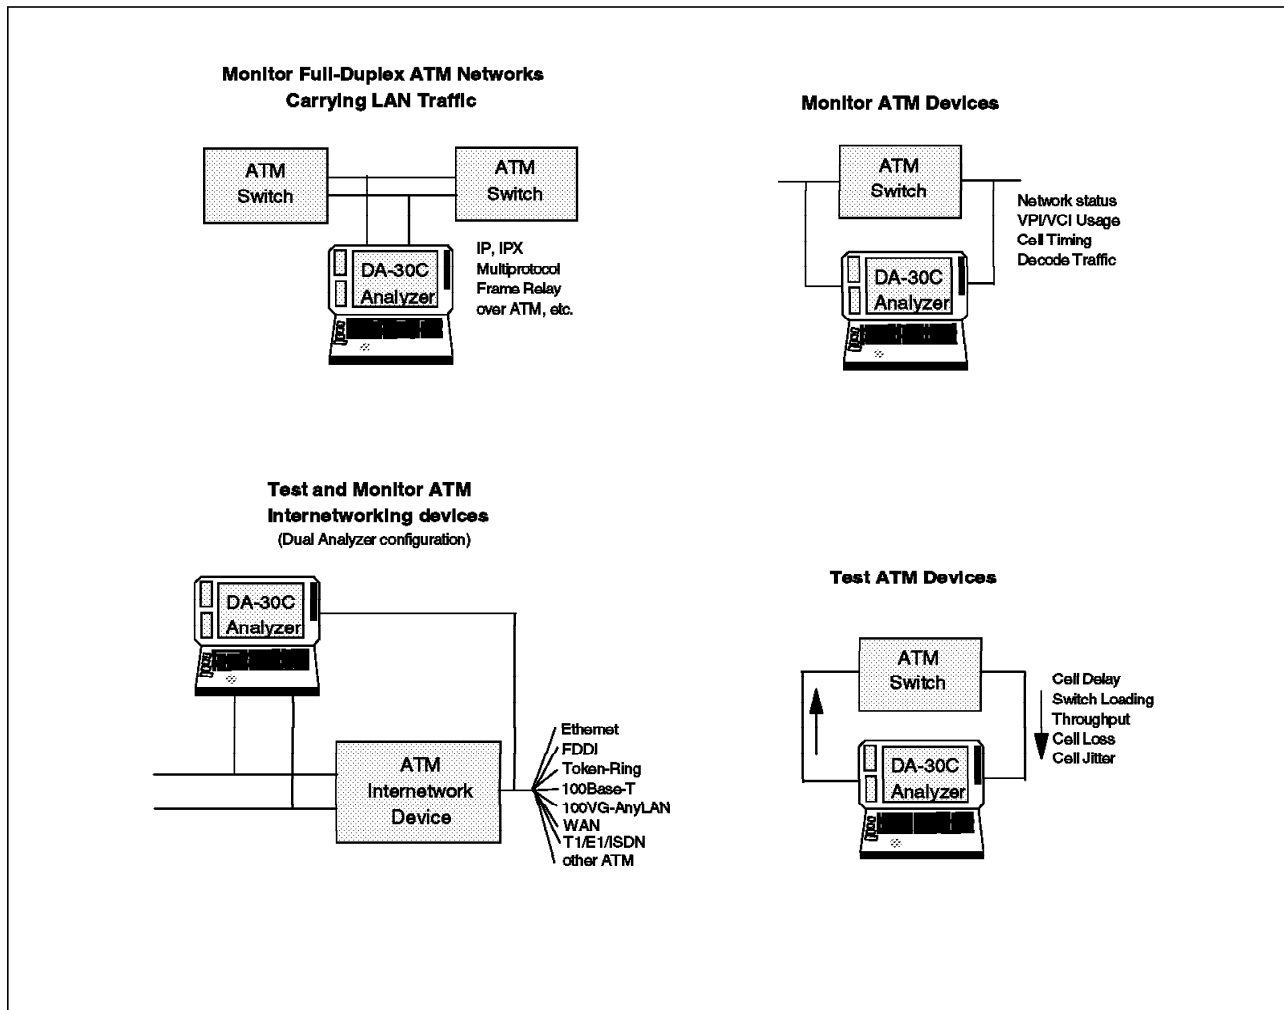


Figure 327. ATM DA-30C Applications

A DA-30C analyzer equipped with the E3/DS3/OC3 ATM analysis packages is uniquely suited to test the services carried by ATM networks. Utilizing the dual-analyzer, real-time reassembly and full protocol decode capabilities, the DA-30C analyzer provides powerful advantages to labs, including the ability to verify the functionality of any LAN traffic (Ethernet, token-ring, FDDI, etc.) running on an ATM link. Furthermore, these analysis packages give users the ability to measure performance across a network.

Appendix G. Special Notices

This publication is intended to help service personnel and sales to troubleshoot or analyze Multiprotocol LAN/ATM Campus Network. The information in this publication is not intended as the specification of any programming interfaces that are provided with the LAN/ATM Campus products. See the PUBLICATIONS section of the IBM Programming Announcement 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 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.

The following document contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples contain the names of individuals, companies, brands, and products. All of these

names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

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.

You can reproduce a page in this document as a transparency, if that page has the copyright notice on it. The copyright notice must appear on each page being reproduced.

The following terms are trademarks of the International Business Machines Corporation in the United States and/or other countries:

AIX	AT
IBM	Micro Channel
NetView	Nways
OS/2	P2P
RISC System/6000	RS/6000
RT	ThinkPad
TURBOWAYS	3090

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 H. Related Publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

H.1 International Technical Support Organization Publications

For information on ordering these ITSO publications see "How to Get ITSO Redbooks" on page 519.

- *Campus ATM Design Guidelines*, SG24-5002
- *Local Area Network Concepts and Products: Adapters, Hubs and ATM*, SG24-4754
- *IBM 8260 As a Campus ATM Switch*, SG24-5003
- *Local Area Network Concepts and Products: LAN Architecture*, SG24-4753
- *Campus ATM Network Management Guideline*, SG24-5006
- *Internetworking over ATM*, SG24-4699
- *IBM Nways RouteSwitch Implementation Guide*, SG24-4881

H.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
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
Application Development Redbooks Collection	SBOF-7290	SK2T-8037
Personal Systems Redbooks Collection	SBOF-7250	SK2T-8042

H.3 Other Publications

These publications are also relevant as further information sources:

- *IBM 8250/8260/8285 Planning and Site Preparation Guide*, GA33-0285
- *8260 Nways Multiprotocol Switching Hub Product Description*, GA33-0315
- *Nways MSS Server Service Manual*, GY27-0354
- *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
- *Events Logging System Message Guide*, SC30-3682-01
- *Nways MSS Server Configuration Guide*, SC30-3821

- *8260 Nways Multiprotocol Switching Hub ATM Control point and Switch module Installation and User's Guide*, SA33-0326
- *IBM 8260 ATM/WAN Installation and User's Guide*, SA33-0397
- *Video Distribution Module Installation and User's Guide*, GA27-4173
- *ATM Campus Introduction, Planning, and Troubleshooting Overview*, GA27-4089
- *ATM 155-Mbps Multimode Fiber Universal Feature Card Planning and Installation Guide* GA27-4156
- *IBM 8285 Nways ATM Workgroup Switch Installation and User's Guide*, SA33-0381
- *IBM 8285 Nways ATM Workgroup Switch SAFETY and SERVICE Catalog*, SA33-0398
- *ATM 4-Port 100 Mbps Module Installation and User's Guide*, SA33-0324
- *Nways 8260 ATM 155 Mbps Flexible Concentration Module* SA33-0357-01
- *8260/8285 ATM 25 MBps Concentration Module Installation and User's Guide*, SA33-0383
- *IBM 8260 ATM 4-Port 155 Mbps Installation and User's Guide*, SA33-0358
- *Nways 8260 ATM TR/Ethernet LAN Bridge Module Installation and User's Guide*, SA33-0361
- *IBM 8260/8285 ATM WAN Module Installation and User's Guide*, SA33-0396

H.4 Performance Information

These are sources of performance information:

- *Factors Influencing ATM Adapter Throughput*, by Andrew Rindos, Steven Woolet, David Cosby, Leonard Hango, Mladen Vouk (IBM Networking Hardware Division) available at URL <http://www.networking.ibm.com/per/perprod.html>
- *The IBM Turboways 155 PCI ATM Adapter: Classical IP and LAN Emulation performance for AIX* (IBM Networking Hardware Division) available at URL <http://www.networking.ibm.com/per/perprod.html>
- *RFC 1323, TCP Extensions for High Performance*, May 1992
- *IBM Performance Monitoring Guide*, SC23-2365-04
- *RS/6000 and Asynchronous Transfer Mode*, SG24-4796
- *Banking on ATM Networking - Real LAN Emulation Interoperability Scenarios*, by David Cosby, Lon Hall, Wes Kinard, Cindy Kueck Young (IBM Corporation)

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
TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET LISTSERV PACKAGE
```

To register for information on workshops, residencies, and redbooks, type the following command:

```
TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ITSOREGI 1996
```

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.htm>). Redpieces are redbooks in progress; not all redbooks become redpieces, and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

How Customers Can Get ITSO Redbooks

Customers may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **Online Orders** — send orders to:

In United States:	IBMMAIL usib6fpl at ibmmail	Internet 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.com/pbl/pbl

- **Internet Listserver**

With an Internet e-mail address, anyone can subscribe to an IBM Announcement Listserv. To initiate the service, send an e-mail note to announce@webster.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.htm>). 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	CRC	Cyclic Redundancy check
ABR	Available Bit Rate	CSMA/CD	Carrier Sense Multiple Access with Collision Detection
ACN	ATM Cluster Number	CTL	Control field (LLC field)
A-CPSW	ATM Control-Point and Switch	DA	Destination Address
AIX	Advanced Interactive Executive	DAAT	Destination Address Association Table
ANR	Automatic Network Routing	DE	Discard Eligibility
APPN	Advanced Peer-to-Peer Networking	DIX	Digital, Intel and Xerox
ARB	All Routes Broadcast	DTL	Designated Transit Lists
ARE	All Routes Explorer	DTR	Data Terminal Ready / Direct Token-Ring
ARI	Address Recognize Information	DXI	Data Exchange Interface
ARP	Address Resolution Protocol	ECC	Error Correction Code
ASCII	American (National) Standard Code for Information Interchange	EDEL	End Delimiter
ATM	Asynchronous Transfer Mode	EFCI	Explicit Forward Congestion Control
AUI	Attachment Unit Interface	ELID	Emulated LAN Identifier
B-ISDN	Broadband ISDN	EMC	Electromagnetic Compatibility
BCM	BroadCast Manager	ERM	Explicite Rate Marking
BOOTP	Boot Protocol (IP)	ESI	End System Identifier
BPDU	Bridge Protocol Data Unit	ETSI	European Telecommunication Standards Institute
Bps	Bytes per second	ELS	Event Logging System
bps	bits per second	FCI	Frame Copied Information
BRI	Basic Rate Interface	FCS	Frame Check Sequence
BUS	Broadcast and Unknown Server	FDDI	Fiber Distributed Data Dnterface
CAC	Call Admission Control	FPGA	Field Programmable Gate Array
CAD	Common ATM Datamover	FTP	File Transfer Protocol
CAM	Content Addressable Memory	Gbps	Gigabits Per Second
CAP	Common ATM Processor	GCAC	Generic Connection admission Control
CBR	Constant Bit Rate	GFC	Generic Flow Control
CCITT	Comite Consultatif International Telegraphique et Telephonique (International Telegraph and Telephone Consultative Committee) now ITU-T	HDLC	High-level Data Link Control
CE	Circuit Emulation	HDTV	High-Definition Tele-Video
CIP	classical IP	HEC	Header Error Check
CLP	Cell Loss Priority	HPR	High Performance Routing
CPCS	Common Part Convergence Sublayer	IBM	International Business Machines Corporation
		IEEE	Institute of Electrical and Electronics Engineers

IETF	Internet Engineering Task Force	MB	MegaBytes
IISP	Interim Inter-Switch Signaling Protocol. (P-NNI phase 0)	Mbps	Megabits per second
ILMI	Interim Local Management Interface	MIB	Management Information Base
INARP	Inverse Address Resolution Protocol	MPOA	MultiProtocol Over ATM
IP	Internet Protocol	MPM	Management Process and Control
I-PNNI	Integrated PNNI	MSS	Multiprotocol Switched Services
IPX	Internetwork Packet eXchange	MTU	Maximum Transmission unit
IRQ	Interrupt Request	NBBS	Networking BroadBand Services
ISA	Industry Standard Architecture	NBMA	NonBroadcast Multiaccess Network
ISDN	Integrated Services Digital Network	NDIS	Network Driver Interface Specification
ISO	International Organization for Standardization	NDPS	NonDisruptive Path Switch
ITSO	International Technical Support Organization	NetBIOS	Network Basic Input/Output System
ITU-T	International Telecommunication Union - Telecommunication	NHRP	Next Hop Resolution Protocol
KB	kilobyte	NHS	Next Hop server
Kbps	kilobits per second	NIC	Network Information Center
LAA	Locally Administered Address	NIX	Network Information Exchange
LAN	Local Area Network	NMS	Network Management Station
LANE	LAN Emulation (ATM Forum)	NNI	Network-to-Network Interface
LE	LAN Emulation (also, LANE)	nrt-VBR	Non-real-Time Variable Bit Rate
LEC	LAN Emulation Client (ATM Forum LANE)	NSAP	Network Service Access Point
LECS	LAN Emulation Configuration Server	NRB	Non Reserved Bandwidth
LES	LAN Emulation Server	OAM	Operations Administration and Maintenance
LIS	Logical IP Subnetwork	OC-n	Optical Carrier level n
LLC	Logical Link Control	ODI	Open Data-link Interface
LNNI	LAN emulation Network Node Interface	OID	Originator IDentifier
LPDU	Logical link Control Protocol Data Unit	OSI	Open Systems Interconnection
LSU	Link State Update	OSPF	Open Shortest Path First
LUNI	LAN emulation User-to-Network Interface	PAR	PNNI Augmented Routing
MAC	Medium Access Control	PC	Personal Computer
MAT	Management Application Transporter	PCR	Peak Cell Rate
MARS	Multicast Address Resolution Server	PCI	Peripheral Component Interconnect
		PCM	Pulse Code Modulation
		PDH	Plesiochronous Digital Hierarchy
		PDU	Protocol Data Unit

PG	Peer Group	SONET	Synchronous Optical Network
PGI	Peer Group Identifier	SR-TB	Source Route Translational Bridge
PGL	Peer Group Leader	SRB	Single Route Broadcast / Source-Route Bridging
PIM	Product Independent Module	SRM	Source Route Manager
PNNI	Private Network-to-Network Interface	SRF	Specifically Routed Frame
PSM	Product Specific Module	SRT	Source Route Transparent bridging
PT	Payload Type	SSAP	Source Service Access Point
PTSE	PNNI Topology State Element	SSCOP	Service-Specific Connection-Oriented Protocol
PTSP	PNNI Topology State Packet	SSCS	Service-Specific Convergence Sublayer
PVC	Permanent Virtual Circuit	SSI	Switch to Switch Interface
PVP	Permanent Virtual Path	STE	Spanning Tree Explorer
QoS	Quality of Service	STM	Synchronous Transfer Mode
RAIG	Resource Availability Information Group	STP	Shielded Twisted Pair / Spanning Tree Protocol
RB	Reserved Bandwidth	SVC	Switched Virtual Circuit
RCC	Routing Control Channel	SVN	Switched Virtual Networking
RIF	Route Information Field	TA	Terminal Adapter
RIP	Routing Information Protocol	TAXI	Transparent Asynchronous transmitter-receiver Interface
RISC	Reduced Instruction Set Computer/cycles	TCP	Transmission Control Protocol
RMON	Remote Monitor	TCP/IP	Transmission Control Protocol/Internet Protocol
rt-VBR	Real-Time Variable Bit rate	TDM	Time Division Multiplexing
RTP	Rapid-Transport Protocol	TE	Terminal Equipment
SA	Source Address	TFTP	Trivial File Transfer Protocol
SAAL	Signaling ATM Adaptation Layer	TRS	Topology and Route Selection
SAAT	Source Address Association Table	TP	Twisted Pair (Wiring)
SAP	Service Access Point	TTRT	Target Token Rotation Time
SAR	Segmentation And Reassembly	UAA	Universally Administered Address
SDEL	Start Delimiter	UBR	Unspecified Bit Rate
SDH	Synchronous Digital Hierarchy	UDP	User Datagram Protocol
SDLC	Synchronous Data Link Control	UFC	Universal Feature Card
SDU	Service Data Unit	ULEC	Unknown LAN Emulation Client
SEAL	Simple and Efficiency Adaptation Layer (AAL5)	UME	UNI Management Entity
SFE	Specific Front End	UNI	User-to-Network Interface
SNA	Systems Network Architecture	UTP	Unshielded Twisted Pair
SNAP	Subnetwork Access Protocol	VBR	Variable Bit Rate
SNMP	Simple Network Management Protocol		

VC	Virtual Channel (ATM) Virtual Connection (Frame Relay) Virtual Circuit (X.25)	VPCI	Virtual Path Connection Identifier
VCC	Virtual Circuit Connection (X.25 and ATM)	VPI	Virtual Path Identifier
VCI	Virtual Channel Identifier	VPL	Virtual path Link (UNI 3.0)
VCL	Virtual Channel Link (UNI 3.0)	VS/VD	Virtual Source/Virtual Destination
VLAN	Virtual LAN	VSS	Viewing ATM Service Statistication
VP	Virtual Path	WAN	Wide Area Network
VPC	Virtual Path Connection	XIWT	cross industry Working Team

Index

Numerics

- 100VG-anyLAN 16
- 8260
 - commands 335
 - LECS address 311, 346
- 8281
 - configuration 442
 - status 436, 455

A

- A-CPSW 467
 - commands 96, 335
- AAL 477
- AAL-0 66
- AAL-1 66
- AAL-2 67
- AAL-3/4 68
- AAL-3/4 SAR and CPCS PDUs 68
- AAL-5 70
- AAL-5 frames (user plane) 72
- AAL-5 SAR and CPCS PDUs 70
- abbreviations 523
- acronyms 523
- adaptive cut-through 246
 - add party 159
- add party failure
- Address authorities 60
- Address formats in ATM 59
- address table 245, 250, 252, 271
- Addresses in PNNI peer groups 60
- Addresses in SSI clusters 60
- aging time 245, 250, 252, 267, 270
 - analyzer 103, 314, 398
- architecture problem 420
- ARP 35
- ATM
 - address authorities 60
 - address formats 59
 - addressing in SSI clusters 46
 - cell format 56
 - characteristics 47
 - commands 104
 - concepts 43
 - configuration sheet 92
 - congestion control 50
 - connection setup 115
 - connection tear-down 116
 - connections 52
 - empty cells 57
 - flow control 50
 - icons 91
 - idle cells 57
 - initial registration 115

ATM (continued)

- layer model 61
- MIB 77
- multicast trees 48
- network structures 43, 46
- networking models 78
- performance 85
- performance optimization 88
- physical interfaces in 62
- quality of service (QoS) 49
- routing of cells 54
- rules 111
- service classes 64
- signalling 58, 115
- switching of cells 54
- tools 95
- user interface 105
- ATM adaptation layer (AAL) 64
- ATM address authorities 60
- ATM connections 52
- ATM layer 64
- ATM layer model 61
- ATM module
 - leds 95, 96
- ATM routing loop
 - reachability 228
- Authorities for ATM addresses 60

B

- bridging 27
 - data transfer 418
 - healthy network 417
 - PD guidelines 418
 - PD methodology 423
 - problem 420
- broadcast 146
- BUS 80

C

- Call establishment and clearing 73
- cause codes
 - information element (IE) 76
 - provided by network management 151
 - value=01 170
 - value=03 162, 173
 - Value=100 160
- Cell discard 50, 57
- Cell format 56
- Cell loss priority (CLP) 57
- Cell routing 54
- Cell switching 54
- classical IP (RFC 1577)
 - address resolution 293

classical IP (RFC 1577) *(continued)*

- ATMARP 81
- ATMARP cache 386
- ATMARP client 320, 321
- ATMARP client (IBM Nways 8274 LAN RouteSwitch) 389
- ATMARP client (RS/6000) 383
- ATMARP server 304, 320, 321, 384
- broadcasts 294
- concepts 80, 294, 303
- configuration 381, 388
- diagnosis methodology 320, 325
- healthy network guidelines 293
- performance optimization 88
- problem determination guidelines 306
- TCP/IP parameter configuration 88
- VCCs 324
- collision domain 10
- Common part convergence sublayer (CPCS) 72
- concepts 78
- Congestion control 50
- Connect and disconnect procedures 73
- Connection Acceptance Control (CAC) 50
- Connection setup and tear-down 73
- Connectivity
- constant bit rate (CBR)
- control point
- CPCS 72
- cross-over cable 239
- cut-through 15, 246

D

- designated transit list (DTL)
 - path selection 203
- duplicate ATM addresses
 - redundancy 214

E

- E-1 frame (2.048 Mbps) 62
- E-3 frame (34.368 Mbps) 62
- ELAN 78
- Empty cells 57
- Emulated LAN 78, 295, 342, 353, 360, 393, 408
- End System Identifier (ESI) 60
- ESI 60
- Ethernet 9, 293
- Ethernet V2 (DIX)
- etherpipes 239

F

- Fast Ethernet 15
- FDDI 22
- Flow control 50
- frame encapsulation 447
- frame format
 - AAL-3/4 SAR and CPCS PDUs 68

frame format *(continued)*

- AAL-5 frames (user plane) 72
- AAL-5 SAR and CPCS PDUs 70
- BPDU with STP 29
- E-1 frame (2.048 Mbps) 62
- E-3 frame (34.368 Mbps) 62
- Ethernet IEEE 802.3 13
- Ethernet v2 11
- FDDI 25
- LANE frames 72
- NetWare's IPX/SPX 39
- Q.2931 75
- RFC 1483 frames 72
- RFC 1577 frames (classical IP) 72
- STM-1 frame (155.52 Mbps) 62
- STS-1 frame (51.84 Mbps) 62
- STS-3c frame (155.52 Mbps) 62
- token-ring IEEE 802.5 18
- frame translation 32, 250

G

- Generic flow control (GFC) 57, 76
- GFC 57, 76
- gigabit Ethernet 16
- group 241, 258

H

- Header error check (HEC) 57
- HEC 57

I

- IBM 8281 426
- ICMP 35
- Idle cells 57
- IISP 46, 166, 182, 204, 210
- ILMI 77
 - tracing 123
- Information Element (IE) 75
- interim interswitch protocol (IISP)
- Interim local management interface (ILMI) 77
- IPX 37, 432

L

- Label swapping 52, 54
- LAN emulation (LANE) 78
- LAN switching
 - broadcast 237
 - connection 238
 - diagnosis methodology 252
 - Ethernet 15
 - problem determination guidelines 247
 - token-ring 21
- LANE
 - address resolution 293
 - ATM Forum-Compliant VCCs 295, 297, 300

LANE (continued)

- broadcasts 294
- BUS 80, 295, 312, 313, 315, 327, 360, 374, 406, 408
- concepts 294, 295
- configuration 346
- data direct VCC 395
- data transfer 299, 393
- diagnosis methodology 310, 325
- frame format 72
- healthy network guidelines 293
- hints and tips 335
- IBM compliant 300
- IBM compliant VCCs 300
- IBM-compliant 329
- initialization 297
- LANE 2.0 294
- LE client 295, 312, 353
- LE_ARP cache 393
- LEC 80
- LECS 79, 295, 298, 301, 311, 313, 314, 327, 342, 372
- LECS well-known address (WKA) 298
- LECS well-known VCC (WK VCC) 298
- LES 80, 295, 312, 313, 315, 327, 360, 374, 408
- problem determination guidelines 306
- proxy LE client 295, 312
- redundancy 327, 408
- VCCs 313

Layer model for ATM 61

LEC 80

LECS 79

LES 80

logical IP subnet 381, 388

M

MIB 40, 493

microcode 370

MSS 293, 332

- bridging 456, 458
- commands 335, 428
- features 301
- process 438
- redundancy 302
- routing 428, 439, 440
- status 343

multi-protocol over ATM (MPOA)

Multicast trees 48, 80

multiprotocol over ATM (MPOA) 294

N

netBIOS 34, 426

network analyzer 314, 398

network monitor

Next Hop Routing Protocol (NHRP) 83

next-hop resolution protocol (NHRP) 302

NHRP client (NHC) 84

NHRP server (NHS) 84

NNI 43, 46

NNI cell format 56

O

OAM 57

Operation and management (OAM) cells 57

P

path selection for PNNI

- shortest path 209
- widest path 207

path selection pnni 205

Payload type indicator (PTI) 57

peak cell rate (PCR) 314, 315

peer group

performance 85, 326, 406, 419

Performance optimization 88

permanent virtual connection (PVC)

- classical IP 305

phase 133

phases

Physical interfaces 62

Physical layer 62

PNNI 46, 58, 183, 230

- migration To 214
- modification To 161

PNNI ATM addresses 60

PNNI group identifier 46

PNNI level ID 46

PNNI peer groups 46

point-to-point-bridging (PToP) 241

Policing 50

PTI 57

Public UNI 47

PVC 320, 321, 380

Q

Q.2110 72

Q.2130 72

Q.2931 72, 75

QoS 49

Quality of service (QoS) 49

R

reachable address

- PNNI 161, 230
- PNNI connected to this UNI port is down, the address 178
- VPC link 195

Reserved VCI values 56

RMON 41, 104

route selection (ATM)

- traces 198

- routeswitch
 - auto-tracker VLAN aging time 245, 267, 277
 - CAM 243
 - commands 255, 335, 428
 - configuration 354
 - hints and tips 254
 - LED 254, 258
 - login console 254
 - mirroring 283
 - MPM 243, 254
 - NSM 254
 - redundant MPM 260
 - route switch manager 258
 - routing 428
 - services 459
 - status 344
 - submenu 255
 - terminal emulation 254
- routing 36
 - data transfer 418
 - healthy network 417
 - PD guidelines 418
 - PD methodology 423
 - problem 421
- Routing of cells 54

S

- SAAL 72, 73, 75
- SAR 72
- SDH 62
- SEAL 70
- Server
- Service classes 64
- Service specific connection oriented protocol (SSCOP) 72
- Service specific convergence sublayer (SSCS) 72
- Service specific coordination function (SSCF) 72
- Signaling
 - AAL (SAAL) 73
 - control plane 73, 75
 - frames 75
 - ILMI 77
 - management plane 76, 77
 - Q.2931 75
 - UNI 75, 76, 77
- Signalling AAL (SAAL) 73
- SMT 41
- SNMP 40
- SONET 62
- SONET lite 63
- source-route bridging 31
- source-route switching 33
- source-route translational bridging 32
- source-route transparent bridging 32
- spanning tree protocol 28, 32, 450
 - MSS 428, 456, 458
 - routeswitch 267, 428, 459

- specifications
- SSCF 72
- SSCOP 72
- SSCS 72
- SSI 46
- SSI ATM addresses 60
- SSI clusters 46
- STM-1 frame (155.52 Mbps) 62
- store and forward 15, 246
- straight through cable 239
- STS-1 frame (51.84 Mbps) 62
- STS-3c frame (155.52 Mbps) 62
- SVC 321, 380
- switched Ethernet
- switched virtual connection (SVC)
 - classical IP 304
- Switching of cells 54
- Synchronous Digital Hierarchy (SDH) 62
- Synchronous Optical Network (SONET) 62

T

- TCP/IP 34
- token-Ring 17, 250, 293
- tokenpipes 239
 - topology and routing services 116
- traces
- transparent bridging 27, 450
- tree 425
- trunking 240

U

- UNI 44, 46, 157, 166
- UNI cell format 56

V

- Virtual channel (VC) 53
- Virtual channel connection (VCC) 53
- Virtual channel indicator (VCI) 53, 54, 57
- Virtual Channel Indicator (VCI), reserved values 56
- Virtual channel link (VCL) 54
- Virtual channel switching 54
- Virtual connections 53
- virtual LAN (VLAN)
 - aging time 245
 - assignment 243, 244, 249, 251, 265, 277
 - broadcast 244
 - leakage 250, 252, 283
 - policy rules 242
 - port rule 267
 - processing 244
 - protocol rule 286
- Virtual path (VP) 53
- Virtual path connection (VPC) 53
- Virtual path connection identifier (VPCI) 53
- Virtual path indicator (VPI) 53, 54, 57

Virtual path link (VPL) 53
Virtual path switching 55
VOID 173, 196, 207
VOID port
VPC 173, 195, 207
VPC link
VPI/VCI swapping 52, 54

W

Web networking informations
well-known address (LECS)
 call to WKA 124, 221
WKA 163, 181

ITSO Redbook Evaluation

Troubleshooting IBM LAN/ATM Campus Networks &titleline2.
SG24-2105-00

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.com>
- Fax this form to: USA International Access Code + 1 914 432 8264
- Send your comments in an Internet note to redbook@vnet.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!)**



This soft copy for use by IBM employees only.

Printed in U.S.A.

SG24-2105-00

