SG24-4755-00

Local Area Network Concepts and Products: Routers and Gateways

May 1996

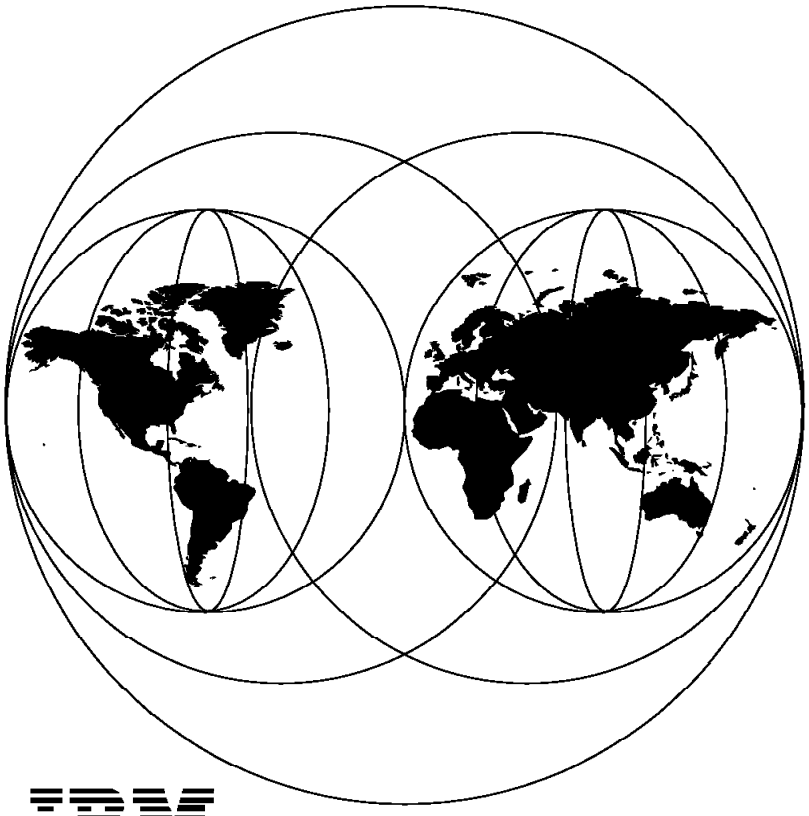

IBM

International Technical Support Organization
Raleigh Center

International Technical Support Organization

SG24-4755-00

Local Area Network Concepts and Products: Routers and Gateways

May 1996

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix A, “Special Notices” on page 265.

First Edition (May 1996)

This edition applies to the most recent IBM LAN products and LAN architectures.

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

Contents

Preface vii
How This Redbook Is Organized vii
The Team That Wrote This Redbook viii
Comments Welcome ix

Chapter 1. LAN Interconnection 1
1.1 LAN End Nodes 1
1.2 Intermediate Nodes 1
1.3 IBM LAN Bridge Programs 2
1.3.1 Bridge Product Positioning 3
1.3.2 Highlights 4
1.3.3 Systems Management 4
1.4 IBM Local Token-Ring Bridge/DOS Version 1.0 4
1.5 IBM Remote Token-Ring Bridge/DOS Version 1.0 4
1.6 IBM LANStreamer Token-Ring Bridge/DOS Version 1.0 5
1.7 IBM LAN Bridge Manager/2 Version 1.0 5
1.8 IBM Frame Relay Token-Ring Bridge/DOS Version 1.0 6
1.9 IBM LAN-to-LAN WAN Program (LTLW) 8
1.9.1 Many-to-One or One-to-Many 10
1.9.2 LU 6.2 Sessions 10
1.9.3 LTLW V1.07 and ELTLW V1.01 11
1.9.4 LTLW Summary 16
1.10 RouteXpander/2 17
1.10.1 Highlights 17
1.10.2 Technical Description 18
1.10.3 Network Management 19
1.10.4 RouteXpander/2 Interoperability 19
1.10.5 Compatibility 19
1.10.6 RouteXpander/2 (RXR/2) Features 20
1.11 IBM AnyNet Product Family 21
1.11.1 IBM AnyNet 21
1.11.2 Customer Requirements 22
1.11.3 Multiprotocol Transport Networking (MPTN) 23
1.11.4 NetBIOS 24
1.11.5 NetBEUI 24
1.11.6 Positioning AnyNet/2 NetBEUI over SNA with LTLW 24
1.11.7 AnyNet/2 Version 2.0.2 27
1.11.8 AnyNet/6000 28
1.11.9 AnyNet/400 29
1.11.10 VTAM Version 4 Release 2 AnyNet Feature 29
1.11.11 AnyNet/2 Sockets Over SNA Gateway Version 1.1.6 30
1.11.12 AnyNet SNA over TCP/IP Gateway for OS/2 31
1.11.13 AnyNet IPX over SNA Gateway for OS/2 32
1.11.14 AnyNet APPC over TCP/IP for Windows 32
1.11.15 AnyNet Advantages 33

Chapter 2. Bridges and Routers 37
2.1 Bridges 37
2.1.1 Simple Bridges 38
2.1.2 Complex Bridges 38
2.1.3 Local Bridges 39

2.1.4 Remote Bridges . . . 39
2.1.5 Bridging Methods . . . 40
2.2 Routers . . . 47
2.2.1 Packet . . . 48
2.2.2 Routing Table Maintenance Protocols . . . 49
2.3 Performance in Routers . . . 50
2.3.1 Performance in a Router Network . . . 51
2.3.2 Relationship Between Response Time, Throughput, and Utilization . . . 54
2.3.3 Performance Metrics . . . 55
2.4 To Route or to Bridge? . . . 59
2.4.1 Router Connections . . . 59
2.4.2 Bridge Connections . . . 59
2.5 IBM 8229 LAN Bridge . . . 61
2.5.1 Identifying the 8229 Front Panel . . . 61
2.5.2 Identifying the Attachment Module . . . 62
2.5.3 Operation in General . . . 66
2.5.4 IBM 8229 Utility Program . . . 68
2.5.5 Details of IBM 8229 Token-Ring to Ethernet Use . . . 69
2.5.6 IBM 8229 Token-Ring to WAN Connectivity . . . 71
2.5.7 Frame Format and Address Conversion . . . 74
2.5.8 8229 Summary . . . 80
2.6 IBM 2210 Nways Multiprotocol Router . . . 84
2.6.1 Models of the IBM 2210 . . . 84
2.6.2 Indicators on the IBM 2210 . . . 86
2.6.3 The Reset Button on the IBM 2210 . . . 87
2.6.4 Networks Supported by the IBM 2210 . . . 87
2.6.5 Accessing the IBM 2210 . . . 88
2.6.6 Software Package . . . 88
2.6.7 MRNS Overview . . . 89
2.6.8 The IBM 2210 As an IP Router . . . 98
2.6.9 Data Link Switching . . . 109
2.6.10 Features and Facilities . . . 112
2.7 IBM 6611 Router . . . 119
2.7.1 Hardware Overview . . . 120
2.7.2 Multiprotocol Connectivity . . . 125
2.7.3 Bridging with IBM 6611 . . . 127
2.7.4 Data Link Switching . . . 140
2.7.5 IBM 6611 Network Processor Enhancements - Release 4 . . . 149

Chapter 3. LAN Gateways . . . 151
3.1 S/390 Open Systems Adapter . . . 152
3.2 3172 Gateway . . . 159
3.2.1 3172 LAN-To-Host Mode . . . 159
3.2.2 3172 WAN-to-Host Mode . . . 165
3.3 3174 Gateway . . . 168
3.3.1 SNA Environment . . . 168
3.3.2 TCP/IP Environment . . . 175
3.4 3745/3746-9x0 Gateway . . . 181
3.4.1 3745/3746-9x0 Models . . . 181
3.4.2 IBM 3746-9x0 Connectivity . . . 183
3.4.3 3745 SNA Environment . . . 184
3.4.4 3745 TCP/IP Environment . . . 186
3.4.5 3746-9x0 SNA/APPN Environment . . . 188
3.4.6 3746-9x0 TCP/IP Environment . . . 190
3.5 RISC System/6000 . . . 192

3.5.1 AIX TCP/IP . . . 192
3.5.2 AIX SNA Server/6000 . . . 193
3.5.3 AIX SNA Gateway/6000 . . . 195
3.5.4 SNA Application Access . . . 196
3.5.5 SNA Client Access for AIX . . . 199
3.5.6 Summary . . . 200
3.6 2217 Nways Multiprotocol Concentrator . . . 200
3.6.1 Benefits . . . 201
3.6.2 Release 2 Enhancements . . . 202
3.6.3 Hardware Description . . . 203
3.6.4 SNA Support . . . 203
3.6.5 LAN Functionality . . . 207
3.7 PC Gateways . . . 210
3.8 PC/3270 Gateway . . . 213
3.8.1 Gateway Status Utility . . . 214
3.9 IBM Communications Server for OS/2 Warp, Version 4.0 (CommServer) 215
3.9.1 SNA Gateway . . . 216
3.9.2 Advanced Peer-to-Peer Networking (APPN) . . . 217
3.9.3 Multiprotocol Support . . . 219
3.9.4 Emulator Support . . . 220

Chapter 4. Remote LAN Access . . . 221
4.1 Remote LAN Access Environments . . . 221
4.1.1 Remote-to-Remote . . . 221
4.1.2 Remote-to-LAN . . . 222
4.1.3 LAN-to-Remote . . . 222
4.1.4 LAN-to-LAN . . . 223
4.2 Remote LAN Access Technologies . . . 224
4.2.1 Remote Control Approach . . . 224
4.2.2 Remote Client Approach . . . 225
4.2.3 Remote Node Approach . . . 226
4.3 IBM LAN Distance . . . 226
4.3.1 Technology . . . 227
4.3.2 Components and Environments . . . 228
4.3.3 Connections, Interfaces and Applications . . . 231
4.3.4 Key Features . . . 232
4.4 IBM 8235 Dial-In Access to LAN Servers . . . 233
4.4.1 8235 System Components . . . 234
4.4.2 Dial-In Access to LAN Servers (DIALs) Client Software . . . 234
4.4.3 IBM 8235 New Features . . . 235
4.4.4 What is Virtual Connection? . . . 239
4.4.5 What is Channel Aggregation? . . . 239
4.4.6 Management Facility . . . 240
4.4.7 8235 Hardware . . . 240
4.4.8 Models Summary . . . 246
4.4.9 Communication Options . . . 246
4.4.10 Supported Protocols . . . 247
4.4.11 Security . . . 255
4.4.12 The Activity Logger . . . 258
4.5 Distributed Console Access Facility (DCAF) . . . 259
4.5.1 Product Positioning . . . 259
4.5.2 Highlights . . . 259
4.5.3 Description . . . 259
4.6 LANHOP/6000 . . . 261
4.6.1 Highlights . . . 261

4.6.2 Description 262
4.6.3 Technical Description 263

Appendix A. Special Notices 265

Appendix B. Related Publications 269
B.1 International Technical Support Organization Publications 269
B.2 Other Publications 269

How To Get ITSO Redbooks 271
How IBM Employees Can Get ITSO Redbooks 271
How Customers Can Get ITSO Redbooks 272
IBM Redbook Order Form 273

Index 275

Preface

Local Area Network Concepts and Products is a set of four reference books for those looking for conceptual and product-specific information in the LAN environment. They provide a technical introduction to the various types of IBM local area network architectures and product capabilities. The four volumes are as follows:

- SG24-4753-00 - *LAN Architecture*
- SG24-4754-00 - *LAN Adapters, Hubs and ATM*
- SG24-4755-00 - *Routers and Gateways*
- SG24-4756-00 - *LAN Operating Systems and Management*

To obtain all four books, order the set SK2T-1306.

These redbooks complement the reference material available for the products discussed. Much of the information detailed in these books is available through current redbooks and IBM sales and reference manuals. It is therefore assumed that the reader will refer to these sources for more in-depth information if required.

These documents are intended for customers, IBM technical professionals, services specialists, marketing specialists, and marketing representatives working in networking and in particular the local area network environments. Details on installation and performance of particular products will not be included in these books, as this information is available from other sources.

Some knowledge of local area networks, as well as an awareness of the rapidly changing intelligent workstation environment, is assumed.

How This Redbook Is Organized

The redbook is organized as follows:

- Chapter 1, “LAN Interconnection”

  This chapter describes LAN interconnection techniques and products available to provide this function.

- Chapter 2, “Bridges and Routers”

  This chapter describes bridges and routers and their functions in the LAN environment.

- Chapter 3, “LAN Gateways”

  The chapter describes LAN gateways and the methods and functions of each type.

- Chapter 4, ″Remote LAN Access″

  This chapter describes remote LAN access technology and products including the 8235 and the IBM LAN Distance product.

The Team That Wrote This Redbook

This redbook was produced by a team of specialists from around the world working at the International Technical Support Organization Raleigh Center.

The advisors for this project were:

Ricardo Haragutchi
International Technical Support Organization, Raleigh

John Parker
International Technical Support Organization, Raleigh

The authors of this document were:

Edmilson Barbosa
IBM Brazil

Ingvar Hyleborg
IBM Sweden

Jefferson da Silva
IBM/GSI Brazil

Klaus Wichmann
ITSO Raleigh

Marcello Belloni Gomes
IBM Brazil

Thanks to the following people for their invaluable advice and guidance provided in the production of this document:

Toshi Shimizu
International Technical Support Organization, Austin

Aroldo Yai
Barry Nusbaum
Donna Fox
Fergus Stewart
Jose Boo
Juan Rodriguez
Mark DeCain
Mohammad Shabani
Robert Macgregor
Stephen Breese
Volkert Kreuk
International Technical Support Organization, Raleigh

Alan Millard
Arthur Bond
Bert Wendle
Carol Carson
Dean Stockwell
Erik Dixon
H. Parrish

Paul Carter
IBM Research Triangle Park, Raleigh NC.

Comments Welcome

We want our redbooks to be as helpful as possible. Should you have any comments about this or other redbooks, please send us a note at the following address:

redbook@vnet.ibm.com

Your comments are important to us!

Chapter 1. LAN Interconnection

This chapter discusses the products and methods used to interconnect LANs of various types. Local and remote bridging as well as routing techniques are discussed. For LAN Gateway functions please refer to Chapter 2 of *Local Area Network Concepts and Products: LAN Adapters, Hubs and ATM*, SG24-4754.

1.1 LAN End Nodes

LAN end nodes attach to a single LAN in an internetwork. End nodes are typically general purpose computing devices such as host computers, servers, and workstations.

End nodes do not provide interconnectivity between networks, nor do they participate in any exchange of information about the topology of an internetwork.

End nodes that only support physical and data link layer protocols may exchange data with other end nodes directly over a network, while those that also support a network layer protocol may use intermediate nodes to access end nodes on remote networks in an internetwork. In the latter case, the end nodes must be configured with internetwork topology information, or be capable of automatically receiving it from intermediate nodes.

1.2 Intermediate Nodes

Intermediate nodes connect two or more networks and allow information to be exchanged between them. Very often they exchange information about network topology or information that allows network reconfiguration in the event of failure. Intermediate nodes provide the *glue* that connects individual networks into an internetwork.

Four types of intermediate nodes are available. They are:

- **Repeaters** that electrically regenerate, retime and forward all packets between the networks to which they are attached.
- **Bridges** that selectively forward data between networks, based on the MAC sublayer destination address or control information in each packet.
- **Routers** that selectively forward data between networks, based on the network layer destination address in each packet.
- **Gateways** that selectively forward packets between different network environments, for example, systems network architecture (SNA) and open systems interconnect (OSI).

Each intermediate node implements different levels of the OSI Reference Model (see Figure 1 on page 2).

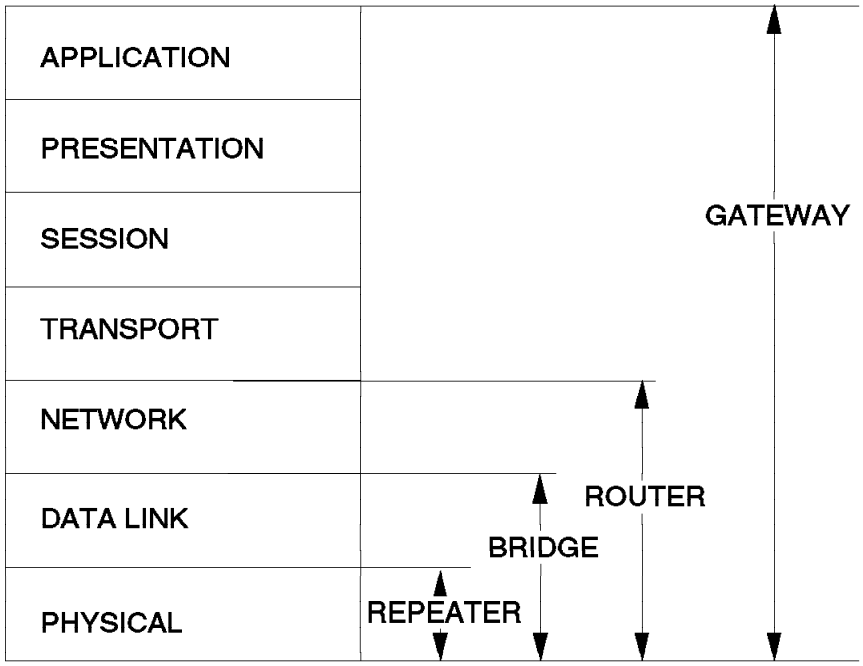

Figure 1. The OSI Model - Repeaters, Bridges, Routers and Gateways

Repeaters implement physical layer standards only. They are transparent to protocols above the physical layer and act as physical layer relays.

Bridges implement physical and data link layer standards. They are transparent to protocols above the data link layer and act as data link layer relays.

Routers implement physical, data link and network layer standards. They are transparent to protocols above the network layer and act as network layer relays.

Gateways are considered application layer relays between network environments. They must implement all seven layers of the OSI Reference Model.

While all four nodes act as intermediate nodes, normally only bridges and routers are referred to as intermediate nodes. Indeed many would disagree with this definition and use the term intermediate node to refer only to routers.

1.3 IBM LAN Bridge Programs

IBM has announced the following five new bridge programs:

- IBM LANStreamer Bridge/DOS Version 1.0
- IBM LAN Bridge Manager/2 Version 1.0
- IBM Local Token-Ring Bridge/DOS Version 1.0
- IBM Remote Token-Ring Bridge/DOS Version 1.0
- IBM Frame Relay Token-Ring Bridge/DOS Version 1.0

The IBM LANStreamer Token-Ring Bridge/DOS Version 1.0 provides high performance throughput when combined with the LANStreamer 32-bit adapter in a local token-ring environment. It also allows the hop count to be increased from the previous limit of seven to a maximum of thirteen.

The IBM LAN Bridge Manager/2 Version 1.0 enables distributed installation, setup and management of IBM Token-Ring bridge program products (IBM 8209 excluded). The IBM LAN Bridge Manager/2 Version 1.0 implements a Bridge Manager station which manages a client/server relationship with existing server and bridge clients. Bridges and their resources can now be managed remotely. Existing network management functions between LAN Network Manager Version 1.1 and IBM bridges are not affected.

The IBM Local Token-Ring Bridge/DOS Version 1.0 replaces the current local bridge program product. This new local bridge provides identical function to the local bridge provided in the IBM Token-Ring Network Bridge Program Version 2.2. The new, lower priced IBM Local Token-Ring Bridge/DOS Version 1.0 provides an efficient, cost effective local bridging solution for connecting and segmenting token-ring networks.

The IBM Remote Token-Ring Bridge/DOS Version 1.0 replaces the current remote bridge program product. A single license now supports both halves of a remote bridge. Remote dial function is included.

This new remote bridge provides enhancements to the remote bridge provided in the IBM Token-Ring Network Bridge Program Version 2.2, including:

- Support for full T1 (1.544 Mbps) or full E1 (2.048 Mbps) line speeds when using the High-Speed Communications Co-Processor/2 Adapter.
- Communication adapter transmit buffer depth is increased from 64 KB to 256 KB. Remote bridge performance is improved through enhanced buffering of bursty traffic.
- Two new bridge filters.

IBM Remote Token-Ring Bridge/DOS Version 1.0 is equivalent to PTF 37463 and IBM Remote Bridge Program 2.2. The new, lower priced IBM Remote Token-Ring Bridge/DOS Version 1.0 provides an efficient, cost effective remote bridging solution for connecting remote token-ring networks.

IBM Frame Relay Token-Ring Bridge/DOS Version 1.0 provides WAN communications to support multiple remote customer locations with a wide range of bandwiths. It provides bridge filters and fully supports LAN Network Manager bridge management functions.

1.3.1 Bridge Product Positioning

Separate local and remote bridge products offer price advantages over the Token-Ring Network Bridge Program Version 2.2. A single IBM Remote Token-Ring Bridge/DOS Version 1.0 program license now installs the primary and secondary halves of the remote bridge pair. If high-speed performance is not required for some LAN segments, then these products are very competitive in the token-ring environment.

If high-speed performance is the major requirement, the LANStreamer Token-Ring Bridge, together with the LANStreamer 32-bit adapter, provide the highest throughput performance attainable in a token-ring environment.

The Bridge Manager saves the duplication of skills that are required to manage and maintain bridges in various locations. It provides the mechanism to distribute code and run the token-ring bridge applications from a central location.

1.3.2 Highlights

- The IBM Local Token-Ring Bridge/DOS Version 1.0 provides an efficient, cost effective local bridging solution for connecting and segmenting token-ring networks.
- The IBM Remote Token-Ring Bridge/DOS Version 1.0 includes both halves and provides enhancements for full T1 line speed support and improved performance through enhanced buffering of bursty traffic.
- The IBM LANStreamer Token-Ring Bridge/DOS Version 1.0 provides high-speed performance when combined with the LANStreamer 32-bit adapter and allows increased hop counts.
- The IBM LAN Bridge Manager/2 provides two major functions, centralized software distribution and bridge management.
- The IBM Frame Relay Token-Ring Bridge/DOS Version 1.0 provides low cost, affordable WAN communication for telecommunication lines and bridge hardware.

1.3.3 Systems Management

The IBM LAN Bridge Manager/2 Version 1.0 allows remote, centralized management of IBM Token-Ring bridges across multisegmented networks. The Bridge Manager is designed to allow LAN administrators to be more productive by automating and centralizing the tasks associated with setting up, installing, modifying and managing the bridge application. All of IBM′s Token-Ring bridge products are fully supported by LAN Network Manager. The Bridge Manager further enhances the networking capability of the bridges running in a PS/2 by providing faster response and increased functionality in a token-ring environment.

1.4 IBM Local Token-Ring Bridge/DOS Version 1.0

The local bridge is a source-routing bridge that connects two adjacent token-rings operating at 4 or 16 Mbps. It supports communications with up to four IBM LAN Network Manager programs. The IBM LAN Network Manager collects information, such as network performance data, and forwards alerts to an IBM NetView host. All other functions supported in the IBM Token-Ring Network Bridge Program Version 2.2 are included in this program.

1.5 IBM Remote Token-Ring Bridge/DOS Version 1.0

The remote bridge is a source-routing bridge which connects distant token-rings via dedicated or multiplexed data communication links with speeds from 9.6 Kbps to 2.048 Mbps, including T1 speeds of 1.544 Mbps. Each remote bridge interfaces to a WAN link via a synchronous modem, DSU/CSU, T1 multiplexer or statistical multiplexer. This program contains both the primary and secondary halves of the remote bridge pair. The remote bridge program can communicate with up to four IBM LAN Network Manager programs, which collect information, such as network performance data and forwards alerts to an IBM NetView host.

The communications adapter transmit buffer depth has been increased from 64 KB to 256 KB, which improves performance through enhanced buffering of bursty traffic. All other functions, including the remote dial application, are the same as those in the IBM Token-Ring Network Bridge Program Version 2.2.

1.6 IBM LANStreamer Token-Ring Bridge/DOS Version 1.0

The LANStreamer bridge is a source-routing high-speed local bridge that connects adjacent token-ring LANs. With the support of the LANStreamer 32 MC adapter, the adapter bottleneck has been eliminated, and the bridge performance is determined by the speed of the PS/2 processor. A high end processor is able to approach media speed performance; however, it can also run on various other processors, and the small frame performance is at least two times greater than that currently achieved on the IBM Local Token-Ring Bridge/DOS Version 1.0 and the IBM Token-Ring Network Bridge Program Version 2.2. This bridge program also supports the full Route Information Field and is enabled for an increased hop count from the previous limit of seven to a maximum of thirteen. This allows addressing of a larger network, but to fully utilize this support, each end station and the LAN Network Manager must also have the increased hop count support. The LANStreamer bridge also provides a ring utilization counter. This function was previously obtained only by adding necessary additional equipment.

The following additional functions are included:

- Frame forwarding and adapter interface
- Configuration parameter end user interface
- Ring status display
- Configuration report server
- Ring parameter server
- Ring error monitor
- Accumulation and display of performance statistics
- Support for frame-forward filtering
- Ring utilization counter

In addition, all previous local network functions provided in Token-Ring Network Bridge Version 2.2 are supported with this program. The LAN Network Manager functions are enhanced with the support of the LAN Bridge Manager/2.

Note

The IBM LANStreamer Token-Ring Bridge/DOS Version 1.0 requires a PS/2 with a 32-bit architecture such as Models 70, 76, 77, 80, 90, and 95.

1.7 IBM LAN Bridge Manager/2 Version 1.0

The Bridge Manager provides remote, centralized management of token-ring bridges across multisegment networks. The implementation uses the concept of manager/agent, where the manager resides in any OS/2 or DOS workstation with access to an IBM OS/2 LAN Server or NetWare server, and the agent resides on the bridge machine. The program includes both the manager and the agent segments. The agent segment is also offered as a distributed feature so that additional licenses are purchased for a multiple bridge network that requires one manager in the server and one agent in each bridge computer. The agent

communicates with the LAN server for software management and distribution, while the network administrator interacts with the manager for software installation, server setup, configuration information and changes, and asset management. The following information is viewed from your workstation:

- DOS version installed
- Bridge program version installed
- Bridge network address
- Token-ring encoded adapter address
- Type of token-ring adapter installed
- Microcode level on the token-ring adapter

The following are the functions of the manager component:

- Manages multiple levels of the bridge program through alias definitions on the LAN server
- Manages and distributes bridge filters
- Reboots the machine remotely (time delay included)
- Displays the bridge hardware and software configuration, together with statistics obtained from the agent activity log
- Provides asset inventory information on the identification and location of each bridge machine
- Provides the capability to mass change the link password
- Provides a bridge installation and setup function for server setup

The following are the functions of the agent component:

- Simplifies the software installation, setup and customization of the bridge machines
- Manages the distribution of all software between the LAN server and the bridge machines
- Communicates with either an IBM OS/2 or Novell NetWare server
- Provides updated information on the machine hardware and software configuration
- Forwards alerts to LAN Network Manager, if installed, that communicates in an IBM LAN server or Novell server environment

1.8 IBM Frame Relay Token-Ring Bridge/DOS Version 1.0

Frame relay provides affordable WAN communications that are tailored to support multiple remote customer locations with a wide range of bandwidths.

With Frame Relay Token-Ring Bridge/DOS, customers can lower costs for the following:

- Telecommunications lines
- Bridge hardware
- Bridge software
- Network administration overhead

It provides source-route bridging and can connect either point-to-point over a leased line or across a frame relay network. It attaches locally to a token-ring LAN and bridges across the WAN communications link to a compatible bridge on a remote token-ring LAN.

The DOS Frame Relay Bridge supports the following four bridge filters and provides filter source code to assist customers with user-written filters.

- NetBIOS filter
- Address filter
- SAP Address filter
- SNAP filter

DOS Frame Relay Bridge fully supports LAN Network Manager bridge management functions (Bridge Server, RPS, CRS, and REM). It forwards alerts and performance statistics to a linked LAN Network Manager. Bridge configuration parameters are changed from the remote LAN Network Manager station.

In addition to IBM′s family of MCA and ISA Token-Ring Adapters, it supports both IBM′s ARTIC family of adapters and the WAC (Wide Area Connector) adapter.

Communications Adapters for PS/2 Micro Channel:

- X.25 Interface Co-Processor/2 - line speeds up to 1.544 Mbps
- High Speed Communication Co-Processor/2 - line speeds up to 2.048 Mbps
- WAC adapter for MCA - line speeds up to 2.048 Mbps

Communications Adapters for ISA Bus Value Point PCs:

- X.25 Interface Co-Processor - line speeds up to 1.544 Mbps
- Realtime Interface Co-Processor - line speeds up to 64 Kbps
- WAC adapter for ISA - line speeds up to 2.048 Mbps

Additional Software Required: RTIC Family Communications Adapters (X.25 ICA and HSCC) require driver support from IBM RTIC DOS Support Program, Version 1.03 or later. There is no additional software requirement for the Wide Area Connector adapter.

A device that attaches to the telecommunications line is also required, such as a DSU/CSU (IBM 5822), IDNX or a Multiplexor.

IBM′s Communications adapters each provide a choice of standard electrical interfaces, such as V.35, X.21, RS232, and RS442/449.

Over frame relay networks, DOS Frame Relay Bridge supports up to *five* concurrent DLCIs. This allows customers to connect multiple remote locations to one another for the most efficient and direct interoffice communications. This eliminates the need for dedicated leased lines between remote locations and the dedicated Remote Bridge hardware required for each line.

The Frame Relay Bridge interoperates with all of IBM′s bridges supporting RFC 1294 (multiprotocol routing and bridging over frame relay).

Other IBM products provide frame relay handler (FRH) function, including IBM 3172 Communications Controller, IBM 3174 (bridging) and the IBM 3745 (INN protocol transport over frame relay using ACF/NCP V6.2 or V7.1).

1.9 IBM LAN-to-LAN WAN Program (LTLW)

The IBM LAN-to-LAN Wide Area Network Program and LAN-to-LAN Wide Area Network Entry programs are software-based router products that are designed to link geographically separate LANs together to carry NetBIOS, IPX and IP traffic between them. When the stations attach to LANs in close geographical proximity, local bridges provide the most flexible approach to network design. Other protocols besides NetBIOS are able to share the common resource of the bridge. When the LANs are too far apart to make use of local bridges, the question arises: what should be done to provide interconnection? Many customers already have a wide area network (WAN) installed, based on either SNA or X.25 switches. A WAN is usually designed to be a shared resource and resulting cost savings may mean that meshing becomes an economic proposition. This approach has other significant benefits, such as providing multiple paths through the network and hence resilience to failure of individual components. Could the LANs be linked over the existing WAN?

A remote bridge is another option. However, it requires the installation of a dedicated telecommunications link between the LAN sites. This link cannot be shared with the SNA WAN facilities and to mesh the various LAN sites may require many remote bridges with consequent costs and management implications for the associated links. The IBM Token-Ring remote bridge implementation does not allow for the link to cross a packet switched data network (PSDN) accessed with the X.25 protocol, a common WAN implementation. The IBM LAN-to-LAN Wide Area Network Program provides an answer to the question by allowing multiprotocol communications between LANs to share an SNA or X.25 WAN. The IBM LAN-to-LAN Wide Area Network Program:

- Allows existing LAN-attached stations to communicate via a WAN.
- Is able to use an SNA backbone (either an SNA subarea or SNA Advanced Peer-to-Peer SNA APPN) network or an X.25 Packet Switched Data Network (PSDN) as the WAN.
- Is an application designed to run on a PC or Personal System/2 machine under OS/2 V2.0 or higher.
- Has been designed to use the OS/2 Extended Edition Communications Manager Advanced Peer-to-Peer Communication (APPC) LU 6.2 interface to send and receive traffic between NetBIOS switches.
- Supports one-to-many and many-to-one NetBIOS switch configurations.
- Provides a graphical and text-based operator interface for planning, configuring, installing, monitoring, error detection and recovery.

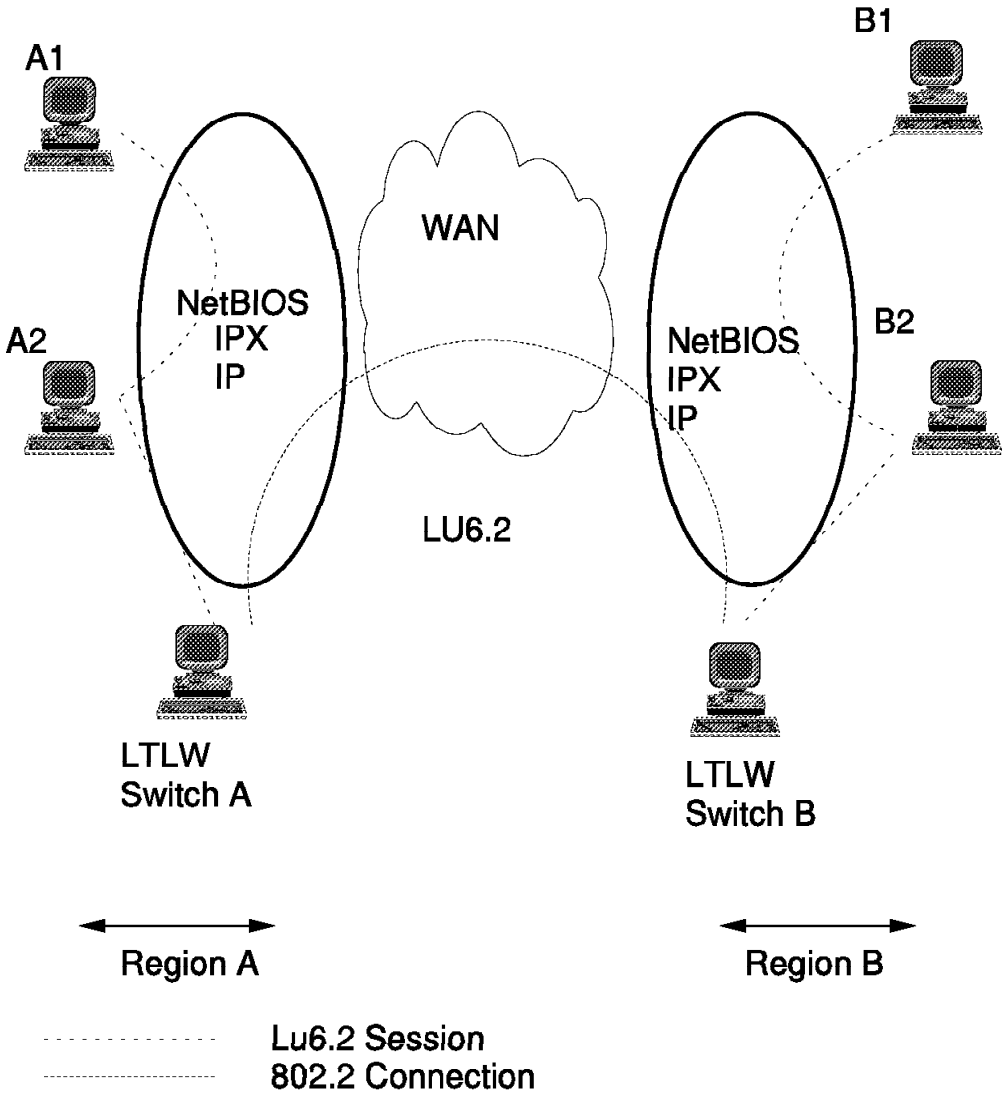

Figure 2. Connection of LANs via WAN Using IBM LTLW

In Figure 2, a NetBIOS application in station A1 is communicating with a partner in station B1. Station A1 might be a DOS LAN Requester accessing a file server in station B1. At the same time, station A2, perhaps a print server, is being accessed by station B2. Both connections are multiplexed onto an APPC LU 6.2 SNA session between the two NetBIOS switches. In fact, there are actually two LU 6.2 sessions and conversations between the switches, though only one is shown. One of the sessions is used for sending data between switch A and switch B, the other for data travelling in the opposite direction.

Figure 2 indicates the LANs, with their respective NetBIOS switches, can be termed a *region.* A region is a self-contained group of LAN segments connected by MAC level bridges. Within a region, station MAC addresses and segment numbers must be unique. When regions are connected with NetBIOS switches, segment numbers and MAC addresses may be duplicated between regions.

This is an important migrational benefit, because when regions are to be connected for the first time, a major renumbering task may not be required. If the regions were to be connected by a remote bridge, then uniqueness would be a requirement, and the regions would be merged into one.

The IBM LAN-to-LAN Wide Area Network Program supports multiple connectivity options. On the LAN side the following are supported:

- IBM Token-Ring
- IBM PC Network (Broadband and Baseband)
- Ethernet

On the WAN side, any communication method that is used by OS/2 Extended Edition 1.2 or higher for advanced program-to-program communication (APPC) is supported. This includes all three LAN types described earlier, as well as the following:

- Synchronous data link control (SDLC), using the standard IBM PC or PS/2 line driver cards. This link could run directly to another NetBIOS switch, to an SNA Communications Controller, to an IBM 9370 or to an IBM AS/400 processor.
- CCITT X.25 using the appropriate PC or PS/2 card and interfacing to a public or private PSDN.

1.9.1 Many-to-One or One-to-Many

LTLW is capable of having LU 6.2 sessions with many other switches. This means that when multiple regions are to be connected, only one copy of the IBM LAN-to-LAN Wide Area Network Program is needed in any region. However, if there are two switches in the same region, both cannot have sessions with one or more switches in another region. If they did, then a parallel path would exist between the two regions and this is not supported.

1.9.2 LU 6.2 Sessions

The purpose of the LU 6.2 sessions is to provide a pipe for communication between the NetBIOS switches. This pipe is used for NetBIOS data transport between one LAN and the other, for setting up the communications path between the NetBIOS applications, and for internal communication between the switches for flow control and error information. Between any two switches, there are two sessions: one for transmitting data, the other for receiving data. The purpose of this implementation is to make the data path between the switches full duplex. The sessions are able to provide reliable data transport over the WAN, as well as being able to use SNA pacing to provide a flow control mechanism. A *pacing function*, using the 802.2 logical link control is also provided between the switches and the stations within their region. This ensures that no part of the connection path can become over-congested with data. The IBM LAN-to-LAN Wide Area Network Program is defined so as to have knowledge of its own LU, its own region, and the LU names of its partners. The region name of the partner is learned when the partner session is established.

The LU sessions can be permanent, started automatically when the program is initialized, or temporary, that is started by operator command. If the partner switch is not active when session initialization is attempted, and the session is permanent, then the session will be retried. If the partner switch is not active

when session initialization is attempted, and the session is temporary, the initiating switch will wait for the partner switch to initiate the session.

The status of the sessions is displayed from the operator's screen, together with operational statistics.

1.9.3 LTLW V1.07 and ELTLW V1.01

These are two new enhanced LAN-to-LAN Wide Area Network programs. LTLW Version 1.07:

- Can communicate with up to 47 other LTLWs at one time
- Can support 500 circuits
- Can support 250 local workstations
- Contains loopback code
- Provides support for NetBIOS / IP / IPX / CID-enabled

Entry LTLW Version 1.01:

- Can communicate with up to 47 other LTLWs at one time
- Can support 30 circuits
- Can support 10 local workstations per protocol
- Contains loopback code
- Provides support for NetBIOS / IP / IPX / CID-enabled

LTLW enables two stations on separate LANs using IBM NetBIOS, IP or IPX protocols to communicate by sending frames across a WAN that connects the two LANs. Each of the LANs must have an LTLW station, and that station must be able to communicate with another LTLW station through the WAN.

NetBIOS sessions use link station connections established at the IEEE 802.2 interface level. IP and IPX are connectionless link layer protocols.

Sessions on the WAN between the two LTLW stations use APPC, the IBM implementation of LU 6.2. The LTLW stations appear to the WAN components as LUs.

The first LTLW station:

- Appends APPC headers to the frames sent to it by the source LAN stations
- Sends the *encapsulated* frames to the second LTLW station (its Partner LU) via an LU-to-LU session that conforms to the APPC protocol

The second LTLW station:

- Removes the APPC headers
- Sends the LAN protocol frames to the target LAN station

The LTLW *Loopback Utility* provides a means to connect NetBIOS, IP and IPX applications to remote LANs without the use of LAN hardware. This is done by using the LTLW *Loopback MAC driver* that simulates an IBM Token-Ring LAN to OS/2. Using the NDIS stack, applications can communicate using IEEE 802.2 protocol with MAC layer support underneath. The LTLW Loopback driver appears to be an IBM Token-Ring MAC driver. When applications send frames

to the Loopback driver, the driver converts the NetBIOS, IP or IPX information in the frame to LTLW information and sends the frame back to the IEEE 802.2 protocol stack. Both LTLW and the application think they are communicating with a local station on the LAN. Even OS/2 Services think that a LAN is attached. Therefore configuring the Loopback driver is the same as configuring an IBM Token-Ring or any other MAC layer driver.

In a stand-alone environment, there is no need to buy a LAN adapter and an Multiaccess Unit (MAU). You would be able to run LTLW and LAN Requester in a workstation that is not LAN-attached and use LTLW to have the requester connect to a LAN Server across the existing WAN.

Testing has been done to support SNMP management of the 8250 via new LTLW support.

How LTLW Routes:

- LTLW uses the IEEE 802.2 stack.
- An LTLW adapter cannot be shared for routed protocol.
- LTLW uses Qualification Exchange Routing (QER).

The LTLW uses the IEEE 802.2 protocol stack to interface with the LAN adapter. Therefore, you must configure the LTLW LAN adapter to use this stack. Furthermore, the LTLW *cannot* share this LAN adapter with other programs. While your workstation uses LTLW to route NetBIOS, IP and IPX frames, other programs cannot be configured to use these protocols on the LTLW adapter.

If you install the IBM TCP/IP for OS/2 product and the LTLW on the same workstation with only one LAN adapter, the LTLW will *not* be able to route IP frames because the TCP/IP for OS/2 product will receive those IP frames from the LAN. In other words, the IBM TCP/IP for OS/2 opens the X′AA′ Service Access Point and will get all frames addressed to the X′AA′ SAP and LTLW will not.

The LTLW maintains two lists of local resources: the *name qualifier list* and the *IP qualifier list*. When an LTLW establishes an LU 6.2 session with a partner LTLW, it exchanges these lists, if applicable, so that each LTLW knows the local resources available through the other LTLW. Those available LTLW resources are associated with the LU 6.2 Session ID that was created by the session establishment between LTLWs. When an LTLW receives frames intended for a certain name and address, it uses the lists to *qualify* the address to determine which partner LTLW should receive it. It then sends the frames to the partner LTLW. Furthermore, it puts the qualification information into memory so that it can quickly send subsequent frames without having to qualify them again. When the partner LTLW receives the frames, it also puts the qualification information into memory. This method of routing is called *Qualification Exchange Routing* (QER).

It is only necessary for you to define your local LTLW resources that will be *accessed* by other stations. This would mean that in most cases only domains and servers need be defined, or IP Network Addresses that represent the IP networks accessible through this local LTLW. The LTLW learns from the requester/server flows where the requester resides. If a resource is not qualified by the local LTLW, the frame will not be sent out across the WAN because the LTLW does not know the resource exists.

NetBIOS Support:

- LLC Type 1 and Type 2 frames are used.
- NetBIOS uses *names*.
- Qualifier list is used for *name query*.
- Qualification information is cached.
- Local LLC is used for acknowledgments and termination of hop count.

The *NetBIOS* protocol uses LLC Type 1 (connectionless) frames and LLC Type 2 (connection-oriented) frames. The protocol uses names for addressing. A requestor broadcasts to a domain using a name to find a resource. The domain controller or server will respond and a few other LLC Type 1 frames with names for addressing will flow. Eventually the stations will initialize a session and begin using LLC Type 2 frames that have session numbers, but no names, for addressing.

The LTLW routes the broadcasts using the *qualifier list*. Most applications restrict NetBIOS resource names to 8 characters. However, the frames on the LAN are 16 bytes. Applications add a prefix and suffix to the user′s destination name. Since the user′s name can be found anywhere in the 16-byte name, the LTLW uses an implied wildcard to determine the qualification of a destination name. Once learned, the LTLW does not qualify the name again, but routes frames directly to the partner using the qualification information in memory.

Note:

Each LTLW is capable of having a total of 100 names, whether they be NetBIOS or IPX, in the name qualifier list.

The LTLW will interact on the behalf of remote partners, acknowledging frames (at the LLC layer) from the LAN. This interaction allows the LTLW to satisfy the IEEE 802.2 timing parameters and hop count restrictions. The LTLW logs the session, known as a *circuit*, and tracks the amount of data transmitted over that circuit.

IP Support:

- Uses LLC Type 1 frames.
- Uses *network address*.
- Uses qualifier list for *ARP* broadcast.
- Caches qualification information.
- Uses termination of hop count.
- Supports PING function through LTLWs.

The *IP* protocol is a connectionless protocol that supports many higher layer protocols such as TCP/IP, UDP and SNMP. Acting as an IP router, the LTLW uses qualifier lists composed of IP network addresses and network masks. Using IP addresses and a mask for the subnetwork, LTLW qualifies IP frames to learn where the destination station resides. If the IP address is an address on a separate network supported by another router, the IP router address is included as part of the IP qualifier to direct LTLW to route those IP frames destined for that IP network to the local router. The LTLW uses the address resolution

protocol (ARP) to locate remote resources and determine the LAN addresses associated with the IP addresses.

IP qualifiers specify what IP resources are locally accessible through the LTLW′s WAN connection. Partner LTLWs receive the list from the local LTLW when they connect over APPC. The LTLWs use the qualifiers to route the frames to the LAN that contains the intended IP address.

Under the IP protocol, each host is assigned a 4-byte address consisting of three fields: network ID, subnetwork ID and user address. A given implementation of the IP protocol may include various combinations of these fields. The LTLW is concerned only with the network ID and the subnetwork ID fields.

An IP qualifier is specified in LTLW as three 4-byte entries, each byte in decimal format and separated with decimal points. The format is:

```
Network address:    ddd.ddd.ddd.ddd
IP router address:  ddd.ddd.ddd.ddd
Mask:               ddd.ddd.ddd.ddd
```

where ddd is a decimal value ranging from 0-255. An address of 0s and all 255s is invalid (if the IP router address is zero, the default is the LTLW′s IP address). The first entry is the host′s network ID (or address). The second entry is the local IP address of a router that supports the host′s network address (default is LTLW′s IP address for the local LAN). The third entry is the subnetwork mask. To determine the subnetwork address, the LTLW performs a logical AND operation on the IP network address and the mask. For example, IP address 192.213.45.77 and Mask 255.255.0.0 yield an IP subnetwork address of 192.213.0.0.

After deriving the subnetwork address, the LTLW compares the result with the subnetwork address yielded by performing a logical AND operation on the qualifier network address and mask. If the subnetwork addresses match, the frame is routed to the appropriate LAN.

When a frame is received from the WAN destined for a station on the local LTLW′s region, the LTLW will determine the physical address on the LAN to route the IP frame to by seeing what IP router address services the network address of the frame. If the IP router address is the LTLW′s, then the LTLW knows the station is on the local LAN and sends the frame directly to the host′s physical address. If the IP router address is not the LTLW′s IP address, the LTLW sends the IP frame to the physical address of the IP router so that the IP router may forward the frame to the proper LAN.

In LTLW, when specifying the IP qualifier, the IP address entry may be only one of any number of stations on the local LAN. By supplying the correct subnetwork mask, any IP frame destined for any station on the local LAN is properly routed. Thus, one IP qualifier will satisfy the address ability of many stations. Stations that are supported through other IP routes on the local LAN may have their subnetworks specified. To properly route to the IP networks on remote links to local IP routers, the IP addresses of the appropriate IP router must be supplied. The LTLW must know the physical LAN address of the IP router in order to send it the proper frames. Specifying the IP address of the IP router assists LTLW in determining the physical address. This allows a LAN connected over LTLW links to address an IP address that exists on the remote end of another router attached to the LTLW LAN.

Although LTLW does not have a function within it to PING another device, PINGs that are ICMP ECHO frames are supported through the LTLW and will use the Qualifier List.

IPX Support:

- Uses LLC Type 1 frames.
- Uses *SAP and RIP*.
- Exchanges SAP and RIP at LTLW session initiation.
- Uses SAP *Server Name*.
- Sends changes *only* when required.
- Qualifier list may be used for large networks.
- Uses termination of hop count.
- An LTLW can have up to 100 name qualifiers.

The *IPX* protocol is a connectionless protocol. IPX relies on applications to advertise their services in order to find addresses. A NetWare network may contain file, print, communications and other types of *servers*. The devices that need the services must know what services are available and where they are located. The server will broadcast, every 60 seconds, the server name, address and server type. Routers also help by passing the servers′ advertisements around the network. These functions are called service advertising protocol (SAP) and routing information protocol (RIP). When a requester wants to access a server, it sends out frames asking where the nearest server is and then how to get to the server that it wishes to attach. Each router on the network must be able to maintain and offer SAP and RIP information from other regions of the network to help local stations address the correct servers.

The LTLW supports the SAP and RIP protocols. It obtains SAP and RIP information from the local LAN when starting up and passes this information to its partner LTLWs over the WAN. Each LTLW sends SAP and RIP information to remote sites that are linked to the LTLW.

The LTLW can store up to 100 SAP entries and 100 RIP entries. To manage a network with more than 100 servers, you must apply the NetWare name qualifier list to the SAP entries. This filters out the local servers that should not be accessed over the WAN LTLW links.

Unlike the process used with NetBIOS and IP, the LTLW supporting IPX does not need to clear the table containing network topology (SAP and RIP) information the moment the information changes. To limit the number of WAN broadcasts needed to update the SAP and RIP entries, the LTLW only sends *changes* to the tables on a periodic basis (usually about once every minute).

You can configure IPX to support various types of LAN protocols including:

- IEEE 802.2 (token-ring or Ethernet) using the X′E0′ SAP
- SNAP (token ring or Ethernet) using the X′AA ′ SAP

The LTLW keeps information about either type of server in its SAP table. The LTLW can provide concurrent access to IEEE 802.2 IPX and SNAP IPX stations.

An LTLW can have up to 100 name qualifiers. This is for *both* NetBIOS and IPX protocols.

Link Startup:

- LTLW maintains resources for up to 47 links to partner LTLWs.
- Of those 47 Links, some are defined as:
  - Autolink
  - Manual
  - Dynamic

A *link* is the connection between an LTLW and one of its partner LTLWs. A link can start in one of the following ways:

- Automatically (Autolink)
  - Link starts when communication function is started.
  - If link fails once to establish, an attempt is made every 2 minutes to restart the link until the link is stopped manually by the operator.
  - If the link is stopped manually, it must be restarted manually. It will then act as an automatic link again.
- Manually
  - Link is started only by the operator after the communication function is started.
  - If the link fails once to establish, it is restarted in response to a request from the partner LTLW (Autolink at Partner LTLW perhaps), or it must be restarted manually.
- Dynamically
  - Link is started when the occurrence of a specific NetBIOS name is addressed on the LAN.
  - When the link becomes idle for 15 minutes (user-configurable), the link is stopped.
  - Dynamic links share common LTLW resources.
  - LTLW uses dynamically defined resources by sharing the available resource with links that are defined as Autolink and Manual. If there are more LTLW Dynamic Links defined than links available then LTLW will service requests on a first-come-first-served basis. The dynamic links defined are cycled through as the resources become available and then used again.

Each link definition specifies the link startup option. The default setting is *Autolink*.

1.9.4 LTLW Summary

The IBM LAN-to-LAN Wide Area Network Program:

- Provides a means for interconnecting LANs across a wide area network to convey NetBIOS, IPX and IP traffic from one LAN to another.
- Allows the WAN to be an SNA subarea network, an SNA APPN network, or an X.25 network. An X.25 network must be capable of supporting qualified logical link control (QLLC - SNA over X.25) switches, which can be connected together by a dedicated link.
- Allows the LAN traffic to share the bandwidth of the WAN with other traffic.

- Uses APPC LU 6.2 sessions as pipes for the interregional LAN traffic.
- Uses the facilities of Communications Manager to support the APPC LU 6.2 communications.
- Can attach to Ethernet, token-ring or PC Network (Broadband and Baseband) LANs.
- Is an OS/2 application, therefore need not run in a dedicated machine.

  Note: If the machine is required to run another NetBIOS application apart from the IBM LAN-to-LAN Wide Area Network Program, the LAN adapter cannot be shared between the applications. Two adapters must be installed.
- Provides a region resolution and filter mechanism to control access to resources across the WAN. The filter is tailorable to customers′ needs.
- Does not support *concatenated* or *parallel* connection between regions.
- Is provided with a set of utilities, allowing operator control, statistics gathering and comprehensive problem determination facilities.
- Is positioned for customers who have made an investment in their WANs and who do not wish to invest in the link costs of a remote bridging solution. The product should be used by customers for inter-LAN NetBIOS, IP and IPX traffic.

1.10 RouteXpander/2

The RouteXpander/2, an Operating System/2 (OS/2) licensed program, extends an existing OS/2 Version 2.0 communicating workstation into a high function, entry level node in bridge and router networks. The RouteXpander/2 source-route bridge and multiprotocol routing facilities transport multiple protocols, including TCP/IP, SNA/APPN and NetBIOS over a single physical link using either a frame relay or point-to-point connection. RouteXpander/2 can be used as a low cost 6611 data link switching feeder node to an upstream 6611 Network Processor. When combined with the new IBM Wide Area Connector adapter, RouteXpander/2 provides high-speed wide area network communication up to 2.048 MB per second. RouteXpander/2 provides an economical, entry-level feeder node that will interoperate with IBM 6611 and industry standard routers. As a licensed program offering, it may be easily added to an existing OS/2 file/print/mailserver or stand-alone workstation, with available capacity. Because RouteXpander/2 presents the appearance of a token-ring LAN to higher-level protocols, existing communications products need not be modified to gain the benefits of frame relay, source-route bridging, and router networks. Existing OS/2 communications protocols, such as SNA/APPN, TCP/IP, and NetBIOS, may be routed or bridged into the backbone network. RouteXpander/2 is well suited to small LANs and stand-alone workstations. Configurations that require off-LAN communication, but that do not need a dedicated bridge or router, can now gain high-speed access to WANs.

1.10.1 Highlights

RouteXpander/2 provides business solutions by:

- Exploiting the frame relay technology and by serving as a feeder node to the IBM 6611 Network Processor and industry-standard router networks.
- Providing OS/2 applications with high-speed, affordable access to enterprise data with minimal consideration of geographic location or processor type.

- Multiplexing multiple protocols over a single physical link thus, reducing network cost.
- Supporting the multiple communications protocol stacks available on OS/2, allowing development of business solutions using the most appropriate applications, without regard for underlying communications protocols.

RouteXpander/2 brings the multiprotocol router, bridging, and frame relay technologies to OS/2 workstations. Among the key benefits of RouteXpander/2′s frame relay support are its ability to multiplex multiple protocols on a single physical link and communicate with multiple destinations (logical links) on a single physical link. By allowing consolidation of multiple communications links and associated equipment into a single link, the customer can reduce network costs and develop new applications using multiple protocols. This may not have been economically feasible when separate networks were required. With the new support for multiple protocols on a single link, stand-alone and LAN applications connected via RouteXpander/2 have more access to resources attached in the customer network. By acting as a feeder node, RouteXpander/2 offers applications access to frame relay and router networks, without requiring a dedicated processor for the connection. RouteXpander/2 communicates with another copy of RouteXpander/2 over a nonswitched, point-to-point link, offering direct LAN-to-LAN multiprotocol communications in situations where a network is not required or not available. Because a dedicated machine is not required, a customer may now implement a multiprotocol wide area network at a lower cost and gain the application benefits of such a network.

1.10.2 Technical Description

A primary purpose of RouteXpander/2 is to provide small LANs and stand-alone workstations with access to frame relay, bridge, and router networks. To do this, RouteXpander/2 takes advantage of communications protocol products running on OS/2, including IBM TCP/IP for OS/2 (TCP/IP) and IBM Communications Manager/2 (SNA/APPN and NetBIOS). These product currently perform routing for TCP/IP and SNA between LAN adapters on the OS/2 workstation. This, in effect, creates a multiprotocol platform whose communications protocol support corresponds to the protocol stacks provided by the protocol products. RouteXpander/2 acts as an NDIS interface to provide a token-ring LAN appearance to the higher-level protocol products. RouteXpander/2 manipulates the LAN header information and converts it to frame relay formats. It then writes the data to a communications driver, such as the IBM Wide Area Connector adapter, through a lower NDIS interface.

RouteXpander/2 acts as an NDIS wedge, providing an upper level NDIS appearance to the protocol products and a lower level NDIS appearance to the communications driver. The technique of emulating a token-ring LAN allows the protocol products to be configured as though they were communicating with another workstation on the LAN when, in fact, RouteXpander/2 is providing access to multiple workstations and applications that are reached through the WAN. In order for routing to occur, a protocol must be routable and the protocol stack must exist on the OS/2 workstation with RouteXpander/2. RouteXpander/2 contains a source-route bridge to handle situations where routing is not possible. For example, in the case of NetBIOS which cannot be routed, the source-route bridge is used and can connect two token-ring LANs.

Because the frame relay device driver is configured as a token-ring LAN, traffic is bridged from the LAN to a WAN. Because it is a source-route bridge, only

token-ring LANs are supported for bridging; Ethernet LANs are not supported for bridging. RouteXpander/2 is designed to route packets, if possible. If routing is not possible, the packet will be bridged. For bridging, RouteXpander/2 interoperates with another copy of RouteXpander/2 and with the native bridge support of the 6611 Network Process, over a frame relay network, or using a point-to-point connection. The frame relay support, in effect, permits the single port bridge to establish bridge connections to a maximum of 200 other bridges in the frame relay network.

1.10.3 Network Management

RouteXpander/2 uses Extended Services for OS/2 to build SNA alerts which may be forwarded to NetView. These alerts are issued to indicate conditions requiring attention such as link and logical link (DLCI) outages. RouteXpander/2 extends the SNMP agent facility of TCP/IP for OS/2 to build and forward traps to an SNMP manager, such as AIX NetView/6000. The bridge and frame relay management information base (MIB) objects are supported. Performance statistics are maintained as defined in those MIBs. The SNMP GET functions are supported for each of these MIBs.

1.10.4 RouteXpander/2 Interoperability

RouteXpander/2 is designed to communicate with any IBM or OEM device that complies with the frame relay specifications and conforms to multiprotocol operation as documented in RFC 1294. RouteXpander/2 supports the following ANSI and CCITT standards for frame relay:

ANSI:

- T1S1.2/91-454
- T1.617-1991 (Annex D)
- T1.618-1991

CCITT:

- COM XVIII-R 47E I.370
- COM XI-R 133-E (Q.933) (Annex A)
- COM XI-R 125-E (Q.922) (Annex A)

RouteXpander/2 also supports the RFCs for Multiprotocol Interconnect over Frame Relay (RFC 1294 - proposed Annex F to ANSI T1.617-1991) and Inverse Address Resolution Protocol (RFC 1293). RouteXpander/2 supports source-route bridging, including spanning tree, as specified in the IEEE P802.1d bridging standard, and the IEEE P802.5M/D6 source-routing supplement. For network management, RFC 1286 (bridging) and RFC 1315 (frame relay) are supported.

1.10.5 Compatibility

Because RouteXpander/2 operates as a device driver, it is insensitive to applications running on the higher-level protocol drivers. Some examples of applications and configurations supported by RouteXpander/2 are:

- TCP/IP application on a LAN workstation communicating with a TCP/IP host application. An FTP application is running on a TCP/IP workstation that is Ethernet-attached to an OS/2 server with RouteXpander/2 and TCP/IP for OS/2. Traffic from the originating workstation is received by TCP/IP for OS/2 in the server workstation, which routes it to RouteXpander/2.

RouteXpander/2 sends the traffic over a frame relay network to a 6611 Network Processor. The 6611, which is token-ring-attached to a 3172 control unit that is channel-attached to an S/390 host, sends the traffic to a TCP/IP FTP application.

- A 3270 Emulator on a LAN communicating with a host application. An OS/2 workstation with the Extended Services or CM/2 for OS/2 3270 emulator is token-ring-attached to an OS/2 workstation running RouteXpander/2. Traffic originating at the 3270 emulator is received by the RouteXpander/2 source-route bridge that bridges the traffic over a frame relay network to a 6611 Network Processor, which incorporates it into data link switching for routing within router networks and subsequent delivery to another 6611. This 6611, token-ring-attached to a 3745 controller that is channel-attached to an S/390 host, sends the traffic to the host 3270 application. Note that Extended Services for OS/2 need not be installed in the RouteXpander/2 workstation for bridging of SNA traffic.
- Other workstations on the originating token-ring can be sending other protocols such as TCP/IP and NetBIOS to the workstation running RouteXpander/2. NetBIOS traffic is bridged to its destination and the TCP/IP traffic is routed to its destination, all over the same physical link to the frame relay network. Note that in both of these examples, the traffic could be sent on a point-to-point connection to a 6611 or directly to another copy of RouteXpander/2 without crossing a frame relay network. Likewise, the traffic could be sent across a frame relay network to another copy of RouteXpander/2 or some other equipment manufacturer router. Remote RouteXpander/2 bridge support interoperates with another copy of RouteXpander/2 and with the native bridge support of the 6611 Network Processor, either over a frame relay network or via a point-to-point connection. RouteXpander/2 can coexist on a token-ring LAN with the IBM Token-Ring Network Bridge Program and the IBM 8229 LAN Bridge, but a remote connection between the two or between RouteXpander/2 and the compatibility mode bridge support of the 6611 Network Processor is not supported. RouteXpander/2 interoperates with the IBM 6611 Network Processor router and is designed to operate with OEM devices that conform to the latest frame relay specifications, including RFC 1294.

1.10.6 RouteXpander/2 (RXR/2) Features

The second port on a WAC adapter is supported by RXR/2. This allows daisy chaining small offices together and will support through traffic on the same physical line.

RouteXpander/2 Ethernet Bridge Support: Ethernet transparent bridging is provided by RXR/2. This bridge is in addition to and independent of the source-route bridge available today. This support permits both local Ethernet LANs and remote WAC-attached lines to be transparently bridged. The IBM EtherStreamer adapter is supported.

Transparent bridging requires that the adapter be run in promiscuous mode.

RouteXpander/2 IPX Router Support: The IPX router support will enable RXR/2 to route IPX traffic from either a token-ring or Ethernet LAN into the frame relay and/or router network to an RFC-1294 compliant router. The RIP and SAP broadcast interval is configurable in this support as well as RIP and SAP filtering.

Support and Service: RXR/2 has a bulletin board service available on both CompuServe and Advantis. This service enables customers to exchange experiences on RXR/2 and to have immediate access to service for the product. Service can be downloaded from either bulletin board.

1.11 IBM AnyNet Product Family

This section describes all the products of IBM AnyNet Family (including AnyNet/2 NetBEUI over SNA), NetBIOS and the Multiprotocol Transport Networking (MPTN).

1.11.1 IBM AnyNet

AnyNet is a versatile software family of access node and multiprotocol gateway products that allow customers to choose the applications that meet the needs of their businesses, regardless of what transport protocol is used in their central or remote sites.

AnyNet products are based on the Multiprotocol Transport Networking (MPTN) architecture, which allows applications to be enabled in mixed protocol networks. The industry standard MPTN solution is part of the Networking Blueprint framework introduced in 1992 by IBM.

AnyNet access node products provide a way for enterprises to run new types of applications on AIX, MVS/ESA, OS/2, OS/400 and Windows end-systems on an existing network. For example, SNA printer applications can run over a TCP/IP network. AnyNet *gateway* products allow enterprises to extend the scope of their SNA or TCP/IP backbone networks to embrace multiprotocol support. AnyNet gateways allow a variety of configuration options for LAN/WAN integration.

The AnyNet products are attractive to customers who:

- Have SNA application solutions that they want to provide on TCP/IP networks
- Have TCP/IP Sockets application solutions that they want to provide on IPX, SNA or NetBIOS networks
- Have NetBIOS application solutions that they want to provide on SNA networks
- Want to allow remote IPX, TCP/IP or NetBIOS branch locations to reuse the connectivity of an existing SNA network
- Want to allow remote SNA locations to reuse the connectivity of an existing TCP/IP network
- Want to consolidate or change network backbones

The products in the AnyNet Family are described in Table 1.

Table 1 (Page 1 of 2). IBM AnyNet Product Family at a Glance

Access Node	Products
APPC over IPX	AnyNet/400 in OS/400 V3R6.0
Sockets over IPX	AnyNet/2
	AnyNet/400 in OS/400 V3R6.0
Sockets over NetBIOS	AnyNet/2

Table 1 (Page 2 of 2). IBM AnyNet Product Family at a Glance	
Access Node	**Products**
NetBIOS over SNA	AnyNet/2 NetBEUI over SNA
Sockets over SNA	AnyNet/MVS VTAM feature
	AnyNet/2
	AnyNet/400 in OS/400 V3R1.0
	AnyNet/6000 AIX SNA Server/6000 feature
	AnyNet for Windows (beta)
APPC over TCP/IP	AnyNet/MVS VTAM feature (also supports SNA over TCP/IP)
	AnyNet/2 (also supports SNA over TCP/IP)
	AnyNet/400 in OS/400 V3R1.0
	AnyNet/6000 AIX SNA Server/6000 feature
	AnyNet for Windows
Gateways	**Products**
IPX over SNA Gateway	AnyNet IPX over SNA Gateway for OS/2
	2217 Nways Multiprotocol Concentrator
NetBIOS over SNA Gateway	AnyNet for OS/2 (beta)
	2217 Nways Multiprotocol Concentrator
Sockets over SNA Gateway	AnyNet/2 Sockets over SNA Gateway
	2217 Nways Multiprotocol Concentrator
IPX over TCP/IP Gateway	AnyNet for OS/2 (beta)
NetBIOS over TCP/IP Gateway	AnyNet for OS/2 (beta)
SNA over TCP/IP Gateway	AnyNet SNA over TCP/IP Gateway for OS/2
	AnyNet/MVS VTAM feature

1.11.2 Customer Requirements

The need for this type of solution arises from the diversity of today's networks. With the growth of networking and local area networks in particular, most large networks now run multiple networking protocols.

Also, many more alliances are being formed that cause customers to seek inter-enterprise network interconnection. To better support multiprotocol networks, MPTN-based AnyNet products provide solutions for:

- Adding new application types independent of the existing networking protocol
- Reducing networking costs by consolidating and simplifying multiprotocol networks while protecting the investment of existing applications
- Extending the reach of applications across multiple networks

The AnyNet product family delivers customer value by:

- Expanding business solutions by enabling new applications, unconstrained by network type
- Providing flexibility in network integration and application enablement

1.11.3 Multiprotocol Transport Networking (MPTN)

MPTN is an architecture that defines IBM′s strategic direction of providing complete independence between distributed application services and the underlying network transports. This architecture has been submitted to X/Open. MPTN end nodes offer support for application programs written to one network transport to execute over nonnative transports.

MPTN gateways allow for the interconnection of unlike networks in a way that is transparent to application programs.

The MPTN architecture, being a general solution to internetworking, overlaps with other specific industry solutions. By itself, for example, the MPTN architecture for OSI over TCP/IP may duplicate the function of RFC 1006, which the Internet Engineering Task Force uses to solve the same problem. NetBEUI over TCP/IP, as defined by MPTN, is vaguely similar to the RFCs 1001/1002. While none of the specific solutions (such as the RFCs) are expected to interoperate or even coexist, all AnyNet products by their very nature avoid these problems.

The introduction of gateways that implement the MPTN architecture allows AnyNet to further its goal of application/network independence. Customers are increasingly able to use the applications they want anywhere in their enterprises without changing their logical networks. Further, the reduction of the number of underlying protocols is increasingly possible as AnyNet expands its product set and encourages other MPTN vendors to do the same.

Industry acceptance of the MPTN architecture continues to grow. In January 1994, X/Open Company Limited published its first MPTN document, *X/Open Guide: Multiprotocol Transport Networking Architecture*. In October 1994, X/Open published three MPTN documents as *X/Open Preliminary Specifications*, opening up MPTN to other vendors and achieving another milestone towards MPTN′s acceptance as an industry standard.

The *MPTN Preliminary Specifications* are available from X/Open and are being used by vendors to implement MPTN-compliant products. In July 1995, Wall Data announced MPTN-based APPC over TCP/IP access node support in a new product, *Rumba Access/400*. This is the first non-IBM product available that uses the X/Open standard for Multiprotocol Transport Networking.

X/Open MPTN documents:

- Access Node - ISBN 1 85912 040 7
- Address Mapper - ISBN 1 85912 039 3
- Data Formats - ISBN 1 85912 043 1

X/Open documents may be ordered by mail or fax:

X/Open Publications, PO Box 96, Witney, Oxon OX8 6PG, England

Tel:+44 (0)993 708731, Fax:+44 (0)993 708732

1.11.4 NetBIOS

NetBIOS is a very widespread LAN network transport. Invented by IBM, NetBIOS grew in popularity because it was very small, and very fast. As LANs became more prevalent, so did NetBIOS. Later, Novell′s IPX came along and replaced NetBIOS as the most popular LAN protocol. NetBIOS is nonroutable, having no concept of net IDs. NetBIOS programs can send and receive data over connections, as well as send and receive connectionless datagrams. In fact, a program may send datagrams not only to another individual partner, but to multiple partners who have all registered under the same group name.

NetBIOS, like all LAN protocols, uses LAN broadcasting for name resolution. A program must register a name with the network before another program can communicate with it. When a program requests a name be added, NetBIOS issues a LAN Name Query message over the LAN asking if anyone already has that name. Typically it does this six times at half-second intervals, although this is configurable. If another machine has this name, it responds to the query. Otherwise, the first program is told it has that name.

1.11.5 NetBEUI

NetBEUI stands for the NetBIOS End User Interface and is the industry-standard programming interface for NetBIOS. It is traditionally tied with the NetBIOS transport, in the same way that sockets are associated with TCP/IP transport and CPI-C is associated with the APPC protocol.

AnyNet/2 NetBEUI over SNA allows NetBIOS programs to communicate with other NetBIOS programs across SNA/APPN networks. If LAN Server and LAN Requester are configured to take advantage of AnyNet/2 NetBEUI over SNA, then the large set of LAN Server application programs can also communicate across SNA networks.

Currently, NetBIOS networks consist of machines on a single LAN. As LANs grow and are bridged together, these networks can grow quite large. However, NetBIOS is a chatty protocol; that is, it issues LAN broadcast messages any time it adds or searches for names on the network. It also must use timers in order to wait for replies to its broadcast inquiries. The timeouts must get larger as the LAN grows and, as in all LAN protocols, fail to work at all in a WAN environment.

A very common customer complaint is the amount of broadcasting that is done on their physical networks. AnyNet/2 NetBEUI over SNA offers its own NetBIOS API that directly uses SNA transport. No LAN broadcasts are ever issued, and the traditional NetBIOS timers no longer exist. AnyNet/2 NetBEUI over SNA uses a centralized directory for its name resolution. In comparing AnyNet/2 NetBEUI over SNA with other choices, the customer must answer the following questions. Do I need to support multiple protocols? Is the AnyNet family comprehensive enough for my needs? What are the technical differences (installation, performance, management) between these products?

1.11.6 Positioning AnyNet/2 NetBEUI over SNA with LTLW

Both products offer the ability to transport NetBIOS traffic over LU 6.2 connections. Obvious differences are that the former follows the MPTN architecture, while the latter is a gateway.

1.11.6.1 Design

Native NetBIOS typically uses LAN broadcasts for name resolution. AnyNet/2 NetBEUI over SNA uses a central directory. AnyNet NetBIOS networks are not at all sensitive to topology; end users may be on the same LAN, remote LANs across WANs, or no LAN at all.

AnyNet/2 NetBEUI over SNA has code that must execute on each machine, but the code may be stored on a remote file server. Communications Manager/2 (CM/2) must also execute on each machine, but in the case of CM/2, the code may also exist on a remote file server. The amount of memory taken up by APPC and AnyNet/2 NetBEUI over SNA is estimated at 2-3 megabytes.

As an AnyNet product, AnyNet/2 NetBEUI over SNA uses MPTN line-flows and can participate in future MPTN networks built up with cascaded MPTN gateways.

LAN-to-LAN over WAN (LTLW) is a gateway that today transports NetBIOS over SNA as well as IPX and IP. A stand-alone box listens on the LAN for NetBIOS names according to a defined template for routing traffic to other LANs. Enterprises enforce naming conventions so that LTLW knows what names are routed to other LANs.

1.11.6.2 AnyNet Gateway

Traffic coming from another station is routed over the SNA network using the function of an AnyNet gateway. This is also exactly what LTLW does today. LLC2 timers do not exist at the gateway.

For many customers, a gateway is required because they have NetBIOS applications on DOS/Windows machines or OS/2 machines with little memory. Another frequently-heard comment is that customers are reluctant to install new software on any machines besides code servers and therefore want a gateway. However, as mentioned earlier, all of the AnyNet code may exist on a remote code server, and with CM/2, all of the CM/2 APPC code may reside there as well.

In other cases, a gateway may not be required at all. The obvious advantages of moving to a single APPC/APPN network include network performance, predictability, stability, single network managers, and APPN LAN and WAN connectivity. Simply, APPC/APPN is becoming recognized as a premier LAN protocol. Customers want the new remote load support built into CM/2. This is an attractive way for customers to spread APPC/APPN out from their WAN backbone into their LANs. The growing family of APPC applications would be available to all machines on these LANs.

AnyNet/2 NetBEUI over SNA would participate in this environment by providing WAN support for NetBIOS applications.

1.11.6.3 Remote Access to Central LAN Machines

Several customers simply have no need at all for any-to-any NetBIOS connectivity across a WAN. One bank realized that having all tellers at all sites able to transfer files or log on to any other machine at any other of 1000+ locations may not be desirable. Such customers typically have districts and regions; their requirement is for someone at a regional office to be able to log on or transfer files with one or a few machines at each remote LAN. The configuration and definition requirements required by gateways to enable true any-to-any connectivity on this scale could be quite costly.

AnyNet/2 NetBEUI over SNA offers large corporate customers the ability to design topology-independent networks that support the popular NetBIOS applications. A typical configuration has at least one machine per LAN in a NetBIOS network that spans the WAN backbone. Users can log on to this network using LAN Server/Requester and gain access to the resources on these remote machines. Users could also send notes and forms to these remote machines using such packages as Lotus Notes and ccMail.

The IBM 8235 also supports remote NetBIOS connectivity. The remote user would load NetBIOS on his PC and then dial into the LAN via the 8235 running on Ethernet or token-ring. The remote user would then be able to access NetBIOS application as though locally attached to the LAN. For a more detailed description of the 8235, refer to Chapter 5 in *Local Area Network Concepts and Products: Adapters, Hubs and ATM,* SG24-4754.

Elimination of Broadcast Storms: AnyNet/2 NetBEUI over SNA is the introduction of the NetBIOS API (NetBEUI) as an alternative API for SNA connectivity. This is to support mission-critical programs that are not written to the preferred CPI-C interface. As stated earlier, this approach completely eliminates LAN broadcasting.

NetBIOS Routability: NetBIOS, as a LAN protocol, has no concept of network identifiers. They simply are not available in the API. However, a considered direction for AnyNet's support of NetBIOS is to introduce net IDs into the API. This would allow for worldwide AnyNet NetBIOS networks that are logically separate and yet able to intercommunicate.

NetBIOS broadcasts are completely eliminated. However, AnyNet must support NetBIOS applications that issue broadcasts and multicasts.

A broadcast is the transmission of a data buffer to all nodes on a LAN. Broadcasts may be issued by NetBIOS applications, but are much more frequently used by NetBIOS to add names to a network and find names on a network. When a NetBIOS application does an ADD_NAME to a network, native NetBIOS would broadcast a message to all nodes on the LAN. NetBIOS would then wait for a predetermined duration for another node to announce that it already has that name. If no response comes, this node considers the name safe for use.

In contrast, AnyNet/2 NetBEUI over SNA, when it receives an ADD_NAME, attempts to register the name with the APPC Name Server. If the registration succeeds, we are done; otherwise, we signal to the application that another node already has that name in use. No broadcasting across the network is ever done for name resolution.

If a NetBIOS application wishes to issue a broadcast, it is treated as a special case of *multicast*. A multicast is the transmission of a data buffer to all nodes registered under a NetBIOS group name. Multicast is invoked by NetBIOS applications.

AnyNet/2 NetBEUI over SNA ships a multicast server. When a node requests that a data buffer be sent to all members in a group, that data buffer is sent to the multicast server. The server finds all of the members in the group by querying the APPC Name Server. It then sends the data buffer to each member. We require the multicast server and APPC Name Server to be on the same machine

for ease of configuration. AnyNet/2 NetBEUI over SNA networks are a star configuration around this node.

An AnyNet/2 NetBEUI over SNA network is, by definition, all nodes that are connected to the APPC Name Server. In this case it obviously does not matter where the nodes are located; they may be on different LANs, they may be across a WAN, or they may even be across an asynchronous line.

The original NetBIOS interface, as defined in the *IBM LAN Technical Reference*, is now often referred to as the *NB30 interface*. AnyNet/2 NetBEUI over SNA supports the NB30 interface. A problem with the NB30 interface is that in OS/2 it cannot support multiple NetBIOS transports. That is, if two vendors provide their own OS/2 NetBIOS NB30 implementations, they cannot coexist in the same machine. Microsoft took the initiative on this issue and defined the LM10 interface. LM10 is not widely-documented but is almost identical to NB30. LM10 sits under NB30, and NB30 now simply maps to LM10. But LM10 has a register function, wherein each NetBIOS transport in a machine states which LAN adapter it will support. So, for instance, a Novell NetBIOS and a Microsoft NetBIOS can now coexist, each driving a different LAN adapter. For AnyNet/2 NetBEUI over SNA, the LM10 interface is supported. This allows coexistence with NTS/2.

1.11.6.4 Management

A major emphasis in all AnyNet products is to make sure the user is told in a considerable way what is going on. AnyNet/2 NetBEUI over SNA offers error reporting and tracing consistent with what a user familiar with CM/2 would expect. Error reporting is done with FFST/2, which is supported by all components of the product. This mechanism logs any unusual or fatal condition into message and error log files. CM/2 ships tools to format and display these logs.

1.11.7 AnyNet/2 Version 2.0.2

AnyNet/2 delivers four multiprotocol combinations:

- SNA over TCP/IP
- Sockets over IPX
- Sockets over NetBIOS
- Sockets over SNA

The following text discusses the key functions of each combination:

- SNA over TCP/IP

  As an access node, APPC applications can communicate over TCP/IP to other AnyNet SNA over TCP/IP or APPC over TCP/IP access node products on AIX, OS/2, OS/400, and MVS/ESA.

  Also as an access node, AnyNet/2 is used in a single gateway configuration to integrate communications over connected TCP/IP and SNA networks. APPC applications can communicate over TCP/IP through an AnyNet SNA over TCP/IP Gateway on OS/2 or MVS/ESA to like APPC applications running on a native SNA network.

  When used with the AnyNet/MVS portion of the VTAM V4R2 AnyNet Feature, SNA emulator and printer applications, as well as APPC applications, can run over a TCP/IP network.

- Sockets over SNA

  As an access node, Sockets applications can communicate over SNA to other AnyNet Sockets over SNA access node products on AIX, OS/2, OS/400, and MVS/ESA.

  Also as an access node, AnyNet/2 is used in a single gateway configuration to integrate communications over connected TCP/IP and SNA networks. Sockets applications can communicate over SNA through an AnyNet/2 Sockets over SNA Gateway to like Sockets applications running on a native TCP/IP network.

- Sockets over IPX

  As an access node, Sockets applications can communicate over SNA to other AnyNet Sockets over SNA access node products on OS/2.

- Sockets over NetBIOS

  As an access node, Sockets applications can communicate over SNA to other AnyNet Sockets over NetBIOS access node products on OS/2.

1.11.7.1 Application Enablement

For SNA over TCP/IP, in addition to running APPC applications such as CICS and DB2/2, customers can now run SNA printer and emulator applications when AnyNet/2 is used with the AnyNet Feature for VTAM V4R2.

For Sockets over IPX, SNA and NetBIOS, all C-based Berkeley Sockets Distribution (BSD) applications can communicate over SNA, including DCE, DSOM, FTP, HOST, NFS, PING, REXEC, RSH, SAP R/3, SNMP, TALK, Telnet, and X-Windows.

1.11.7.2 Customer Scenario: AnyNet/2 SNA over TCP/IP

A business with remote branch offices, spread over thousands of miles, depends on its information network to consolidate vital operations and financial data. To control hardware, line, and network management costs, the network service providers have selected to standardize on a single network protocol. The branch local area networks use TCP/IP and a TCP/IP backbone for host communications.

Workers in branch locations are accustomed to and want to continue to use their existing data transfer applications, which are LU 6.2-based, to update daily financial reports. Rewriting the existing applications to a Sockets (TCP/IP) interface is an expensive programming project and is considered if the benefits outweigh the costs.

AnyNet/2 SNA over IP satisfies the need to use existing LU 6.2 applications across the TCP/IP backbone network without any change to the application.

1.11.8 AnyNet/6000

AnyNet/6000 APPC over TCP/IP and AnyNet/6000 Sockets over SNA are features of AIX SNA Server/6000 Version 2 Release 1.1.

1.11.8.1 APPC over TCP/IP Feature

As an access node, AnyNet/6000 APPC over TCP/IP allows existing TCP/IP networks to add APPC or CPI-C applications without adding a separate SNA network. AnyNet/6000 APPC over TCP/IP can communicate with other AnyNet APPC over TCP/IP access node products on AIX, MVS/ESA, OS/2, and OS/400.

Also as an access node, AnyNet/6000 is used in a single gateway configuration to integrate communications over connected SNA and TCP/IP networks. APPC applications can communicate over TCP/IP through the AnyNet SNA over TCP/IP Gateway on MVS/ESA or OS/2 to matching APPC applications running in a native SNA network.

1.11.8.2 Application Enablement

Any APPC or CPI-C application such as CICS/6000 or DB2/6000 can communicate across a TCP/IP network using this support.

1.11.8.3 Sockets over SNA Feature

Sockets over SNA allows existing sockets applications to run over any SNA network, avoiding the cost of adding a separate TCP/IP network. This function is of particular interest to companies wishing to integrate the numerous UNIX sockets applications for the AIX platform into their existing SNA networks.

As an access node, Sockets applications can communicate over SNA to other AnyNet Sockets over SNA access node products on AIX, OS/2, OS/400, and MVS/ESA.

Also as an access node, AnyNet/6000 is used in a single gateway configuration to integrate communications over connected TCP/IP and SNA networks. Sockets applications can communicate over SNA through an AnyNet/2 Sockets over SNA Gateway to like Sockets applications running on a native TCP/IP network.

Application Enablement: Some of the key sockets applications supported are file transfer protocol (FTP), network file system (NFS), NetView/6000, Telnet, RLOGIN, and AIX DCE/6000.

1.11.9 AnyNet/400

APPC over TCP/IP and Sockets over SNA multiprotocol combinations are built in AnyNet/400 functions shipped with the AS/400 base operating system, OS/400 Version 3 Release 1.

1.11.9.1 Application Enablement

The following applications will work in an AnyNet environment:

- Sockets over SNA: FTP, LPR/LPD, SMTP, and SNMP
- APPC over TCP/IP: Any LU 6.2 application, including CICS/400, SNA/DS, DB2/400, and 5250 Display Station Passthrough

1.11.10 VTAM Version 4 Release 2 AnyNet Feature

The AnyNet feature for VTAM Version 4 Release 2 on MVS/ESA provides significant new functions over the VTAM Version 3 Release 4.2 feature. The AnyNet/MVS multiprotocol functions are:

- SNA over TCP/IP
- SNA over TCP/IP Gateway

- Sockets over SNA

Highlights of the VTAM AnyNet feature for VTAM V4R2 include the following:

- Broader application support is available. Now, all SNA LU types are supported for communications over TCP/IP networks. In the AnyNet feature of VTAM V3R4.2, only APPC (LU 6.2) is supported over TCP/IP networks. In the AnyNet feature of VTAM V4R2, SNA printers and emulators are also supported over TCP/IP networks. This means that the customer′s current investment in the emulator and printer programs are significantly extended through availability of those applications to end users in TCP/IP networks without any changes to the applications required and with the same end-user appearance.
- Sockets over SNA can now run over subarea and APPN local and wide area networks.

  In general, AnyNet is related to APPN in that it extends even further the benefits of APPN by allowing additional application types, such as Sockets or NetBIOS applications, to communicate over APPN networks. AnyNet increases the number of applications that can communicate over APPN networks.

1.11.10.1 Application Enablement

For Sockets over SNA, with a new update to the AnyNet/MVS feature of VTAM V4R2, MVS Open Edition (OE) Sockets applications (for example, Distributed Computing Environment (DCE) remote procedure call (RPC) applications) can now run over subarea SNA and APPN networks in addition to TCP/IP networks. DCE RPC applications are Sockets-based and traditionally run over TCP/IP networks. This update delivers integrated OE Sockets support, meaning that the customer can choose either SNA/APPN or TCP/IP at the OE level and then all OE Sockets applications will implement that transport.

So on the host, these Sockets applications can run over SNA networks: DCE, MVS Open Edition, NFS, PING, and X-Windows.

The VTAM V4R3 AnyNet/MVS feature will deliver converged OE Sockets support, meaning that the customer has the ability to dynamically choose either the SNA/APPN or TCP/IP transport based on which network is supporting the target endpoint.

For SNA over TCP/IP, all SNA applications are enabled to run over TCP/IP. Example applications include CICS/ESA, DB2, IMS/ESA, NetView Distribution Manager, CICS OS/2, DCAF and TSO.

Of course, user-written or other vendor-socket and SNA applications can work in addition to those listed in this section.

1.11.11 AnyNet/2 Sockets Over SNA Gateway Version 1.1.6

The AnyNet/2 Sockets over SNA Gateway allows matching TCP/IP Sockets applications to communicate across connected TCP/IP and SNA networks. No changes are required to the SNA and TCP/IP networks other than adding the AnyNet gateway. Several configuration choices are configurable:

- A *single gateway* enables Sockets applications running over TCP/IP to communicate through the gateway to like applications running on AnyNet Sockets over SNA access node software on AIX, MVS/ESA, OS/2, or OS/400.

- *Two gateways* are used to connect Sockets applications on remote TCP/IP LANs across an SNA network.
- *Parallel gateways* increase the number of Sockets ;i1.parallel gateways connections that can be supported simultaneously over integrated SNA and TCP/IP networks.

Sockets applications can take advantage of SNA′s networking features.

- Cost-effective bandwidth utilization.
- Predictable response time.
- Traffic prioritization - allows the association of class of service (COS) and priority, allowing existing SNA traffic to meet response times while adding Sockets traffic.
- Data compression - reduces the amount of data being exchanged between partners, improving response time and providing higher data rates at lower cost.
- APPN dynamics.

AnyNet/2 Sockets over SNA Gateway is available in three separately priced gateways. The 20, 100, and 250 connection sizes support small, medium, and large networks. The customer chooses the configuration and corresponding price that suits current networking needs. The size of the AnyNet Gateway is changed with a size increase kit that provides a password. Customers do not have to reinstall or reconfigure the gateway.

1.11.11.1 Customer Scenario: AnyNet/2 Sockets over SNA Gateway

The AnyNet/2 Sockets over SNA Gateway provides companies and institutions, such as banks, with a software solution to manage remote branch locations over their SNA networks. Customers can install an SNMP-based management product such as IBM′s NetView/6000 and gather problems and alerts from the remote LANs through the SNA network back to a central site.

1.11.12 AnyNet SNA over TCP/IP Gateway for OS/2

The AnyNet SNA over TCP/IP Gateway for OS/2 provides SNA application connectivity across SNA and IP networks. It is a new member of the AnyNet product line. Now customers can interconnect SNA and TCP/IP networks, without impacting applications, and reduce network costs at the same time.

With this gateway, end users on SNA networks can now access and communicate with other SNA networks across private TCP/IP networks or across the worldwide Internet. Several configuration choices are configurable:

- A *single gateway* enables APPC applications on SNA systems to communicate to like applications running on AnyNet APPC over TCP/IP or SNA over TCP/IP access node software on AIX, MVS/ESA, OS/2, and OS/400. And SNA printers and emulators can communicate from an AnyNet/2 Version 2 SNA over TCP/IP access node through the gateway to matching SNA applications.
- *Two gateways* are used to connect any APPC applications, such as CICS and DB2, running on SNA networks, over a TCP/IP network.
- *Parallel gateways* increase the number of SNA sessions that can be supported simultaneously over integrated SNA and TCP/IP networks.

AnyNet SNA over TCP/IP Gateway for OS/2 is available in four separately priced gateway sizes. The low-priced entry size meets the needs of small remote offices by providing support for 20 local connections. The 100, 250, and 500 connection sizes support small, medium, and large networks. The size of the AnyNet Gateway is changed with a size increase kit that provides a password. Customers do not have to reinstall or reconfigure the gateway.

1.11.13 AnyNet IPX over SNA Gateway for OS/2

The AnyNet IPX over SNA Gateway for OS/2 provides IPX application connectivity across IPX LANs and SNA networks. It is a new member of the AnyNet product line. Now customers can interconnect IPX LANs and SNA networks, without impacting applications or network performance, and reduce network costs at the same time.

In a *two gateway* configuration, end users on IPX LANs can access and communicate with other IPX LANs worldwide across SNA networks. The gateway is used to connect any client/server application that runs on a NetWare Network Operating System, such as Lotus Notes or NetWare Management Services, over an SNA network. The gateway fully protects the SNA backbone by filtering IPX broadcasts.

IPX applications can take advantage of SNA′s networking features through:

- Cost-effective bandwidth utilization.
- Predictable response time.
- Traffic prioritization - allows the association of class of service (COS) and priority, allowing existing SNA traffic to meet response times while adding IPX traffic.
- Data compression - reduces the amount of data being exchanged between partners, improving response time and providing higher data rates at lower cost.
- APPN dynamics.

AnyNet IPX over SNA Gateway for OS/2 is available in four separately priced gateway sizes. The low priced entry size meets the needs of small remote offices by providing support for 20 local connections. The 100, 250, and 500 connection sizes support small, medium, and large networks. The size of the AnyNet Gateway is changed with a size increase kit that provides a password. Customers do not have to reinstall or reconfigure the gateway.

1.11.14 AnyNet APPC over TCP/IP for Windows

The AnyNet APPC over TCP/IP for Windows access node allows CPI-C applications to run over a TCP/IP network. End users on Windows workstations in TCP/IP networks can now access CPI-C applications as well as TCP/IP applications. Now customers can administer a single TCP/IP network while running both TCP/IP and CPI-C applications. Network complexity and costs are reduced, without impacting applications. This product runs with IBM and other vendor TCP/IP software that uses the Windows Sockets application program interface (API).

In an *single network configuration*, end users on Windows workstations in TCP/IP networks can now access CPI-C applications in hosts, workstations, and

minicomputers running AnyNet access node software on MVS/ESA, OS/2, OS/400, AIX and Windows platforms.

In a *multiple network configuration*, this new Windows access node product works with the AnyNet SNA over TCP/IP Gateway on OS/2 or MVS/ESA to connect TCP/IP and SNA networks. CPI-C applications running natively in an SNA network communicate through the AnyNet gateway to similar applications running on AnyNet APPC over TCP/IP for Windows access node software in an adjacent TCP/IP network.

1.11.15 AnyNet Advantages

With the IBM AnyNet Product Family of access nodes and multiprotocol gateways, customers can access new types of applications through their existing network protocol, or they can consolidate multiple networks while protecting the investment of existing applications.

When AnyNet is used to run an application over a networking protocol that it wasn′t designed to run over (a nonnative protocol), the application data transfer can take advantage of features of the underlying network protocol.

1.11.15.1 Advantages for Sockets Applications Running over SNA

1. **Performance:** TCP/IP Sockets applications benefit from the steady throughput and predictable response time of SNA networks. Performance tests on AIX/6000, OS/2, OS/400 and MVS have shown that when file sizes are 8 KB or greater, Sockets applications running over SNA may outperform Sockets applications running over native TCP/IP.
2. **Traffic prioritization:** The configuration of AnyNet Sockets over SNA allows the association of class of service (COS) and priority for well known TCP/IP applications such as FTP and Telnet. A specific mode name is associated with an application, based on port number, so Telnet traffic is configured to have a higher priority than FTP traffic.
3. **Administrative simplification - no need to subnet:** One Socket over SNA subnet is sufficient to assign host addresses across an enterprise, regardless of LAN/WAN topology. Using Sockets over SNA, IP addresses get mapped to LU Names, and since LU Names are topology independent, the IP subnet no longer is unique to a bridged LAN. There may be administrative reasons to have multiple Sockets over SNA IP subnets, but they aren′t required for each LAN.
4. **IP address location independence:** A user can move their Sockets over SNA access node from one location to another without needing a new IP address. This is a side benefit of item number three and would benefit laptops.
5. **Session-level security:** This allows verification of the identity of the partner node before session setup using a verification function.

   It′s difficult in native IP to only allow specific users to attach to a server. Routers are used to block one subnet from reaching another, but it′s not easy to do on an individual host-by-host basis (and requires coordination with the router owner). Using session-level security a user can ensure that only those IP users that should have access to a machine do have access.
6. **Data compression:** This reduces the amount of data being exchanged between partners, thus improving response time and reducing traffic over the network.

1.11.15.2 Advantages for NetBIOS Applications Running over SNA

1. **Communications Manager/2 connectivity options:** Using SNA networks to transport NetBIOS application data enables connectivity between NetBIOS applications anywhere on the WAN. When NetBIOS applications are run natively over NetBIOS, they can only communicate within a bridged LAN environment. A NetBIOS application running on Communications Manager/2 can connect to a WAN in the following ways:
   - Synchronous data link control (SDLC)
   - Integrated-services digital network (ISDN)
   - SNA phone modem support
   - Frame relay (with RouteXpander/2)
   - LAN connections such as Ethernet and token-ring
2. **Session-level security:** This verifies the identity of the partner node before session setup using a verification function.

   You may want to use this security for dial-in users as NetBIOS has no inherent security.
3. **SNA Data compression:** This reduces the amount of data being exchanged between applications, thus improving response time and reducing traffic over the network.
4. **Elimination of NetBIOS protocol broadcasts:** NetBIOS over SNA eliminates the broadcasting that is built into the NetBIOS protocol, easing traffic on networks. It cannot eliminate broadcasting done by applications, but it directs those broadcasts only to the machines that will actually receive them instead of forcing every machine on the network to read and discard each broadcast message.
5. **Performance:** NetBIOS applications benefit from the steady throughput and predictable response time of SNA networks. Elimination of the NetBIOS protocol's usual wait for acknowledgment after sending each record can improve throughput even more compared to NetBIOS (on the same speed data links).
6. **Control of network resources:** Each NetBIOS over SNA connection between applications is a separate SNA session that can be managed using all the traditional SNA network management facilities for resource control.
7. **End-to-end timers eliminated:** NetBIOS over SNA replaces the usual NetBIOS end-to-end timers (that can cause problems with very large bridged LANs) with traditional SNA reliable connections.
8. **Support for large datagrams:** Native NetBIOS datagrams are 512 bytes or the length of the transmit buffer. NetBIOS over SNA will support datagrams from 512 bytes up to 30 KB.

1.11.15.3 Advantages for SNA Applications Running over TCP/IP

1. **Non-disruptive session rerouting:** Link failures don't result in session failures.
2. **Connectivity:** TCP/IP can give you access to the worldwide Internet. This would allow SNA applications on AnyNet to have wider connectivity.

1.11.15.4 Advantages for IPX Applications Running over SNA

1. **Elimination of IPX protocol broadcasts:** IPX over SNA Gateway protects the SNA backbone by filtering IPX broadcasts and caching names.
2. **Traffic prioritization:** The configuration of AnyNet IPX over SNA Gateway for OS/2 allows the association of class of service (COS) and priority to gateway-to-gateway links. A specific mode name is associated with all traffic going between a gateway pair.
3. **Data compression:** IPX over SNA reduces the amount of data being exchanged between partners, thus improving response time and providing higher data rates with a lower cost network.

Chapter 2. Bridges and Routers

An internetwork design must ensure the correct choice of bridges and routers as intermediate nodes, and the exploitation of their differing capabilities. Bridges and routers must be fully understood in order to do this.

The following sections provide an overview of the capabilities of these devices and descriptions of some products.

2.1 Bridges

Bridges act as data link layer relays between LANs.

Figure 3 shows that a bridge implements the physical and data link layers of the OSI Reference Model.

A bridge participates as a device on the networks to which it is attached, exchanges information with devices on those networks, and forwards information between the networks selectively through the MAC address.

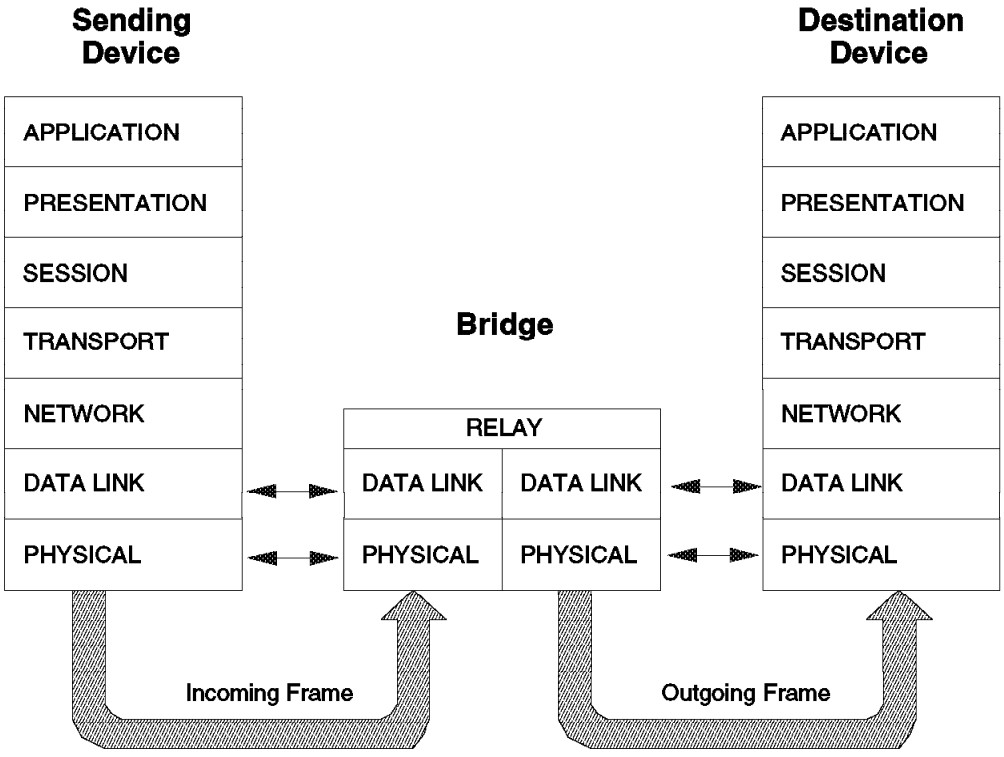

Figure 3. Bridge Implementation

There are four types of bridges and they are classified by their hardware and software capabilities.

2.1.1 Simple Bridges

Simple Bridges consist of two or more linked network interfaces connecting local area networks. Bridges interconnect separate local area networks (LANs) by relaying data frames between the separate MAC (media access control) entities of the bridged LANs.

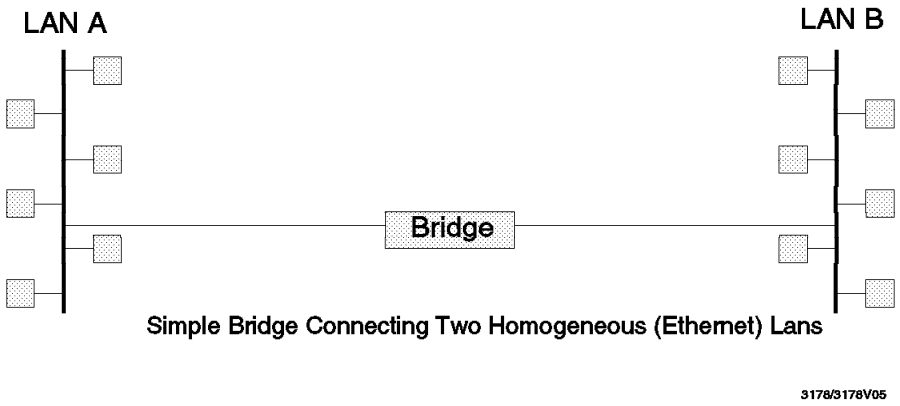

LANs

Figure 4. Example of Simple Bridge Connecting Two Homogeneous (Ethernet)

The main functions of a simple bridge may be summarized as follows:

- The bridge reads all data frames transmitted on LAN A and receives those addressed to LAN B. Simple bridges make no changes to the content or format of the data frames that they receive. They also do not encapsulate frames with any additional headers. Most simple bridges contain routing addressing and routing intelligence. At a minimum, the bridge must know which addresses are on each connected network so that it can know which frames to pass on.
- The bridge retransmits the data frames addressed to LAN B using the MAC protocol for that LAN. Bridges should have enough buffer space to meet peak data traffic demands since data frames may arrive faster than the bridge can transmit them.
- The bridge does the same for LAN B-to-LAN A data frame traffic.

2.1.2 Complex Bridges

Complex Bridges carry out more sophisticated functions than simple bridges. These functions may include the bridge maintaining status information on the other bridges. This information includes the communication path cost as well as the number of hops required to reach each connected network. Periodic exchanges of information between bridges update all bridge information. These types of exchanges allow for dynamic routing between bridges.

Complex bridges can also modify frames and recognize and transmit packets from different LAN technologies (for example, token-ring and Ethernet). In this case the bridge is sometimes referred to as a translational bridge.

The Adaptive Source-Routing Transparent (ASRT) Bridge is the IBM 2210′s implementation of bridge technology. The ASRT Bridge is a collection of software

components capable of several of the bridging options just described and more. All of these functions are explained in greater detail later in this chapter.

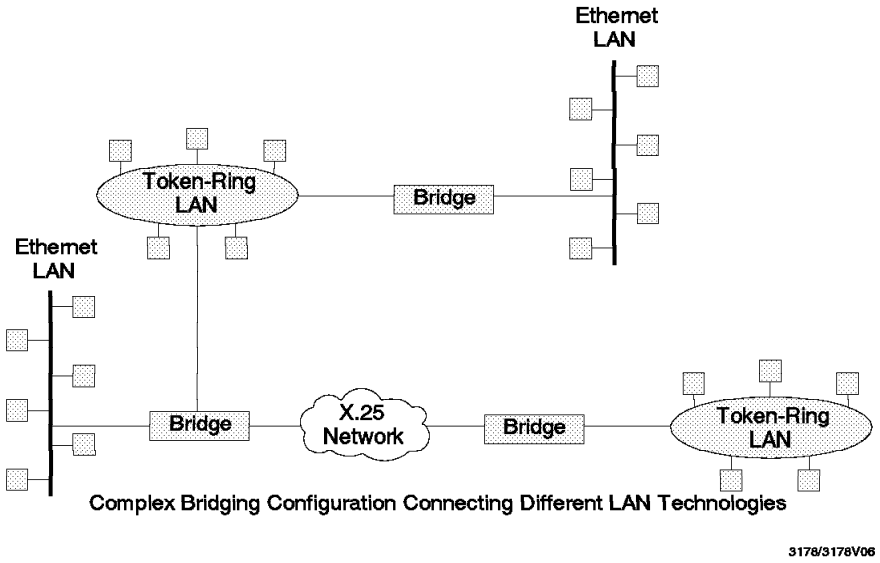

Figure 5. Example of Complex Bridge Connecting Different LAN Technologies

2.1.3 Local Bridges

Local Bridges provide connections between several LAN segments in the same geographical area. An example of this would be a bridge used to connect the various LANs located in your company's main headquarters.

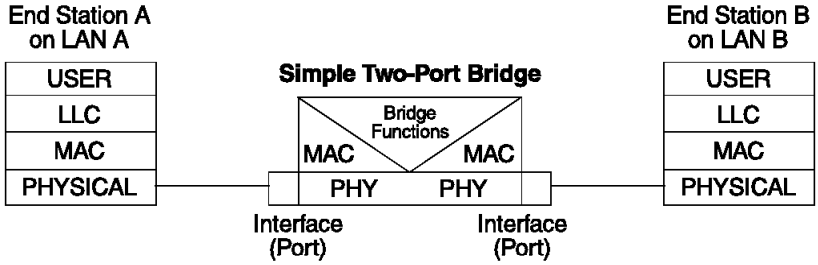

Figure 6. Example of Local Bridge Connecting Two LANs

2.1.4 Remote Bridges

Remote Bridges connect multiple LAN segments in different geographical areas. An example of this would be bridges used to connect LANs located in your company's main headquarters to LANs in other branch offices around the country. Because of the geographical differences, this configuration moves from a local area network configuration to a wide area network (WAN) configuration.

Remote bridges can differ from local bridges in several ways. One major difference is found in the speed in which data is transmitted. WAN connections

are generally slower than LAN connections. This difference in speed can make quite a difference when running time-sensitive applications. Another difference is found in the physical way that remote and local bridges are connected to LANs. In local bridges, the connections are made through local cabling media (for example, Ethernet). Remote bridge connections are made over the serial lines.

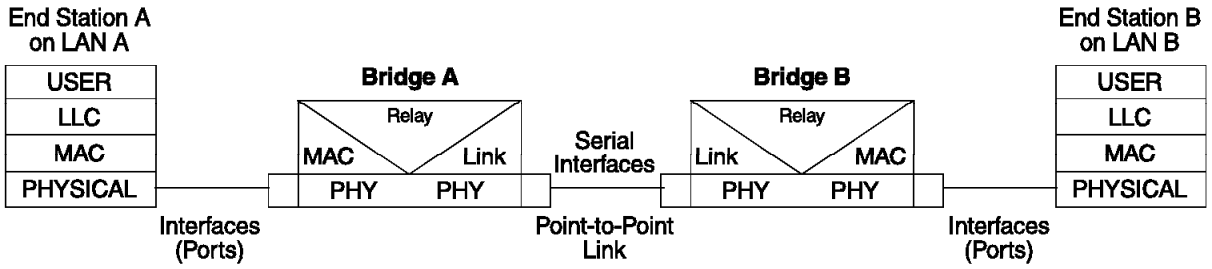

Figure 7. Example of Remote Bridge over a Point-to-Point Protocol (PPP)

2.1.5 Bridging Methods

There are two primary methods of bridging:

- Transparent bridging (mainly used with Ethernet LANs), also called spanning tree bridging (STB)
- Source-route bridging (SRB) (used in 802.5 LANs)

Then, from these two primary methods of bridging, there are other methods listed as follows:

- Source-route transparent bridging (SRT)
- Source-route - translational bridging (SR-TB)
- Tunnel bridge (IP encapsulation)

All of these bridging methods are supported by the IBM 2210.

In the following topics, we provide a summary description of these bridging methods.

For more information about the bridging methods, please refer to *Nways MRNS V1R2 Protocol Configuration and Monitoring Reference*, SC30-3680.

2.1.5.1 Transparent Bridging (STB)

A *transparent bridge* is also called a spanning tree bridge (STB).

Transparent bridging is normally used to connect LAN segments. It is specified in the ISO 8802-1 standard.

This form of bridging could also be used for connection of token-ring LAN segments, although this is not common.

Transparent bridging is based on the principle that a sending device can transmit a frame to a receiving device on a LAN network without having any knowledge of the location of, or the path to, that receiving device.

Transparent bridges within a network are responsible for forwarding the frame to the correct destination, making the determination of whether a frame should be forwarded based on MAC sub-layer destination address.

Transparent bridges achieve this by building and maintaining a *filtering database* that acts as a *forwarding table* for received frames. They build their database by copying all frames from the LANs to which they are attached and learning the location of devices by inspecting the MAC sub-layer *source address* in each received frame.

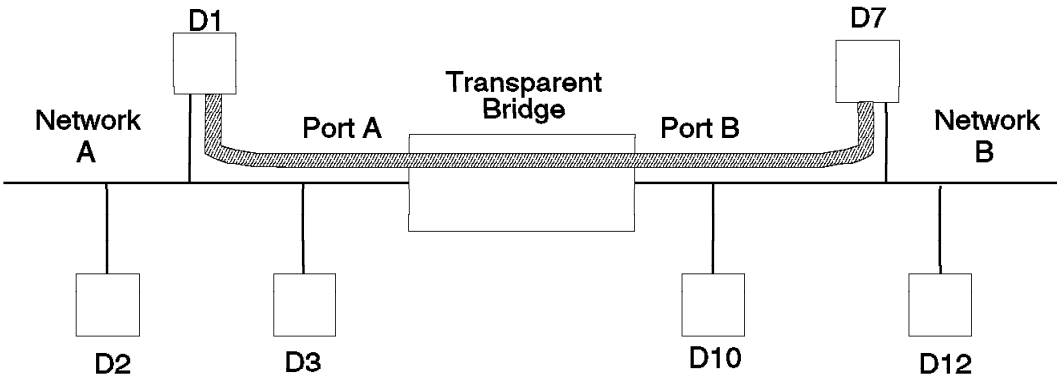

Filtering Database

Address	Port
D1	A
D2	A
D3	A
D7	B
D10	B
D12	B

Figure 8. Transparent Bridging

Figure 8 illustrates how a transparent bridge will build up its filtering database. When the bridge receives a frame from device D1 on port A, it learns that D1 is reached via the LAN on port A. Similarly, if a frame arrives from device D7 on port B, it learns that D7 is reached via the LAN on port B.

For each new source address the bridge sees on the LAN, it adds an additional entry in its database. In time a full picture is built up of all devices on the two LANs and via which port they are reached.

The bridge uses its filtering database to determine if an incoming frame should be forwarded or discarded. This is done by examining the MAC sub-layer *destination address* of each frame and comparing it to the list of addresses in the filtering database:

- If the destination address is not in the database, the frame is forwarded on each port except the receiving port.

- If the destination address is in the database and the frame was received on a port associated with the address, the frame is discarded.
- If the destination address is in the database and the frame was received on a port not associated with the address, the frame is forwarded to the associated port for this destination address in the database.

Transparent bridges require that there be only a *single* active path between any two LANs in an internetwork. This requirement is to ensure that frames do not loop in such a way that they are seen on both ports of a bridge. If this happens, the bridge will be unable to forward the frames correctly to their destination.

Transparent bridges support and use spanning tree protocol, which ensures a loop-free topology between all the transparent bridges within the network.

2.1.5.2 Source-Route Bridging (SRB)

Source-route bridging is implemented by IBM and compatible bridge products for use over token-ring LAN segments.

Source-routing requires a sending device to specify the path that should be taken by a frame across an internetwork, rather than allowing the decision to be made by individual bridges. To do this a sending device must determine the best path to a destination and include it in all frames to that destination. The best path to a destination is found using a discovery process, one implementation of which is described in this section.

Figure 9 on page 43 shows how the routing information field is used to define a route through an internetwork between end nodes D1 and D7.

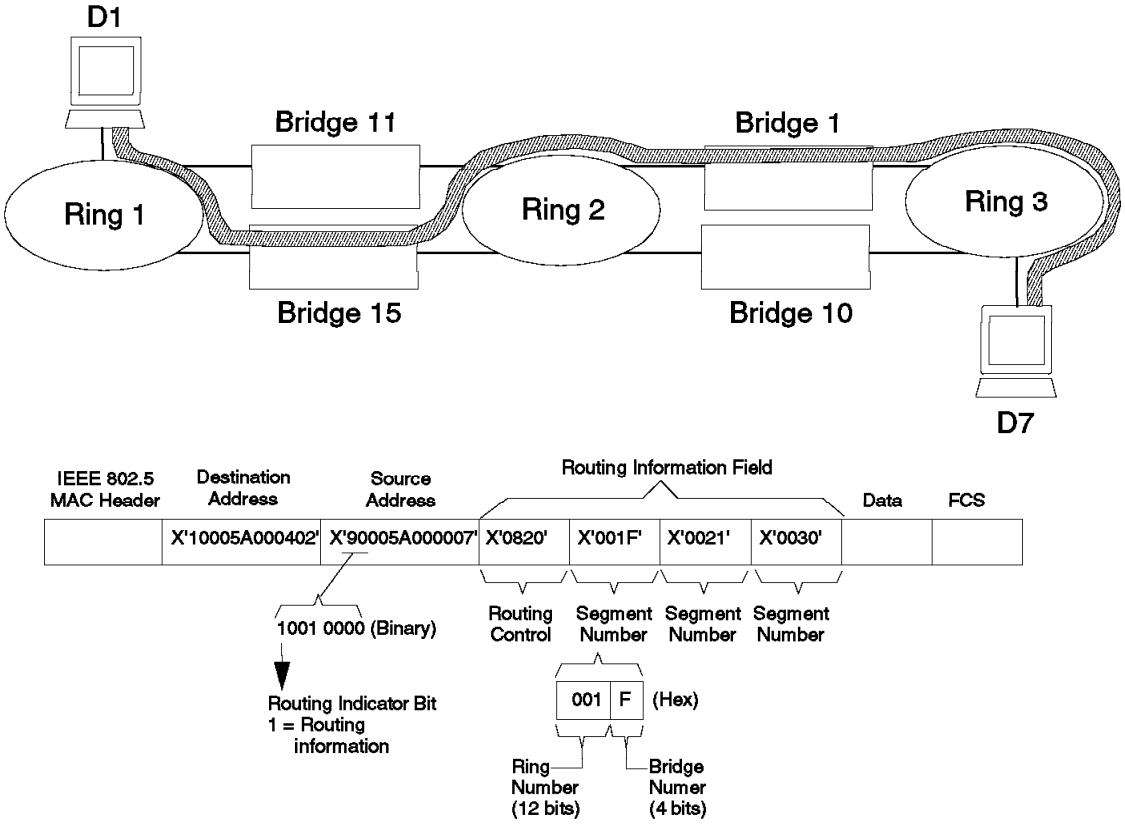

Figure 9. Source-Route Bridging

A sending device sends a discovery frame to the intended destination device marked single-route broadcast. Bridges in a token-ring internetwork should be configured using the token-ring spanning tree algorithm to permit only one path for single-route broadcast frames between devices. The destination device should therefore receive only a single copy of the discovery frame.

The destination device responds to the discovery frame with a discovery response frame, marked all-routes broadcast. This will contain the most significant bit (the route information indicator, also called RII) set in the source MAC address field, and an entry in the routing information field (RIF). This will initially contain zero in the bridge number field, and the number of the networks to which the destination device is attached in the segment number field.

The discovery response frame, because it is marked all-routes broadcast, will pass through all bridges on its way back to the original sending device. Each bridge that the frame passes through must insert its bridge number and LAN segment, and hence the frames that return to the original sending device contain the routes they have taken through the bridged internetwork.

The routing information field can currently only hold data for about seven bridges and eight LAN segments. If a frame is received by a bridge with this field full, it is discarded. This limits the number of bridge *hops* in the network to seven, and consequently the maximum size of source-route bridged internetworks.

The original sending device therefore receives one or more discovery response frames. These frames contain routing control and bridge and LAN segment numbers in their routing information fields. The routing control field indicates the number of bridge/LAN segments in the routing information field and also the maximum frame size that is supported by the route.

The sending device can now select the best route to use through the internetwork to reach the destination device. Current implementations select the route in the *first* received discovery response frame (the fastest path at the time of the discovery process), although the architecture allows route selection based on other criteria, for example, maximum frame size supported by the route.

2.1.5.3 Source-Route Transparent Bridging (SRT)

The IEEE 802.1 committee identified the need for source-route bridges to interoperate with transparent bridges in the same internetwork.

A *source-route transparent bridge (SRT)* standard has been defined to achieve this goal.

The principle behind SRT bridges is very simple. An SRT bridge inspects all received frames and looks for the presence of the routing information indicator (RII) and the routing information field (RIF). If these fields are present, the SRT bridge uses them and acts as a source-route bridge. If not, the SRT bridge operates in transparent bridge mode and forwards frames based on their MAC sublayer destination address and its associated entry in the filtering database.

The source-route transparent bridge does not allow source-route bridge devices to communicate with transparent bridge devices. SRT bridge is the capability for its interfaces to understand both source-route bridging and transparent bridging devices. But an SRT bridge will never translate source-route bridge frames into transparent bridge frames, and vice versa.

Figure 10 on page 45 shows you that frames with RII=1 are forwarded with RII=1, and frames with RII=0 are forwarded with RII=0.

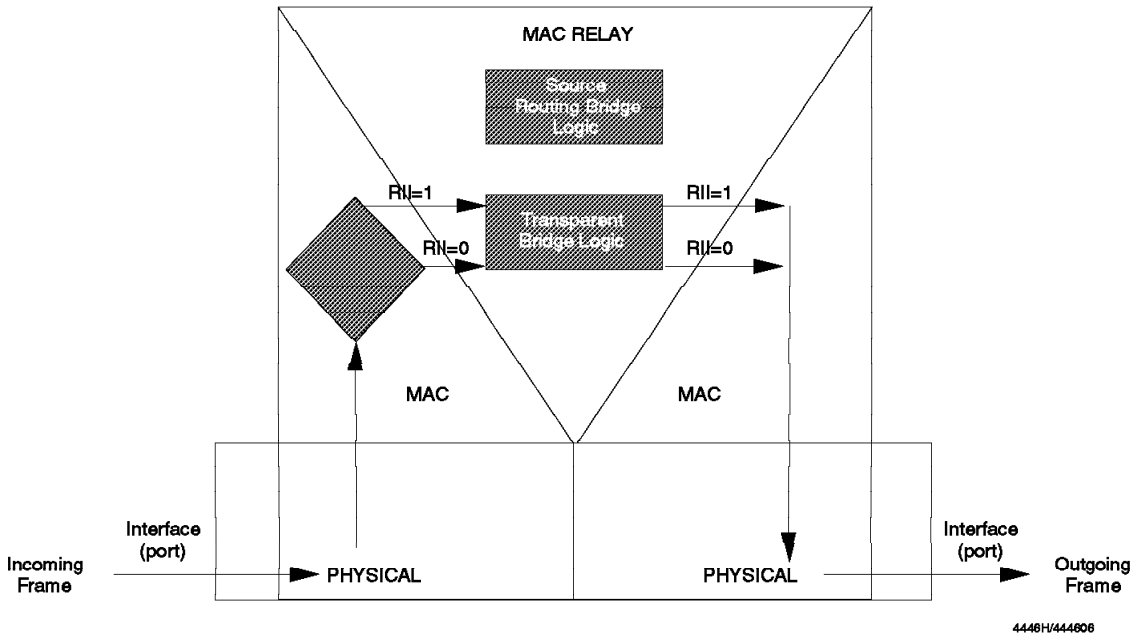

Figure 10. Source-route Transparent Bridging

2.1.5.4 Source-Route - Translational Bridge (SR-TB)

Source-Route - Translational Bridge (SR-TB) is not an ISO standard definition. However, more and more bridges are implementing the SR-TB because of the need to interconnect source-route bridge domain with transparent bridge domain.

The goal of the source-route - translational bridge is to translate source-route bridge frame into a transparent bridge frame, and vice versa.

The SR-TB bridges have to change the MAC layer protocol from (or to) Ethernet protocol to (or from) token-ring protocol. Actually, regarding the ISO bridge definition, this translation does not belong to a bridge. But it is implemented in a lot of bridges, in order to be able to interconnect source-route bridge domain and transparent bridge domain regardless of the protocol of the upper layer.

Figure 11 on page 46 shows you that the SR-TB allows an SRB device with RII=1 to communicate with an STB device (RII=0).

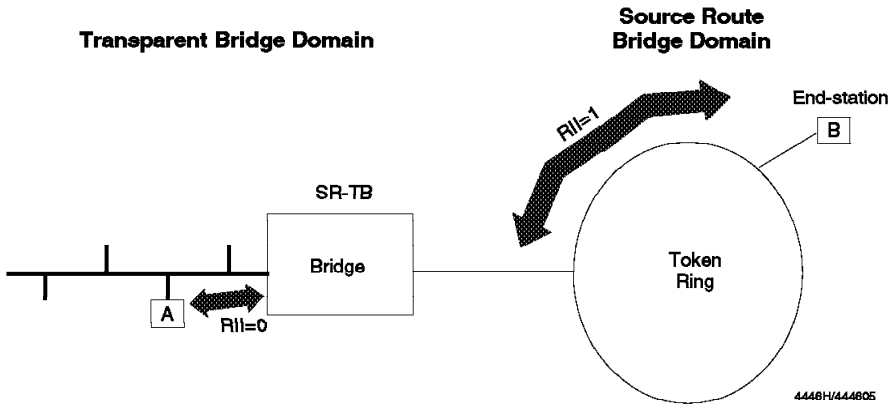

Figure 11. Source-Route - Translational Bridging

2.1.5.5 Tunnel Bridge

The tunnel bridge allows source-route bridge domains or transparent bridge domains to communicate across an IP network.

The tunnel bridge receives bridged frames from its source-route bridge or transparent bridge domain. The frames are encapsulated into IP datagrams that are sent to the destination IP address. These IP datagrams are routed in the IP network as are other IP datagrams, with the IP rules.

The destination IP address is actually another bridge implementing the tunnel bridge feature. This target bridge removes the IP envelope from these IP datagrams making them source-route bridge or transparent bridge frames. Then the target bridge sends these frames to its source-route bridge domain or transparent bridge domain in the same way that other bridged frames are sent.

Figure 12 on page 47 shows you an example of the tunnel bridge implementation.

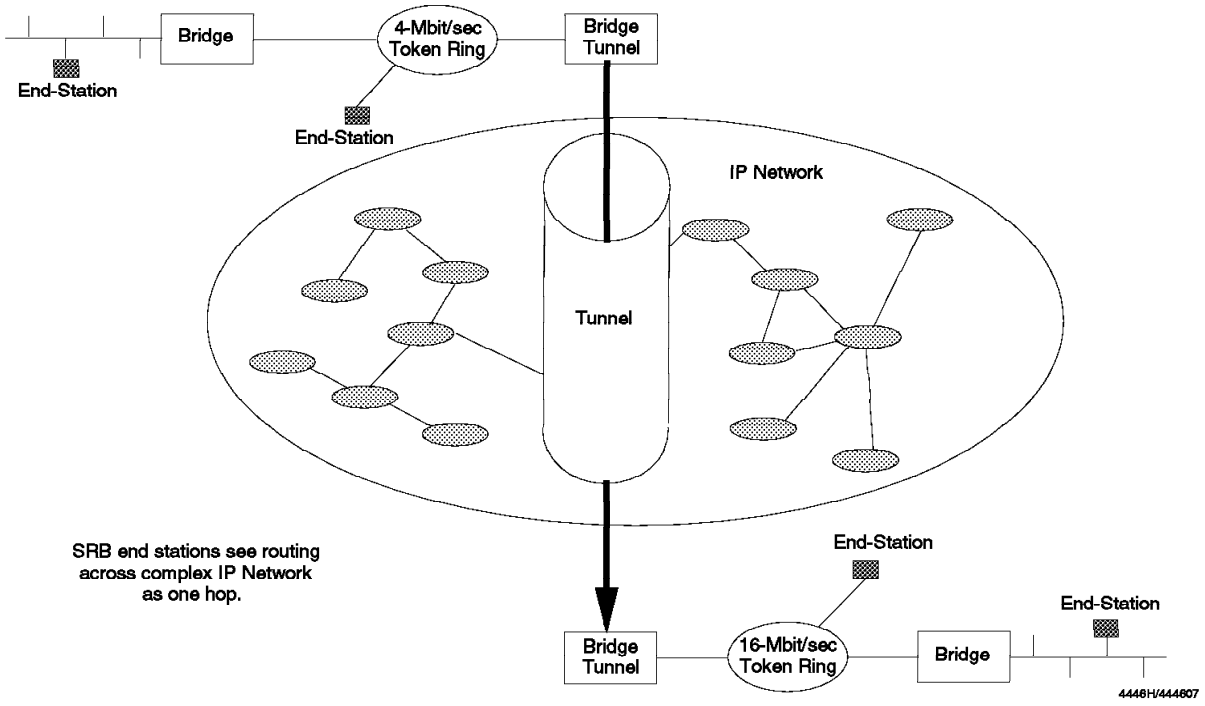

Figure 12. Tunnel Bridging

With tunnel bridging, as far as the source-route bridge is concerned, the IP network is seen as a single LAN segment, regardless of the complexity of the IP network. Then it adds only one hop to cross this IP network.

The number of hops from the source device to the source IP tunnel bridge, plus one hop to cross the IP network, plus the number of hops from the destination IP tunnel bridge to the destination device, must not exceed the 7-hop count limitation of the source-route bridge implementation.

2.2 Routers

Routers act as network layer relays between networks.

IBM offers two Nways router product solutions:

- IBM Nways 6611 Network Processor
- IBM 2210 Nways Multiprotocol Router

While bridges are normally restricted to connecting LANs within an internetwork, routers have the capability of connecting networks of different types (for example: point-to-point, multi-access broadcast and multi-access non-broadcast).

Figure 13 on page 48 shows that routers implement physical, data link and network layers of the OSI Reference Model.

A router participates as a device on each attached network and exchanges information with devices on those networks. These are *end node* capabilities.

A router has the additional capability of exchanging information with other routers and with end nodes on remote networks as long as they support the same network layer protocol.

Therefore, routers are protocol-dependent, unlike bridges.

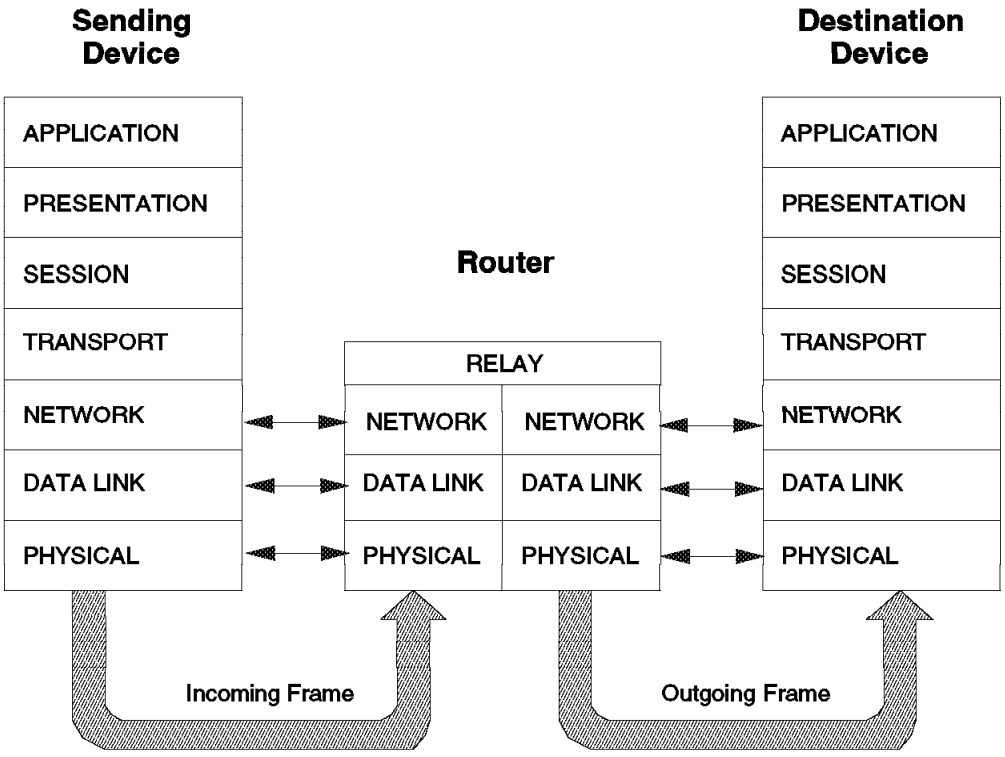

Figure 13. Router Implementation

2.2.1 Packet

Packets (also known as *frames*) are defined as chunk data which has been packaged, addressed, and sent into the network towards its destination much as a letter is placed into an envelope, addressed, and dropped into a mailbox.

The data may represent an entire message or a segment of a message, depending on media or device limitations on the amount of information that may be *enveloped*.

Routers forward packets based on information held in the network layer headers of packets they receive on their attached networks. Network layer headers contain address information (typically consisting of a network identifier and a host identifier) and associated control to allow the packets to be routed to their correct destination. Packets that have no network layer header, or that have a network layer header for a protocol not supported by a particular router, are discarded.

There are a number of network layer protocols, each of which has its own network layer header format for addressing and control. Some of these protocols are vendor proprietary, while others are public domain or ISO standards.

More information on protocols can be found in Chapter 4 of *Local Area Network Concepts and Products: LAN Architecture*, SG24-4753.

Examples of network layer protocols include:

- Internet Protocol (IP) defined in RFC 791
- Internet Packet Exchange (IPX) developed by Novell
- Internetwork Datagram Protocol (IDP) developed by Xerox
- Datagram Delivery Protocol (DDP) developed by Apple

A router must be configured with the network layer address of each of its network connections. When it receives an incoming frame with a compatible network layer header, it determines whether the destination address is on the same network. If it is, then the frame is discarded. Otherwise, the router forwards the frame to the destination device (if on a network attached to the router), or to the next router in the path to the destination device.

In order to do this, a router must maintain routing tables containing information about the next router in the path to every reachable destination in the internetwork. The process for doing this includes two stages:

- It must acquire route information.
- It must determine the best routes to insert into its routing table.

Route information can be acquired by manual configuration (these are called *static* routes), or can be learned automatically from other routers using routing table maintenance protocols (these are called *dynamic* routes).

2.2.2 Routing Table Maintenance Protocols

Routing tables are held by routers to define paths to destinations in an internetwork. They are normally created and maintained by *routing table maintenance protocols*.

Routing table maintenance protocols that maintain routing tables for IP networks are normally referred to as *IP routing protocols.* This term is used throughout this section.

Each routing table maintenance protocol uses a different algorithm to determine when new routes are available and what the best routes are that should be added to the routing table. These algorithms compare routes on the basis of some measurement of *distance* or *cost* associated with a route.

Routes consist of three elements: a destination, the identity of the next hop in the path to the destination and the distance or cost of the complete route to the destination. The distance or cost of the route is referred to as a *metric*.

In some protocols the metric is purely the number of hops to the destination; in others, it is a measure of the number of seconds to reach the destination. Modern protocols allow the cost of each network link to be set during network design. In this case the network link cost is normally an indication of the speed

of the network, although it could be used to indicate other types of cost such as the lease cost of a telecommunications circuit.

Each protocol suite (such as TCP/IP, DECnet, AppleTalk) has its own routing table maintenance protocol. Some, such as TCP/IP, offer alternates.

The routing table maintenance protocols for each suite are:

TCP/IP

- Routing Information Protocol (RIP)
- Open Shortest Path First (OSPF)
- HELLO
- Exterior Gateway Protocol (EGP)
- Border Gateway Protocol (BGP)
- Static Routes

NetWare

- Routing Information Protocol (RIP)

Xerox Network Service (XNS)

- Routing Information Protocol (RIP)

AppleTalk

- Routing Table Maintenance Protocol (RTMP)

DECnet

- DECnet phase IV

Banyan VIrtual NEtworking System Protocol (VINES)

- VINES Routing Update Protocol (VINES RTP)

TCP/IP, NetWare and XNS all have routing information protocols. All derive from the XNS routing information protocol, but all are *different*.

2.3 Performance in Routers

Performance has become one of the *hot* topics in the multiprotocol router market today. Therefore, performance results must provide sufficient information so valid comparisons are made for the network that is to be built. The definition of performance in a multiprotocol router network has required revision as the router network has evolved. When the device was a single protocol, single media machine, performance could easily be reported in frames per second. This speed specification was a sufficient performance rating given the environment was simple and well understood. As multiple media interfaces and multiple connectionless protocols were added, performance was presented as maximum speed for the specific protocol or the generic maximum throughput for the box. This led to some confusion since the type of frames and environmental modifiers were not specified.

In today′s network, where the router is used for LAN interconnection and WAN transport for connectionless and connection-oriented protocols, performance reporting must be standardized to allow for comparison in both the bridged and routed environments. This standardization must result in information valid for a

production environment where multiple protocols are enabled in the router, frames of router exchange information flow, and filters are activated.

Many vendors still quote total theoretical box throughput for their boxes as sufficient input for the router decision. These reports often come without the specification of the environment in which the performance test was made. The test bed may have no likeness to production environments as the test may be chosen to optimize for maximum speed. Without a controlled environment to record the speed, these recordings are not easily compared to other vendors′ performance information. To offset this, the Internet Engineering Task Force (IEFT) formed the Benchmark Methodology Working Group (BMWG) to establish guidelines for performance testing. Terminology, test environment (frame size, router information exchange, single and multiple protocol mixture, etc) and, in the future, packet content are specified to give the marketplace data that are compared and contrasted when making purchase decisions.

The Benchmark Methodology Working Group (BMWG) of the Internet Engineering Task Force (IEFT) defined the following RFCs:

- Benchmarking Terminology (RFC1242) - Available
- Benchmarking Methodology - Draft
- Benchmarking Methodology: Test Frame Formats - Draft

In spite of the attention performance receives, it is only one of several important selection criteria and must always be weighed with the other critical determining factors in the network, such as:

- Required functions and protocols supported
- Reliability and availability
- Network management capability
- Quality and level of service provided by vendor
- Total cost of ownership (hardware, software, service and support)

2.3.1 Performance in a Router Network

A customer perceives the network performance by the response time seen for the application. The IBM 2210 and 6611 Network Processor are some examples of the many contributors to response time, and their performance in the customer environment is really defined by their contribution to the overall system response time.

Performance varies by protocol and packet size in the router network. Kevin Tolly, in an article in *Network World* (Aug. 10, 1992), states that “traffic on a typical corporate internet consists of larger frames, such as those used in Novell, Inc. NetWare or Microsoft Corp. LAN Manager file transfers. Measuring performance using 64-byte packets is largely of academic interest, since virtually no applications use such small packets.”

Performance information then must be specified in such a way that expectations for a given LAN or WAN interface are factored into the network design. WAN speed is often the limiting factor in throughput or response time measurements. Traffic patterns of the protocols being supported are also a factor when designing a network to support the necessary response time for an application.

Another performance consideration that is added to the aggregate of the performance and capacity is the overhead of router-to-router exchanges. As each protocol is added, a new flow of exchanges between the routers must occur. Network management is another internal flow that must be given consideration when determining the necessary box performance to support the system performance requirements.

Routers were designed with connectionless-oriented protocols in mind. Each data packet finds its way through the network with the information available in its header. Path information is maintained via router information exchanges outlined in protocols such as TCP/IP′s RIP and OSPF. Broadcasts are used to learn locations of other end stations, services, and zones. Today′s connection-oriented protocols are encapsulated to take on these characteristics. Performance under this encapsulation may be more comparable from the end stations point of reference, such as response time, as each implementation provides unique function in its support of connection-oriented protocols ranging from spoofing to broadcast and flow control.

Performance is limited by the hardware on which the function is being tested. Each interface must be tested with the protocols and function supported on it. After looking at the card level performance, box level performance information becomes one of the entities in the overall network performance information. Only then can performance be estimated for the network and bottlenecks pinpointed. End-to-end response time is the last of the performance information that can be gathered.

The primary metrics used to characterize the performance of a device are:

- **Responsiveness** is the time interval between an input and the corresponding output.
- **Productivity** is the amount of work done during a time interval.
- **Utilization** is the ratio of *busy time* to elapsed time.

Each is used to describe some aspect of how well a device processes jobs (units of work to be performed by the system). For routers, a job is usually synonymous with a *routed packet*.

2.3.1.1 Responsiveness

Response time is defined as time between the arrival and departure of a given job. The term implies a non-deterministic arrival rate sampled over time and typically refers to the *turn-around time* experienced by the end user of a system as depicted in Figure 14 on page 53. This is calculated as the time from the generation of a request (when the user presses the Enter key) to the time that the entire response is received. Note that total response time perceived by the user will include workstation (WS) latencies.

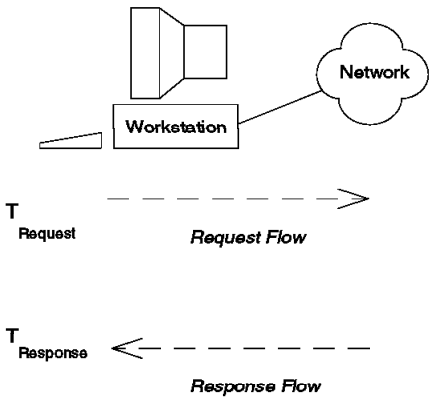

Figure 14. Interactive Transaction Flow

Clearly, total response time (RT) perceived by the end user is as follows:

$$RT_{User} = T_{Response} - T_{Request}$$

Latency is the delay a device imposes on the flow of data. This delay is a combination of software execution time (path length), hardware delays, and queueing (waiting) time due to contention for shared resources (processor, memory, bus, and so on).

Unloaded-box latency is simply a latency measurement for a single job with no other jobs in the system. Without additional work demands on the system, job latency will not include delays resulting from contention for various resource services; therefore, unloaded box latency is a *minimum* latency value.

These concepts are applied to a situation you experience in your drive to work each morning. If you leave home at 5 AM, your drive-time (latency) may approach its minimum theoretical value. On the other hand, leaving home at 7:30 AM puts you in a larger pool of cars contending for shared resources (roads) and you experience greater queueing delays (such as red lights) as well as other *slowdowns*.

2.3.1.2 Productivity

Throughput is the term generally used to describe mean throughput rate, which is an index of productivity defined as the number of jobs (which may be one of several measures of work) processed per unit time. Throughput rate is obtained by counting jobs leaving the system during a specified time interval. For those readers mathematically inclined, throughput may be stated as follows:

$$TP = \sum_{j=1}^{n} D_j / t$$

where **n** is the total jobs, **Dj** is the quantity of data associated with each job **j**, and **t** is the time interval.

When viewed within a system, device throughput is often controlled by an external apparatus, as one or more external limiters (for example, host, application, LAN, and remote link) impact performance with their bandwidth limitations.

Capacity generally refers to the theoretical maximum number of jobs processed per unit time and is independent of such external limiters as a LAN or remote link. The capacity of a resource (for example, a microprocessor) is calculated by dividing measured throughput by the utilization of that resource. However, it is important to use performance measurements for a moderately utilized resource to avoid inclusion of additional work demands imposed by other system components as their bandwidth limitations are approached. The resulting value is useful in determining reserves for future system expansion and possible performance enhancements.

2.3.1.3 Utilization

Utilization is the ratio of time a resource is busy to total elapsed time. It is a dimensionless quantity useful in identifying system *bottlenecks*. A resource that is always busy (no idle time) is said to be 100% utilized.

Customers are primarily concerned with high-level resource utilizations, such as router, bridge, host, and teleprocessing link utilizations. Engineers also have an interest in low-level utilizations of internal processors, buses, buffers, etc.

2.3.2 Relationship Between Response Time, Throughput, and Utilization

As the load on the system increases, there is a commensurate increase in throughput with changes in response time remaining imperceptible until some critical load, sometimes called the *knee value*, is reached. Moving beyond this knee value, two changes become apparent:

1. The slope of the throughput curve decreases, levels off, and may even begin to drop.
2. Response time increases dramatically.

This phenomenon is depicted in Figure 15 on page 55.

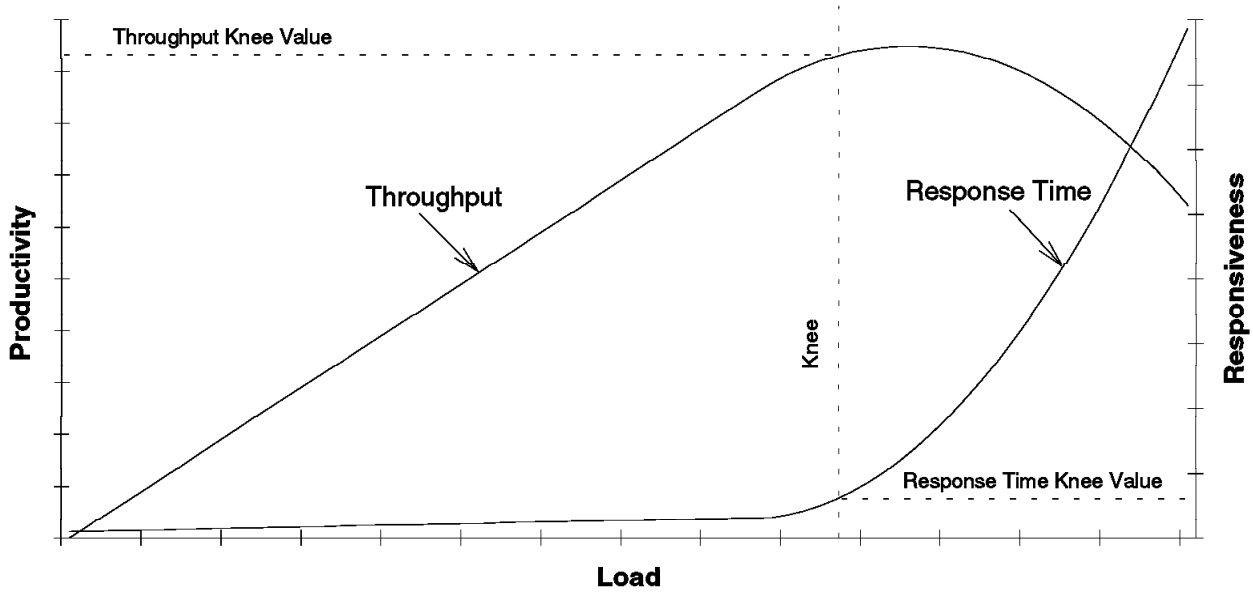

Figure 15. Example of Throughput and Response Time Curves

So what causes the knee in the response-time and throughput curves? To find an answer, we first need to perform a utilization analysis of internal resources comprising the device (processors, memory, buses, etc.) to identify the resource least able to perform its share of the work. At some threshold utilization, this resource exhibits signs of stress by not keeping pace with the load. Beyond this threshold, so many tasks are vying for a share of the resource that the work expended switching between them becomes an increasingly significant component of the total work demand. This limiting resource then establishes the *maximum useable capacity* of the device.

2.3.3 Performance Metrics

Now that the performance basics parameters are defined, we can see how these parameters must be verified in routers, when we are choosing anyone to implement a internetwork environment.

2.3.3.1 Router Latency

When considering a *store-and-forward* component of a router in a network system, the term *packet latency* describes the interval of time from *last byte in* to *first byte out* for a given packet. For interactive workloads, latency is a measure of the contribution to total system response time by a component within the system. Similarly, total system response time is the aggregation of the individual device latencies encountered by a job over the duration of a transaction, as shown in Figure 16 on page 56.

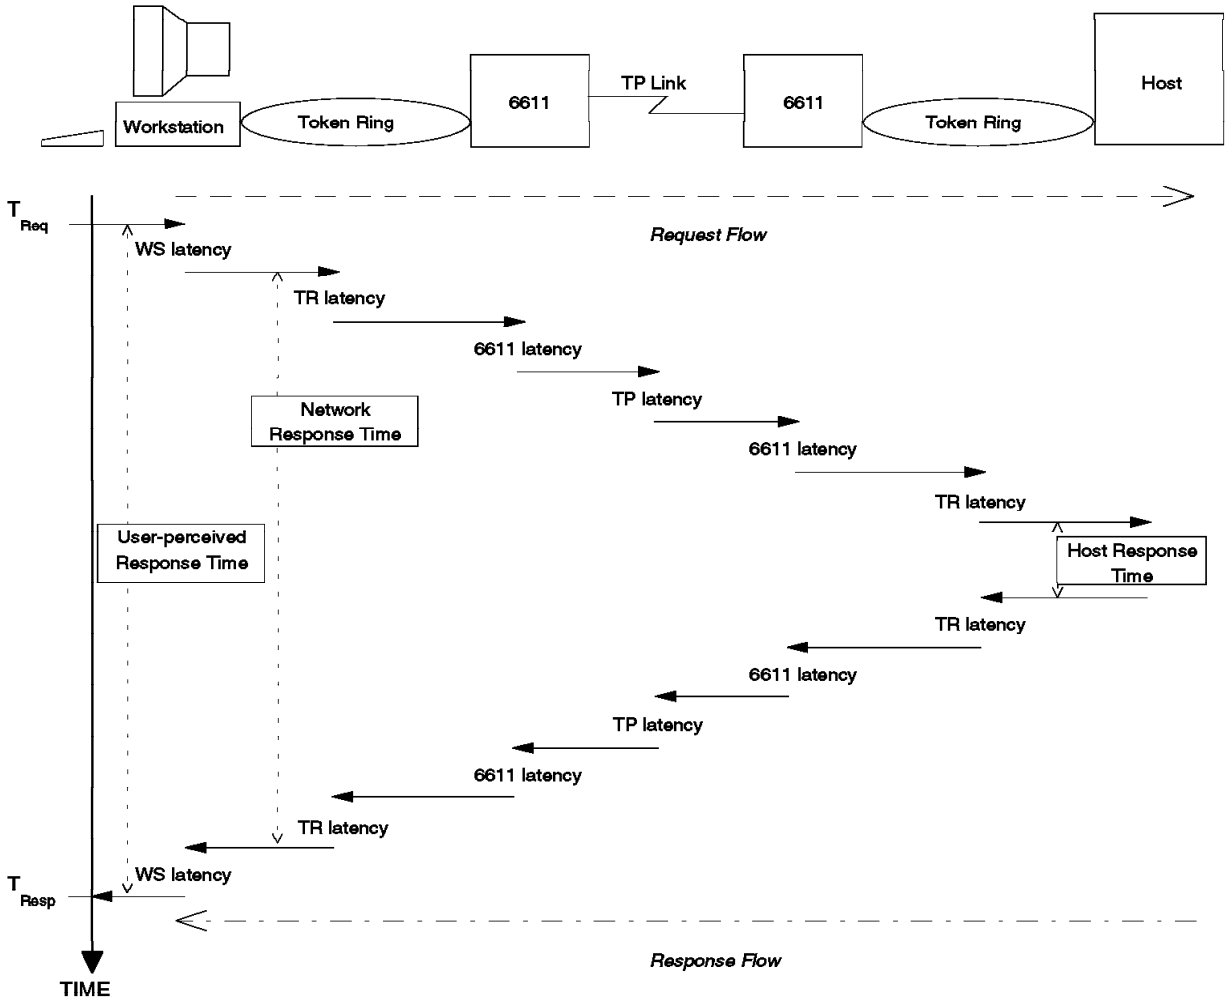

Figure 16. Contributors to Total System Response Time for Interactive Transactions Workloads

2.3.3.2 Router Capacity

When comparing router performance for interactive (small packet) workloads, throughput is generally expressed in thousands of packets per second (Kpps). Of course, throughput in Kpps is meaningful only if packet sizes are known. *When not stated, packet size is often assumed to be 64 bytes.*

On the other hand, when evaluating a router's file transfer (large frame) performance in a specific configuration, throughput may be stated best in terms of millions of bits per second (Mbps) to illuminate the router's ability to consume media bandwidth. For example, if the media is a 50 Mbps full duplex link and router throughput is 25 Mbps full duplex, we instantly understand that the router can utilize 50% of the link capacity.

When considering *connection-oriented protocols*, those involved may express throughput in terms of packets per second or transactions per second (tps). You should be careful to understand the definition of a transaction (there is no standard), which may include one or more packets *and* their respective acknowledgements. For example, an interactive request packet, the

request-acknowledgement packet, the response packet, and the response-acknowledgement packet may be lumped together as a single transaction. On the other hand, a vendor might choose to specify throughput in packets per second (including acknowledgement packets) to yield an inflated throughput specification.

Maximum raw throughput, however, does not provide sufficient information. We need to understand router behavior under worst-case scenarios, with the definition of worst-case depending on the type of media attachments.

The following are three metrics commonly used to characterize the effects of media bandwidth and packet loss on router performance:

- **Rated Throughput**

  For a given packet size, rated throughput is the maximum rate at which packets are forwarded without packet loss. This is probably the most interesting throughput measurement, since packet loss is extremely undesirable. As in the water-pump analogy, we want to know the maximum throughput rate without *spillage*.

- **Maximum Load Throughput**

  What happens if the media delivers packets faster than the router can handle them? For a given packet size, maximum load throughput is the throughput attained when the source media offers frames at its maximum rate, irregardless of packet loss.

- **Maximum Throughput**

  For a given packet size, maximum throughput is the highest throughput achieved across the range of input rates supported by the media, irregardless of packet loss. The rate difference between offered frames and forwarded frames is called the *Frame Loss Rate* (or, in keeping with the water-flow analogy, the spillage rate).

Figure 17 on page 58 illustrates these concepts for a fictitious router (packet size is constant).

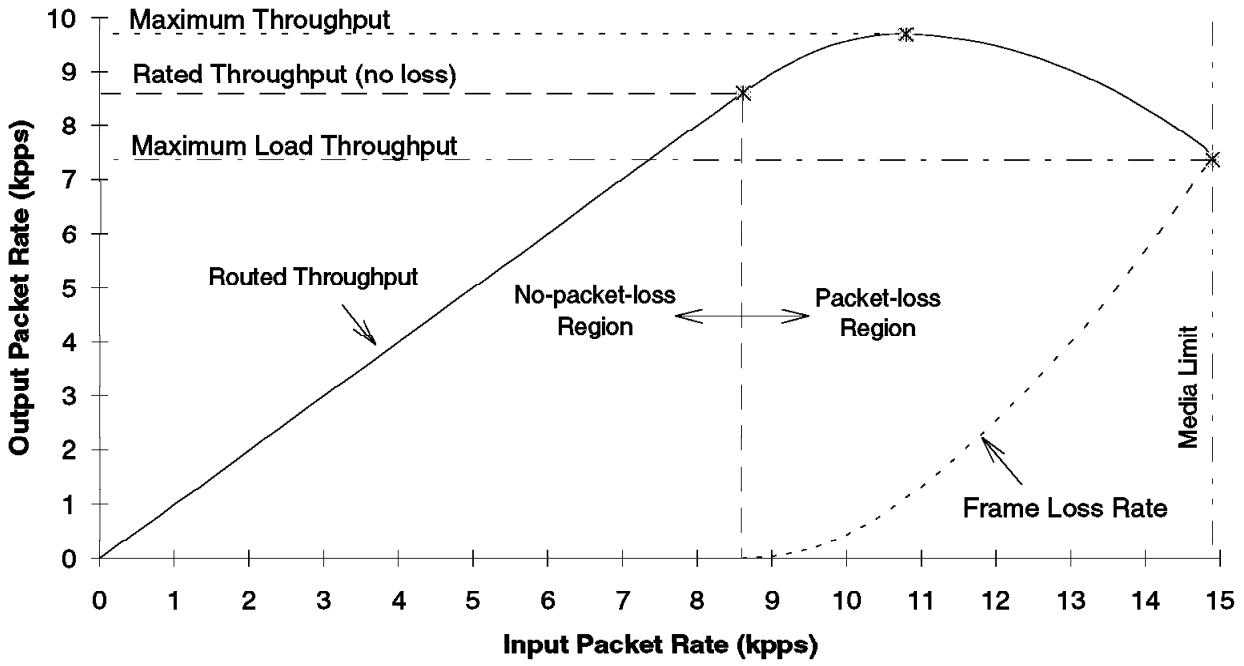

Figure 17. Example Router Throughput Curve

You might expect output to increase with input until 100% utilization is achieved and then to remain constant as input increases further. However, an overloaded router must allocate some resource to each discarded packet, causing the rate of successfully routed packets to drop as the input packet rate increases beyond the rated value.

2.3.3.3 Router Utilization

Low-level utilization concepts applied to router internals are interesting primarily to engineers and performance analysts; performance metrics is most interesting to customers who relate to response time and throughput. However, when viewed at the box level, a utilization specification indicates how hard a device must work to achieve a specified throughput. In certain scenarios, utilization may give an indication of the reserves available for future network expansion.

Customers are very interested in *link* utilization. Since links are often leased, it is important to the customer to maintain link utilization as high as possible to avoid paying for unused bandwidth.

2.3.3.4 Packet Processing Time

Packet Processing Time (PPT) is a performance metric unique to the network environment. It is defined as the amount of time that a processor devotes to processing a single packet, and is given by:

$$PPT = \frac{\text{CPU UTILIZATION}}{\text{PACKET THROUGHPUT}}$$

For most router designs, PPT is almost constant over practical ranges of utilizations and packet sizes.

2.4 To Route or to Bridge?

The above subsections outline the main characteristics of bridges and routers. In many situations it may be possible to use either within an internetwork. This subsection, however, attempts to highlight the key differences between the two types of intermediate nodes to help answer the key question: to route or to bridge?

2.4.1 Router Connections

Router Connections: Connecting at Layer 3 with a router allows connectivity and path selection between end stations located in distant geographical areas. Because of the variety of network and subnetwork configuration options available to you in large networks, connecting LANs through the network layer is usually the preferred method.

- Transparent bridges use only a subset of network topology at any one time as only a single path can exist at any time between two points in a bridged internetwork. Routers, however, can use the best path that exists between source and destination and can readily switch paths as better ones become available.
- Bridges can reconfigure to take into account changed topologies, but the process is much slower than for routers, which can adapt instantaneously. When bridges change topology, it is common for link layer timeouts to occur and sessions to be terminated. Routers are always the termination point for link layer connections, and hence, any router-to-router configuration changes do not affect link layer operation.
- Bridges offer no protection against large volumes of broadcast packets. In forwarding broadcast packets, bridges are only carrying out their normal function, but in doing so can impact internetwork performance and function. This is a particular problem with remote bridges where broadcasts have to traverse interbridge serial links.
- Bridges have to drop packets that are too large for their attached networks. Routers, because they support a network layer, have the capability of fragmenting packets to accommodate networks with a smaller maximum packet size.
- Bridges have no capability to provide congestion feedback to other bridges or to end nodes. This can lead to the need to discard packets with consequent impact on end system performance. Routers provide congestion feedback using the capabilities of the network layer protocol.
- Bridges do not distinguish between packets they forward. They offer no possibility, therefore, of prioritizing traffic based on protocol type.

2.4.2 Bridge Connections

Bridge Connections: Connecting at Layer Two with a bridge provides connectivity across a physical link. This connection is essentially *transparent* to the host connected on the network.

Note: Source-routing bridges are not considered completely transparent. See 2.1.5.2, “Source-Route Bridging (SRB)” on page 42.

The link layer maintains physical addressing schemes (vs logical at Layer 3), line discipline, topology reporting, error notification, flow control, and ordered delivery of data frames. Isolation from upper layer protocols is one of the advantages of bridging. Since bridges function at the link layer, they are not concerned with looking at the protocol information that occurs at the upper layers. This provides for lower processing overhead and fast communication of network layer protocol traffic. Because bridges are not concerned with Layer 3 information, they may also forward different types of protocol traffic (for example, IP or IPX) between two or more networks (as routers do).

Bridges can also filter frames based on Layer 2 fields. This means that the bridge may be configured to accept and forward only frames of a certain type or ones that originate from a particular network. This ability to configure filters is very useful for maintaining effective traffic flow.

- Bridges are *plug and play* and require little expertise to install.
- Bridges require little administrative overhead. Once installed they generally function with minimum attention.
- Bridges let you isolate specific network areas giving them less exposure to major network problems.
- Bridges eliminate node limitation. Local network traffic is not passed on to all of the other connected networks.
- Bridges are truly multiprotocol. They forward all packets and protocols irrespective of whether the protocols are routable or not.
- Bridges are generally transparent to end systems. Routers, however, require end systems to be configured with router addresses, or to run a routing protocol passively. These options may be impractical.
- Bridges generally have superior price/performance to routers.

To summarize the above, *bridges* are best used to provide convenient local connections within single site networks, possibly where there is limited technical and administrative support for the products. They are particularly suitable for environments where it is undesirable or impractical to configure end systems for operation with routers, or where the protocols in use are mainly non-routable.

Routers, on the other hand are functionally more robust, not suffering from the technical shortcomings of bridges such as LLC timeouts, broadcast susceptibility, possible packet loss and poor congestion control. They are able to support large network configurations both in loading and interface terms. They do, however, require more technical and administrative support. Routers, therefore, are most suitable for multi-site production networks, where a fully trained support staff is available.

Bridge routers (or *brouters*) are, of course, replacing conventional bridges and routers in the marketplace; this is because they combine the functions and hence the benefits of dedicated bridges and routers. Subject to costs, they are suitable as standard all-purpose intermediate nodes for most internetwork requirements. The IBM 6611 Network Processor and IBM 2210 are examples of these devices.

2.5 IBM 8229 LAN Bridge

The new IBM 8229 is a hardware bridge that gives a more powerful, cost-effective way to operate and manage local token-ring LANs connected to various types of networks. The 8229 is the successor to the IBM 8209, long a popular choice for interconnecting a token-ring LAN with another token-ring or an Ethernet LAN. It retains the outstanding features of the 8209 and adds a number of its own, including operation at media speed, connection of a token-ring LAN to a WAN, and alternative management methods. The result is a bridge which combines outstanding performance and greater value.

2.5.1 Identifying the 8229 Front Panel

The EIA 232 data terminal element (DTE) port is a 25-pin male EIA 232 connector for loading operating software to the FLASH memory contained on the main logic board. This port is marked DTE, and is configured for DTE operation at 9600 bps with an 8-data-bit, no-parity, 1-stop-bit character format.

Power Light-Emitting Diode (LED): When green, this indicates that power is available to the 8229.

Green Status LED: Indicates that the 8229 has successfully completed its basic tests and is ready for operation.

Yellow Status LED: Indicates that the 8229 has detected an internal fault as part of the basic tests and is inoperative. The fault code appears as a numeric display.

Numeric Display: A 2-digit numeric display indicates the current status of the diagnostics in progress, or the fault code in the case of a detected fault.

Hardware Reset: A recessed reset button is accessible on the front panel of the 8229.

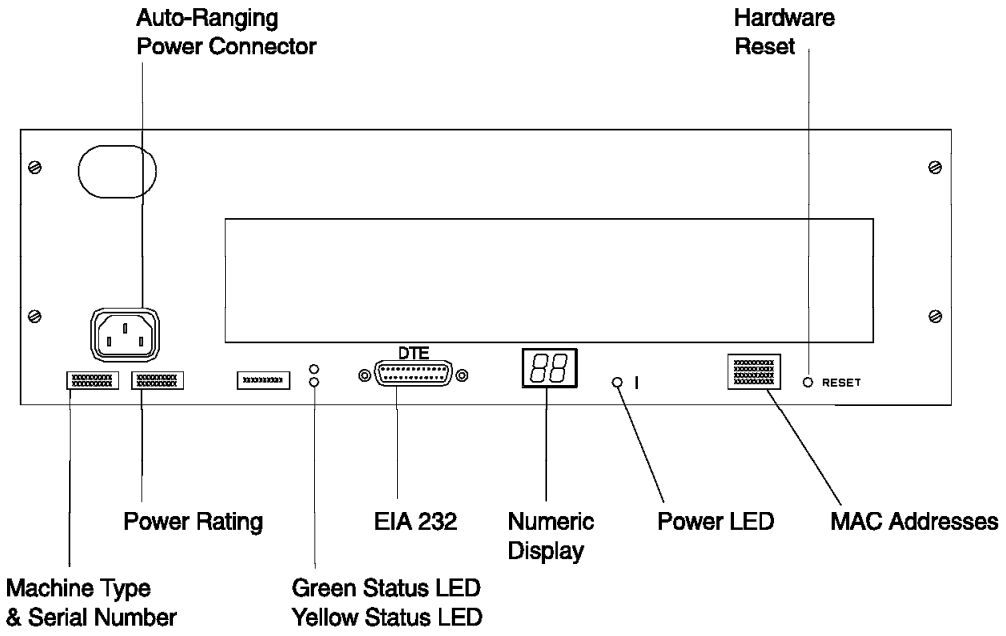

Figure 18. IBM 8229 Front Panel

2.5.2 Identifying the Attachment Module

Port Status Indicators: Each attachment module has a pair of status LEDs for each port. The green status LED, when lit, indicates that the internal tests for the port have been successfully completed; the yellow status LED, when lit, indicates a detected internal fault.

LAN Activity Indicators: Each attachment module indicates outbound activity for each port. The green activity LED indicates that the 8229 is successfully connected to the respective LAN and that traffic is being forwarded by the 8229 from that network. See Figure 19 on page 63 to Figure 22 on page 64 for illustrations of the attachment modules.

2.5.2.1 Using the 8229 Attachment Modules

To use the 8229, it is necessary to have one or two attachment modules. There are four types of attachment modules:

- Single-port token-ring LAN attachment module
- Dual-port token-ring LAN attachment module
- Ethernet LAN attachment module
- Wide Area Network (WAN) attachment module

The *single* and *dual* descriptions refer to the number of port groups that an attachment module has. The single-port token-ring attachment module has one port group consisting of two ports:

- An RJ-45 port for unshielded twisted-pair (UTP) connection
- A port for a 9-pin D-Shell male connector (for STP)

The dual-port token-ring attachment module has two port groups; each port group consists of the same ports as described above.

Only one port within each port group can be used at one time.

It is possible to use one or two attachment modules. The attachment modules plug into the front of the 8229. The 8229 comes with default settings and is pre-configured to work in *one* of the following recommended environments:

- **Token-ring to token-ring:** Uses the dual-port token-ring attachment module or two single-port token-ring attachment modules.
- **Token-ring to Ethernet:** Uses the single-port token-ring attachment module or one port of the dual-port token-ring attachment module *and* the Ethernet attachment module.
- **Token-ring to token-ring WAN:** Uses the single-port token-ring module or one port of the dual-port token-ring attachment module *and* the WAN attachment module.

2.5.2.2 Token-Ring Connection

The physical interface to the token-ring LAN is a 9-pin D female connector (Type 1) used for STP connection or an 8-pin RJ-45 connector used for UTP connection. The single-port token-ring module has two ports (in one group), but only one port can be used at any one time. The dual-port token-ring module has four ports (in two port groups), but only one port of each port group can be used at any one time.

The activity LED is lit only when the 8229 is transmitting outbound on the ring.

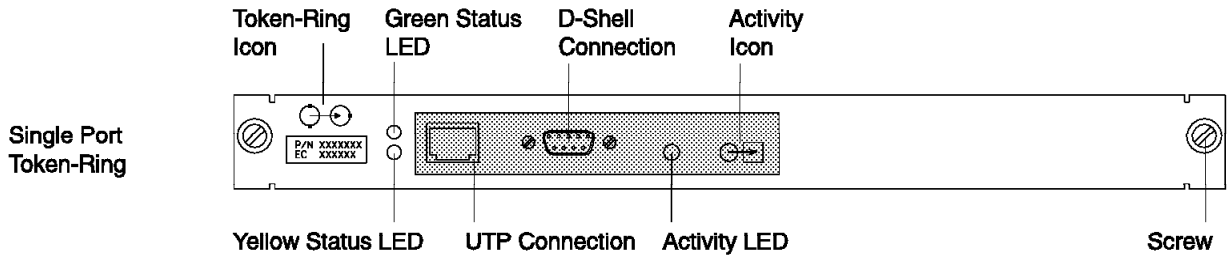

Figure 19. Single-Port token-ring Attachment Module

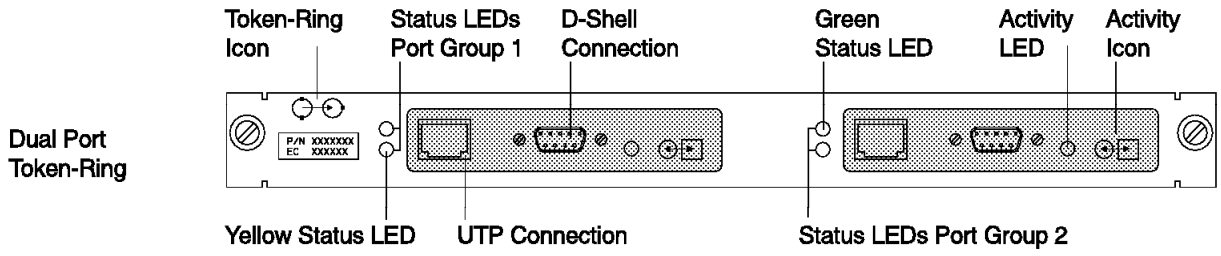

Figure 20. Dual-Port token-ring Attachment Module

2.5.2.3 Ethernet Connection

The physical interface to the Ethernet LAN is a 15-pin female D connector (AUI). The RJ-45 connector supports an Ethernet 10Base-T connection (UTP). The Ethernet attachment module has one port group with two ports, but only one port can be used at one time.

The activity LED indicates transit or outbound traffic.

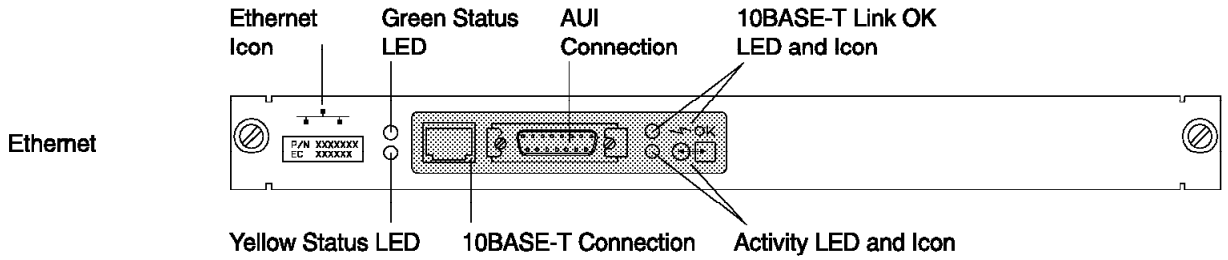

Figure 21. Ethernet Attachment Module

2.5.2.4 WAN Connection

The physical interface to a WAN is a 26-pin D female connector. The single WAN module has one port.

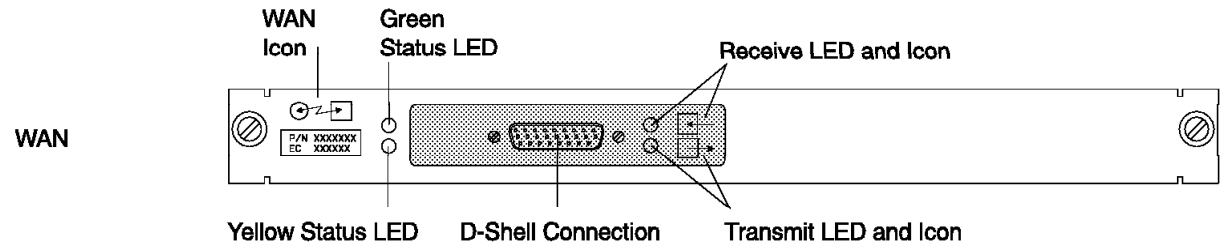

Figure 22. WAN Attachment Module

2.5.2.5 Locating Ports 1 and 2

The locations of Port 1 and Port 2 depend on which LAN attachment modules are installed in the 8229. In general, the first token-ring port on the bridge is Port 1. When you are using two single-port token-ring attachment modules, the port for the lower module is Port 1 and the port for the upper module is Port 2.

When you are using one single-port token-ring attachment module and either an Ethernet attachment module or a WAN attachment module, the 8229 will detect the first operational token-ring port and regard that port as Port 1; Port 2 is the Ethernet or WAN attachment module port, as shown in **A** and **D** in Figure 23 on page 65.

When you are using a single-port token-ring attachment module and a dual-port token-ring attachment module, Port 1 is defined as the first port found, starting at the bottom left and proceeding to the bottom right and then to the top left. Port 2 is defined as the second port found. This configuration is shown in **C** and **F**. Using a dual-port token-ring attachment module with a single-port token-ring

attachment module is not recommended because one of the three available ports is not used.

B and **E** show a dual-port token-ring module being used by itself.

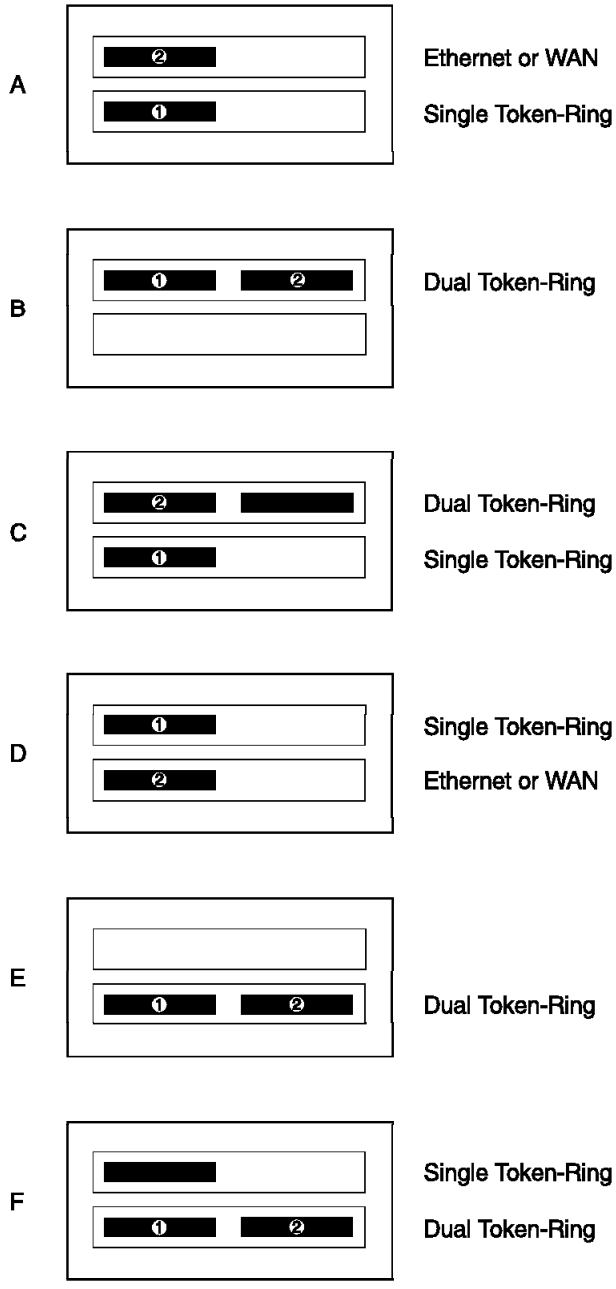

Figure 23. Port Defaults for Several Possible Configurations. In this diagram, the number 1 represents Port 1 and the number 2 represents Port 2.

2.5.3 Operation in General

This section defines the characteristics of the 8229 that are independent of the attachment modules installed.

2.5.3.1 Programs Supplied with the 8229

Two sets of programs are supplied on the diskette for management and configuration of the 8229:

- The *IBM 8229 Utility Program* allows you to change bridge parameters for LLC management only.
- The *LDBRG* program allows you to load code, filters, and SNMP configuration parameters over the network regardless of whether you are using LLC or SNMP.

The EIA 232 Program allows parameters and code to be specified and loaded through the EIA 232 port. This program is resident in the 8229 and is accessible through a terminal that is configured for TTY or VT100 emulation and is connected to the EIA 232 port. The EIA 232 Program is also known as the *out-of-band* utility because it is accessed by the EIA 232 port rather than through a workstation connected to the LAN.

2.5.3.2 Overview of SNMP

SNMP is used in a TCP/IP environment to allow network devices, such as the 8229, to be managed from a remote site over a network. SNMP is organized to use management information bases (MIBs). Some of the variables in these files are set to act as flags; others serve to store data. Within the SNMP design, the two key software entities are the manager and the agent. These two communicate with one another using the SNMP protocol. In general, managers make requests to the agents and collect information from them. The agents respond to the requests made by the managers. Both agents and managers consult the MIBs when they need information that is stored there.

From the SNMP point of view, the 8229 is an agent and a network management program, such as NetView/6000, is the manager. Both the 8229 and the network manager access the same MIBs. The network manager must load with the three private MIBs that are unique to the 8229 before it can communicate with the 8229.

2.5.3.3 Operational Software

Operational software for the 8229 is maintained in the FLASH within the 8229. This software is updated or replaced by use of a personal computer and the EIA 232 port (which is labeled *DTE*) on the front panel of the 8229. File transfer for software load is performed by using the XMODEM protocol, supported by a variety of available ASCII terminal emulation software products (not included on the Operational Software Diskette provided with the 8229).

You can also load the software over the LAN using the IBM 8229 Utility Program or the LDBRG program contained on the diskette that came with the 8229.

There are two operational modes: *full operational mode* and *minimal operational mode.*

2.5.3.4 Full Operational Mode

Full operational mode is achieved after the 8229 has gone through initialization without encountering a problem. This is the normal operational mode.

In full operational mode, the 8229 can support several operations in addition to being able to use the EIA 232 port. When the 8229 is running in operational mode and is configured for SNMP, the Utility Program can no longer be linked to the 8229. The SNMP bridge manager, such as NetView/6000, takes over the management of the bridge.

2.5.3.5 Minimal Operational Mode

In the event that the operational software is corrupt, or if the operational software is valid but does not match the hardware configuration, the 8229 will progress following power-on to minimal mode.

2.5.3.6 Using the 8229 Without Changing the Software

Now you should determine whether you can run the 8229 without changing the default software or the default values for the bridge parameters. You can probably avoid changing the software if the following conditions apply to you:

- Your bridge hardware matches your physical network. For example, you connect a token-ring network to a token-ring network with an 8229 that has two single-port token-ring attachment modules or one dual port token-ring attachment module.
- You do not need to change the bridge configuration parameters.
- You do not plan to use simple network management protocol (SNMP).

2.5.3.7 When You Have to Change the Software

You need to access and change the 8229 software if any of these situations apply to you:

- You need to upload new operational software. You must do this whenever you change the function of the 8229.

  For example, suppose the 8229 is using two single-port token-ring attachment modules to connect two token-ring segments. You remove one of the single-port token-ring attachment modules and replace it with an Ethernet attachment module, to connect the 8229 to an Ethernet segment. In this case, you must upload new operational software.

  You must also upload new operational software if you want to use SNMP in your network.
- You need to upload a filter program other than the default filters that are automatically provided in the 8229 operational software.

2.5.3.8 Overview of the Parameters Set by the Switches

You can set some most commonly used parameters using switches on the attachment modules. In fact, the following two frequently changed parameters cannot be changed by the utility programs or the EIA 232 program, and must be changed using the switches:

1. Token-ring speed
2. Connector type

Other parameters that are changed by using the switches include the following:

1. LAN segment number
2. Bridge mode (for the Ethernet attachment module)
3. Bridge number
4. Parameters for the WAN attachment module:
   - Primary or secondary bridge
   - Largest frame size
   - Wrap check

In many cases, the default switch settings will meet the needs of your network.

Important: The default switch setting for network speed on both the single and the dual token-ring attachment module is 16 Mbps.

2.5.3.9 LAN Management Programs That Support the 8229

The IBM LAN Network Manager Program, Versions 1.0 and 1.1, provides LAN management support to the 8229.

The IBM LAN Network Manager Version 1.1 provides full token-ring and Ethernet network management and configuration support to the 8229, providing all of the functions that the Utility Program provides, with the exception of loading operational software and user-defined filter programs.

When the 8229 is loaded with operational software to support SNMP management, use an SNMP Manager, such as NetView/6000.

2.5.4 IBM 8229 Utility Program

The IBM 8229 Utility Program uses the LLC Type 2 LAN management methodology and runs on a workstation attached to the same network as the 8229. The workstation can access the 8229 by being on the same LAN segment as the 8229 or by accessing the 8229 over other bridges from a different LAN segment.

The Utility Program enables you to:

- Display and change advanced configuration parameter values (such as values for single-route broadcast, spanning tree, and filter parameters).
- Display or change the LAN segment number (for either LAN segment connected to the 8229) to a value other than the value set by the attachment module switches.
- Display or change the bridge number and the bridge name. The configuration parameter for the bridge number overrides the bridge number set by the attachment module configuration switches.
- Display and change the status of the LAN management servers.
- Load a filter program into the 8229.
- Load operational software into the 8229.

2.5.4.1 Differences in the Utility Program for Token-Ring and Ethernet

Most of the functions of the Utility Program are the same for token-ring to token-ring configuration and for token-ring to Ethernet configuration. These are the functions of the Utility Program for token-ring to Ethernet configuration:

- Bridge definition
- Linking the Utility Program to the bridge
- Configuring the bridge. This choice brings up a menu that offers the following options:

  Loading bridge parameters

  Loading static addresses into the Ethernet static address database

  Loading mapped addresses

  Retrieving and reading the Ethernet database

  Loading user-defined filter programs

  Loading operational microcode
- Querying the bridge profile
- Defining a bridge password using System Definition
- Unlinking the Utility Program from the bridge
- Shutting down the Utility Program

These functions are the same as for the token-ring to token-ring configuration, except for some of the choices for configuring the bridge and some of the Configure Bridge panels. You reach the configuration choices when you choose Option 4, *Configure bridge*, from the main menu and then choose Option 1, *Bridge parameters*, from the Configure Bridge Menu.

These are the parameter choices that are specific to Ethernet:

- *Static entries*
- *Address Mapping*
- *Read data base*

These choices bring up panels that enable you to enter values for the following parameters:

- Static addresses for the Ethernet static address database
- Mapped addresses

2.5.5 Details of IBM 8229 Token-Ring to Ethernet Use

The 8229, when configured with a token-ring attachment module and Ethernet attachment module, provides a means to transfer frames between stations on a token-ring LAN and stations on Ethernet Version 2 and IEEE 802.3 LANs, performing the appropriate protocol conversion and format manipulation. The 8229 supports either a 4 or 16 Mbps token-ring LAN and a 10 Mbps Ethernet LAN.

The 8229 supports two LAN connections, one token-ring and one Ethernet.

A simplified connection diagram using the 8229 is illustrated in Figure 24 on page 70.

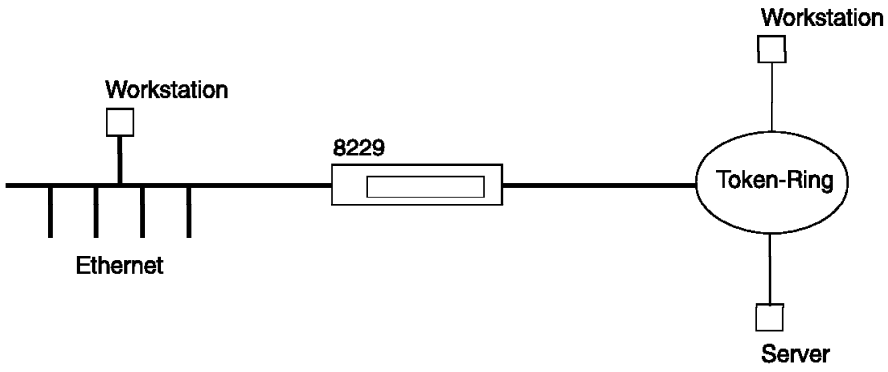

Figure 24. Typical Connection Using the 8229 with an Ethernet Module

2.5.5.1 Supported LAN Types and Protocols

The 8229 with the Ethernet attachment module installed provides:

- Access protocol support (token-passing) for attachment to an IBM Token-Ring Network segment
- Access protocol support (CSMA/CD) for attachment to an Ethernet Version 2 or IEEE 802.3 Ethernet segment
- Novell IPX protocol support to enable communication across the 8229 between stations that use Novell NetWare on the LAN segments
- Access protocol conversion and frame format conversion for each LAN connected to it

Versions of Ethernet Supported:

The following are the CSMA/CD LAN types commonly called Ethernet:

- Ethernet Version 1
- Ethernet Version 2
- IEEE 802.3

The 8229 supports both Ethernet Version 2 and IEEE 802.3 Ethernet. IEEE 802.3 Ethernet is based on, and coexists with, Ethernet Version 2.

Ethernet Version 1 does not coexist with either Ethernet Version 2 or with IEEE 802.3, and is not supported by the 8229. The two Ethernet versions have different end-of-transmission states (half-step versus full-step) and electrical common mode characteristics. Both Ethernet Version 2 and IEEE 802.3 use the half-step end-of-transmission state; Ethernet Version 1 uses the full-step state.

Differences between Ethernet Version 2 and IEEE 802.3 Ethernet exist at protocol levels, requiring the 8229 to support two different modes of frame format conversion between the IBM Token-Ring Network segment and the Ethernet LAN segment. The 8229 provides token-ring to Ethernet Version 2 frame format conversion and token-ring to IEEE 802.3 Ethernet frame format conversion.

2.5.5.2 Protocols Supported

Other differences between Ethernet Version 2 and IEEE 802.3 Ethernet exist at protocol levels above the physical layer and are accommodated in the configurable modes of operation of the 8229.

The 8229 supports the use of SNA, NetBIOS, and TCP/IP on both Ethernet Version 2 and IEEE 802.3 Ethernet. The design of the 8229 does not preclude the use of any protocols that adhere to the industrial protocol standards for token-ring and Ethernet LANs. However, only SNA, TCP/IP, NetBIOS and Novell NetWare IPX protocols have been tested.

2.5.5.3 8229 Modes

In mode 1, the 8229 performs frame format conversion between Ethernet Version 2 and IBM Token-Ring Network frames. In mode 2, the 8229 performs frame format conversion between IEEE 802.3 Ethernet and IBM Token-Ring Network frames.

2.5.6 IBM 8229 Token-Ring to WAN Connectivity

The 8229 (with a token-ring attachment module and a WAN attachment module) connects a token-ring operating at 4 or 16 Mbps to a WAN. This configuration is called a *split bridge* because two bridges are physically present but the two bridges make one connection. The telecommunications link stretching between the two halves of the bridge is considered as if it were part of the bridge. One half of the bridge is designated as primary and the other half as secondary. Figure 25 shows this configuration.

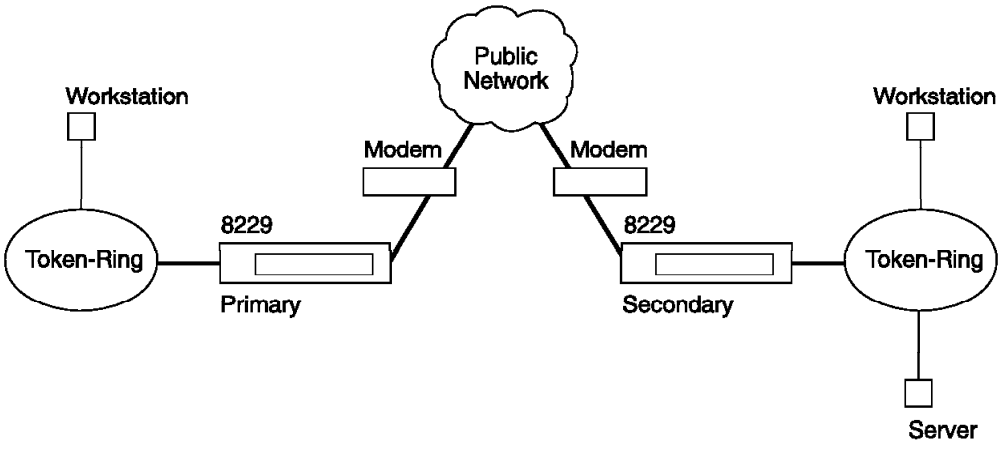

Figure 25. Typical Use of IBM 8229 with WAN Modules

The 8229 WAN Attachment Module communicates with one of the following bridges at the other end of the WAN:

- IBM Remote Token-Ring Bridge/DOS Version 1.0 or IBM Remote Token-Ring Bridge Program Version 2.2 running on a workstation
- Another 8229 with a WAN attachment module

Note:

In this configuration, the 8229 supports only the two versions of the IBM Remote Token-Ring Bridge stated in the previous list: IBM Remote Token-Ring Bridge/DOS Version 1.0 or IBM Remote Token-Ring Bridge Program Version 2.2.

The Remote Token-Ring Bridge Program is always the primary half of the split bridge. In this case, LAN Network Manager must be used to configure the Remote Token-Ring Bridge Program. The parameters set for the primary half of the bridge are used by the secondary half of the bridge, which is the 8229.

If you want to clear a user-defined filter in the 8229, perform one of the following steps:

1. Press the hardware reset button (the parameters in NVR are not controlling the bridge, so erasing them does no harm).
2. Use LDBRG to load the null filter file, FILTER3.X. FILTER3.X erases the user-defined filter.

Important: The Utility Program must be linked to the primary half of this split bridge.

Once the Utility Program is linked to the primary half of the bridge, using the Utility Program is the same for the split bridge composed of two 8229s as it is for a token-ring to token-ring bridge.

2.5.6.1 Supported Protocols

The 8229 functions as a local source-routing MAC level bridge, performing frame forwarding without modification to the contents of the information field within the frame. As such, all upper level protocols supported on the token-ring LAN are supported by the 8229.

The 8229 supports an LLC2 LAN management scheme only, in this configuration. On the telecommunications link, the frame formatting and bridge-half communication is compatible with the split bridge.

2.5.6.2 Supported Network Topologies

The 8229 supports communication in all valid token-ring configurations supported by existing personal computer-based IBM Token-Ring bridge programs.

The WAN port supports serial data rates up to 2.048 Mbps with timing supplied by the network interface equipment. Operation is supported for dedicated path connections; specific network topologies are a function of the carrier service and are functionally transparent to the 8229, as long as acceptable line quality is provided.

2.5.6.3 Filters

Each bridge uses its own filter selections. The 8229 uses the following selections:

- Address filter
- Range filter
- User-defined filter

The IBM Remote Token-Ring Bridge Program uses a user-defined filter.

The bridge-half filter decision is made after data is taken off the token-ring LAN and before the data is sent to the WAN connection.

2.5.6.4 Coexistence with the IBM Remote Token-Ring Bridge Program

You can have a split bridge composed of a workstation with IBM Remote Token-Ring Bridge Program at one side of the WAN and an 8229 at the other side. In this configuration, the 8229 is always the secondary half of the split bridge and will support the following functions:

- Bridge testing
- Configuring data
- Network status
- Path trace
- Performance counters
- Communication status
- Shutdown verification

You cannot use the Utility Program with this configuration; you must designate the workstation with the IBM Remote Bridge Program Version 2.0 as the primary half of the bridge and set the parameters on that bridge, using the LAN Network Manager Program. The 8229 accepts the parameters set for the primary half of the bridge, that is, the parameters for the IBM Remote Token-Ring Bridge Program.

To load and set a user-defined filter program when the 8229 is the secondary half of the split bridge, use LDBRG. To clear a user-defined filter when the 8229 is the secondary half of the split bridge, perform one of the following steps:

1. Press the hardware reset button.
2. Use LDBRG to load FILTER3.X, the null filter. This file erases the user-defined filter program from the secondary bridge.

2.5.6.5 8229 WAN Equipment and Supplies

To install and operate the 8229 module at data rates from 9.6 Kbps to 2.048 Mbps, the following equipment is needed for each bridge half:

- One 8229 (with one token-ring attachment module and one WAN attachment module)
- A DCE device that provides attachment to the telecommunication link
- A modem and its attaching cables, which are compatible with your network (see Table 2)
- A data service unit (DSU)

Table 2 (Page 1 of 2). Electrical Interface, Required Cable, and Line Data Rates

Electrical Interface	Supporting Cable	Line Data Rates	Part Number
V.24	EIA 232-C Attachment Interface Cable	9.6 Kbps to 19.2 Kbps. See note (1).	02F9477

Table 2 (Page 2 of 2). Electrical Interface, Required Cable, and Line Data Rates			
Electrical Interface	**Supporting Cable**	**Line Data Rates**	**Part Number**
V.35	V.35 Interface Cable	9.6 Kbps through 2.048 Mbps	02F9494
X.21	X.21 Interface Cable	9.6 Kbps through 2.048 Mbps. See note (2).	02F9493

Note:

(1) Shorter cables are required to run above 19.2 Kbps.

(2) Part number 02F9493 currently only runs at a maximum throughput of 256 Kbps.

2.5.7 Frame Format and Address Conversion

This section describes the various types of frame format conversion that the 8229 can do. Frame conversion is required when the bridge connects a token-ring network to Ethernet. It also describes how it does the address conversion, which you may have to do manually for some protocols.

2.5.7.1 Address Conversion

The bit order of the 48-bit (12-digit) IEEE adapter (MAC) address is reversed between the token-ring and Ethernet LANs.

As an example of converting an address, refer to Figure 26, which shows how token-ring address 4000 3174 0001 is converted into a bit-inverted Ethernet address.

The same process is used to convert an Ethernet address to a token-ring address.

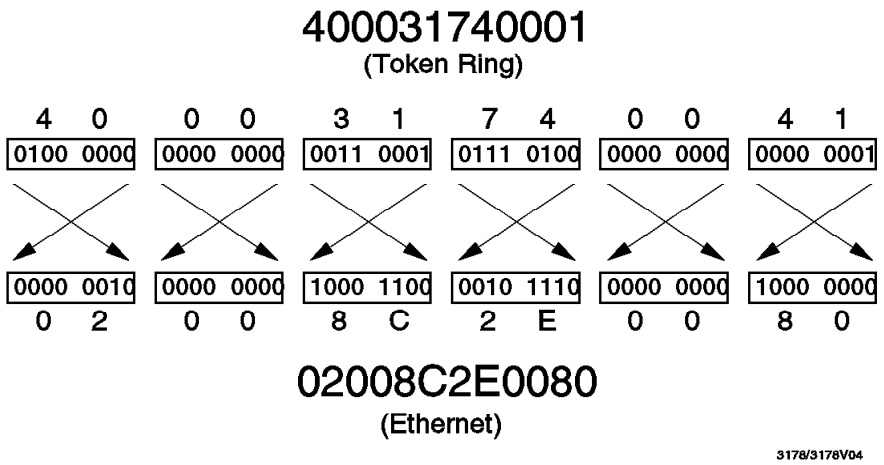

Figure 26. An Example of Token-Ring to Ethernet Address Conversion

Some protocols that use this address may not adjust for the inverted bits when communicating between LANs. Therefore, you must manually convert the address bit order. You may need to do manual conversion if you are:

- Defining a locally administered address for an Ethernet port. The address must be specified in the IEEE 802.5 format in the bridge configuration. (The 8229 does not do bit inversion on addresses of recognized IPX frames when IPX support is enabled.)
- Isolating a problem with a protocol.
- Tracing frames between the networks.

Use the following procedure to convert an address. An example of Address Conversion Chart is also shown to you.

1. Write the 12-digit address on the Address Conversion Chart.

   Separate the 12 digits into pairs. Use the first digit of each pair as the row coordinate and the second digit as the column coordinate.

2. Locate a bit-order inverted pair in Table 3 for each pair you wrote on the worksheet.

3. Combine the 6 pairs from the table into the converted 12-digit address and write them on the Address Conversion Chart and you will have the new address.

Note: This procedure is valid for token-ring to Ethernet conversion or vice-versa.

Table 3. Address Conversion Table

2nd Char (Col.) → 1st Char (Col.) ↓	**0**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**A**	**B**	**C**	**D**	**E**	**F**
0	00	80	40	C0	20	A0	60	E0	10	90	50	D0	30	B0	70	F0
1	08	88	48	C8	28	A8	68	E8	18	98	58	D8	38	B8	78	F8
2	04	84	44	C4	24	A4	64	E4	14	94	58	D4	34	B4	74	F4
3	0C	8C	4C	CC	2C	AC	6C	EC	1C	9C	5C	DC	3C	BC	7C	FC
4	02	82	42	C2	22	A2	62	E2	12	92	52	D2	32	B2	72	F2
5	0A	8A	4A	CA	2A	AA	6A	EA	1A	9A	5A	DA	3A	BA	7A	FA
6	06	86	46	C6	26	A6	66	E6	16	96	56	D6	36	B6	76	F6
7	0E	8E	4E	CE	2E	AE	6E	EE	1E	9E	5E	DE	3E	BE	7E	FE
8	01	81	41	C1	21	A1	61	E1	11	91	51	D1	31	B1	71	F1
9	09	89	49	C9	29	A9	69	E9	19	99	59	D9	39	B9	79	F9
A	05	85	45	C5	25	A5	65	E5	15	95	55	D5	35	B5	75	F5
B	0D	8D	4D	CD	2D	AD	6D	ED	1D	9D	5D	DD	3D	BD	7D	FD
C	03	83	43	C3	23	A3	63	E3	13	93	53	D3	33	B3	73	F3
D	0B	8B	4B	CB	2B	AB	6B	E8	1B	9B	5B	DB	3B	BB	7B	FB
E	07	87	47	C7	27	A7	67	E7	17	97	57	D7	37	B7	77	F7
F	0F	8F	4F	CF	2F	AF	6F	EF	1F	9F	5F	DF	3F	BF	7F	FF

Address Conversion Chart: The following form is used for converting addresses.

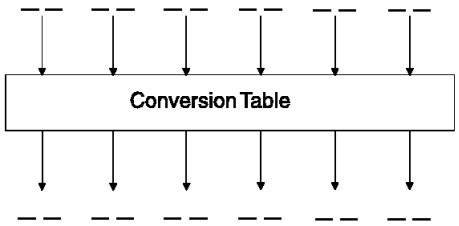

Figure 27. Example of Address Conversion Chart

2.5.7.2 Frame Format Conversion

The 8229 has two main types of frame format conversions:

- Bit inversion

  The bit order of the bytes for the destination and source address fields is inverted by the 8229 as part of the copy process, when the bit order of these fields is reversed between the two LAN types.

- Frame header manipulation

  Address, control, routing, and length information is copied to or deleted from a frame before the frame is forwarded, to provide the fields required by the destination LAN type.

The 8229 provides the following frame format conversions:

- Token-ring to Ethernet Version 2.0
- Ethernet Version 2.0 to token-ring
- Token-ring to IEEE 802.3 Ethernet
- IEEE 802.3 Ethernet to token-ring
- ARP MAC address bit-inversion
- RARP MAC address bit-inversion
- Token-ring to Ethernet Version 2.0 for LLC-based protocols
- Ethernet Version 2.0 to token-ring for LLC-based protocols

Token-Ring to Ethernet Version 2.0 Conversion: This conversion runs on the 8229 in mode 1.

The conversion from a token-ring frame to an Ethernet Version 2 frame for TCP/IP operation is represented in Figure 28 on page 77. In this conversion, the routing information (RI) and the destination service access point (DSAP), source service access point (SSAP), control (CONT), and protocol ID contained in the subnetwork access protocol (SNAP) header are extracted from the token-ring frame and discarded. The destination address (DA), source address (SA), and information field (TYPE and INFO) are copied into an Ethernet frame and sent to the Ethernet LAN.

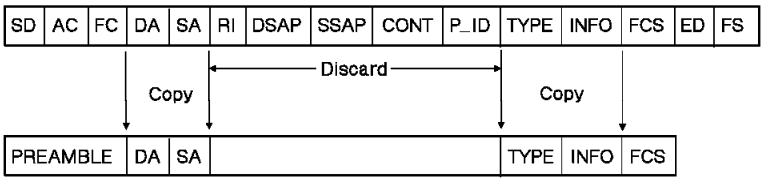

Figure 28. Token-Ring to Ethernet V2 Frame Conversion

In this conversion process, the bit order of the bytes for the destination and source address fields is inverted by the 8229 as part of the copy process because the bit order of these fields is reversed between the two LAN types.

Ethernet Version 2.0 to Token-Ring Conversion: This conversion runs on the 8229 in mode 1.

This conversion is the reverse of the token-ring to Ethernet conversion. In the conversion from an Ethernet frame to a token-ring frame, the destination address (DA), source address (SA), and information fields (TYPE and INFO) are copied into the respective fields of a token-ring frame. Before sending the frame to the token-ring LAN segment, the 8229:

- Retrieves the source-routing information associated with the token-ring destination address and inserts the information into the frame
- Inserts the fixed hexadecimal values AA AA 03 (representing the DSAP, SSAP, and control fields) into the frame
- Inserts a protocol ID of hexadecimal 00 00 00 into the frame

(The SNAP header consists of the protocol ID and the TYPE fields.)

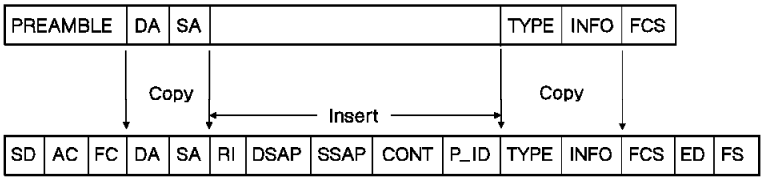

Figure 29. Ethernet V2 to Token-Ring Frame Conversion

In this conversion process, the bit order of the bytes for the destination and source address fields is inverted by the 8229 as part of the copy process because the bit order of these fields is reversed between the two LAN types.

Token-Ring to IEEE 802.3 Ethernet Conversion: This conversion runs on the 8229 in mode 2.

The conversion from a token-ring frame to an IEEE 802.3 Ethernet frame is represented in Figure 30 on page 78. In this conversion, only the routing information (RI) is extracted from the token-ring frame and discarded. The destination address (DA), source address (SA), and 802.3 INFO data is copied into an IEEE 802.3 frame, a LENGTH field calculated by the 8229 is inserted into the frame, and the frame is sent to the IEEE 802.3 Ethernet LAN.

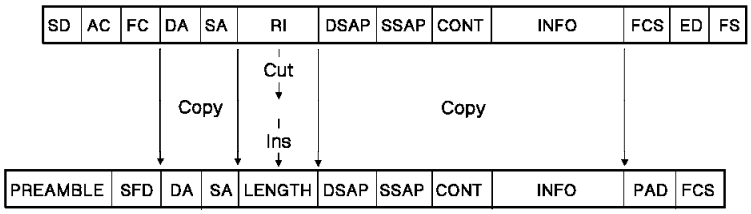

Figure 30. Token-Ring to IEEE 802.3 Ethernet Frame Conversion

In this conversion process, the bit order of the bytes for the destination and source address fields is inverted by the 8229 as part of the copy process because the bit order of these fields is reversed between the two LAN types.

IEEE 802.3 Ethernet to Token-Ring Conversion: This conversion runs on the 8229 in mode 2.

This conversion is the reverse of the token-ring to IEEE 802.3 Ethernet conversion. In the conversion from an IEEE 802.3 frame to a token-ring frame, the destination address (DA), source address (SA), and IEEE 802.3 Info data is copied into the respective fields of a token-ring frame. The Length field is discarded. The 8229 then retrieves the source-routing information fields associated with the token-ring destination address and inserts these fields following the source address before sending the frame to the token-ring LAN.

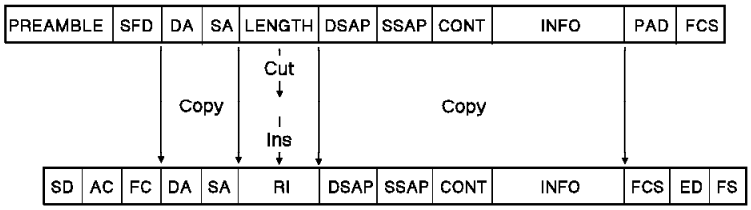

Figure 31. IEEE 802.3 Ethernet to Token-Ring Frame Conversion

In this conversion process, the bit order of the bytes for the destination and source address fields is inverted by the 8229 as part of the copy process because the bit order of these fields is reversed between the two LAN types.

ARP Conversion: The address resolution protocol (ARP) is used to determine a target station's hardware network address when only the station's internetwork or protocol address is known in a TCP/IP environment. The mechanism used is a broadcast frame containing the protocol address of the desired LAN station. The station having the protocol address responds to the broadcast packet with a message containing its hardware address within the information field of the frame.

The ARP frames are uniquely identified by X'0806' in the TYPE field in the frame header.

Like the source and destination addresses of the frame, the bit order of the address within the information field of the ARP is inverted between the token-ring and Ethernet Version 2 or IEEE 802.3 Ethernet LANs.

The ARP conversion process is shown in Figure 32 on page 79.

RARP Conversion: RARP uses the same format as ARP and operates in a similar manner except that it is used by a station that requests its protocol address from a network server. The RARP frames are uniquely identified by X′8035′ in the TYPE field in the frame header.

Like the source and destination addresses of the frame, the bit order of the addresses within the information field of RARP is inverted between the token-ring and Ethernet Version 2 or IEEE 802.3 Ethernet LANs.

The RARP conversion process is shown in Figure 32.

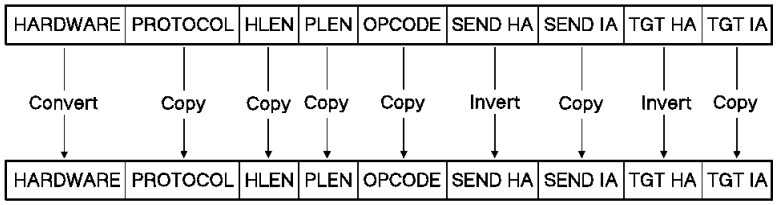

Figure 32. ARP and RARP Conversion

Token-Ring to Ethernet Version 2.0 Conversion for LLC-Based Protocols: A method for transporting logical link control (LLC) based protocols on Ethernet Version 2 is supported by the PC-RT and OS/2 Extended Edition (OS/2 EE). The Ethernet type value of X′80D5′ is used to indicate an LLC-based protocol.

The conversion from token-ring to Ethernet for LLC-based protocols is done by the 8229 when the following conditions exist:

- The 8229 recognizes the destination station as communicating in Ethernet Version 2 format (mode 1).
- The Forward LLC Traffic function is enabled in the 8229 configuration. (The default is enabled.)
- The DSAP in the token-ring frame is contained in the 8229 SAP table. The default values in this table are hexadecimal 00, 04, 08, F0, F4, and FC.

When the conditions for the LLC-based protocols exist, the conversion is done and the frames are forwarded as shown in Figure 33; otherwise, the conversion is not done.

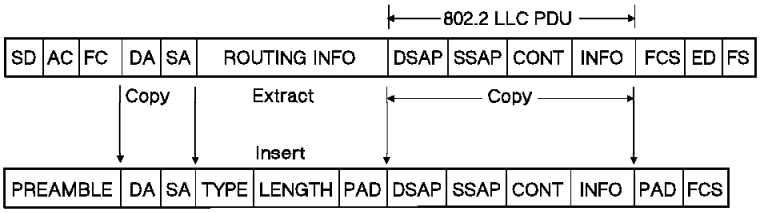

Figure 33. Token-Ring to LLC-on-Ethernet Frame Conversion for LLC-Based Protocols

In this conversion, the routing information (RI) is extracted from the token-ring frame and is discarded. The destination address (DA) and source address (SA) and LLC protocol data unit (PDU) are copied into the Ethernet frame.

The LLC type field (X′80D5′), the length of the PDU, and the pad characters are inserted into the Ethernet frame and sent to the Ethernet LAN.

Ethernet Version 2.0 to Token-Ring Conversion for LLC-Based Protocols: This conversion is the reverse of the token-ring to Ethernet conversion for LLC-based protocols. In this conversion, if a frame is received from an Ethernet station with a type value of hexadecimal 80D5, the conversion process is done as shown in Figure 34.

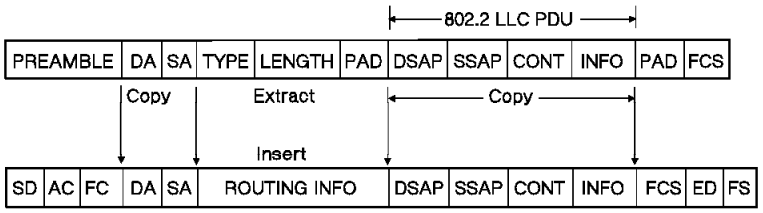

Figure 34. LLC-on-Ethernet to Token-Ring Frame Conversion for LLC-Based Protocols

In this conversion, the type, length, and pad fields are extracted and discarded. The LLC PDU and destination address (DA) and source address (SA) fields are copied into the token-ring frame. The 8229 then retrieves the routing information (RI) associated with the token-ring destination address, inserts it into the token-ring frame, and sends the frame to the token-ring LAN.

2.5.8 8229 Summary

Faster speed: The 8229 achieves faster frame forwarding rates through the use of a high-speed 50 MHz 486 SLC forwarding processor, an internal 32-bit bus and 2 MB of system memory. Your operating speeds are improved with the 8229 media speed between local LANs and up to T1/E1 (2.048 Mbps) between local and remote token-rings. In addition, the 8229 uses high-performance technology such as source-routing for token-ring and transparent bridging for Ethernet. Faster transmission of information translates directly into more productive users on the network.

The IBM 8229 LAN Bridge interconnects a token-ring network with another token-ring network or with an Ethernet V2 or IEEE 802.3 LAN.

Systems and workstations with compatible protocols such as IPX, TCP/IP, SNA, NetBIOS or IEEE 802.2 can communicate concurrently across this connection. The 8229 handles all necessary conversions to route information between dissimilar LANs.

More flexibility: The 8229 is designed to give you meaningful choices in the way you operate and manage your local area networks.

You have a choice of connections: link a token-ring segment with another local token-ring segment, with a remote token-ring via a WAN, or with a local Ethernet LAN.

You also have choice of management protocols: LAN network manager (LNM) or simple network management protocol (SNMP).

Remote Token-Ring bridging allows customers to interconnect their networks over a WAN into one larger manageable network independent of geographic location. Remote users can choose a distributed management strategy or be managed centrally from one location.

More value: The 8229 bridge can help you extend the value and productivity of your current network. At the same time, it gives you new speed, greater reliability and enhanced network management support capabilities while providing for future enhancements with the new flash memory feature. Combine all this with the low price, and the 8229 stands out as the ideal solution for your bridging needs.

Table 4. Bridging Architecture Enhancements

Feature	Benefits
Added connectivity	• Lets you link one token-ring segment to another token-ring LAN, to an Ethernet LAN or to a remote token-ring via a WAN • Permits management of multiple LANs from a central site, reducing the need for local management resources or travel to remote locations
Exclusive IBM 486 SLC processor	• Increases performance to media speed, delivers information faster • Boosts user productivity, reduces operating costs
New flash memory	• Simplifies changing or upgrading bridge capability • Contributes to management of remote rings from central site
Enhanced management	• Enables choice of SNMP or LNM - whichever best fits your needs; simplifies and streamlines network management • Permits management in a network with other network devices without requiring two management vehicles • Provides out-of-band management via RS-232-C port for selecting SNMP parameters as well as uploading new operating code, configuration data or filters in flash memory
Improved exterior design	• Saves space and speeds installation using standard mounting racks • Simplifies connections and monitoring by locating all cable connections ports and status LEDs on front panel • Speeds installation with front-mountable modules

2.5.8.1 Easy to Install

You don′t need special tools and you don′t need special training. For most installations, the 8229 can be used without changing any of the preset configuration parameters. All power, cable connections, and indicators are conveniently located on the front panel. That means you can be up and running in about 10 minutes. If changes to flash memory are necessary, add another 10 minutes.

If your network requires other functions or options, you can easily select them through hardware switches or configuration functions of the included Utility Program or the LAN management program. You may also write your own filters to reduce forwarding of unnecessary traffic and prevent degradation of LAN service. It′s one more way 8229 flexibility helps you improve LAN performance.

2.5.8.2 Performance Specifications

Performance of the 8229 depends on frame size and inter-frame arrival rate. The values, below, represent base case values measured under laboratory conditions meant to reflect user environments. The actual performance of a particular configuration may be different.

The filtering rate on the Ethernet 802.3 port is approximately 14,400 frames per second (64-byte frames). The filtering rate on the token-ring port is in excess of 25,000 frames per second (64-byte frames).

For 64-byte token-ring frames and 2-way traffic across a single bridge, the 8229 frame forwarding rate is approximately 23,000 packets per second (pps). This rate is based on normal traffic generated by a 20-station ring. Translational bridging forwarding rates are significantly improved over the 8209 bridge.

2.5.8.3 What You Get

IBM 8229 bridges are preloaded with LAN Network Manager software and are shipped with the following:

- Power cord
- Rack mount kit
- Bridge manual
- Operational software; maintenance and flash code utilities

2.5.8.4 Attachment Module

In addition to the bridge, you will need to order the appropriate attachment modules (refer to Table 5 on page 83).

- For the Model 001, either two single-port token-ring LAN modules or one dual-port token-ring LAN module
- For Model 002, one single-port token-ring LAN module and one Ethernet LAN module
- For Model 003, one single-port token-ring LAN module and one single-port WAN module

Users may later choose to change modules. You can convert from Model 001 to another model by downloading the appropriate operating software into a flash memory (microcode upload) and installing the appropriate attachment modules. If you already use a dual-port token-ring attachment module and wish to change models, you may continue to use it while adding either an Ethernet attachment module (to convert to Model 002) or a WAN attachment module (to convert to Model 003). One port on the dual-port token-ring attachment will become inactive, while the other port builds the bridge to the other function being implemented with the model change. Multiport functions are not supported.

Attachment modules are shipped with appropriate token-ring or Ethernet wrap plugs.

Table 5. Attachment Module Alternatives

8229 Base	Alternative	Single-port Token-Ring Module	Dual-port Token-Ring Module	Ethernet Module	WAN Module*
Model 001**	A	2			
	B		1		
Model 002		1		1	
Model 003		1			1

Note:

* To use the WAN attachment, one IBM cable must be ordered (see the WAN attachment cables table below).

** Select A or B.

Table 6. WAN Attachment Cables (select One)

Electrical Interface	Maximum data rate
RS-232-C V.24/28	19.2 Kbps
X.21 serial	256 Kbps
V.35 serial	2.048 Mbps

Table 7 (Page 1 of 2). IBM 8229 Bridge at a Glance

Machine type	Model 001, PN 73G4761
	Model 002, PN 73G4770
	Model 003, PN 73G4771
Product description	Model 001: Token-ring to local token-ring LAN
	Model 002: Token-ring to local Ethernet LAN
	Model 003: Token-ring to remote token-ring WAN
Attachment modules	Single-port token-ring card for all models
	Dual-port token-ring card for Model 001
	Single-port Ethernet card for Model 002
	Single-port WAN card for Model 003
Standards	Complies with IEEE 802.1, 802.3 and 802.5
Management	Choice of LAN Network Manager or SNMP. LNM is preloaded; SNMP is downloadable from utility disk via a LAN-based PC or through the RS-232-C serial port
SNMP MIBs	SNMP MIB II (RFC 1213), Standard Bridge MIB (RFC 1286), IBM Enterprise MIB extension, Surrogate token-ring MIB
	For token-ring-Ethernet connection, also supports Ethernet-like MIB (RFC 1398), IBM Token-Ring MIB and 8229 Product MIB
Protocols	For token-ring to Ethernet: supports IPX TCP/IP, SNA, NetBIOS or IEEE 802.2; supports Ethernet Spanning-Tree protocol.
	For token-ring to token-ring: supports all valid token-ring protocols
Physical specifications	Width: 17.5″ (444.5mm)
	Depth: 14″ (355.6 mm)

Table 7 (Page 2 of 2). IBM 8229 Bridge at a Glance	
	Height: 5.25″ (133.4 mm)
	Weight: 25.1 lbs (11.4 kg)
Operating environment	Temperature: 50° to 104°F (10°to 40°C)
	Relative humidity: 8 to 80%
	Electrical power: 0.115 KVA
	Noise level: < 40 dB
Installation	Base unit may be rack mounted
	Attachment modules are installed by customer
Warranty	One year

2.6 IBM 2210 Nways Multiprotocol Router

This section provides an overview to the IBM 2210, including a description of the hardware, and an overview of the software package. Further information is found in the *IBM 2210 Nways Multiprotocol Router Maintenance Information*, SY27-0345 and the *IBM 2210 Nways Multiprotocol Router Planning and Setup Guide*, GA27-4068.

2.6.1 Models of the IBM 2210

The IBM 2210 is available in several models, based on the types of networks you want to support.

IBM withdrew the Models 121, 122, 123, 124, 125 and 126. Models 121, 122, 123 and 124 had one LAN port, two serial connections, 2 MB Flash and 4 MB DRAM and were replaced with the Models 12T and 12E. Models 125 and 126 had one LAN port, two serial connections, 2 MB Flash and 4 MB DRAM and were replaced with the Models 127 and 128.

Table 8 on page 85 shows the different models and the offerings of the IBM Nways Multiprotocol Routing Network Services that are available.

The only differences between some of the models is the amount of flash memory and DRAM. Flash memory is used to store a compressed version of the router′s software while DRAM memory provides the working memory for the router programs and the router network tables.

Note: Flash memory is not able to be upgraded on the 12x models of the IBM 2210.

You can add an additional 4 MB of flash memory to the 14T and 24x models of the IBM 2210 by replacing the installed flash memory with an 8 MB Memory Expansion Feature. This upgrade provides a total of 8 MB of flash memory for those models.

If you want to maintain multiple copies of software for various releases, you may want to consider a model with 4 MB of flash memory.

IBM 2210′s DRAM provides the working memory for the router programs and the router network tables. The amount of required DRAM in an IBM 2210 is

determined by the size and complexity of the network the IBM 2210 must support.

You can upgrade the DRAM on all models of the IBM 2210 to a maximum of 16 MB using IBM's 16 MB Memory Expansion Feature.

Certain models of the IBM 2210 support ISDN. You cannot use one of the standard WAN ports for ISDN. Software support for ISDN must be ordered separately.

Table 8. IBM 2210 Models

Hardware						Software		
Model	**LAN**	**No. of WANs (See Note)**	**Flash Memory**	**DRAM**	**ISDN**	**Base**	**Additional Routing**	**ISDN**
12T	Token-Ring	2	4 MB	4 MB		x	x	
12E	Ethernet	2	4 MB	4 MB		x	x	
127	Token-Ring	2	4 MB	4 MB	x	x	x	x
128	Ethernet	2	4 MB	4 MB	x	x	x	x
14T	Token-Ring	4	4 MB	8 MB	x	x	x	x
24T	2 Token-Ring	4	4 MB	8 MB	x	x	x	x
24E	2 Ethernet	4	4 MB	8 MB	x	x	x	x
24M	1 Token-Ring, 1 Ethernet	4	4 MB	8 MB	x	x	x	x

Note: The standard WAN ports on the IBM 2210 will support any of these physical interfaces:

- EIA RS 232-D/V.24
- V.35
- V.36
- X.21

The ports of the different models are shown in Figure 35 to Figure 38 on page 86. The two models shown in each figure differ only in the amount of DRAM and flash memory they contain, as described above.

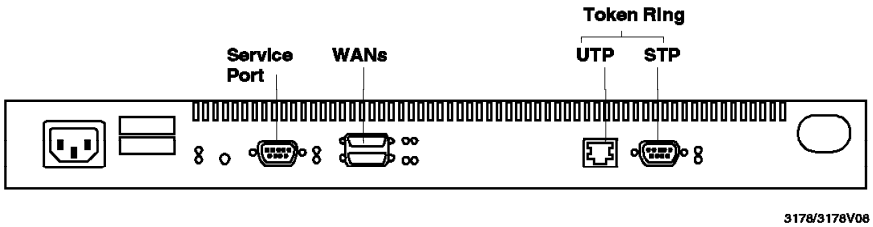

Figure 35. Model 12T

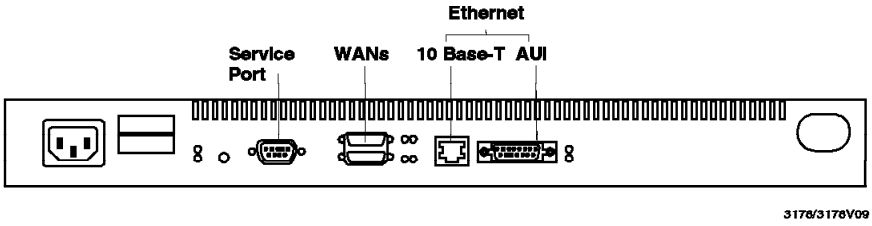

Figure 36. Model 12E

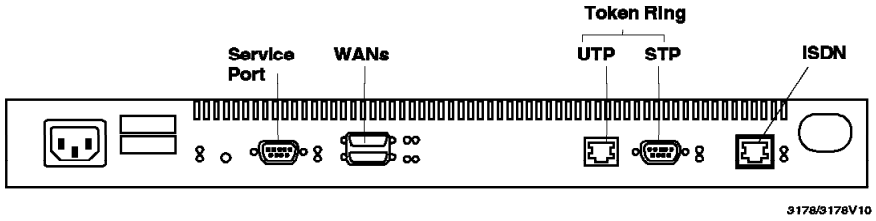

Figure 37. Model 127

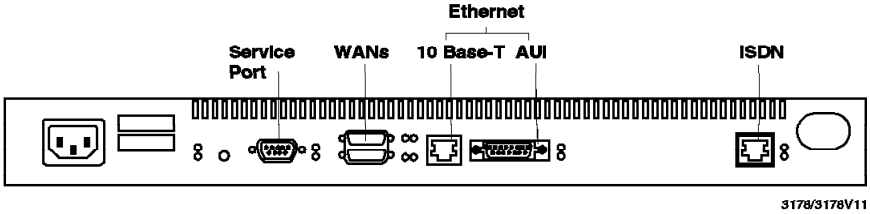

Figure 38. Model 128

2.6.2 Indicators on the IBM 2210

The IBM 2210 has green and amber LEDs that indicate the status of the system and of individual ports. Green indicates normal operation; amber indicates a problem.

The LEDs are on both the front and the back of the IBM 2210, so you can place it with either side facing forward. This is shown in Figure 39 on page 87 and Figure 40 on page 87.

Note: The figures shown are for Model 12T. The port LEDs are specific to each model.

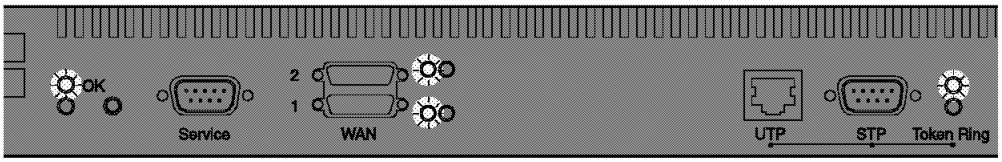

Figure 39. LEDs on the Port Side of Model 12T

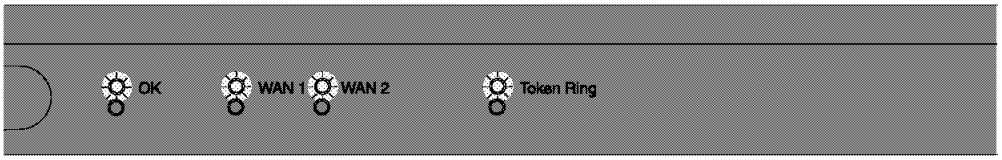

Figure 40. LEDs on the Side Opposite the Ports for Model 12T

2.6.3 The Reset Button on the IBM 2210

If you press the reset button, it will re-load the operational code. Also, if you press this button within 10 seconds of powering on, the 2210 will enter the extended power-on self-test (POST). Extended POST allows you to test memory more extensively than POST.

The reset button on the IBM 2210 is recessed to prevent accidental activation and is shown in Figure 41.

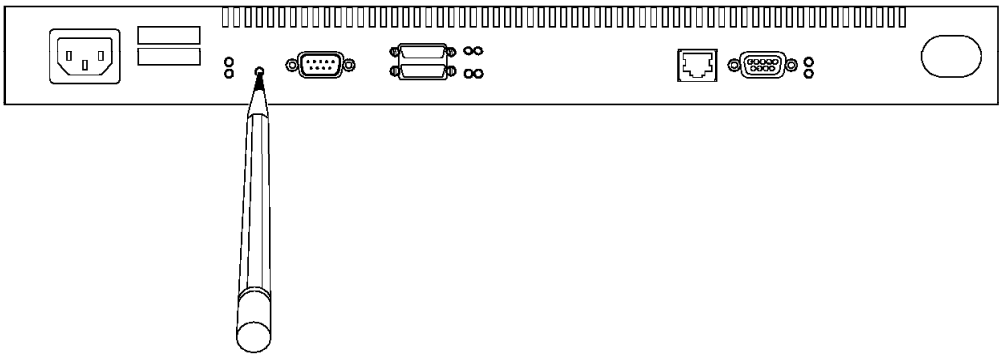

Figure 41. Reset Button on the IBM 2210

2.6.4 Networks Supported by the IBM 2210

The IBM 2210 supports the following LAN connections:

- Token-Ring (IEEE 802.5) with STP or UTP connection
- Ethernet (IEEE 802.3) with AUI or 10Base-T connection

Every IBM 2210 supports the following serial connections:

- EIA 232D/V.24

- V.35
- V.36
- X.21

Note: RS449 is also supported, using the V.36 cable available for the IBM 2210.

In addition to these serial connections, you can order optional support for ISDN.

2.6.5 Accessing the IBM 2210

You can access the IBM 2210 using the following methods:

- An ASCII terminal (or emulator) attached directly to the service port
- An ASCII terminal (or emulator) attached via a modem to the service port
- A Telnet session

2.6.5.1 Local Access

You can access the IBM 2210 locally through its service port, using an ASCII terminal or emulator. The DEC VT100 terminal is supported, as well as devices that are configured to emulate it. The settings should be:

- No parity
- 8-bit word length
- 1 stop bit
- 300 bps-38.4 Kbps bit rate

The IBM 3101, 3151 and 3161 Display Stations are also supported. For further information on these, please refer to *The IBM 2210 Nways Multiprotocol Router Planning and Setup Guide*, GA27-4068.

2.6.5.2 Remote Access

You can access the IBM 2210 remotely using either Telnet or a terminal attached to the service port via a modem.

The modem must use asynchronous operation and support the AT command set. The modem connected to the IBM 2210 must be set to auto-answer mode.

2.6.6 Software Package

Nways Multiprotocol Routing Network Services (MRNS) is the software that runs on the IBM 2210 and it comes as a base package, plus two separately orderable packages - one containing support for additional routing protocols, and the other containing the ISDN support. The protocols supported by each package are:

- Base offering
  - TCP/IP over point-to-point (PPP), frame relay, and X.25
  - Bridging over PPP
    - source-routing bridge (SRB)
    - transparent bridge (TB)
    - source-routing transparent bridge (SRT)
    - source-routing - translational bridge (SR-TB)

  - SNA/DLSw over PPP, frame relay, X.25, and SDLC
  - Bandwidth reservation for PPP
- Additional Routing Protocols Feature
  - Internetwork Packet Exchange (IPX) over PPP, frame relay, and X.25
  - AppleTalk over PPP
- ISDN Feature
  - Supported over IP, IPX, AppleTalk, SRB, TB, SRT, and SNA/DLS

2.6.7 MRNS Overview

This section provides an overview of the Nways Multiprotocol Routing Network Services (MRNS) software for the IBM 2210. It includes descriptions of the boot process, the user interface and the event logging system (ELS). Further information can be found in the *Nways MRNS Software User's Guide.*

The Nways MRNS is the software that supports the IBM 2210. The Nways MRNS has three components:

- The code that provides the routing, bridging, data link switching, and SNMP agent functions for the IBM 2210.
- The configuration program, which offers a graphical user interface that allows you to configure the IBM 2210 from a workstation.
- A monitoring system that allows you to perform network management, problem determination, and configuration.

2.6.7.1 Boot Files and Boot Processes

The IBM 2210 does not have a hard drive like the 6611 Network Processor, so it needs another method to load its operating system (referred to here as the boot file).

The boot file can be loaded (booted) from the following sources:

1. Flash memory (referred to as the Integrated Boot Device (IBD).
2. An external server which supports the TFTP server function. This could be another router which supports the TFTP server function (such as another IBM 2210).
3. The console port using ZModem.

Note: The IBM 2210 is delivered preloaded with a boot file in the IBD.

The IBM 2210 has a boot configuration database which holds information on all available boot files. Each entry in the database contains the location of the server host where the boot file resides and the path, file name, and a timeout value for the boot file. You can add entries to the database by using the following command:

```
Boot Config>add boot-entries
```

On startup, the IBM 2210 will normally load itself with the boot file stored in the IBD, but can use the boot configuration database to obtain a copy from a TFTP server should this boot file become corrupted or unusable.

The IBM 2210 may also use the Boot Protocol (BOOTP) to obtain its boot file, and uses the BOOTP client function to do so. The IBM 2210 will use the BOOTP protocol to learn its own IP address and the location (TFTP server) from which the boot file is obtained. It will then use TFTP to load the boot file from the TFTP server.

In order to cause the IBM 2210 to act as a BOOTP client, the interfaces over which the BOOTP packet should be broadcast are indicated by using the following command:

```
Boot Config>add bp-device
```

Note:

When the IBM 2210 obtains its boot file at boot time from an external source, it loads the boot file into executable memory, but does not save a copy in the IBD. If you want to move a copy into the IBD, you need to issue the following commands:

```
Boot Config>Copy Config
```
or

```
Boot Config>TFTP get
```

Both commands use the TFTP protocol. The only difference is the format in which you specify the location of the file to be transferred.

The IBM 2210 does not allow you to initiate a transfer from another device to the IBM 2210, so you will need to start the transfer from the router operator′s console.

The ZModem boot allows you to load router code through the console port, using an ASCII terminal emulator package that supports the ZModem protocol. To load the code via this method you enter:

```
>zb
```

The > prompt is the Boot prompt which is accessed by entering Ctrl+C while the IBM 2210 is reloading.

Your ZModem software documentation will explain the commands required to start the upload.

2.6.7.2 IBM 2210 Configuration

The configuration process customizes the IBM 2210 for the network in which you intend to run it, and the physical equipment being used. The configuration file may be created via the Nways MRNS Configuration Program and then transferred to the IBM 2210, or via commands entered at the operator console.

The configuration data resides in IBM 2210′s non-volatile RAM (NVRAM) and is combined with the boot file when the IBM 2210 is restarted or reloaded, creating the operating environment of the IBM 2210. NVRAM is the only place from which the IBM 2210 will obtain the configuration information during a restart or reload.

Reloading the IBM 2210 causes the router to reload the boot file into RAM. At the same time, it customizes the operating environment using the configuration file on NVRAM.

To reload the IBM 2210 you issue the `Reload` command from the OPCON prompt.

Restarting the IBM 2210 doesn't cause the router to reload the boot file. It simply takes the configuration file on NVRAM and feeds it into the operating environment.

To restart the IBM 2210 you issue the `Restart` command from the OPCON prompt.

Changes made from the operator console configuration process (CONFIG) are immediately saved in NVRAM and, in most cases, will take effect once the IBM 2210 is restarted or reloaded. However, there are a few changes which will take effect immediately without the need to restart or reload.

Changes made from the operator console monitoring process (GWCON) take effect immediately. However, once the router is restarted or reloaded, these changes are lost. This facility could be useful if you wish to test some changes prior to making them permanent.

NOTE: The parameters which are changed from the GWCON process are a subset of the parameters which can be changed from the CONFIG process.

The Nways MRNS Configuration Program may also be used to configure the IBM 2210. The Nways MRNS Configuration Program runs under AIX, OS/2 and Windows and uses a GUI interface. When configuring via the Nways MRNS Configuration Program, you create a configuration file on the workstation which can be saved in two formats:

- An archive format which is stored in the workstation configuration database, and is readable by the Nways MRNS Configuration Program
- A 2210-readable format for transferring to the IBM 2210 via TFTP

Note: The 2210-readable format cannot be reloaded into the Nways MRNS Configuration Program, so it is highly recommended that you save an archive copy before creating and sending a 2210-readable file to the router. The 2210-readable file must be manually transferred to the IBM 2210 using one of the following commands:

- `The Boot Config>Copy Config` or
- `The Boot Config>TFTP get` or
- `The >zc`

If you choose to create your configuration on the IBM 2210 console, then you should save a copy of it on an external server in case the NVRAM fails or the file is corrupted. You do this with the commands:

`Boot Config>Copy Config` or

`Boot Config>TFTP put`

The `>zc` command allows you to load a configuration file via the console port, using an ASCII terminal emulator that supports the ZModem protocol.

To access the > prompt, you need to enter Ctrl+C while the router is reloading.

Your ZModem software documentation will explain the commands required to start the upload.

2.6.7.3 MRNS User Interface

You access the Nways MRNS user interface through an ASCII console or emulator, as mentioned in 2.6.5, “Accessing the IBM 2210” on page 88.

By default, when you connect to the IBM 2210 you will not be required to enter a user ID or password, and you will have access to all router functions and commands. However, for security reasons you may want the users to enter a user ID and password when they connect to the router.

2.6.7.4 The Event Logging System (ELS)

ELS is a monitoring system that manages messages logged as a result of router activity. Using ELS commands, you can configure the system such that you only see the messages you need to. ELS uses the concepts of *subsystem, event number, message text, logging level, and group* to help you manage the messages you see.

Subsystem is a predefined name for a router component, such as an interface or protocol. For example, IP is the subsystem name for the IP protocol, and TKR is the subsystem name for the token-ring interface.

The ELS Config process is accessed by issuing the `Config>event` command.

You can obtain a complete list of the subsystem names by issuing the `ELS Config>list subsystem` command. The output shows the subsystem name, the number of events for the subsystem, and a description of the subsystem.

Event number is a predefined number assigned to each message within a subsystem. You can obtain a complete list of events for a particular subsystem by issuing the `ELS Config>list subsystem` *subsys* command, where *subsys* is the name of the particular subsystem you are interested in.

For example:

```
ELS Config>list udp
```

will list all possible events in the UDP subsystem. The output shows the event number, the logging level and the message text.

Message text is the actual text related to the event that has occurred and is used along with the subsystem and event number when the message is displayed by the MONITOR process. *Logging level* is a predefined category that each event will belong to, and which indicates the importance of the event. Note, whenever you use the `ELS Config>list subsystem` *subsys* command to list all the events within a subsystem, the logging level for each event is displayed.

Group is a user-defined collection of events that is given a name. A group can consist of events from different subsystems and of different logging levels. Once you have created a group, you can use the group name to manipulate the events in the group as a whole.

The *Nways MRNS Event Logging System Messages Guide* also contains a complete list of all events for all subsystems and includes the logging level for each event.

2.6.7.5 The IBM 2210 Configuration Program

The IBM Nways Multiprotocol Routing Network Services Configuration Program allows you to perform a complete configuration of an IBM 2210 Nways Multiprotocol Router. The Configuration Program is run on a workstation and has a graphical user interface.

Before using the Configuration Program you must perform an initial configuration on the 2210 to allow you to transfer these settings across to the IBM 2210 Router.The minimum requirement is that IP Routing is enabled to use the Trivial File Transfer Protocol (TFTP) or IP and SNMP are enabled to use the Communication option within the configuration program.

An Overview of the IBM 2210 Configuration Program: The IBM 2210 Configuration Program consists of two main windows:

- The Navigation Window
- The Configuration Window

The Navigation Window displays a directory tree, consisting of the various components that you can configure.

To select any particular configuration screen, click the left mouse button on the item you are interested in. The Configuration Window will now display the configuration screen you have selected.

Help is available for each field within a panel. You may access the help by pressing PF1.

If the field requires you to enter a value, be sure you press CR (Enter/Return) after entering your value. If you don′t do this, the value may not be saved.

Hardware and Software Requirements: The following hardware and software are required to run the Configuration Program on the RISC System/6000 workstation:

- IBM AIX 3.1.5 or higher with Transmission Control Protocol/Internet Protocol (TCP/IP) enabled

  Note: AIX 4.0 and higher is not supported.
- IBM AIXwindows
- 16 MB of memory
- A 3.5-inch diskette drive that can read and write
- 1.44 MB formatted diskettes
- 10 MB of available space on the fixed disk drive
- A graphics display that supports 640x480 resolution and 16 colors or gray scales
- A mouse

The following hardware and software are required to use the Configuration Program on a PS/2 workstation using an Intel 80386 or higher processor or a compatible system that has an Intel 80386 or higher processor.

For workstations running the Microsoft Windows program you need:

- IBM DOS 3.3 or higher, MS-DOS 3.3 or higher
- Microsoft Windows 3.1 or later versions
  - Win32s, included with the MRNS Configuration Program diskettes
  - WinSock 2.0 DLL (included with Win32s)
- TCP/IP application that uses WinSock 2.0 (This is only required for using the Configuration SEND function)
- 8 MB of memory
- 3.5 inch diskette drive that can read and write 1.44 MB formatted diskettes
- 10 MB of available space on the fixed disk drive
- A graphics display that supports 640x480 resolution and 16 colors or gray scales
- A mouse

For workstations running the IBM Operating System/2 (OS/2) Program, you need:

- OS/2 2.1 or later, including Warp
- IBM TCP/IP 1.2.1 or OS/2 or later (this is only required for using the Configuration SEND function)
- 10 MB of memory
- 3.5 inch diskette drive that can read and write 1.44 MB formatted diskettes
- 10 MB of available space on the fixed disk drive
- 10 MB of available swapper disk space on the swapper fixed disk drive partition
- A graphics display that supports 640x480 resolution and 16 colors or gray scales.

  Note: There is a known problem when running the Configuration Program on Warp. A selection of 65535 colors will prevent the program logo from displaying.
- A mouse

Anonymous FTP Site for the IBM 2210: IBM has established an anonymous FTP site for providing information and configuration program updates (and in the future other program updates) relating to the 2210.

The host name for the anonymous FTP site is nways.raleigh.ibm.com. If you have trouble resolving this name, the IP address is 192.35.236.5. After connecting to the machine specify anonymous as the user ID and your E-mail address as your password. Check the README file on the anonymous FTP site in the /pub directory for the latest information.

The subdirectories where the Configuration Programs reside are as follows:

- /pub/config/2210/1.2.0.0/GA/diskettes for the diskette images
- /pub/config/2210/1.2.0.0/GA/runtime for the RISC System/6000 files

2.6.7.6 IBM Nways Multiprotocol Routing Network Services Release 3 - Enhancements

The MRNS Configuration Program Release 3 supports configuration for all the functional enhancements for Nways Multiprotocol Routing Network Services Releases 1 and 2, and in addition offers the following:

- **Support for the new 2210 Models 14T, 24T, 24E and 24M.**

  There are many packages of the MRNS Release 3 to support these new 2210 models or those currently available.

- **Local LAN-to-LAN bridging support**

  With the addition of multiple LAN connectivity on the new models, there is the obvious need for local bridging support. Users may configure LAN-to-LAN and LAN-to-WAN bridging using any of the following as appropriate:

  - Transparent bridging (TB)
  - Source-route bridging (SRB)
  - Source-route transparent bridging (SRT)
  - Source-route - Translational bridging (SR-TB)

- **AIW Version 1 DLSw for SNA, and NOW NetBIOS, support (RFC 1795 compliant)**

  MRNS′s DLSw is now compliant with RFC 1795, referred to as the AIW Version 1 DLSw. MRNS′s DLSw will still interoperate with the DLSw implementation in MRNS V1 R1 and R2 for SNA traffic but not for NetBIOS (prior releases support NetBIOS only via bridging).

- **EasyStart, automatic configuration capability**

  The goal of EasyStart is to eliminate the need for local initial configuration, essentially creating a “plug and play” installation.

  EasyStart allows network download of initial router configuration. When the system starts, and there is no configuration information, EasyStart attempts to obtain it from a network server. If EasyStart fails, the fall back is to use the local ASCII console.

  Once the initial configuration is retrieved from the network, the system is automatically restarted to cause the new configuration parameters to take effect.

- **Data Compression over Point-to-Point Protocol (PPP)**

  Support has been added for the draft standard PPP Compression Control Protocol and, currently, for a single data compression engine:

  - *Deflate - LZ77*

    PPP data compression is negotiated by PPP at link open time; the algorithm(s) used and the preference order can be set on pre-interface basis (once additional algorithms are introduced), to allow for control of the (substantial) memory usage of compression dictionaries (about 80 KB per direction with Deflate, 24 KB per interface with Stacker, over 90 KB per direction with BSD, and 64 KB per direction with Predictor).

    PPP data compression can be used over any supported PPP interface, and can be used at the same time as Bandwidth Reservation (BRS will operate on data before compression is applied). When compression is in

use, all data that passes over the interface is compressed. The impact of attempting to compress already compressed traffic varies according to the algorithm in use.

The compression achievable varies greatly according the traffic. Using the Calgary Corpus standard of binaries, text files and image files, the Deflate algorithm achieves a ratio of 2.08:1. This compares to the following other algorithms:

- *Stacker-LZS: 1.82:1*
- *BSD Compress-LZW: 2.235:1*
- *Predictor: 1.67:1*

- **LAN Network Manager (LNM) support**

  The 2210 / MRNS LNM support is a source-route (SR) bridging option that enables LAN Network Manager agents on the 2210 bridge. The LNM function supports the following LNM agents:

  - *Configuration Report Server (CRS)*

    The CRS agent collects and reports MAC ring topology changes to the IBM LNM application. It will send out CRS MAC requests to query the status of other ring stations when requested by the LAN Network Manager.

  - *Ring Error Monitor (REM)*

    The REM agent collects MAC error reports from ring stations. When thresholds are exceeded, REM forwards error information to the LAN Network Manager.

  - *Ring Parameter Server (RPS)*

    The RPS agent services MAC requests from ring stations for ring parameter information and informs the LAN Network Manager of ring insertions.

- **National ISDN-1, AT&T #5 ESS and Nortel′s DMS-100 (US and Canada) supported on the 2210 ISDN Models 127 and 128**

  The North American ISDN support is provided in Release 3 on the 2210 ISDN Models 127 and 128. With this support, users can attach the 2210 ISDN BRI port to one of the following:

  - *AT&T #5 ESS switch*
  - *Nortel′s DMS 100 switch*

- **WAN Re-Route**

  The WAN Re-Route function is an enhancement to the IBM 2210 Multiprotocol Routing Network Services (MRNS) software. It allows the activation of a alternate network interface when a primary interface fails. WAN Re-Route is more flexible than the standard WAN Restoral feature (WRS) currently provided because the alternate link may have a different termination point than the primary link. It uses the dynamic routing abilities of the different routing protocols (IP RIP, IP OSPF, IPX RIP, etc.) or bridging protocols to find alternate paths through the new network topology. It also allows the backup of all DLC types, that is frame relay, PPP and X.25, whereas WRS supports PPP links only.

- **SNMP Enhancements**

As new functions are added to the MRNS, additional SNMP support is also necessary to ensure comprehensive network management capability. With Release 3, expanded SNMP MIB support has been added for SDLC links, LLC, BRS and the enhanced DLSw functions.

- **IBM MRNS Configuration Program - Release 3 Enhancements**

  The Release 3 MRNS Configuration Program enhancements include the following changes of the Release 2 support:

  - InARP support for IP, IPX and AppleTalk.
  - Ability to retrieve a configuration file from a 2210 and display its parameters.
  - Ability to create an ASCII flat file for printing purposes. The ability to import an ASCII file, verify contents and subsequently send to a 2210 is not yet available.
  - Drag and drop of certain lists.
  - Enhanced validation of file parameters.

- **Additions to the Additional Routing Protocol Package:**

  - DECnet IV over PPP, frame relay (FR) and X.25 (2210 to 2210)**

    Release 3 includes support over PPP data links as well as MRNS Release 2.

  - DECnet V / OSI protocols over PPP, FR and X.25

    The Digital Network Architecture (DNA) Phase V packet forwarder provides packet forwarding for 2210 routers in accordance with the Phase IV and Phase V router specifications of the DECnet protocol family. This allows a router to connect to systems using DECnet software (DNA Phase IV and Phase V network protocols) on different physical networks.

  - Banyan VINES over PPP

    Support of BVCP (Banyan VINES Control Protocol), over frame relay and X.25 (2210 to 2210) was initially offered in MRNS Release 2 and continues with Release 3. With MRNS Release 3, support of PPP data links is also provided. Because PPP is a non-proprietary protocol, the BVCP addition allows 2210 routers to interoperate with other vendor routers which abide by RFC 1763. Another advantage of the BVCP implementation is that we can expand VINES supports to any media that supports PPP.

- **Optional Switch for Filtering Nonbridged Packets (Inbound Only)**

  The switch is stored in SRAM and new user interface commands have been added to allow the customer to specify whether or not the non-bridged packers are filtered.

  A MAC Filtering/Bridging Switch for Non-Bridged Packets has been inserted which allows the user to select whether non-bridged packets are filtered or not.

  The filtering of non-bridged packets will only occur when the following conditions are met:

  - Bridging is enabled.
  - For inbound packets only. That is, packets coming from a LAN segment and not from a WAN interface.

– When the switch is set to allow filtering of non-bridged packets and when the filter parameter indicate the packet should be filtered.

MRNS, together with the IBM 2210 Nways Multiprotocol Router, provides users with a broad range of networking products and services for high-speed, integrated, manageable, and open networks. The 2210 Nways Multiprotocol Router connects local-area and wide-area networks to form a physically integrated network that transports multiple networking protocols between applications speaking the same protocol.

Plans are in process to eliminate sending copies of backup media diskettes since the desired software package is preloaded on the 2210. Instead, current licenses provide instructions on how to retrieve a copy of the code via Internet access to the MRNS Code Server.

Note:

* AIW is the APPN Implementers Workshop who support the DLSw Related Interest Group (RIG) that evolved the RFC 1795 standard.

** DECnet IV over FR and X.25 (2210 to 2210) was introduced in a PTF to MRNS Release Manufacturing and Delivery (ISMD) as well being preloaded / shipped with current MRNS Release 2 orders.

2.6.8 The IBM 2210 As an IP Router

The IBM 2210 supports three dynamic routing protocols. All these three routing protocols can run simultaneously on the IBM 2210.

The IP dynamic routing protocols supported by IBM 2210 are:

- Routing Information Protocol (RIP)
- Open Shortest Path First (OSPF)
- Exterior Gateway Protocol (EGP)

Additionally, the IBM 2210 implements IP multicasting routing protocols MOSPF and DVMRP.

The IBM 2210 supports ARP Subnet Routing (RFC 1027), also known as Proxy-ARP, and static routing.

This section describes the IP routing implementation on the IBM 2210.

The IBM 2210 implements the following IP functions:

- **IP**

  This is an unreliable and connectionless delivery mechanism which defines the IP datagram and specifies the delivery of these datagrams across the underlying network.

- **ICMP**

  Internet Control Message Protocol is used to report errors and provide information about unexpected circumstances. It includes support of Echo Request/Reply messages (known as PING), redirect messages (to direct a host to use another hop) and Source-Quench messages (used for congestion control).

- **TCP**

  Transmission Control Protocol is the connection-oriented protocol that allows the reliable stream delivery of data across a network from a TCP module on one machine to a TCP module on another machine.

- **Telnet**

  A simple remote terminal protocol that allows a user at one site to establish a TCP connection to a Telnet server at another site.

- **UDP**

  User Datagram Protocol provides a mechanism that allows application programs to send datagrams to other application programs.

- **SNMP**

  Simple Network Management Protocol is used to monitor IP routers and the network to which they attach.

- **TFTP**

  Trivial File Transfer Protocol is a simple file transfer protocol which runs on top of UDP.

- **BOOTP**

  The Bootstrap Protocol is used by diskless machines to learn their IP address and the location of the boot file and boot server.

2.6.8.1 General IP Parameters

When planning to use the IBM 2210 as an IP router, there are a number of IP parameters that you may configure regardless of the routing protocol used in your IBM 2210. These parameters are:

- **Internal IP address**

  You may assign an internal IP address to the IBM 2210. The internal address belongs to the router as a whole, and not to a particular interface. This address is always reachable as long as one interface on the router is active. This address is also used by Data Link Switching (DLSw) feature.

- **Router ID**

  You may also assign a router ID to your IBM 2210. This is the default IP address used in various kinds of IP traffic originated from the router. For example, it is used as the IP source address in PING, TFTP or Traceroute packets.

- **Routing table size**

  Each IBM 2210 has a routing table which contains the dynamic routing information known by your router. Each entry in the routing table is 64 bytes, and by default, the routing table size is 768 entries.

  You may change the number of entries in the IP routing table based on the requirements of your network.

- **Router cache size**

  The IBM 2210 uses a routing cache which contains the recently routed destinations. The router will reference the cache first before using the routing table. The minimum and default size for the router cache table is 64 entries. However, you may change the router cache size based on your requirements.

- **IP broadcast format**

  IBM 2210 allows you to specify the format that is used by your IBM 2210 when broadcasting packets out on a specific interface. On doing so, you must specify the *style* and the *fill-pattern* used.

  The *style* parameter can be either *local-wire* or *network*.

  When you specify local-wire for the style, the router will use the broadcast address of either 255.255.255.255 or 0.0.0.0. The former is used if you have specified the fill-pattern to be 1, and the latter is used with the fill-pattern of 0.

  When you specify network for the style, the router will send the broadcast messages that begin with the network and the subnetwork portion of the IP address of the interface. The host portion of the broadcast messages are either all 1s or 0s depending on the value specified for the fill-pattern parameter.

  Note: When receiving messages, the IBM 2210 recognizes all forms of the IP broadcast addresses regardless of the settings of these parameters.

- **Reassembly size**

  You can configure the size of the buffers that are used for the reassembly of the fragmented IP packets received by the router.

  By default, IBM 2210 uses buffer of 12000 bytes.

  You can configure a route to a default gateway and the cost of reaching that default gateway. Normally, the default gateway is a router which has a more routing information about the network.

- **Default subnetwork gateway**

  In a subnetted network, you can configure a separate default gateway and the cost of reaching it, for each subnet network.

  All the packets detained for unknown subnets of a known subnetted network are forwarded to the subnetwork′s default gateway.

- **IP access control**

  The Access Control system allows the IBM 2210 to determine which packets are to be forwarded and which packets are to be discarded. For more information, refer to 2.6.8.10, “Access Controls” on page 106.

2.6.8.2 Interface Address Assignments

When you assign IP addresses to the router, you must note the following:

- You must assign at least one IP address to an interface. A hardware interface does not accept or send IP packets unless it has at least one IP address.
- It is possible to assign more than one IP address to an interface.
- You must specify an IP address together with its subnet mask.

Note

Serial lines do not *need* addresses. Such lines are called un-numbered and can be configured without IP addresses, but you must still enable them for IP traffic using the following command:

```
IP Config>Add address 1 0.0.0.1
```

Using un-numbered serial lines has some restrictions which are documented in information APAR II08361.

2.6.8.3 RIP Implementation in IBM 2210

The following must be considered when configuring RIP for your IBM 2210:

- Only the network portion, as defined by a mask, is entered into the routing table.
- Masks are not sent in RIP broadcasts.
- Maximum number of hops is 15 and a hop count of 16 indicates infinity.
- Destination entries timeout after three minutes.
- RIP updates are sent every 30 seconds.
- Variable length subnet masks are not supported.
- RIP is not supported across X25 circuits.
- Split horizon is always used.
- Poison reversed may be enabled for individual interfaces.
- The 2210 does not accept host-routes in RIP updates.

RIP Interoperability with 6611 Network Processor: To use RIP between the 6611 Network Processor and the IBM 2210 you need to take the following into consideration:

1. The broadcast address type used by the IBM 2210.

   The 6611 only recognizes *Local-wire* broadcasts. In our case, testing with V1R3 of MPNP, we found that both filling types are accepted. So broadcasting to 255.255.255.255 or 0.0.0.0 are both accepted by the 6611.

2. IBM 2210 does not accept host IP routes.

   The 2210 does not accept host-routes in a RIP response. The 6611 will advertise only the host address (not the network address) for the attached neighbors using point-to-point protocol (PPP).

3. The RIP version configured for 6611 Network Processor.

   6611 Network Processor can be configured to use either RIP Version 1 or RIP Version 2. IBM 2210 only supports RIP Version 1. Therefore, when using RIP between the IBM 2210 and the 6611 Network Processor, the 6611 must be configured to use RIP Version 1.

2.6.8.4 OSPF Implementation

OSPF implementation sets the OSPF router ID to the address of the first OSPF interface appearing in the router's configuration. However, you may change the router ID using the configuration commands from the ASCII console or the General panel in the IP subdirectory of the Nways MRNS Configuration Program.

Note

When you change the router ID of your IBM 2210, the link state advertisements originated by the router before the router ID change may persist in the network for as long as 30 minutes. This may cause an increase in the size of link state database.

The OSPF implementation in the IBM 2210 provides support for TOS-based (Type Of Service) routing for TOS 0 only.

IBM 2210 provides support for *simple password*, allowing for the authentication of the link state advertisement received from the other routers. To provide authentication, you must do the following:

1. Specify authentication type 1 when you define the OSPF area.
2. Specify the authentication key to be used when you configure the OSPF parameters for each interface.

You can import routes learned from other protocols (EGP, RIP or static routes) into the OSPF domain when the OSPF router is configured as an AS boundary router. An OSPF router can also originate a default route into the area. For these purposes you need to enable *AS boundary routing*.

OSPF and Non-Broadcast Networks: If the IBM 2210 is connected to a non-broadcast multiaccess (NBMA) network and is eligible to become the *designated router*, you need to provide the router with the information to find its OSPF neighbor(s). You can achieve this by performing the following tasks:

- Define the interface to the NBMA network as *non-broadcast*.
- Specify the IP address of the OSPF neighbor(s) on the NBMA network.
- Configure your IBM 2210 to become the Designated Router.

In a *star* frame relay network with only 2210s, you can use the *OSPF point-to-multipoint frame relay* enhancement. Refer to Figure 42 on page 103 for an example of a *start* or *partially meshed*. This type of network is also known as a *spoke and hub* network.

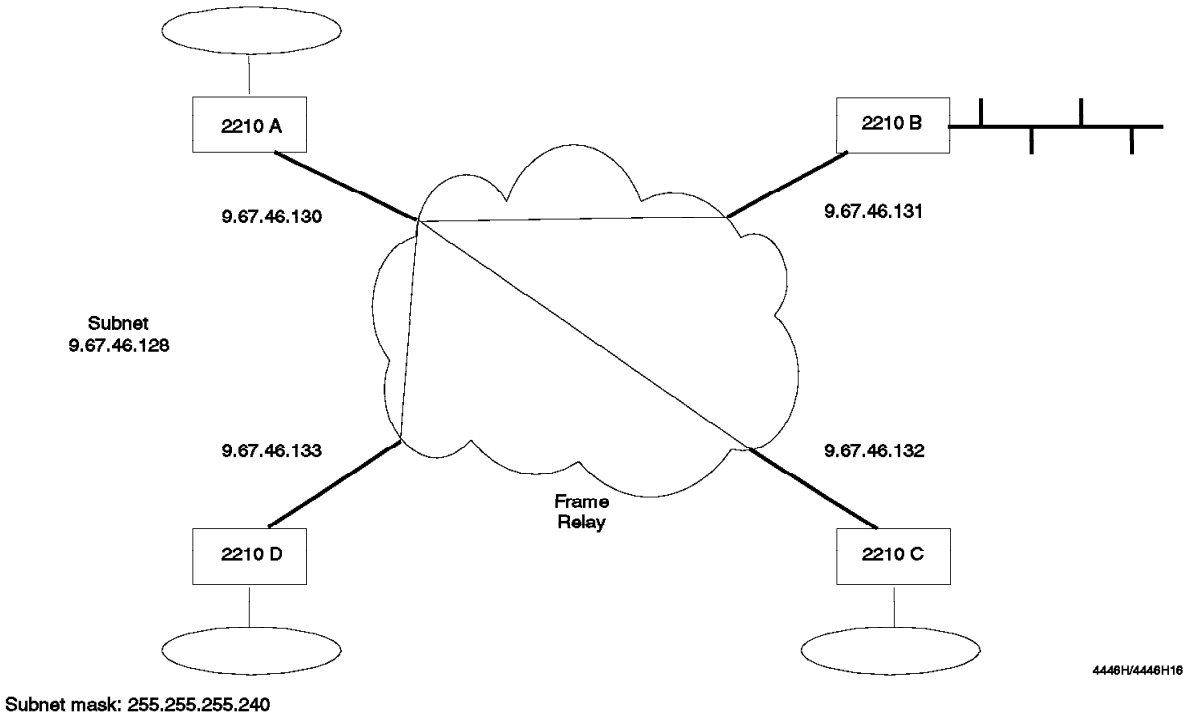

Figure 42. OSPF Point-to-Multipoint Frame Relay

Using the OSPF point-to-multipoint frame relay enhancement provided by IBM 2210, you may now assign a single IP subnet to an entire frame relay cloud and thus a single IP address to each frame relay interface of the router. In this case you only need to specify the OSPF neighbor at one side of each DLCI. In configuring such a network, you need to perform the following tasks:

1. Assign an IP address to the frame relay interface.
2. Enable OSPF on this interface.
3. Define the OSPF neighbor on one side of each DLCI (PVC).
4. To prevent one of the *spokes* from becoming the designated router, specify a *router priority* of 0 for the spokes and anything else but 0 for the hub router.

Note

In this type of OSPF configuration environment,it is not necessary to use the `set non-broadcast` command for each interface. By *not* using this command the router will determine that you intend to use the OSPF point-to-multipoint frame relay enhancement.

OSPF Interoperability with 6611 Network Processor: There are no specific OSPF considerations for connecting the IBM 2210 to the 6611 Network Processor when using OSPF.

Concerning frame relay, OSPF and 6611 interoperability, two scenarios were tested, scenario A and B:

- A: A fully meshed frame relay network with two 2210 routers and one 6611.

- B: A partially meshed frame relay network in a star configuration where the 6611 is the hub and the 2210 routers are the spokes.

Scenario A: Below, the steps concerning frame relay and OSPF are summarized, including the 6611 basic definitions:

- Assign an IP address to the 2210 frame relay interface.
- Enable OSPF and assign the interface to be an OSPF interface.
- Specify the interface as *non-broadcast*.
- Specify the 6611′s IP address as your OSPF neighbor on that interface and make it eligible to become the designated router.

On the 6611:

- Assign an IP address to the 6611 frame relay interface.
- Specify this interface as *fully meshed*.
- Enable OSPF and assign the interface to be an OSPF interface.

The interface type on the 2210 is multispecifying nonbroadcast multiaccess (NBMA).

Scenario B: The differences are summarized in the steps below:

- Assign an IP address to the 2210 frame relay interface.
- Enable OSPF and assign the interface to be an OSPF interface.
- Specify the 6611′s IP address as your OSPF neighbor on that interface and make it eligible to become the designated router.

On the 6611:

- Assign an IP address to the 6611 frame relay interface.
- Specify the DLCIs with their destination IP address as point-to-point links.
- Enable OSPF and assign both interfaces, represented by the IP destination address, as OSPF interface.

The interface type on the IBM 2210 is point-to-multipoint. Using this configuration, the spoke routers can still reach each other via the hub. The 6611 will take care of the routing between the spokes.

2.6.8.5 MOSPF

Multicasting is already used within OSPF. OSPF packets are sent to a standard multicast IP address of 224.0.0.5.

The 2210 extends this mechanism by implementing Multicast OSPF (MOSPF). When you enable the multicast forwarding capability, for each interface you can specify the following:

- Enable multicast forwarding on the interface
- Enable the forwarding of multicast packets as unicast or multicast
- Configure the IGMP polling interval
- Configure the IGMP local database timeout

The MOSPF function is used by the IBM 2210 for DLSw and IP Tunneling. Both implement client/server groups and peer groups for partner definitions.

DLSw uses a base multicast address of 225.0.1.0 for client and peers and an address of 225.0.65.0 for servers. The last octet of this address is used to identify the DLSw group number of the client/server group or peer group.

The IP bridge tunnel uses 224.168.0.0 as a base address for client/server groups as well as for peer groups. In this case the last two octets are used to identify a group.

Within this implementation it is also possible to manually change these addresses and to *join* or *leave* a multicast group specifying its IP address.

2.6.8.6 DVMRP

Distance-Vector Multicast Routing Protocol (DVMRP) allows you to define IP tunnels between MOSPF domains and a DVMRP domain/router. You can configure an IBM 2210 to use DVMRP, and define interface(s) to use it.

2.6.8.7 EGP Implementation

EGP implementation includes the following:

- You can configure the set of routes you want to exchange with a particular neighbor by using the *interchange flags* and the *interchange tables*. In addition, you can select the *cost* you want to assign to a route.
- An EGP router may advertise itself as the default router via its IGPs (OSPF and RIP). This is called *originating default*. For information about specifying as a default router, refer to the sections 2.6.8.3, “RIP Implementation in IBM 2210” on page 101 and 2.6.8.4, “OSPF Implementation” on page 102.

EGP Interoperability with 6611 Network Processor: There are no specific EGP considerations when connecting the IBM 2210 to 6611 Network Processor.

2.6.8.8 Static Route Implementation

You can define a static route for:

- Default gateway

  Packets are routed to the default gateway when the destination cannot be found in the routing table.

- Default subnet gateways

  If you are using subnetted networks, you can define a separate default gateway for each subnetted network.

- Static network/subnet routes

  For each destination that is to have a fixed route, you can define a static route.

2.6.8.9 IP Filters

You can use IP filters to prevent forwarding of the packets for a network or subnet. This includes distribution of routing information about these networks.

2.6.8.10 Access Controls

The access control system allows you to be much more specific in filtering IP traffic. You can control access to particular classes of IP addresses and services by controlling source and destination IP addresses, IP protocol number and port numbers for the TCP and UDP protocols.

When you enable access control and add an entry to the list, all the IP packets originated, forwarded, or received by the router are checked against the access control list. The following rules apply to this checking mechanism:

- For each packet received, the headers are compared to all the specified fields in each entry in the list.
- If the entry matches the packet and the entry is *inclusive*, the packet is forwarded.
- If the entry matches the packet and the entry is *exclusive*, the packet is discarded.
- If there is no match with the entries in the access control list, the packet is discarded.
- Each entry has an IP address as well as source and destination IP address.
- Each IP address is logically *ANDed* with the mask and compared to the address in the entry.
- A mask of 255.255.255.255 matches only the resulting address itself.
- A mask of 0.0.0.0 and the resulting address of 0.0.0.0 is a *wildcard*, and matches any IP address.
- Each entry may have an optional IP protocol number range. A range of 0 to 255 will match to all IP packets (within the address range).
- Each entry may have an optional port number range for UDP or TCP headers.

This implication of the above rules is that if you want to make one exclusion, you need to add inclusion(s) for all the other IP traffic you want to be forwarded by the router.

2.6.8.11 BOOTP Implementation

The IBM 2210 implements the *Boot Process (BOOTP) Client* function and the *Boot Process (BOOTP) Relay Agent* also known as *BOOTP Forwarder*. The 2210 may use the BOOTP client function to obtain its boot file (refer to 2.6.7.1, “Boot Files and Boot Processes” on page 89). It may also be configured to forward BOOTP requests to a BOOTP server.

The 2210 cannot act as a *BOOTP server*. You need a host running the *BOOTP* daemon. A BOOTP server contains a file that lists all the BOOTP clients that this server is responsible for, their associated IP addresses, and the location and name of their boot files.

The following is a summary of BOOTP process:

1. The BOOTP client copies its MAC address into a BOOTP packet (based on UDP) and broadcasts it onto the LAN.
2. If the BOOTP client and server are not on the same network, a local BOOTP relay agent will receive the request from the client, and route it to its defined

BOOTP server(s) or to the next BOOTP relay agent and route to the BOOTP server.

3. The BOOTP server receives the request and tries to match the MAC address with one in its list. If it finds a match, it will send a BOOTP reply with the client's IP address, subnet mask, and BOOTP server name. If the BOOTP client and server are not on the same network, the BOOTP reply may go through relay agent(s) to reach the client. In this case, the relay agent will receive a BOOTP reply, adds an entry to its ARP table and forwards the reply to the client.
4. The client uses the information that is contained in the reply to initiate a TFTP request to the TFTP server to download the boot image.

You need to assign two parameters when you define the router as a BOOTP forwarder (relay agent):

- The maximum number of hops you want the BOOTP request to go through. This is not the number of of IP subnetworks, but the number of BOOTP relay agents needed to get the server from the client (and vice versa).
- The number of seconds you want the client to retry before the BOOTP request is forwarded. BOOTP uses a technique of *timeout and retransmission*. When a client sends a BOOTP request, it starts a timer. If it does not receive a response before the timer expires, it retransmits the request. This process will be repeated the number of times that you have specified.

2.6.8.12 Telnet Implementation

To allow you to access the ASCII console interface remotely, the IBM 2210 implements the Telnet function. It allows you to have 5 Telnet sessions: two servers (inbound to the router), and three clients (outbound from the router).

From a Telnet session to the IBM 2210 does not provide you with any indication of which router your are logged-in to. You may determine the router by displaying the configuration information of the router. Alternatively, you may use (Ctrl+Break) to access the Telnet command mode. You can then issue the `status` command to display the IP address of the station that you are connected to as well as the current terminal mode.

2.6.8.13 SNMP Implementation

SNMP (simple network management protocol) runs on top of Users Datagram Protocol (UDP) and is used for monitoring and managing IP hosts in an IP network. SNMP enables network hosts, running vendor-supplied software, to read and modify some of the router's operational parameters. In this way, network management is established for the IP community. The software that processes the SNMP requests from the network management hosts runs on the IBM 2210 and is called an *SNMP agent*.

The following is the various aspects of the SNMP that you need to consider when configuring the SNMP for your IBM 2210.

Authentication: In SNMP you can define a *community*. The SNMP community is simply a group of nodes that share network management information. The community is established at configuration time.

The community allows you to define the IP address of the SNMP management station that is allowed to access the information in the SNMP agent's Management Information Base (MIB). It allows you to define a *community name* in accessing the MIB. The community name is used as an authentication scheme that prevents unauthorized users from learning information about an SNMP agent or modifying its characteristics. By defining an authentication scheme, you can provide security in your network management system.

Note: If no IP address is defined for the SNMP manager in your community table, any IP station that provides the correct community name will be able to access the MIB in the SNMP agent.

MIB Support: The operational parameters or variables are defined by an MIB (Management Information Base). The standard MIBs supported by IBM 2210 are described in Appendix D of *The Nways MRNS Protocol and Monitoring Reference*.

For each community name, you can specify which MIB or which part of an MIB can be accessed by the members of that community. To do so, you must first add one or more MIB Object IDs (the identification of a MIB item) to a *view*, creating a *sub-tree*. Then, you assign a view to a community.

Traps: SNMP agents can create *trap messages*. These are unsolicited messages that are sent from the router to an SNMP manager in response to a router or network event or condition, such as a router reload or network down. The IBM 2210 provides two types of traps which can be enabled or disabled separately for a specific community name:

- General traps

  These traps are defined by the RFCs and allow the router to send the traps asynchronously to the SNMP Manager in case of a specific event. There are six general traps defined:

  - Link-up
  - Link-down
  - Cold start
  - Warm start
  - EGP neighbor loss
  - Authentication failure

- Enterprise specific traps

  These traps are specific traps which can be generated by event logging system (ELS) messages. You can use the ELS *trap* command to enable sending of messages or groups of messages via an SNMP trap. To enable this to be forwarded by the SNMP agent of your router, you need to enable the trap type *enterprise*. However, the SNMP manager must support these enterprise traps because they are specific to the IBM 2210.

2.6.8.14 TFTP Implementation

The IBM 2210 implements the *TFTP client* function and the *TFTP server* function. The client function allows you to send or receive configurations or boot images to and from a TFTP server. The server function is implemented to provide other routers with a boot image or a configuration file. This implementation allows multiple, simultaneous file transfers between the router′s nonvolatile configuration memory (NVCNFG), the Integrated Boot Device (IBD), and remote hosts. Refer to 2.6.7.1, “Boot Files and Boot Processes” on page 89 for more information about the boot mechanism.

The TFTP implementation does not allow you to use PUT or COPY to transfer files to another router.

When a router acts as a TFTP server, transfers are transparent to the user. Use the ELS message log to view the transfers in progress. To view all TFTP messages, go to the ELS prompt of the GWCON and issue the following commands:

```
+ event
ELS>display subsystem tftp all
```

You can view the messages either by using the following command which displays the messages on the CONFIG console:

```
* divert 2 0
```

or, use the following command to view the messages on the MONITOR console:

```
* talk 2
```

2.6.8.15 ARP Subnet Routing

The IBM 2210 implements Proxy-ARP router function. When the router is configured for ARP subnet routing, it will reply by proxy to to the ARP requests for destination which are reachable via the 2210′s interfaces.

2.6.9 Data Link Switching

This section provides a brief overview of data link switching (DLSw) and discusses configuration of data link switching on the IBM 2210.

2.6.9.1 Data Link Switching Overview

DLSw is designed to facilitate integration of SNA traffic into a multiprotocol network. DLSw functions include:

- Transporting of SNA in a multiprotocol routed backbone
- Dynamic re-routing in the wide area network
- Reliable delivery of SNA traffic
- Termination of LLC acknowledgements on the LAN segments
- Broadcast traffic control through the WAN
- LAN and WAN control for congestion and data flow

DLSw uses IP encapsulation of SNA as its transport vehicle across the internetwork. To supply the reliability SNA requires in the internetwork, DLSw

uses Transmission Control Protocol (TCP) flows between edge-node routers (those routers joining the LAN segments to the IP portion of the network).

DLSw routers establish TCP connections to other DLSw routers using ports 2065 and 2067. Port 2065 is a read port on which all DLSw information is received, and port 2067 is a write port from which all DLSw information is sent.

DLSw also uses a technique known as DLC termination, or spoofing, to minimize T1 timer expirations and to keep acknowledgements isolated to the local LAN segment.

Spoofing is the process that acknowledges receipt of the frame on the local LAN segment by masquerading as the destination end station. Spoofing keeps the receiver ready and/or supervisory poll frames from leaving their subnet media. Therefore, it ensures local media response speeds to acknowledge layer 2 timers (T1 timers for example) and lessens the bandwidth overhead requirements in the WAN.

2.6.9.2 DLSw on the IBM 2210

The DLSw function of the IBM 2210 supports the interconnection of SNA devices attached to either a LAN (token-ring or Ethernet), or an SDLC multipoint non-switched line.

As a prerequisite for DLSw, if the IBM 2210 supports LAN-attached SNA devices, it must be configured to support source-route bridging on the token-ring interface, or transparent bridging on the Ethernet interface.

A DLSw virtual segment number also needs to be configured for IBM 2210s implementing DLSw. This virtual segment must be the same for all IBM 2210s participating in the DLSw function. This is to ensure that the end stations both see the TCP/IP network as one token-ring.

SNA devices attached to an IBM 2210 via SDLC multipoint non-switched lines are each assigned a token-ring locally administered address (LAA), service access point (SAP) and SNA XID (Exchange ID). These will be used by the IBM 2210 to represent such devices to other SNA devices that are using the DLSw function as if they are attached to a token-ring LAN. SDLC-attached devices can have SNA connections with token-ring and/or Ethernet-attached devices connected to the same IBM 2210.

SNA devices attached to an IBM 2210 establish connections with SNA devices attached to other IBM 2210s as if they are on the virtual segment.

SNA devices attached to an IBM 2210 via LAN segments establish connections with SNA devices attached to the same IBM 2210 via SDLC as if they were on the virtual segment.

Data Link Switching Supported Topology: There are two types of data link switching:

- Local data link switching
- Remote data link switching

In local DLSw, the data link switching function is performed within a single IBM 2210. In remote DLSw, stations attached to two or more IBM 2210s communicate across an IP network using DLSw.

Local Data Link Switching: Local DLSw allows communication between a token-ring or Ethernet-attached SNA device and an SDLC secondary PU2.0 or PU2.1 station that is link attached to the IBM 2210.

With Version 1 Release 2 of the IBM 2210 Nways MRNS software, both PU2.0 and PU2.1 link stations can coexist over SDLC lines at the same time.

The LAN-attached device is locally attached to the same IBM 2210 or attached to a remote LAN which is bridged to your IBM 2210.

Each SDLC-attached PU2.0 or PU2.1 device is assigned a MAC and SAP address and will appear to the other SNA devices as if it is attached to a token-ring LAN on your IBM 2210. Local DLSw converts SDLC frames to LLC2 frames. The encapsulated SDLC frames are passed to the DLSw function which will in turn use source-route or transparent bridging function to deliver them to the LAN-attached device.

Remote Data Link Switching: SNA stations attached to an IBM 2210 via a token-ring, Ethernet or SDLC connection can establish sessions with other SNA stations which are attached to a remote IBM 2210 or 6611 Network Processor via a token-ring or an Ethernet connection. The connection between the two IBM 2210s or between the IBM 2210 and the 6611 Network Processor is over an IP network which can include OEM routers which support compatible IP functions such as RIP or OSPF. Note that only the two routers connected to the end stations must be enabled for DLSw. DLSw function is not required in the routers which might exist between the two edge-node routers.

The DLSw in the IBM 2210 encapsulates the SNA frames in a TCP/IP datagram and delivers the encapsulated frames to its partner over the IP network.

Remote DLSw supports:

- SDLC to LAN over WAN

  SDLC frames are converted into LLC2 frames. This allows a link-attached SDLC secondary device to communicate with a LAN (token-ring and Ethernet) attached device.

- LAN to LAN over WAN

  Remote DLSw allows communication between SNA devices attached to token-ring or Ethernet networks. Remote DLSw can convert frames between the token-ring and Ethernet allowing token-ring and Ethernet-attached devices to communicate to each other using DLSw.

DLSw Using MOSPF: The IBM 2210 supports use of the DLSw Group Membership function to allow it to dynamically discover its DLSw partners, instead of having to manually configure the partner addresses. This feature utilizes the Multicast OSPF (MOSPF) function, which is described in 2.6.8.5, “MOSPF” on page 104.

The DLSw Group Membership defines two types of group:

- Client-to-server
- Peer-to-peer

Client-to-server groups have members that are designated either a client or a server. Server routers only form DLSw connections with client routers. This group type is used for subarea SNA connections. *Peer-to-peer* groups have

members that are all designated peers. All members of a peer-to-peer group will form DLSw connections with all other members of the group. This group type could be used for APPC connections.

DLSw group membership will only work between routers that support it, so a combination of group membership and preconfigured DLSw partner definitions may be required in your network.

2.6.10 Features and Facilities

This section describes the different features provided by the IBM 2210, the Bandwidth Reservation (BRS), the MAC Filtering (MCF), and the WAN Restoral (WRS) also called Dial Backup. It describes also some facilities provided by the IBM 2210 such as the dial-on-demand, NetBIOS name caching, and NetBIOS filtering.

2.6.10.1 Bandwidth Reservation (BRS)

In this section, we explain what is the Bandwidth Reservation feature, we show the Bandwidth Reservation configuration commands, and a scenario of Bandwidth Reservation is provided.

Introduction to Bandwidth Reservation (BRS): The Bandwidth Reservation feature allows allows you to reserve part of the bandwidth on the link for a specific traffic type.

Note:

- For Version 1 Release 1 of the Nways MRNS software for the IBM 2210, Bandwidth Reservation (BRS) is supported only over PPP serial links and applies to outbound traffic only.
- For Version 1 Release 2 of the Nways MRNS software for the IBM 2210, Bandwidth Reservation (BRS) supports the point-to-point protocol, frame relay, and dial circuits (ISDN and V.25 bis). Again this applies to outbound traffic only.

Figure 43 on page 113 shows specific data streams assigned to a part of the WAN bandwidth.

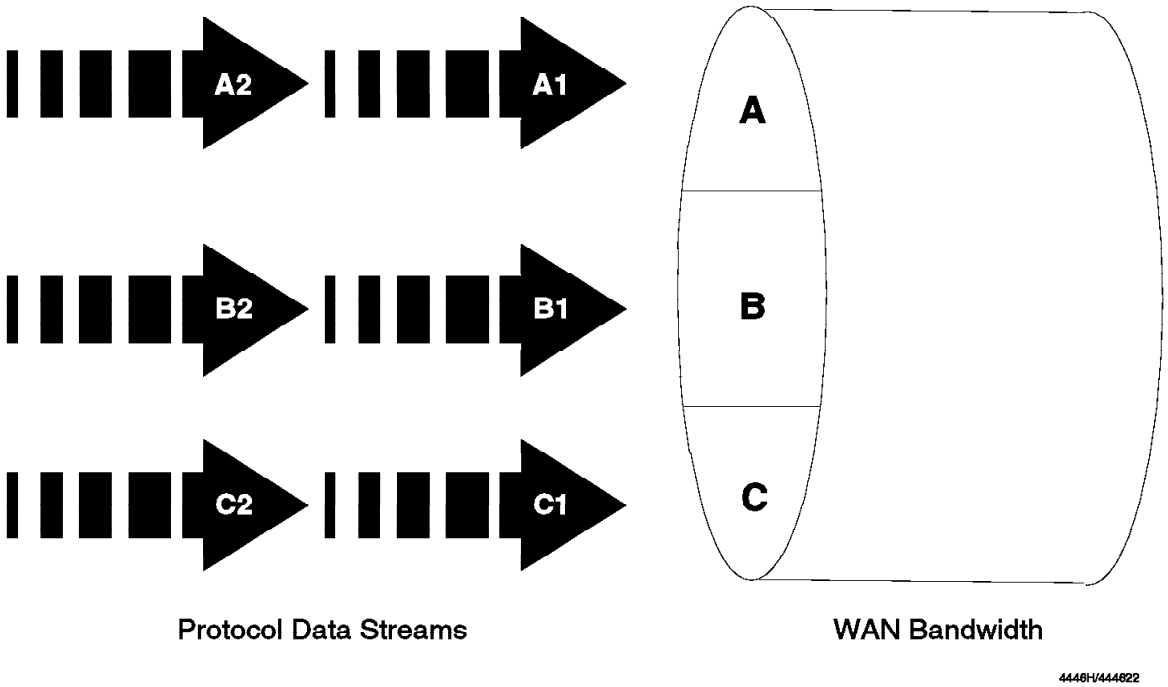

Figure 43. Bandwidth Reservation

First of all, you assign a name to a percentage of the bandwidth. This is called *class name*.

Note: All the names of the classes are case sensitive.

By default, there are two classes of names that you can neither delete nor change. You are just allowed to change their percentage of the bandwidth. These two classes by default are:

- *LOCAL* with 10% of the bandwidth by default
- *DEFAULT* with 40% of the bandwidth by default

The total of all the percentages of all the classes defined must not exceed 100%.

The reserved percentages are the guaranteed minimum slice of the bandwidth for the network connection. If the network is operating at full capacity, the messages from a specific traffic class can only be transmitted as long as they don't use more bandwidth than allocated for that class. If the rate of the messages exceeds the reserved bandwidth, the messages are held until other bandwidth transmissions have been satisfied.

In the case of light traffic on the network, packet streams can use bandwidth exceeding their allowed minimum (up to a maximum of 100% of the bandwidth) if there is no other traffic.

When you assign a class to a type of traffic, you will must also assign the priority class of this traffic within its class. There are four priority classes:

- Low
- Normal

- High
- Urgent

For example, a traffic assigned with class *DEFAULT* and priority *urgent*, will be delivered faster than a traffic assigned with class *DEFAULT* and priority *normal*.

The priority setting within the bandwidth class has no effect on other bandwidth classes. That is, none of the bandwidth classes has priority over the others.

Note: If no priority is assigned within a class, the default priority is *normal*.

After defining the class names, you may assign these classes to the following traffic types:

- The *DEFAULT* traffic class

  The *DEFAULT* traffic class is used by all the traffic that is not assigned to a specific class. By default, the DEFAULT traffic class uses the class *DEFAULT*, with the default class priority *normal*.

- The *protocols* (IP, ARP, IPX, ASRT, APL or AP2)

  For the protocols, you can assign a specific class and priority for each of the following protocols:

  – IP
  – ARP (with ASCII console only)
  – IPX
  – ASRT (Means bridged traffic)
  – APL (AppleTalk phase 1)
  – AP2 (AppleTalk phase 2)

  Note: The ARP protocol is not currently available on Nways MRNS Configuration program. You must customize it via Nways MRNS program on ASCII console.

- The *filter* (RLOGIN_IP, TELNET-IP, NetBIOS, SNA Bridged, SNPM-IP, DLSw-IP, MULTICAST-IP, TUNNELING-IP and SDLC-IP)

  For the filters, you can assign a specific class and priority for each of the following filters:

  – RLOGIN_IP
  – TELNET-IP
  – NetBIOS (bridged NetBIOS traffic)
  – SNA (bridged SNA traffic)
  – SNMP-IP
  – DLSw-IP (SNA traffic via DLSw)
  – MULTICAST-IP
  – TUNNELING-IP (with ASCII console only)
  – SDLC-IP (with ASCII console only)

  The *TUNNELING-IP* filter and the *SDLC-IP* filter are not currently available on Nways MRNS Configuration program. You must customize them via Nways MRNS program on ASCII console.

- Five *TAGs* (from MAC filtering on bridged traffic only)

  You can assign a specific class and priority for the following tags defined by the MAC Filtering (MCF) feature:

  – TAG1
  – TAG2
  – TAG3
  – TAG4
  – TAG5

  Note: The TAG number is assigned to a bridged traffic with the MAC filtering features.

2.6.10.2 WAN Restoral (WRS)

This section provides a description of the WAN Restoral feature and its configuration commands. A scenario of how to configure WAN Restoral on the IBM 2210 is also provided.

Introduction to WAN Restoral (WRS): The WAN Restoral (WRS) feature, which is also called Dial Backup feature, allows you to back up a primary leased PPP serial link with a switched V.25 bis PPP serial link.

Note

Backing up of frame relay or X.25 serial link is not supported. WAN Restoral only supports backing up of PPP leased serial link.

The WAN Restoral feature is supported over every routed protocol (IP, IPX, AppleTalk and DLSw) and for every bridging method, including tunnel bridge.

The backup switched line supported by this feature is over V.25 bis modem. In a future release, the WAN Restoral with the backup serial line over ISDN serial line will be provided for IBM 2210 models 127 and 128.

When the IBM 2210 detects the loss of connectivity on the primary PPP serial link, it automatically dials the configured phone number to establish the dial connection via the V.25 bis modem.

There is only one remote phone number configured in the IBM 2210. This must be the phone number of the same remote IBM 2210 which is reached via the primary serial link.

When the switchover from the primary link to the backup link occurs due to the failure of the primary link, the whole set of protocols configured on the primary leased PPP serial link will be automatically switched over to the switched V.25 bis serial link. All the protocols, IP, IPX, AppleTalk, DLSw and all the bridging methods will survive the switchover to the switched V.25 bis serial link.

When the IBM 2210 detects that the primary PPP serial link has come back up, it automatically drops the V.25 bis dial connection and restores all the protocols to use the primary leased PPP serial connection.

Figure 44 on page 116 shows the typical configuration of a network using WAN Restoral.

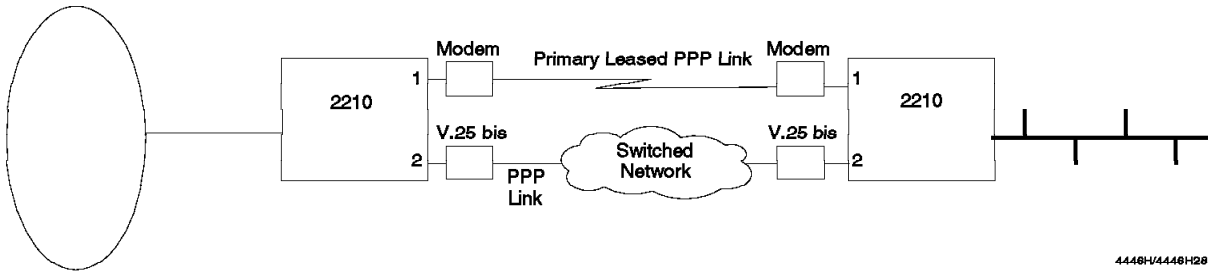

Figure 44. Typical Implementation of WAN Restoral

To be able to use the WAN Restoral, both 2210s at each end of the primary serial link must be customized for WAN Restoral.

To configure a 2210 to use WAN Restoral, you must customize one of its serial interfaces with PPP link, and the other serial interface as a dial interface using V.25 bis modem with PPP encapsulation method.

Since this feature is not supported by the 6611 Network Processor, the only possible way to use this feature in a network that includes 6611 Network Processor is shown in Figure 45. In this configuration, the IBM 2210 could detect the primary link failure and dial the 6611 Network Processor over the backup link.

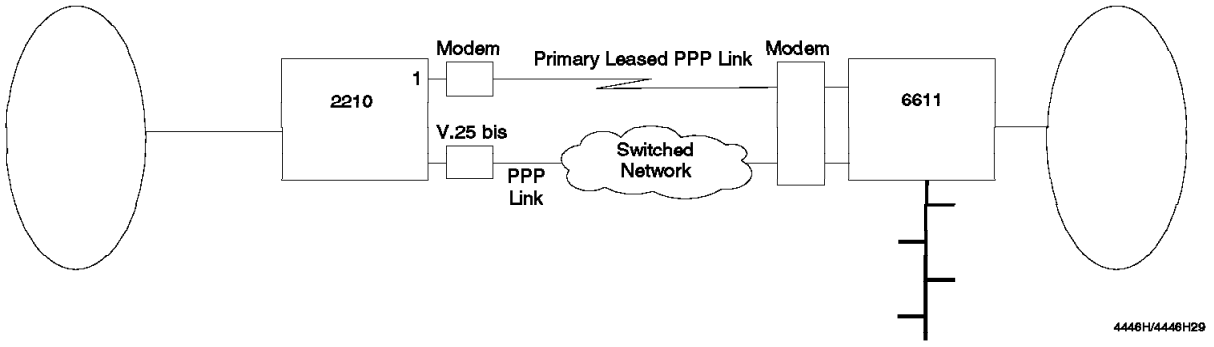

Figure 45. Possible Interoperability of WAN Restoral with 6611 Network Processor

2.6.10.3 Dial-on-Demand

This section provides a description of the dial-on-demand facility. It shows the dial-on-demand configuration commands, and provides an example scenario of the dial-on-demand configuration.

Introduction to Dial-on-Demand: The dial-on-demand facility is designed for remote sites that do not need to be connected to the central site all the time, but just when there is some data to be sent.

When the IBM 2210 detects that a packet needs to be sent over the switched network to a remote IBM 2210, it automatically dials the customized phone number to establish the dial connection via the V.25 bis modem.

You could customize several phone numbers in the IBM 2210, and map each remote phone number to a specific protocol address (IP or IPX address). However, note that only one connection to a remote site is allowed at any single point in time. This means that if there is already a connection to a remote site, you cannot send any packets to another remote site. In this case, you must wait until the first connection is terminated before trying to reach the second remote site.

To use the dial-on-demand facility, you must configure all the parameters of the desired protocol (IP or IPX) on the corresponding virtual dial-circuits and not on the physical V.25 bis interface.

When the IBM 2210 detects that no more packets are required to be sent over the switched interface for a certain lapse of time (idle time), the switched line is automatically dropped and the V.25 bis modem becomes available.

Note that when you customize a serial interface as a dial interface using V.25 bis modem with PPP encapsulation method, the other physical serial interface is able to be used for anything else at the same time. Also, both 2210s at each end of the primary serial link must be customized for dial-on-demand.

Note

It is recommended that you allow only one site to issue outbound calls, and the other site should allow inbound calls only. This will prevent dial collision in case both sides want to call each other at the same time. However, this is not a requirement and you can enable both sides for both inbound and outbound calls. In this case, you must be aware that if the IBM 2210s want to call each other at the same time, the V.25 bis modems will loop with DIALING:, then BUSY, then DIALING, then BUSY, etc. This will be repeated until one side decides to no longer send data to the other side. Then the switched link will be activated from the other side.

For IP routing over dial-on-demand, it is recommended that you customize static routes. This prevents the IBM 2210 from establishing the connection for each routing table update which is sent by the dynamic routing protocols.

If there is DLSw customization over a dial-on-demand circuit, be sure to not enable the Keepalives parameter. By enabling this parameter to verify that the remote DLSw partner is alive, the dial-up connection would remain active permanently.

IPX does not provide static routing, therefore, you are advised to specify large RIP and SAP update intervals to ensure that the dial-on-demand circuits are not frequently established as a result of the frequent RIP and SAP messages in an IPX environment.

Note

Dial-on-demand cannot be used to provide additional bandwidth over a switched serial interface in case of over utilization of the bandwidth of a primary leased serial interface.

The dial-on-demand facility is only supported over:

- TCP/IP (including DLSw and Tunnel Bridge)
- IPX protocol

Note

Dial-on-demand is not supported for any bridging methods, except for tunnel bridge method which is actually using IP protocol over the serial links.

The dial-on-demand is only supported over a switched V.25 bis PPP serial link.

Figure 46 shows you a typical drawing of a dial-on-demand network.

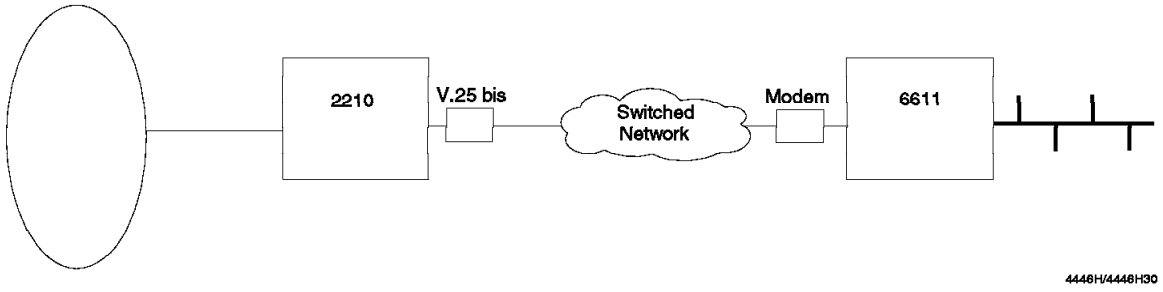

Figure 46. Typical Implementation of Dial-on-Demand

This facility is not supported by the 6611 Network Processor; therefore, the only possible way of using this feature in a network which includes 6611 Network Processor is shown in Figure 47 on page 119. In this configuration, the IBM 2210 could dial the 6611 Network Processor when it has data to send to the 6611. But, if the switched link is not up and the 6611 Network Processor has to send data to the IBM 2210, it must wait until 2210 establishes the call and this will happen when the IBM 2210 has data to send to the 6611 Network Processor.

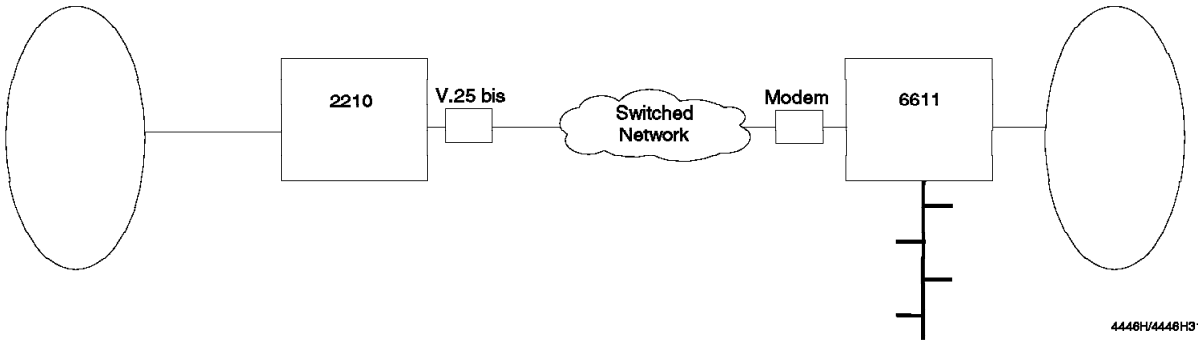

Figure 47. Possible Interoperability of Dial-on-Demand with 6611 Network Processor

2.7 IBM 6611 Router

This section provides a summary of the hardware and functions of the IBM 6611 Network Processor when used with the IBM Multiprotocol Network Program.

Further information on the IBM 6611 Network Processor hardware can be found in the *IBM 6611 Network Processor - Installation and Service Guide*.

Further information on the functions provided by the IBM 6611 Network Processor when used with the Multiprotocol Network Processor can be found in the *IBM 6611 Network Processor - Introduction and Planning Guide*.

The IBM 6611 uses its bridging, routing and data link switching functions to receive and transmit multiple protocols from one LAN to another. The Multiprotocol Network Program provides the necessary configuration functions to support each protocol. The 6611 is not a gateway and therefore requires the end stations that want to communicate with each other to use the same protocol. The data link switching function encapsulates SNA and NetBIOS frames into an IP datagram for transport over a WAN. With all other protocols it uses the packet or frame format prescribed by that protocol to route or bridge that protocol. Each of the adapters has their own high-performance processors and are called *peer-capable* adapters. Except in the case of data-link switching, the adapter processors eliminate the need to pass packets to the system processor enabling faster system performance and packet transfer.

The Multiprotocol Network Program collects and stores status information about the IBM 6611 connections. Performance and other data are stored in its MIB variables. Traps are sent to the SNMP manager for events that occur in the network and router itself. The SNMP manager can then retrieve MIB information to help with problem determination.

The 6611 supports local or remote access and control via the System Manager component of the Multiprotocol Network Program. This program allows you to set passwords, run software and hardware diagnostics, view statistics and error logs and shut down the 6611. Access can be via a local or remote interface.

2.7.1 Hardware Overview

There are three main user components that make up the 6611:

- The IBM 6611′s Family
- The Multiprotocol Network Program (MPNP)
- The System Manager

There were many modifications of IBM 6611′s Family, as described below:

- New 6611 Model 120 configurations
- New 6611 Model 125
- New 6611 Models 145 and 175 replacing Models 140 and 170 respectively
- New adapters

2.7.1.1 Model 120 Enhancements

The following is a complete list of the Model 120 fixed configuration which will be available. The new configurations are:

- Four SDLC ports / two multi-interface serial ports
- One token-ring port and one Ethernet port
- Two token-ring ports
- Two Ethernet ports

The existing configurations are:

- One token-ring port and four SDLC ports
- One Ethernet port and four SDLC ports
- One token-ring port and one X.25 port
- One Ethernet port and one X.25 port
- One token-ring port and two multi-interface serial ports
- One Ethernet port and two multi-interface serial ports

The benefits with these changes are:

- **Expanded Configuration Options**

  These key new configurations will allow the 6611 to be used as a local bridge, both between like media as well as between disparate media. When used in conjunction with MPNP V1R3′s new Translational Bridging function, the 6611 Model 120 can now provide translational bridging between token-ring and Ethernet LANs.

- **Current Configurations Enhanced**

  The existing Model 120 configurations have been replaced by new configurations which utilize the new 6611 adapters, providing the improved performance and increased connectivity previously described. Even though the new 6611 adapters increase the number of ports per adapter, the Model 120s will still be limited to the same number of ports as today. In other words, if a combination adapter is used to achieve a configuration that is currently available today on the Model 120, then the second adapter slot will not be used.

For example, the one token-ring port and two multi-interface serial ports Model 120 configuration will now be handled by one adapter. The performance of the new Model 120 will be equivalent to the old Model 120 with the two adapters.

The Model 120 configurations involving a four-port SDLC adapter or an X.25 adapter will use both slots of the Model 120. The other configurations will use the new adapters.

IBM 6611 Model 120 is positioned for the small or remote office with *two* LAN attachments.

2.7.1.2 IBM 6611 Model 125

This open, two-slot model complements the Model 120′s fixed configuration offerings. This versatile new model provides the following benefits:

- **Flexible configurations**

  The Model 125 can support any of the wide range of new 6611 adapters up to a maximum of eight ports. In many instances, the Model 125, coupled with the new multiport and combination adapters, can support a configuration which previously required a four-slot Model 140, representing a significant savings.

- **Future flexibility**

  Unlike the Model 120, which is available only in fixed configurations that cannot be changed after installation, the Model 125 gives customers the ability, in the future, to change adapters as their network configuration needs change.

  Adapters ordered for a Model 175/145 can be installed and used successfully in a Model 125. This allows flexibility in using adapters as network needs change.

  Note:

  Please be aware that adapters ordered for a Model 125 cannot be used in a Model 175/145. If a Model 125 adapter is installed in a Model 175/145, the adapter is marked as invalid at IPL time. When a configuration is attempted to be loaded into the 6611, the configuration will be invalid since the adapter is invalid.

- **Full function**

  While the Model 125 is a relatively small box in terms of the number of adapters supported, it is supported by the same software as the larger 6611 models with no restriction on the available functions.

IBM 6611 Model 125 is also targeted at the small or remote office but it can handle up to *three* LANs and a couple of WANs.

2.7.1.3 IBM 6611 Models 145 and 175

As replacement models for the current Models 140 and 170, the Models 145 and 175 were designed to offer improvements in packaging and usability while maintaining the same external interfaces. In this way, customers can capitalize on the improvements provided while investing a minimal amount of time familiarizing themselves with the new models. The IBM 6611 Models 145 and 175 use the same physical environment.

These new four and seven-slot models support any mix of the new adapters and offer the following benefits:

- **Rack mount options**

  There are two rack mount features available for the Models 145 and 175. One is a set of brackets that attaches to the sides of the box and permit installation on any industry-standard 19-inch, two or four-rail open or closed rack (including the IBM 9309). This enables optimal use of space in wiring closets and machine rooms.

  If faced with installing a 6611 in an area which is densely populated with equipment, or is in a hard-to-reach location, customers may want to consider the sliding shelf feature. This exceptionally sturdy steel cantilevered shelf mounts on any industry-standard 19-inch rack and is equipped with a recessed handle which enables the shelf to be easily pulled forward, extending to a depth of 27 inches. When the 6611 is placed on the shelf, the user has full range of access to all sides of the machine, significantly simplifying installation and removal of adapters or other maintenance activities. The 6611 can be screwed onto the shelf, and the rubber feet sit in holes on the shelf to prevent slippage of the shelf.

- **Customer setup**

  The new models of the 6611 are designed to support customer setup, further streamlining the installation process. The new adapter features also support customer setup on the new models, making any future configuration changes easier to accommodate and schedule.

- **Space savings**

  The seven-slot Model 175, like the four-slot Model 145, is designed for horizontal installation on a rack or used stand-alone on a table or desktop. This represents a considerable space savings compared to its predecessor, the Model 170, which could be installed only in a vertical position. The Model 175 is also considerably lighter, weighing only 42 pounds fully populated, compared to the Model 170′s maximum weight of 88 pounds.

- **Usability improvements**

  To enable easier access for attachment of an ASCII display or SCSI tape drive for diagnostics or service, the S1 service port and SCSI port have been moved to the front of the box. This makes cabling between the devices easier, as well as reduces the risk of disturbing an installed adapter cable or power cord.

  A cable management bracket is provided as a standard feature for both the Model 145 and 175. This bracket mounts on the rear of the box to provide strain relief for adapter cables, as well as improve cable management by allowing each cable to be dressed through an individual opening.

- **External interfaces preserved**

  Although the packaging of the new models has changed, the interfaces customers use have remained the same as the predecessor models. The Model 145 and 175 use the same three-character display on the operator panel for information and error codes, support the function-rich System Manager for diagnostic and management tasks, and utilize the easy-to-use 6611 Configuration Program for initial and subsequent configurations. Use of these common configuration and management tools across the product line simplifies network operation and management, and protects customers′ investment in training and support resources.

- **Scalability**

  In the event a change in a customer's network configuration causes his requirements to exceed the capacity of his installed Model 145, a Model Upgrade is available to convert the Model 145 to a Model 175, enabling the use of three additional adapter slots.

  As network needs change, the adapters from Models 175/145 can be moved to another Model 175, 145 or 125. This allows flexibility in using adapters as network needs change.

  Note:

  Please be aware that adapters ordered for a Model 125 cannot be used in a Model 175/145. If a Model 125 adapter is installed in a Model 175/145, the adapter is marked as invalid at IPL time. When an attempted configuration is to be loaded into the 6611, the configuration will be invalid since the adapter is invalid. Also, the old adapters for the Models 140 and 170 will not work in the new Models 145 and 175.

IBM 6611 Model 145 is suitable for building a backbone in a location with a number of connections. It can handle 8 LANs for 16 serial connections.

IBM 6611 Model 175 is the largest 6611 model, which provides seven adapter slots that can support the connection of a maximum of 14 LAN ports or 28 WAN ports or a combination of LAN and WAN ports, each at less than their maximum capacity. Thus, IBM 6611 Model 175 is a solution for large regional headquarters and campuses.

2.7.1.4 New Adapters

The new adapter features apply to all models of the 6611. These adapters include a new processor and twice the memory as the previous 6611 adapters. The following are the benefits of the new adapters:

- **Increased port density**

  New LAN adapters, which offer either two token-ring or Ethernet ports are now available; a new WAN adapter is added which provides four serial ports. This doubles the number of LAN and WAN ports previously available for the 6611.

- **LAN/WAN combinations**

  In addition, two new combination adapters are introduced, each offering one LAN port (either token-ring or Ethernet) plus two WAN serial ports on a single adapter. This allows maximum flexibility while preserving adapter slots in all models.

- **Improved performance**

  In general, the new adapters perform better than the old adapters. A four-port serial adapter can fully load four serial lines at T1 speeds. At E1 speeds, the four-port serial adapter performs better than two of the old two-port serial adapters. A token-ring serial combination adapter can handle all the traffic that, previously, could be handled by two adapters (a token-ring and a two-port serial adapter). In the case of an Ethernet serial combination adapter, if the serial interfaces are heavily used with small frame sizes, there is a slight reduction on the Ethernet maximum throughput due to the processing power being shared with the serial interfaces.

The 6611PERF package on MKTTOOLS provides in-depth information on performance. Your IBM account representative will be able to provide you with a copy of this document.

- **Increased connectivity**

  All new adapters with multi-interface serial ports, including the new combination adapters, can support any of the following physical interfaces on any port, including a mix of different interfaces per card:

  – CCITT V.35 - at speeds from 9600 bps to 2.048 Mbps

  – CCITT V.36 - at speeds from 9600 bps to 2.048 Mbps

  – EIA 422/449 - at speeds from 9600 bps to 2.048 Mbps

  – EIA 232/CCITT V.24 - at speeds from 4800 bps to 19.2 Kbps

  – CCITT X.21 - at speeds from 4800 bps to 2.048 Mbps

  Selection of the interface is determined by the adapter cable. So, if a change in the network interface equipment is required in the future, only a new cable is needed to switch interfaces.

- **Investment protection**

  These adapters are all supported on the Model 140 and 170 as well as the new models. This enables customers with installed 6611s to exploit the versatility and performance improvements of these new adapters without requiring an investment in a new platform.

The following is a list of all of the types of adapters which will be available for any 6611 Model (note that Model 120 is available only in fixed configurations). Different adapters must be ordered depending on whether you're putting the adapters in a Model 125 or a Model 145/175. The new adapter types are:

- Four-port multi-interface serial adapter
- Two-port token-ring network 16/4 adapter
- Two-port Ethernet adapter
- Multi-interface serial/token-ring combination adapter
- Multi-interface serial/Ethernet combination adapter
- Two-port multi-interface serial adapter (new, reduced cost)
- One-port token-ring network 16/4 adapter (new, reduced cost)
- One-port Ethernet adapter (new, reduced cost)

The existing adapters are:

- Four-port SDLC adapter
- X.25 adapter

Note

The four-port SDLC adapter and the X.25 adapter are unchanged. The new processor and double the memory used by the new adapters are not applicable to the four-Port SDLC and X.25 adapters.

2.7.2 Multiprotocol Connectivity

The IBM 6611 Network Processor provides routing of the network layer protocols used by the following protocol suites:

- Internet Protocol (IP)
- Novell NetWare Internetwork Packet Exchange (IPX)
- Xerox Network Systems (XNS) Internet Transport Protocol
- DECnet Phase IV and DECnet Phase IV-Prime
- AppleTalk Phase 2
- Banyan VIrtual NEtworking Systems (VINES)

2.7.2.1 Communication Adapter Features Supported

The communication adapter features supported for each of the protocols that can be routed by the IBM 6611 Network Processor are summarized in Table 9.

Table 9. IBM 6611 Adapter Ports and Supported Protocols

Adapter Ports:	**Ethernet•**			**Token-Ring**		**Serial**			**SDLC**	**X.25**
Standard:	**Version 2**	**IEEE 802.3**		**IEEE 802.5**						**CCITT X.25**
Framing → Protocols ↓	**Type**	**LLC**	**SNAP**	**LLC**	**SNAP**	**PPP**	**Frame Relay**	**Token-Ring Bridge Program**	**SDLC**	**X.25**
IP	X		X		X	X	X	X		X
XNS	X	X	X	X	X	X	X	X		
IPX•	X	X	X	X	X	X	X	X		X
AppleTalk			X		X	X	X	X		
DECnet	X				X	X	X	X		
Banyan VINES	X		X	X	X	X	X	X		
SNA•	X		X	X		X	X	X	X	X
APPN•	X		X	X		X	X	X		X
NetBIOS•	X		X	X		X	X	X		X
Source-route Bridging				X	X	X	X	X		
Transparent Bridging	X	X	X			X	X			
Translational Bridging	X	X	X	X	X	X	X	X		

Note:

• Also supports native Novell 802.3 for IPX.

• To run APPN, DLSw must be configured. APPN also requires that DLSw or IP be configured for APPN network nodes to communicate across a WAN.

• For local DLSw of SNA, the configuration of IP is not required. For remote DLSw of SNA and NetBIOS, IP must be configured on the link between DLSw session partners.

All the protocol suites that are supported for a communication adapter feature can be used concurrently across the same communication adapter interface. For example,an interface on the Multi-Interface Serial Adapter can be configured to support the transport of TCP/IP, NetWare, XNS, DECnet and AppleTalk protocol suites concurrently.

This is possible because the data link protocols used by the communication adapter features that support multiple protocol suites provide a mechanism for distinguishing between the various protocol suites sharing the same communication interface.

For example, the PPP data link protocol uses a 2-byte protocol code within each frame to distinguish between protocol suites sharing the same communication interface.

Note: The communication adapter features supported for the TCP/IP protocol suite can also be used to support the transfer of information that originates from nodes that use either the SNA or the NetBIOS protocol suites. This is achieved using the IBM 6611 Network Processor data link switching function which encapsulates the SNA or NetBIOS protocols inside the TCP protocol. This is described further in 2.7.4, “Data Link Switching” on page 140.

2.7.2.2 Routing Table Maintenance

The IBM 6611 Network Processor uses separate routing tables for each of the protocol suites it supports. That is, there is one routing table for each protocol suite supported by the IBM 6611 Network Processor.

For the DECnet, XNS, NetWare, AppleTalk and Banyan VINES protocol suites, their routing tables are maintained using the corresponding routing table maintenance protocol dynamically. For example the XNS protocol suite uses XNS RIP (Routing Information Protocol) for this purpose.

For the TCP/IP protocol suite, several routing table maintenance protocols can be used either singularly or in combination to maintain the single TCP/IP routing table. Additionally, *static routes* can be manually defined during configuration of the IBM 6611 Network Processor.

The TCP/IP routing table maintenance protocols supported by the IBM 6611 Network Processor are:

- Interior protocols used within an autonomous system:

  TCP/IP RIP (Routing Information Protocol)
  Hello
  OSPF (OSPF)

- Exterior protocols used between autonomous systems:

  EGP (Exterior Gateway Protocol)
  BGP (Border Gateway Protocol)

2.7.2.3 Filtering

The IBM 6611 Network Processor multiprotocol routing function provides a very comprehensive filtering capability. There are three types of filtering provided:

1. Filtering based on protocol suite

   The routing of each supported protocol suite can be selectively disabled or enabled for each IBM 6611 Network Processor. That is, each IBM 6611 Network Processor can be configured to either ignore (filter) or route each of the supported protocol suites.

   For example, a IBM 6611 Network Processor can be configured to ignore the token-ring segments DECnet protocol suite, and only route the TCP/IP, XNS, AppleTalk and NetWare protocol suites. Frames received by the IBM 6611 Network Processor that are identified as DECnet will be discarded, and

frames received that are identified as either TCP/IP, XNS, AppleTalk or NetWare will be routed.

2. Filtering based on communication interface

   If the routing of a particular protocol suite is enabled for a IBM 6611 Network Processor, it can be selectively disabled or enabled for each communication interface. That is, each communication interface can be configured to either ignore or route a particular protocol suite.

   For example, a IBM 6611 Network Processor that is enabled for routing the TCP/IP protocol suite, can be configured to ignore the TCP/IP protocol suite on one of its communication interfaces, and only route the TCP/IP protocol suite on the remaining communication interfaces.

3. Filtering based on network layer address

   For each protocol suite the IBM 6611 Network Processor provides additional filtering capabilities that allow the enabling or disabling of routing based on network layer addresses. These filters are either specific to a particular communication interface or global to all communication interfaces.

   The specifics of these filters vary between protocol suites as each protocol suite uses a different form of network layer addressing.

2.7.3 Bridging with IBM 6611

The 6611 supports routing and three types of bridging:

- **Source-route bridging**

  Source-route bridging is used on the 6611 to bridge frames between token-ring LANs.

- **Transparent bridging**

  Transparent bridging is used on the 6611 to bridge frames between Ethernet LANs.

- **Translational bridging**

  Translational bridging allows you to bridge frames between token-ring and Ethernet LANs.

The following topics provide a brief description of bridging with 6611.

2.7.3.1 Source-Route Bridging

Source-route bridging is used to interconnect networks at the data link layer of the OSI reference model. Source-route bridging involves forwarding MAC frames based on information in the MAC header. A frame is passed from bridge to bridge until it reaches the final destination.

A bridge examines each frame to determine whether it is destined for the bridge itself or for another device. The bridge uses data from its tables or information in the frame header to determine whether the frame should be forwarded to another device. Source-route bridging depends on the device that sends the frame (the source) to indicate, within the frame, the complete route to the final destination. The route is a sequence of identifiers for the bridges and rings along the path from the source to the destination device.

Unlike a router, a bridge does not examine the network protocol header that is imbedded in the data field of the MAC frame. The bridge is unaware of the

network protocol information in the data field. Consequently, a bridge is sometimes referred to as protocol independent.

The 6611 can be configured to provide *local* or *remote* bridge functions.

Local Bridge Function: A single 6611 can be used to interconnect multiple token-rings that are directly attached to the 6611. Figure 48 illustrates this local bridge function.

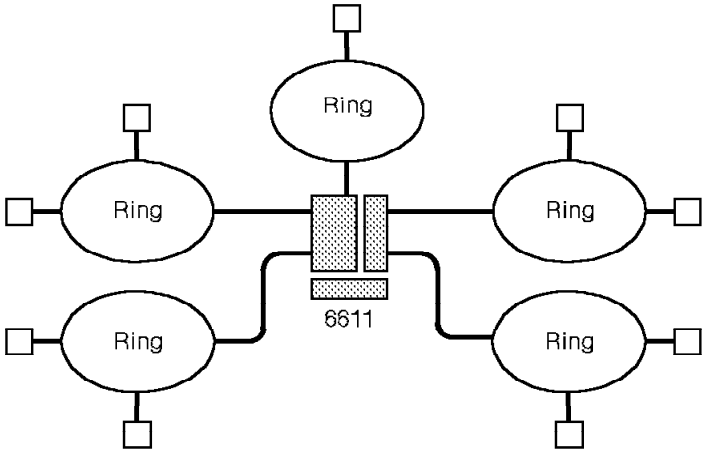

Figure 48. Local Source-Route Bridge Function

Each token-ring segment is attached to the IBM 6611 Network Processor using a IBM 6611 Token-Ring Network 16/4 Adapter. IBM 6611 Network Processor can be used to interconnect two or more token-ring segments across an intervening frame relay network or telecommunication link.

The IBM 6611 Network Processor when used as a source-route bridge can forward three types of frames:

All-Routes Broadcast

When the IBM 6611 Network Processor receives an all-routes broadcast frame on one of its token-ring interfaces, it copies the frame to all the other IBM Token-Ring Network segments to which it is attached. In doing so it updates the RI (Routing Information) field of each copy of the received frame with its bridge number, and the segment number of the destination token-ring segment. The RI field is also updated with the source segment number if it is not already present within the RI field.

Single-Route Broadcast

When the IBM 6611 Network Processor receives a single-route broadcast frame, it only copies the frame to the other token-ring segments if the corresponding interface has been enabled for the forwarding of single-route broadcast frames. Each interface can either be manually or automatically configured for the forwarding of single-route broadcast frames. The RI field for each copy of the received frame is updated in the same manner as for all-routes broadcast frames.

Non-Broadcast with Routing Information Field

When the IBM 6611 Network Processor receives a non-broadcast frame that contains an RI field it will forward the frame if the next entry in the RI field contains the bridge number of the IBM 6611 Network Processor and the segment number of a segment attached to the IBM 6611 Network Processor.

The IBM 6611 Network Processor is able to participate in the automatic configuration of the single-route broadcast function using the spanning tree algorithm with other source-route bridges that support this capability.

Remote Bridge Function Between 6611s: Two 6611s can be used to interconnect two or more token-rings across an intervening frame relay network or telecommunications link. Figure 49 shows two sample configurations that use this remote bridge function. The function is sometimes called *native mode bridging*, to distinguish it from the remote bridge function described below.

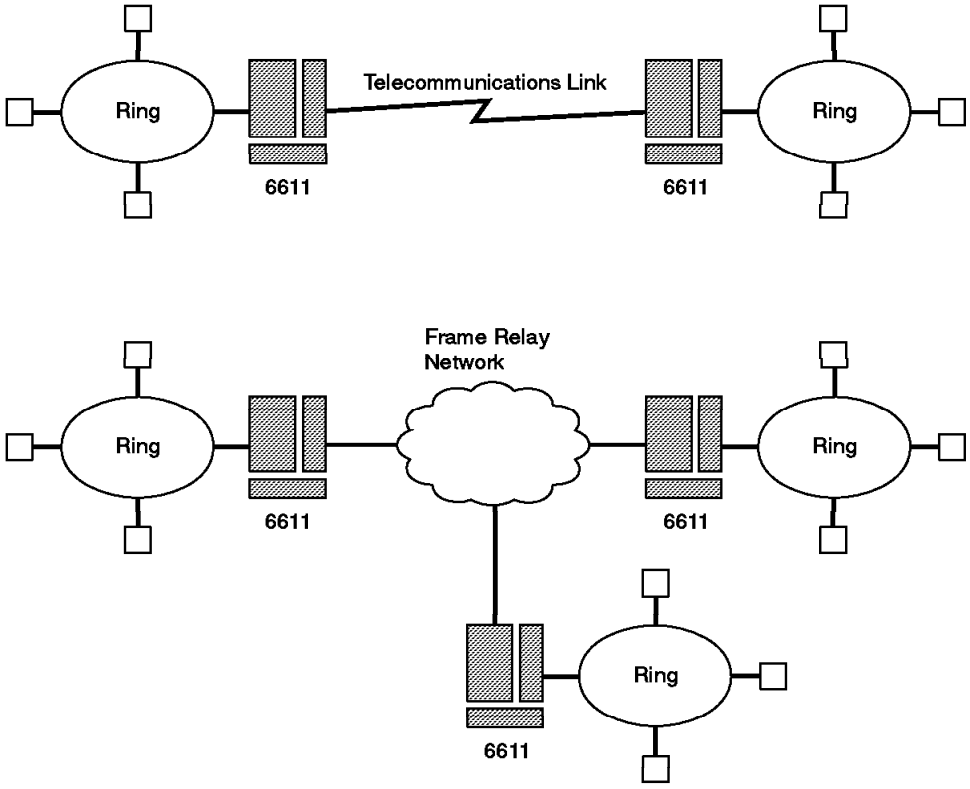

Figure 49. Remote Source-Route Bridge Between 6611s

Each token-ring segment is attached to a IBM 6611 Network Processor using a IBM 6611 Token-Ring Network 16/4 Adapter. The remote connections between each IBM 6611 Network Processor can utilize the 2 multi-interface serial ports, and can use either the PPP or frame relay data link protocols.

Each connection between IBM 6611 Network Processors can be either:

- A point-to-point communication facility such as the T1 or E1 services provided by many common carriers. Such a connection would use PPP data link protocols.

- A DLC (Data Link Connection) across a frame relay service. Many DLCs can share the same physical interface to a frame relay service using a unique DLCI (Data Link Connection Identifier) to distinguish between each DLC. This allows a IBM 6611 Network Processor to establish connections with many other IBM 6611 Network Processors using a single physical interface to a frame relay service.

The bridge number assigned to the IBM 6611 Network Processor will be used not only for bridging with remote token-ring segments attached to other IBM 6611 Network Processors, but also for local bridging and remote bridging with PS/2s.

Remote Bridge Function Between a 6611 and a PS/2: The IBM 6611 supports a remote bridging between a 6611 and a PS/2 workstation running either the IBM Token-Ring Network Bridge Program, Version 2.2, or the IBM Remote Token-Ring Bridge/DOS, Version 1.0.

Figure 50 shows a sample configuration using this remote bridge function. The function is sometimes called *compatibility mode bridging*. In this configuration, the 6611 functions as the primary half of the bridge and the Bridge Program functions as the secondary half of the bridge. A telecommunications link connects the 6611 to the PS/2 workstation running the bridge program. The devices communicate using a proprietary protocol.

Note:

The proprietary protocol used on the telecommunications link is referred to as the *LAN Bridging Protocol* within the 6611 library.

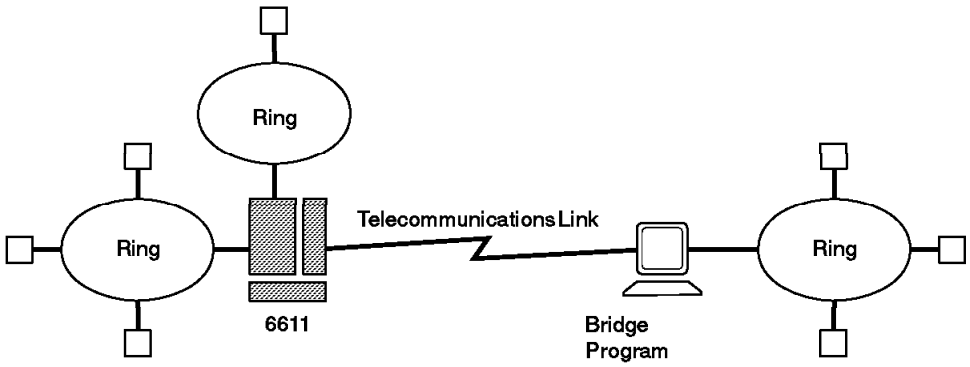

Figure 50. Remote Source-Route Bridge Between a 6611 and a PS/2 Workstation Running a Bridge Program

Token-ring segments are attached to the IBM 6611 Network Processor using the IBM 6611 Token-Ring Network 16/4 Adapter. Remote connections between IBM 6611 Network Processors and PS/2s utilize point-to-point protocol (PPP), and can be attached to the IBM 6611 Network Processor using the 2 multi-interface serial ports.

The bridge number assigned to the IBM 6611 Network Processor will be used not only for bridging with remote token-ring segments attached via PS/2s, but also for local bridging and remote bridging with other IBM 6611 Network Processors.

Additionally, one of the token-ring segments locally attached to the IBM 6611 Network Processor must be selected to become the *designated ring*. All the

PS/2 remote bridges connected to a IBM 6611 Network Processor are logically bridged to the designated segment. An example of how to use a *designated ring* is shown on Figure 51 on page 131.

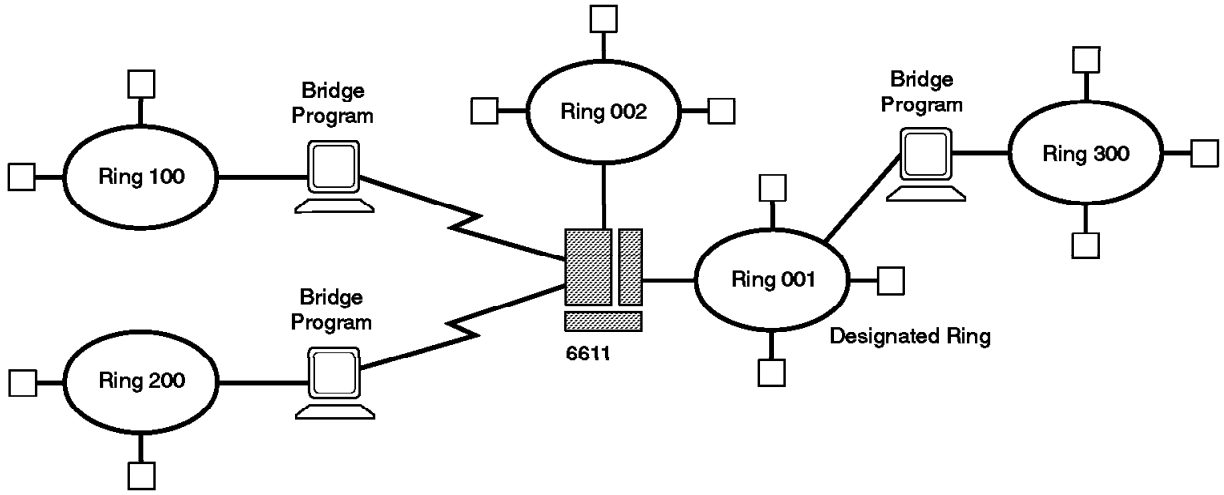

Figure 51. Remote Source-Route Bridge and the Designated Ring

Note: Frames transported by the IBM 6611 Network Processor between token-ring segments other than the designated segment, do not appear on the designated segment. Instead they are processed entirely within the IBM 6611 Network Processor. However, the designated segment number does appear in the RI field of frames transported to or from remote token-ring segments attached to PS/2 remote bridges.

Filtering: The IBM 6611 Network Processor source-route bridging function provides a very comprehensive filtering capability.

Filters can be configured for each communication interface that participates in source-route bridging. This includes interfaces on both the IBM 6611 Token-Ring Network 16/4 Adapter and the Multi-Interface Serial Adapter when remote source-route bridging is used.

For each communication adapter interface, both inbound and outbound filters can be configured. Inbound filters act upon frames received by the IBM 6611 Network Processor across the communication interface. Outbound filters act upon frames scheduled for transmission by the IBM 6611 Network Processor across the communication interface.

There are five types of filters which can be configured for each interface. With the exception of the hop count filter, each type can be configured separately for inbound and outbound operation. The five filter types available are:

Hop Count This filter can be used to process frames that have more than an allowable number of hops in their RI (Routing Information) field.

MAC Address This filter can be used to process frames that are to or from specific MAC (Media Access Control) addresses.

Source SAP This filter can be used to process frames that contain a specific source SAP (Service Access Point).

SNAP Value This filter can be used to process frames that contain a specific SNAP header. SNAP headers exist in frames that have source and destination SAP values of X′AA′.

Segment Number This filter can be used to process frames that contain a specific origin segment number within the RI (Routing Information) field.

Each type of filter only acts upon either single-route broadcast, or all-routes broadcast frames, or both. Each type of filter can be set to operate in one of two modes:

- Include only frames which match the filter characteristic (not used by the hop count filter). This is *permit* mode.
- Exclude only frames which match the filter characteristic (always used by the hop count filter). This is *deny* mode.

With the exception of the hop count filter, each type of filter provides the capability for multiple values to be filtered concurrently, and a mask capability allows a range of values to be specified with a single entry. Only those bits set in the mask are used for comparisons between the value specified and the frame being processed by the filter.

All five types of filters can be used concurrently if required. With the exception of the hop count filter, each type of filter can be individually enabled or disabled.

Notes:

Use of the SNAP value filter requires that the corresponding source SAP filter also be enabled. For example, to use the outbound SNAP value filter for an interface, the outbound source SAP filter for the same interface *must* also be enabled. No SAPs need be defined for the source SAP filter if only the SNAP value filter is required.

The hop count filter can be effectively disabled by setting the hop count value to *7 (seven)* which is the maximum hop count possible in token-rings.

To illustrate how multiple filters work together, consider the following scenario where outbound source SAP, outbound ring number and hop count filters are used concurrently for a token-ring interface. The filter settings are listed in Table 10.

Table 10. Example Filter Settings

Filter Type	Mode	Value(s)
Outbound Source SAP	Deny	X′AA′ X′F0′
Outbound Ring Number	Permit	X′100′ X′200′ X′300′
Hop Count	Deny	2

For a frame to pass through the interface for which these filters are enabled it must meet *all* of the following criteria:

1. It must have a source SAP that is not X′AA′ or X′F0′. For example, a frame with a source SAP of X′04′ would meet this requirement, whereas a frame with a source SAP of X′F0′ would not.

2. It must contain an origin segment number of X′100′, X′200′ or X′300′. For example a frame with a routing information field of X′100 1 300 0′ would meet this requirement, whereas a frame with a routing information field of X′400 1 300 0′ would not.
3. The routing information field must contain 2 hops or less. For example a frame with a routing information field of X′100 1 200 1 300 0′ would meet this requirement, whereas a frame with a routing information field of X′200 1 800 1 100 1 300 0′ would not.

2.7.3.2 Transparent Bridging

Transparent bridging, like source-route bridging, is a method used to interconnect networks at the data link layer. The 6611 supports Ethernet transparent bridging, as defined in the IEEE standard for Media Access Control Bridges (802.1D).

In source-route bridging, the device sending a frame discovers the preferred route to a destination device and that route is included within the frame transmitted by the sending device. In transparent bridging, a sending device transmits frames without regard for the location of a destination device. The bridges in the network are responsible for forwarding each frame to its proper destination.

Transparent bridges receive all frames transmitted on the LAN segments to which they are attached, and examine the source and destination addresses of each frame. By examining the source address of a frame, the bridge learns the port and LAN segment associated with a sending device. This information is stored in a routing table or *filtering database* and is used to make future decisions about how to forward frames. By examining the destination address of a frame arriving on a port, the bridge determines if the frame should be forwarded to another port or discarded (the destination device and sending device, in this case, are on the same side of the bridge). Each adapter maintains its own filtering database.

Transparent bridges, like source-route bridges, do not examine the network protocol header imbedded in the data field of the MAC frame. The bridge is unaware of the network layer protocols, and bridges all frames independently of these protocols.

The 6611 can be configured to provide the following transparent bridge functions:

- Local bridging
- Remote bridging

Local Bridging Function: A single 6611 can be used to interconnect multiple Ethernet LANs that are directly attached to the 6611. Figure 52 on page 134 illustrates this local bridge function.

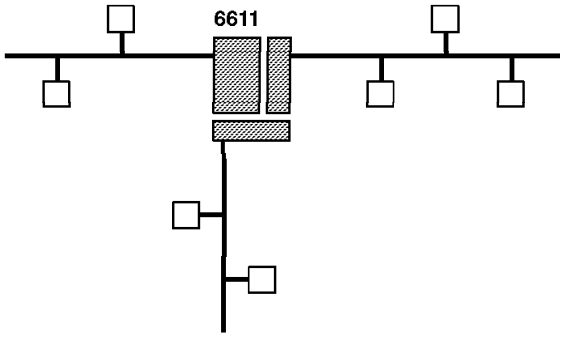

Figure 52. Local Transparent Bridge Function

Remote Bridging Function: Two 6611s can be used to interconnect two or more Ethernet LANs across an intervening frame relay network or telecommunications link. Figure 53 on page 135 shows several 6611 configurations using the remote bridge function. As indicated in the figure, Ethernet and token-ring frames can be transported over the same telecommunications link or frame relay connection.

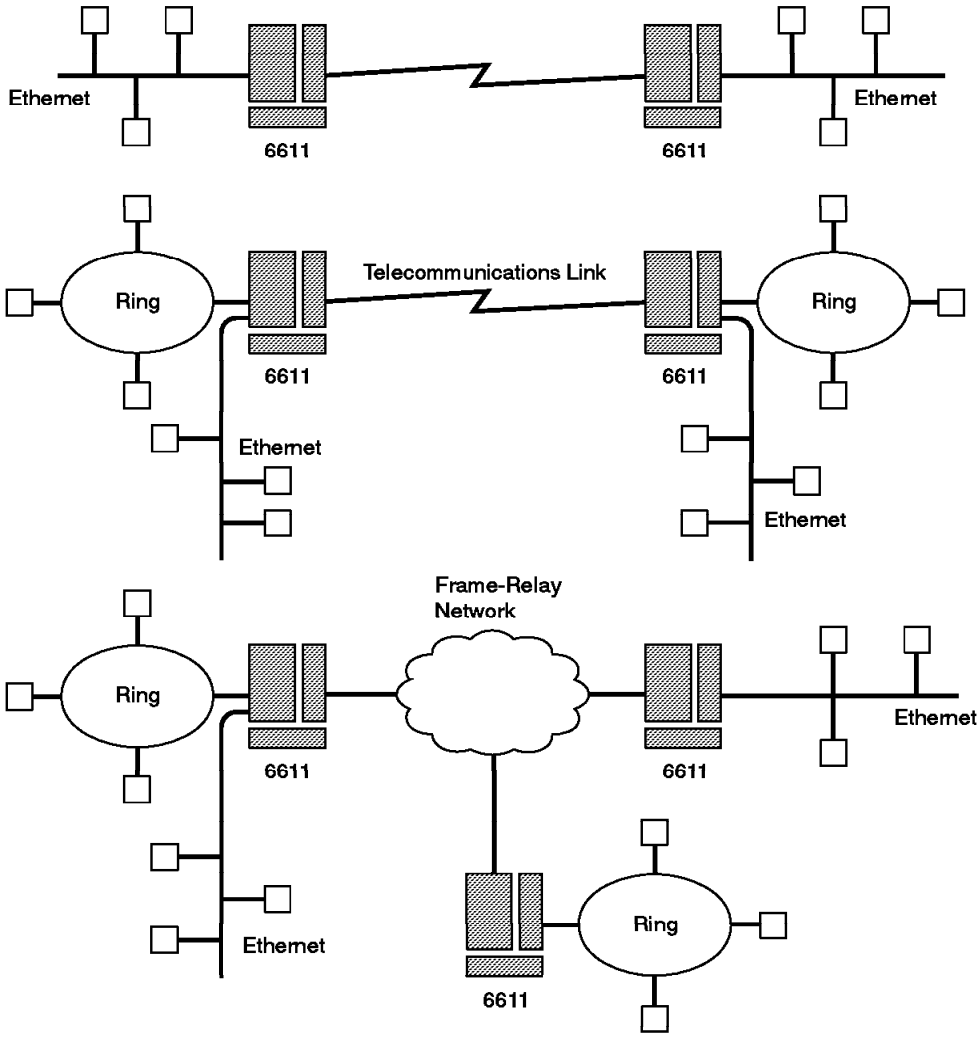

Figure 53. Remote Transparent Bridge Function

2.7.3.3 Translational Bridging

On the 6611, token-ring ports can be configured to support source-route bridging, and Ethernet ports can be configured to support transparent bridging. Because each LAN type uses a different frame format and bridging technique, token-ring and Ethernet LANs cannot be interconnected without providing a method of translation. *Translational bridging* is the method used on the 6611 to bridge frames between these different LAN types. Translational bridging, as implemented on the 6611, is sometimes referred to as *source route transparent bridging* (SRTB or SR-TB).

When you configure the 6611 node as a translational bridge, it operates in the following manner:

- If the source and destination ports for a frame use the same bridging technique, the frame is bridged between the ports without translation.
- If the source and destination ports for a frame use different bridging techniques, the translational bridge converts the frame into the format required for the destination LAN, and bridges the frame.

Frames in IEEE 802.5 format (for token-ring LANs) will be converted to either Ethernet Version 2.0 or IEEE 802.3 format as required by the destination Ethernet LAN. Ethernet frames will be converted to IEEE 802.5 format as required.

To a device on a token-ring LAN, the 6611 translational bridge appears as a source-route bridge. To a device on an Ethernet LAN, the translational bridge is functionally transparent. To enable it to interconnect token-ring and Ethernet LANs, the translational bridge maintains two address databases, as follows:

- The Ethernet database contains the source addresses for stations detected on Ethernet LANs and the frame format that each station uses for data transmission (Ethernet V2.0 or IEEE 802.3).
- The token-ring database contains the source addresses and routing information for stations on token-ring LANs that have forwarded frames to Ethernet LANs.

Notes:

- The translational bridging function on the 6611 is compatible with functions provided by the IBM 8209 and 8229 LAN Bridge products.
- The 6611 does not support *source-routing transparent (SRT) bridging*, which combines source-route bridging and transparent bridging techniques into a single bridging method for token-ring LANs.

The 6611 translational bridge can be configured to provide the following bridge functions:

- Local bridge function
- Remote bridge function between two 6611 translational bridges
- Remote bridge function between a 6611 translational bridge and a 6611 source-route bridge or transparent bridge
- Remote bridge function between a 6611 translational bridge and a PS/2 workstation running either the IBM Token-Ring Network Bridge Program Version 2.2, or IBM Token-Ring Network Bridge/DOS Version 1.0

Local Bridging Function: A single 6611 can interconnect multiple token-ring and Ethernet LANs that are directly attached to the 6611. Figure 54 on page 137 illustrates this local bridge function.

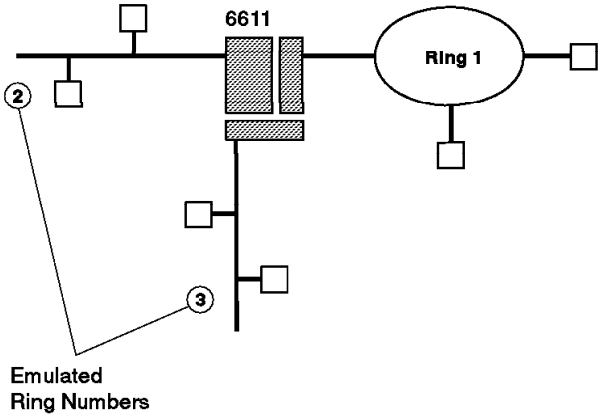

Figure 54. Local Translational Bridge Function

Remote Bridging Function Between 6611s: Two 6611s can be used to interconnect token-ring and Ethernet LANs across an intervening frame relay network or telecommunications link. Figure 55 on page 138 shows two sample configurations that use this remote bridge function. The recommended method for connecting two 6611 translational bridges is to configure *dual mode bridging* on each end of the serial link. When you configure dual mode bridging, bridged frames are translated only if the source and destination LANs require different MAC frame formats.

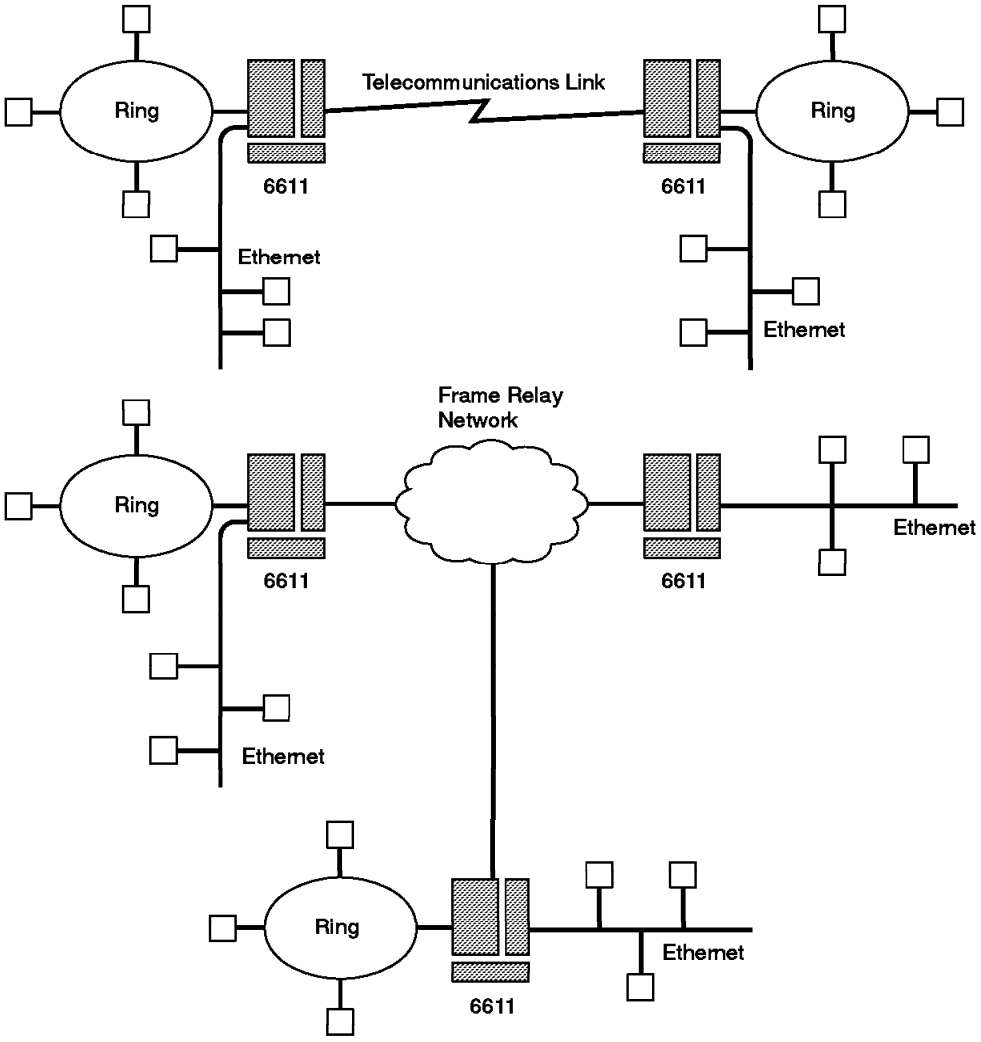

Figure 55. Remote Bridging Function Between 6611 Translational Bridges

Remote Bridging Function Between a 6611 Translational and Non-Translational Bridge: A 6611 translational bridge can be connected to a 6611 source-route bridge or transparent bridge across an intervening frame relay network or telecommunications link. The LANs attached to each bridge can communicate across the WAN connection. Figure 56 on page 139 shows a sample configuration that uses this remote bridge function.

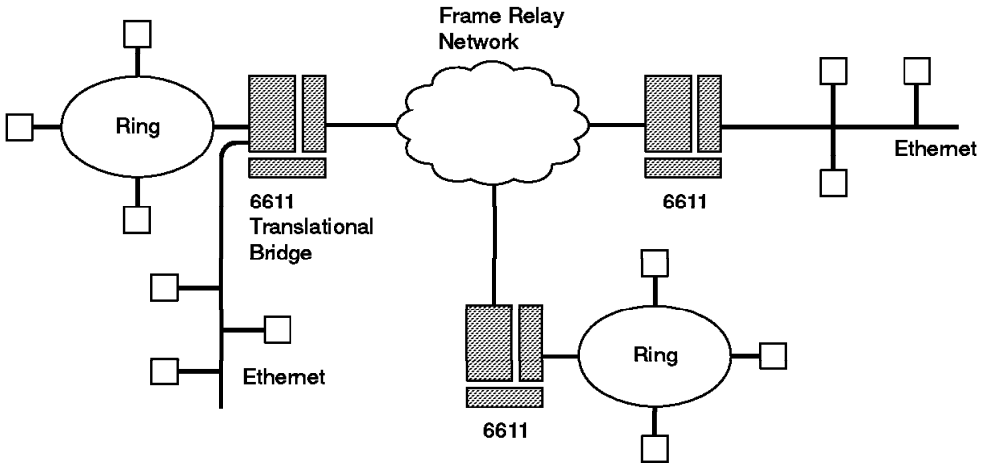

Figure 56. Remote Bridging Function Between a Translational Bridge and Non-Translational Bridge

Remote Bridging Function Between a 6611 Translational and a PS/2: On remote bridging between a 6611 translational bridge and a PS/2 workstation running either the IBM Token-Ring Network Bridge Program Version 2.2, or the IBM Remote Token-Ring Bridge/DOS Version 1.0, the frames can be bridged between 6611 ports configured for source-route, transparent, or dual mode bridging and the PS/2 workstation running the bridge program.

Figure 57 on page 140 shows a sample configuration using this remote bridging function. The function is sometimes called *compatibility mode bridging*. In this configuration, the 6611 functions as the primary half of the bridge, and the bridge program functions as the secondary half of the bridge. A telecommunications link connects the 6611 to the PS/2 workstation running the bridge program. The devices communicate using a proprietary protocol.

Note:

The proprietary protocol used on the telecommunications link is referred to as the *LAN Bridging Protocol* within the 6611 library.

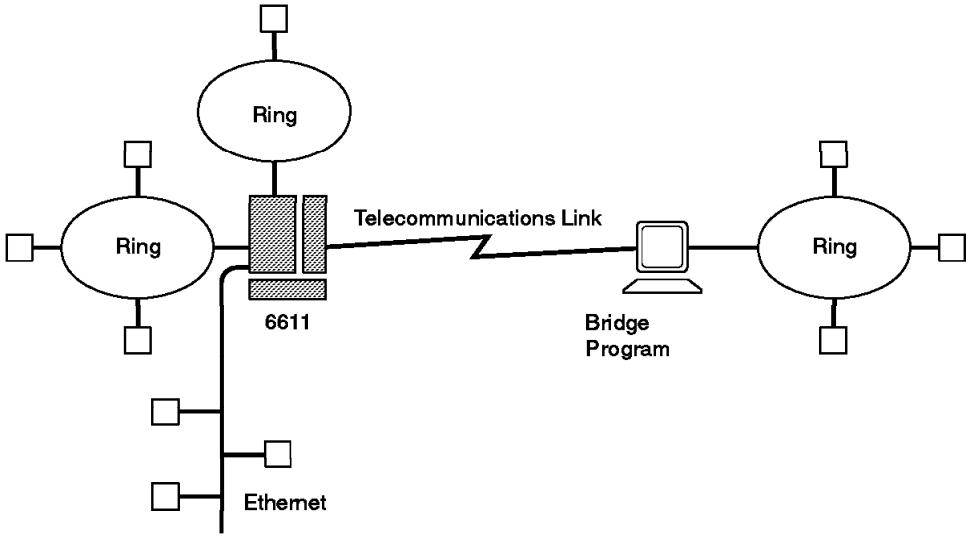

Figure 57. Remote Bridging Function Between a Translational Bridge and a PS/2 Workstation Running a Bridge Program

2.7.3.4 Coexistence with Other IBM Bridge Products

The IBM 6611 Network Processor can coexist with other bridges, such as the IBM 8209 or IBM 8229 and the IBM Personal System/2, using the IBM Token-Ring Network Bridge Program Version 2.2. This includes support for automatic single-route broadcast configuration using the spanning tree algorithm.

However, the IBM 6611 Network Processor does not implement the following functions provided by other IBM bridge products:

- RPS (Ring Parameter Server)
- REM (Ring Error Monitor)
- CRS (Configuration Report Server)
- LRM (LAN Reporting Mechanism)
- LBS (LAN Bridge Server)

As a consequence, there are some limitations when using IBM LAN Network Manager to manage interconnected token-rings that incorporate IBM 6611 Network Processor-based bridges.

2.7.4 Data Link Switching

DLSw is a method of transporting SNA and NetBIOS frames.

The DLS function provides the capability to integrate the transport of the NetBIOS and SNA protocol suites with the other protocol suites that can be routed by the IBM 6611 Network Processor.

Devices that make use of the DLS function are configured as if they were directly attached to each other via a single data link or data link network.

In reality these devices only have a direct data link or data link network connection to an IBM 6611 Network Processor. The IBM 6611 Network Processor

then transports information received on the data link or data link network connection to another IBM 6611 Network Processor. This second IBM 6611 Network Processor has a direct data link or data link network connection with the ultimate destination device.

The two data links or data link networks that are connected via the DLS function need not be the same type of data link or data link network. For example, an SNA device attached via an SDLC data link to a 6611 Network Processor can use the DLS function to connect to an SNA device attached via a token-ring network data link network.

The DLS function uses the TCP transport layer protocol (part of the TCP/IP protocol suite) to implement a transport network between IBM 6611 Network Processors. This transport network can comprise many intermediate nodes, data links and data link networks, if required, through the use of the IP network layer protocol (also part of the TCP/IP protocol suite).

Note

Intermediate nodes in the transport network used to connect IBM 6611 Network Processors that are providing the DLS function do not have to be IBM 6611 Network Processors, provided that they can support the IP network layer protocol.

A TCP connection is automatically established between each pair of IBM 6611 Network Processors that are participating in the DLS function across the TCP/IP transport network. To support the establishment of these TCP connections, each IBM 6611 Network Processor is configured with the TCP/IP network addresses of the other IBM 6611 Network Processors participating in the DLS function.

It is possible to configure an IBM 6611 Network Processor to accept incoming DLS TCP connections from other IBM 6611 Network Processors without explicitly configuring the other IBM 6611 Network Processors. This may reduce the amount of configuration effort required to set up complex DLS environments. However, at least one of the two IBM 6611 Network Processors participating in each DLS TCP connection must be configured with the TCP/IP network address of the other IBM 6611 Network Processor.

The communication adapter features that can be used with the DLS function fall into the following four categories:

- Those that support direct data links to SNA devices
- Those that support direct data links to NetBIOS devices
- Those that support indirect data links to token-ring devices (both SNA and NetBIOS) via a remote source-route bridge configuration
- Those that support connection to the TCP/IP transport network used to interconnect IBM 6611 Network Processors that provide the DLS function

The DLS function incorporates several features to reduce the need to send data across the TCP/IP network that interconnects the IBM 6611 Network Processors participating in the DLS function.

The key feature is the *cache* in which each IBM 6611 Network Processor maintains a table of remote SNA and NetBIOS devices along with the IBM 6611

Network Processor that is able to reach that remote device through the fastest path. Each IBM 6611 Network Processor constructs its cache dynamically by sending queries to other IBM 6611 Network Processors only when needed. The cache can be preloaded with default entries when the IBM 6611 Network Processor is configured to further reduce the need for queries to be sent to other IBM 6611 Network Processors.

An age out timer is used to remove old cache entries after a a period of time. The timeout used by the age out timer can be set when the IBM 6611 Network Processor is configured.

Note

At the time of writing, the cache used by the DLS function could only be used to locate the MAC addresses of remote SNA and NetBIOS devices. As a consequence, NetBIOS requests to locate particular NetBIOS names were copied to all interfaces enabled for DLS on all IBM 6611 Network Processors that participate in the DLS function. However, it is intended that the cache be used to locate NetBIOS names of remote NetBIOS devices. This would dramatically reduce the number of NetBIOS broadcasts that flow across the TCP/IP network that interconnects all IBM 6611 Network Processors participating in the DLS function.

To explain how data link switching is implemented in the 6611, we define two types of data link switching: *local data link switching* and *remote data link switching*. In local data link switching, the data link switching function is performed within a single 6611. In remote data link switching, stations attached to two or more 6611s communicate across an IP network using data link switching. The following topics summarize the features of the two types of data link switching.

There are several differences in the operation of the DLS function for SNA and NetBIOS devices. For this reason each will be described separately on 2.7.4.3, “SNA Data Link Switching” on page 147 and 2.7.4.4, “NetBIOS Data Link Switching” on page 149.

For more information about DLSw networking considerations, see Chapter 4 of *Local Area Network Concepts and Products: LAN Architecture*, SG24-4573.

2.7.4.1 Local Data Link Switching

Local data link switching is used for SNA transport only. It supports communication between a LAN-attached SNA device and a synchronous data link control (SDLC) secondary station that is link-attached to the 6611. The LAN-attached SNA device may be on a LAN directly attached to the 6611, or it may be on a remote LAN that is joined to the 6611 by one or more bridges.

The SDLC secondary station must be a physical unit (PU) type 2.0 or 2.1 and must be operating in normal response mode. During configuration of the 6611, the secondary station is assigned a MAC sub-layer address so that it appears to other network devices to be on a LAN.

Local data link switching converts SDLC frames to IEEE 802.2 LLC type 2 frames. Bridging is used to transport the converted frames (SNA frames encapsulated in a MAC sub-layer frame) to a directly attached LAN or to the next bridge in the path of an interconnected LAN. The local data link switching function does not

convert token-ring MAC sub-layer frames to Ethernet MAC sub-layer frames. However, a route to an interconnected LAN may contain a bridge, such as an IBM 8209 or 8229 LAN Bridge, that converts token-ring MAC sub-layer frames to Ethernet MAC sub-layer frames. A technique called spoofing is used to send acknowledgments to the source station from the 6611 to which the source station is attached, instead of from the destination station.

When configuring local DLSw, configuration of DLSw partners and IP routing is optional.

A sample local data link switched network is shown in Figure 58 on page 144.

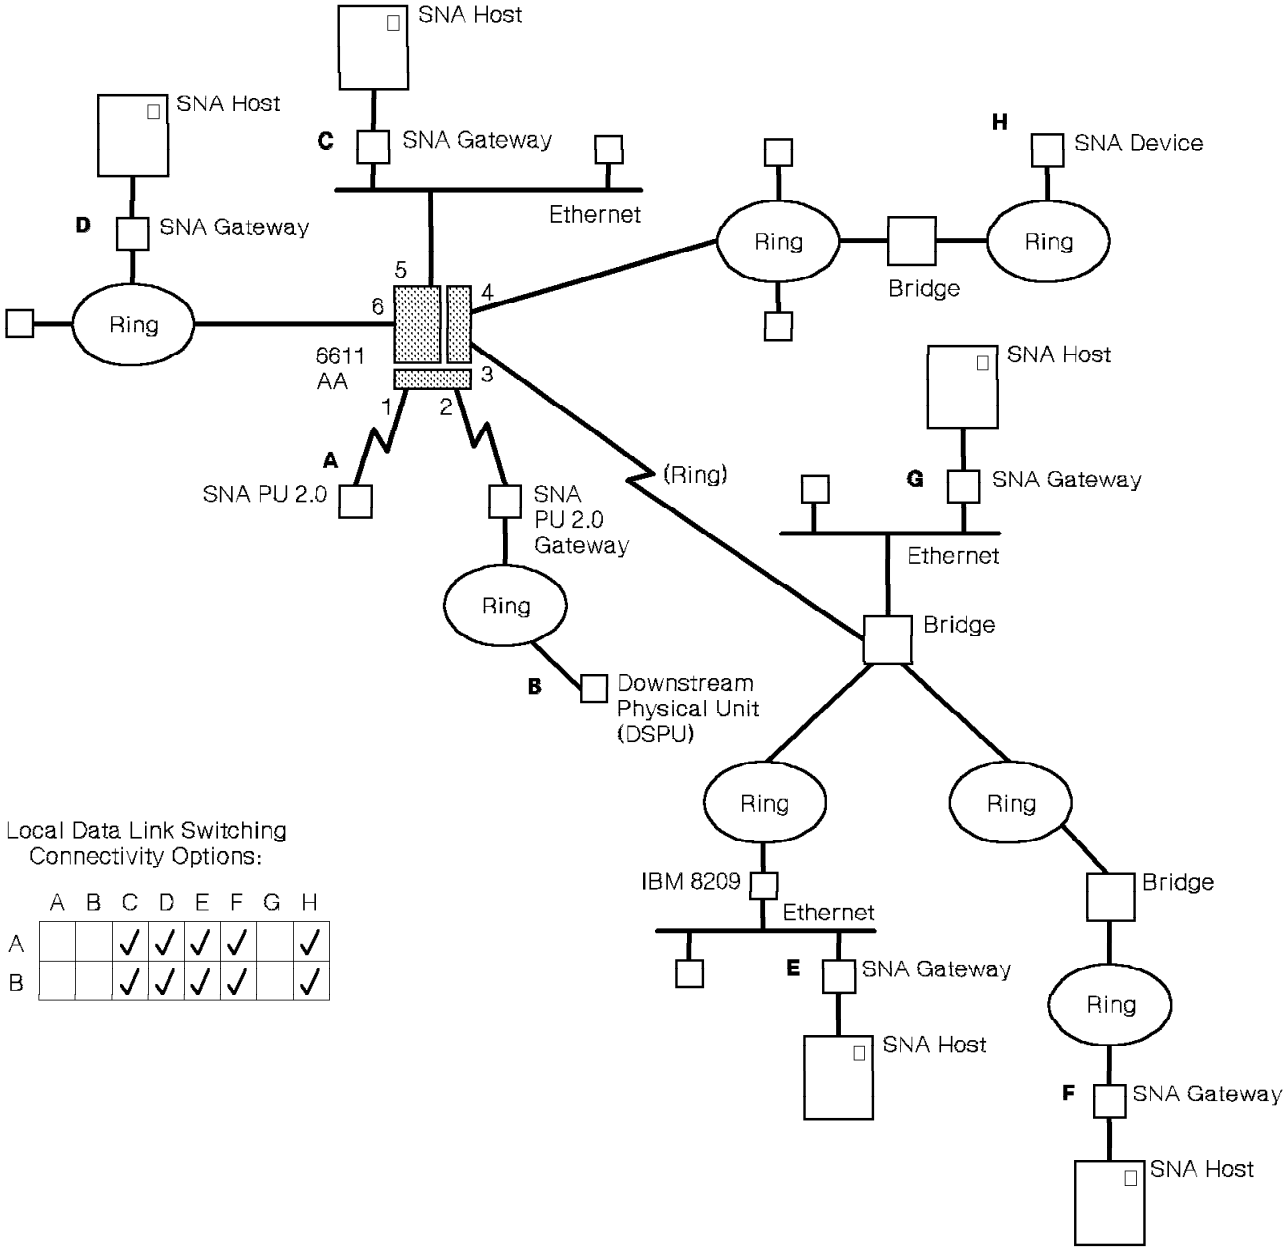

Reference	Configuration Item	Node-Level or Port-Level Configuration
AA	6611	Source-route bridging, transparent bridging, DLSw for SNA
1	SDLC port	SDLC, SNA
2	SDLC port	SDLC, SNA
3	Serial port	Source-route bridging, DLSw for SNA
4	Token-ring port	Token-ring, source-route bridging, DLSw for SNA
5	Ethernet port	Ethernet, transparent bridging, DLSw for SNA
6	Token-ring port	Token-ring, source-route bridging, DLSw for SNA

Figure 58. Sample Local Data Link Switched Network

2.7.4.2 Remote Data Link Switching

Remote data link switching is used for both SNA and NetBIOS transport. An SNA or NetBIOS station attached to a 6611 uses remote data link switching to communicate with an SNA or NetBIOS station attached to another 6611. SNA stations may be link-attached or LAN-attached to the 6611s; NetBIOS stations must be LAN-attached. The 6611s, called *partners*, must be configured for data link switching. The partners communicate with each other across an IP network.

- SDLC-to-LAN communication across a WAN

  Remote data link switching performs SDLC-to-IEEE 802.2 type 2 conversion. This permits a link-attached SDLC secondary station to communicate with a LAN-attached SNA device.

- LAN-to-LAN communication across a WAN

  Remote data link switching supports communication between SNA or NetBIOS stations on token-rings and Ethernets. Remote data link switching can convert token-ring MAC sub-layer frames to Ethernet MAC sub-layer frames, and conversely, so that devices on token-rings and Ethernets can communicate with each other.

The 6611s communicate with the SNA and NetBIOS stations using IEEE 802.2 LLC type 2. The LLC connections are terminated at the 6611s. Spoofing is used to send acknowledgments to the source station from the 6611 to which the source station is attached, instead of from the destination station. This reduces traffic on the WAN.

The hop count for source-route bridging is also terminated at the 6611s. Thus, the source station may be up to 7 hops from the first 6611 in the path and the receiving station may be up to 7 hops from the last 6611 in the path.

For transport between the data link switching partners, the SNA or NetBIOS frames are encapsulated in IP datagrams. The partners communicate with each other using TCP. The route between two partners can contain IP routers that are not 6611s, as long as they are compatible with the 6611. The 6611s in an IP route between partners must be configured for IP routing, but they need not be configured for data link switching.

A sample remote data link switched network is shown in Figure 59 on page 146.

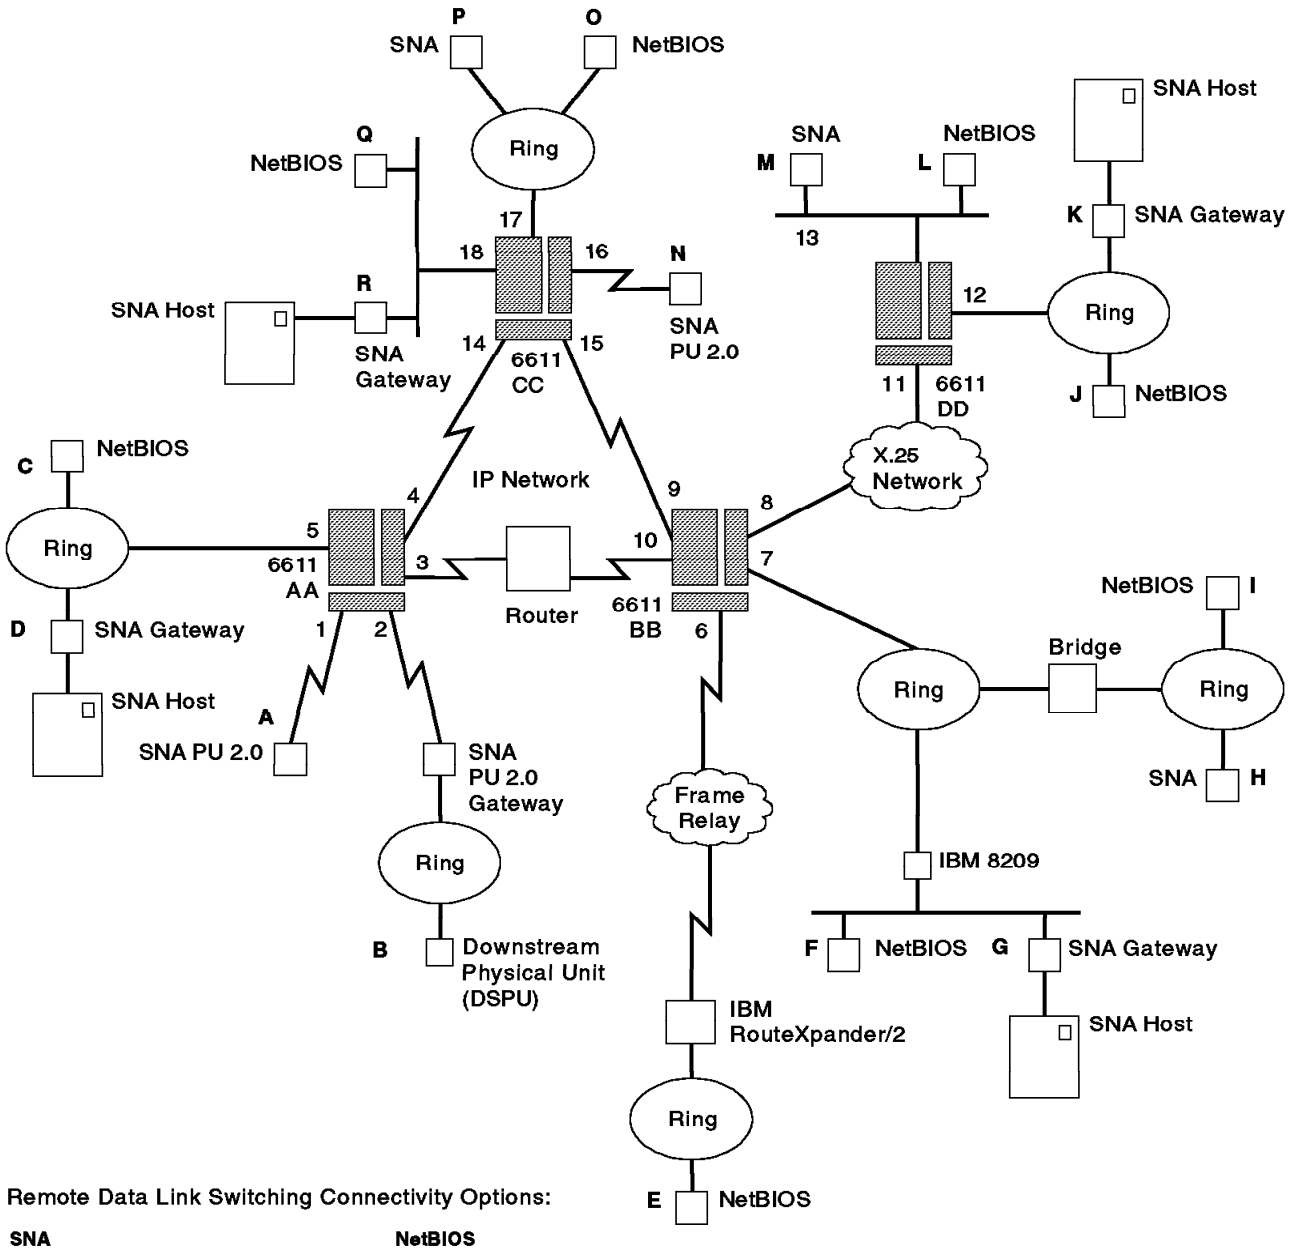

Remote Data Link Switching Connectivity Options:

SNA

	R	P	M	N	K	H	G	D	B
A	✓				✓		✓		
B	✓				✓		✓		
D	✓	✓	✓	✓	✓	✓	✓		
G	✓	✓	✓	✓	✓				
H	✓	✓	✓		✓				
K	✓	✓		✓					
N									
M	✓	✓							
P	✓								

NetBIOS

	Q	O	L	J	I	F	E
C	✓	✓	✓	✓	✓	✓	✓
E	✓	✓	✓	✓			
F	✓	✓	✓	✓			
I	✓	✓	✓	✓			
J	✓	✓	✓				
L	✓	✓					
O							

Figure 59. Sample Remote Data Link Switched Network

The node-level and port-level configurations for the 6611s in Figure 59 are summarized in Table 11 on page 147.

Table 11. Configuration of the Sample Remote Data Link Switched Network

Reference	Configuration Item	Node-Level or Port-Level Configuration
AA	6611	OSPF, source-route bridging, DLSw for SNA and NetBIOS
BB	6611	OSPF, source-route bridging, IP over X.25, DLSw for SNA and NetBIOS
CC	6611	OSPF, source-route bridging, transparent bridging, DLSw for SNA and NetBIOS
DD	6611	OSPF, source-route bridging, IP over X.25, DLSw for SNA and NetBIOS
1	SDLC port	SDLC, SNA
2	SDLC port	SDLC, SNA
3	Serial port	PPP, IP
4	Serial port	PPP, IP
5	Token-ring port	Token-ring, source-route bridging, DLSw for SNA and NetBIOS
6	Serial port	Frame relay, source-route bridging, DLSw for SNA and NetBIOS
7	Token-ring port	Token-ring, source-route bridging, DLSw for SNA and NetBIOS
8	X.25 port	X.25, IP
9	Serial port	PPP, IP
10	Serial port	PPP, IP
11	X.25 port	X.25, IP
12	Token-ring port	Token-ring, source-route bridging, DLSw for SNA and NetBIOS
13	Ethernet port	Ethernet, transparent bridging, DLSw for SNA and NetBIOS
14	Serial port	PPP, IP
15	Serial port	PPP, IP
16	SDLC port	SDLC, SNA
17	Token-ring port	Token-ring, source-route bridging, DLSw for SNA and NetBIOS
18	Ethernet port	Ethernet, transparent bridging, DLSw for SNA and NetBIOS

2.7.4.3 SNA Data Link Switching

The DLS function supports the interconnection of SNA devices attached to either a token-ring or an SDLC multipoint non-switched line. A typical example of the use of the DLS function for SNA devices is illustrated in Figure 58 on page 144.

As a prerequisite for the DLS function, each participating token-ring segments IBM 6611 Network Processor that supports token-ring-attached SNA devices, must be configured to support source-route local bridging on all token-ring interfaces used with the DLS function.

Note

Local bridging will be used in preference to the DLS function to provide connections between token-ring-attached SNA devices that are connected to the same IBM 6611 Network Processor via different token-ring segments.

Each IBM 6611 Network Processor participating in the DLS function must also be configured with a *virtual segment number*. This virtual segment number must be the same for all IBM 6611 Network Processors participating in the DLS function.

Additionally, SNA devices attached to a IBM 6611 Network Processor via an SDLC multipoint non-switched line are assigned a token-ring LAA (locally administered address), SAP (Service Access Point) and SNA XID (Exchange ID). These will be used by the IBM 6611 Network Processor to represent such devices to other SNA devices that are using the DLS function.

Note

A single hop is used in the RI (Routing Information) field to reach an SNA device accessible via the DLS function from a token-ring segment directly attached to a IBM 6611 Network Processor. Therefore, SNA devices can be, at most, six hops from an IBM 6611 Network Processor to reach SNA devices accessible via the DLS function.

The DLS function only supports the attachment of SNA devices via SDLC multipoint lines that are of PU (Physical Unit) type 2.0. The attachment of PU type 2.1 devices is not supported unless they provide a PU 2.0 compatibility mode. The attachment of PU type 4 devices (such as the IBM 3745 Communications Controller) is not supported either.

There are two consequences of this:

1. SDLC-attached devices cannot establish connections with other SDLC-attached devices. This is because SNA PU type 2.0 devices cannot directly communicate with each other as peers.
2. SDLC-attached devices can only support a single connection to another SNA device attached to a token-ring. The other SNA device will usually be a PU type 4, such as the IBM 3745, or a PU type 5.

DLSw SNA Traffic Prioritization: This function was implemented in the Multiprotocol Network Program Version 1 Release 3 (MPNP). It can be defined as a method that allows SNA frames to have adequate priority over NetBIOS frames. It applies to the DLSw traffic from all the ports on the 6611. Additional priority can be given to SNA frames by a two-pronged approach as follows:

1. **SNA/NetBIOS Ratio (Bias)**

   The user can specify the ratio of how many SNA frames are to be sent per NetBIOS frame. Valid SNA/NetBIOS ratio settings are from 0 to 9. If the ratio is set at 9, nine SNA frames will be transmitted on the link per NetBIOS frame. The frames are selected from the DLSw data stream preserving the order of the frames.

   There is no capability that allows NetBIOS frames to have priority over SNA frames. This function is for increasing the priority for SNA traffic.

2. **NetBIOS Frame Size Reduction**

   NetBIOS tends to send frames as large as the transport mechanism will allow, while SNA tends to send very small frames. Often this can lead to NetBIOS using most of the transport's bandwidth. The NetBIOS largest frame size option allows users to force the frames to be broken into segments. In other words, NetBIOS will be forced to use smaller frames, thus allowing SNA Bias to have a more predictable effect. The choices of the valid largest allowed NetBIOS frame in bytes are 2052, 1500 and 516.

2.7.4.4 NetBIOS Data Link Switching

The DLS function supports the interconnection of NetBIOS devices attached to either a token-ring or a CSMA/CD (Carrier Sense Multiple Access/Collision Detection) LAN using either DIX Ethernet V2 or IEEE 802.3 frame formats. A typical example of the DLS function for NetBIOS devices is illustrated in Figure 59 on page 146.

NetBIOS devices on token-rings are handled in a similar way to SNA devices on token-rings. That is, remote NetBIOS devices will appear as if they are on the DLS virtual segment.

NetBIOS devices on CSMA/CD LANs cannot be handled in a similar way to that used for SNA devices on token-rings. Instead, the ability of NetBIOS to dynamically bind a MAC address to a NetBIOS name is exploited.

From the perspective of NetBIOS devices on CSMA/CD LANs, all remote NetBIOS devices appear as if they have the MAC address of the 6611 Ethernet Adapter. This is possible because the NetBIOS protocol discovers the MAC address of other NetBIOS devices using broadcast frames sent to the NetBIOS functional address.

2.7.4.5 Estimating DLSw Storage Requirements

Developing a DLSw configuration requires careful design and planning for efficient utilization of available system resources. To assist you in planning your configuration and determining your 6611 memory needs, IBM provides a storage estimating tool called the IBM 6611 Storage Estimate EXEC. For information on this tool, contact your IBM marketing representative and ask for the Internetworking Marketing Specialist for your trading area.

Memory expansion features are available if additional memory is required for the 6611. An 8 MB memory expansion (feature code 4008) is available on Models 125, 145, and 175. A 16 MB memory expansion (feature code 4016) is available on Models 145 and 175. The 16 MB memory expansion for Models 140 and 170 is available by RPQ 8Q1414.

2.7.5 IBM 6611 Network Processor Enhancements - Release 4

There are many enhancements that will be available on IBM 6611 - Release 4 that we can emphasize:

- High Performance Routing (HPR), with the following features:
  - Automatic Network Routing (ANR) is a sophisticated new source-routing method that delivers unmatched price/performance for mission-critical data.
  - Rapid Transport Protocol (RTP) allows safe reroute of data around failed links or notes.

  - Adaptive Rate-Based (ARB) provides superior flow and congestion control.
- Dependent LU Requester (DLUR) which enables dynamic configuration of dependent LUs.
- Enhanced Priority Queueing Support with three new HPR data queues enrich the 6611′s priority queueing scheme.
- FR Boundary Access Node (BAN) that provides the ability to bridge token-ring and Ethernet SNA traffic directly to a FEP (3745) without frame conversion by DLSw router.
- Frame Relay RFC 1490 is a standard that specifies how SNA and multiprotocol LAN traffic can be natively and efficiently encapsulated in frame relay frames for transport across a wide-area network.
- ITU-T LMI Support via Frame Relay - ITU-T Q.9333 Annex is a standard that defines means of status and the notification of outage for frame relay PVC.
- DLSw V1 Compliance RFC 1795 is an industry-standard method for transmitting SNA and NetBIOS traffic across a TCP/IP wide area network.
- Support for RFC 1027; Transparent Subnetting which enables the 6611 to act as a transparent subnet ARP gateway.
- Support for RFC 1542; BOOTP which enable the 6611 to act as an BOOTP relay agent. Also allows the 6611 to act as relay agent for hosts RFC 1534.
- 2210 EasyStart that allows the IBM 6611 Network Processor to act as a BOOTP relay agent for 2210s which needs to download its initial configuration information from the network.
- IPX Filtering enhancements with New IPX RIP filters that allows a network administrator to filter inbound and outbound RIP filters using network numbers ranges; one filter can be applied to all ports.
- Fast IPL Time for the 6611 Network Processor has been significantly improved.
- Auxiliary Power Shutdown restricts shutdown of UPS
- System Manager Enhancements where several new enhancements to the System Manager function are provided.
- Up to 32 MB of Memory Upgrade for M125 enable Customers to order additional memory (up to 32 MB) for their 6611 Model 125 using Feature Code 4008.
- DASD Size Enhancement where new models of 6611 will begin shipping with larger hard drives.
- OS/2 & DOS/WIN Configuration Transfer Support for sending configurations through the network using TCP/IP socket connection to a 6611.
- Multiple Retrieve Function that provides the ability to retrieve configurations files from multiple routers for configuration updates.

Chapter 3. LAN Gateways

Commonly, the term gateway is used to mean a device that handles interconnection and transfer of data between computers in different network architectures, as shown in Figure 60.

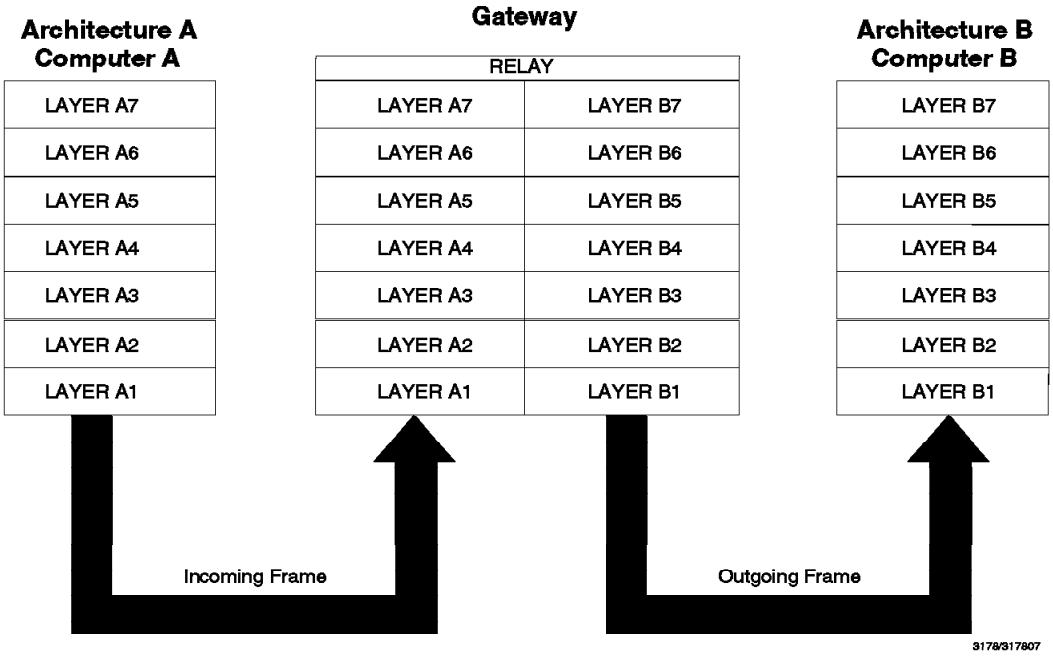

Figure 60. Gateway Implementation

However, the term gateway can be used in a generic sense that is applicable to any layer; a data link layer gateway is called a bridge and a router is the network layer gateway. Thus, a gateway is required when an entity using one protocol needs to communicate with another entity using a different protocol. This is only when the addressing structure of the two entities is inconsistent, or when a connection is required between independent networks.

In this chapter, the various IBM LAN gateways are discussed. A LAN gateway refers to a combination of hardware and software, which provides interconnection between host-based entities and LAN-based entities. The entities may be either programs or physical devices.

A LAN gateway may be located inside, attached or remote to the host.

As many LAN gateways provide interconnection between SNA hosts and LAN devices, there are some SNA concepts that are important to know in order to understand the gateway′s functionality. These concepts are presented in Chapter 4 of *Local Area Network Concepts and Products: LAN Architecture*, SG24-4753.

3.1 S/390 Open Systems Adapter

The Open Systems Adapter 1 (OSA 1) provides token-ring, Ethernet and FDDI interfaces on the S/390 Parallel Enterprise Servers and Parallel Transaction Servers, and the ES/9000 711 and 511 based models. OSA supports SNA, APPN, TCP/IP, and IPX protocols. In addition to integrated, direct LAN connectivity, OSA performs TCP/IP protocol processing (TCP/IP offload), supports the LAN Resource Extension and Services/MVS (LANRES/MVS) program product, and the LAN File Services/ESA (LFS/ESA) program product for NFS clients. The Open Systems Adapter, as a complementary member of the IBM family of communications products, is an integral part of the S/390 initiatives of open systems, client/server, and reduced cost of computing.

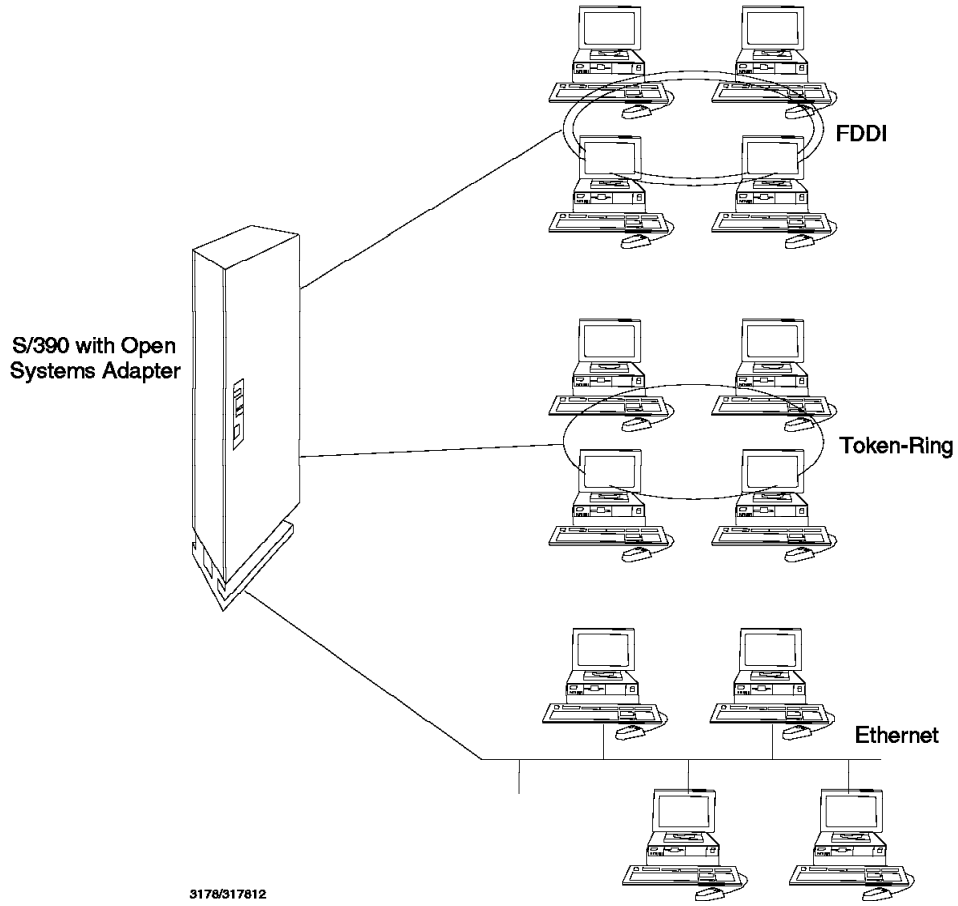

Figure 61. Direct Connection between Server and LAN

The Open Systems Adapter 2 (OSA 2) delivers a lower-priced alternative for S/390 Parallel Enterprise Server Models R2 and R3 running the SNA/APPN and TCP/IP networking protocols.

3.1.1.1 Description

OSA is made of three basic sections: S/390 Input/Output (I/O) Subsystem Interface, intelligent engine for protocol processing and application support, and industry-standard LAN interfaces.

S/390 I/O Subsystem Interface: OSA contains two channel engine modules. These are the same engines that go on S/390 channel adapters. One of these engines runs channel Licensed Internal Code (LIC) and the other engine runs control unit-like LIC. This makes the adapter look like a channel/control unit combination to MVS and VM. No application code changes are required, so a customer's investment is protected. OSA looks like another channel type to MVS and VM. This I/O Subsystem Interface is the same one used by S/390 ESCON and Parallel channels.

Intelligent Engine: OSA uses an IBM manufactured 486 66 MHz engine to run the OSA Application Platform that provides the TCP/IP protocol offload, LANRES, and LFS functions. The OSA Application Platform is a derivative of Novell's NetWare program product.

LAN Interfaces: OSA supports 4 and 16 Mbps token-ring, 10base5 Ethernet, and ANSI X3T9.5 FDDI connections.

IBM intends to enhance the network connectivity capabilities of the S/390 Open Systems Adapter 2 by providing support for ATM as explained in Chapter 5 of *Local Area Network Concepts and Products: LAN Adapters, Hubs and ATM*, SG24-4754.

3.1.1.2 Packaging

By combining the LAN interface, I/O subsystem interface, and intelligence on one adapter, the OSA eliminates some of the data moves required when these functions are on separate adapters.

The use of *off the shelf* components to combine these functions onto one adapter drives the size of this version of the adapter, which causes it to be packaged in a feature frame. Some possible configurations are described in *Planning for the System/390 Open Systems Adapter Feature*, GC23-3870.

With the introduction of OSA 2, the size of the package has been reduced. OSA 2 can be plugged into any I/O slot in a CEC (Central Electronic Complex) or I/O expansion cage.

3.1.1.3 OSA 1 - Configurations

There are three media types available with OSA. Both the Ethernet and token-ring feature consist of two cards, one card for interfacing with the S/390 channel subsystem and one card for five LAN ports. The FDDI feature has one dual-ring port on one card.

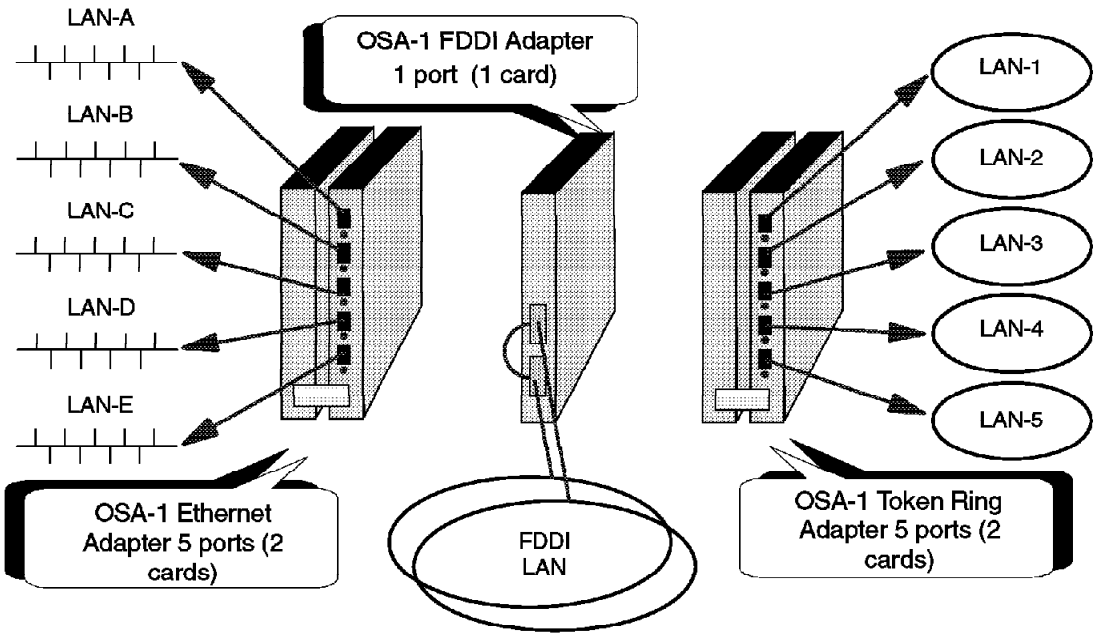

Figure 62. OSA1 Adapter Ports Configuration

The ES/9000 9021 711-based models can support one OSA frame per side; each frame can support two OSA cages and each cage has eight slots.

The ES/9000 9121 511-based models can support one OSA frame per side; each frame can support one OSA cage and each cage has eight slots.

The 9672 CMOS-based processors can support one to two OSA cages depending on the I/O configuration. Each cage in a CMOS processor has nine slots.

3.1.1.4 OSA 2 - Configurations

Three media types are available with OSA 2. The ENTR (Ethernet and token-ring) has two ports that are self-defining when the connector is plugged in. Each port can be defined as either Ethernet or token-ring. The FDDI feature supports a single or dual-ring station.

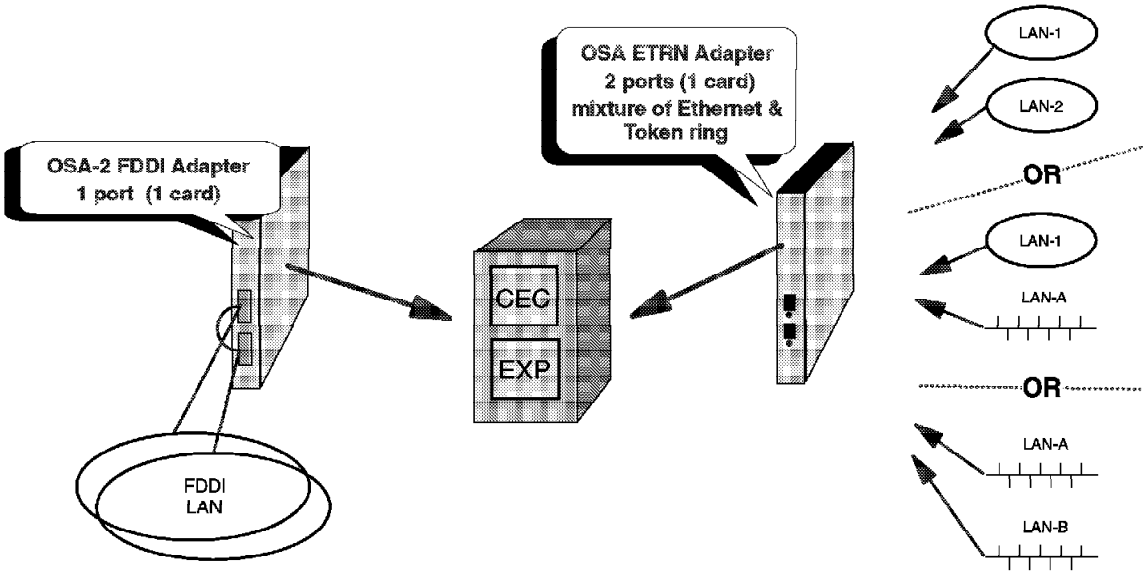

Figure 63. OSA2 Adapter Ports Configuration

The OSA 2 feature can be inserted into any I/O slot in a CEC or I/O expansion cage, up to a maximum of twelve features.

The OSA 2 is supported on the S/390 Parallel Enterprise Server models R2 and R3 only.

Table 12. OSA Hardware Support

Media and Processors	**OSA 1**	**OSA 2**
FDDI	1 port	1 port
Token-Ring	5 ports	2 ports mixed
Ethernet	5 ports	2 ports mixed
ATM	NO	S.O.D
9021 711-based	YES	NO
9121 511-based	YES	NO
9672 E0x, P0x, Rx1	YES	NO
9672 R2,R3	YES	YES

3.1.1.5 Software Compatibility

OSA is supported by MVS/ESA V4R3 or above, and VM V2R1.0. For utilization with ACF/VTAM V4R2 for MVS, a PTF is required.

3.1.1.6 Scenarios

OSA can be used in the following scenarios.

SNA Environment: OSA supports the traditional hierarchical subarea environment. In SNA networks, OSA functionality is similar to the IBM 3172 Interconnect Controller and is defined in VTAM as an XCA Switched Major Node.

For SNA (and for APPN), OSA uses Link Services Architecture (LSA), which terminates the LAN logical link control (LLC) and passes the packets to VTAM for SNA and APPN processing. OSA provides a high-speed SNA access from the LAN to the VTAM mainframe, supporting VTAM to VTAM flows, VTAM to NCP flows and VTAM to device (PU2.X) flows.

APPN Environment: Advanced Peer-to-Peer Networking (APPN) was designed to address the dynamic client/server (C/S) environment. OSA does not differentiate between SNA hierarchical flows to VTAM as a PU5, and APPN flows to VTAM as an APPN network node or APPN end node. OSA is not an APPN network node and thus does not provide dynamic APPN routing, but rather provides static routing to VTAM based on the SAP address and adapter port. Dynamic routing is provided by VTAM or other APPN network nodes in the network.

TCP/IP Passthrough: TCP/IP for MVS V2R2.1 and TCP/IP for VM V2R3 in Passthrough mode. Passthrough mode is when the protocol processing of the IP layer and above is done by TCP/IP for MVS/VM. In this case, we have a higher performance than in the offload mode.

TCP/IP Offload: S/390 CPU utilization can be reduced by performing the TCP/IP protocol processing on the adapter (offload). OSA support for TCP/IP offload will be part of TCP/IP for MVS V3R1.

LFS/ESA: OSA enables large scale file serving for TCP/IP NFS clients via the LAN File Services/ESA (LFS/ESA) program product. This also allows customers to consolidate file servers onto the S/390 platform, reducing the need for intermediate servers. When using LFS/ESA for MVS, there is a cooperating application on OSA that provides file caching for NFS requests, returning CPU cycles for use by applications and clients.

LANRES/MVS: OSA provides large scale disk serving for NetWare clients via the LAN Resource Extension and Services/MVS (LANRES/MVS) program product. The OSA/LANRES combination enables customers to consolidate disk servers onto the S/390 platform. This reduces the requirement for intermediate servers between clients and the mainframe.

3.1.1.7 LPAR Support

Processor Resource/Systems Manager (PR/SM) Logical Partition Support. Each LAN port on an Open Systems Adapter can be shared across multiple logical partitions (LPARs) when OSA is supporting SNA, APPN, TCP/IP passthrough, and LANRES. A LAN port is dedicated to an LPAR when running in TCP/IP offload or LFS/ESA modes. The maximum number of LPARs supported by OSA is the same as the maximum number supported on the System/390 processor. Additional configuration and application concurrency information is available in *Planning for the System/390 Open Systems Adapter Feature*, GC23-3870.

3.1.1.8 OSA Support Facility for MVS/ESA

This is an MVS/ESA application that acts as a configuration and operations management tool for OSA. It is accessed via a non-dedicated OS/2 workstation. It provides the following functions:

- Simple method to download the software to the OSA
- Customize OSA
- Log incidents and commands issued to OSA
- Gather configuration definition information and statistics from the system dynamically

3.1.1.9 OSA 1 and OSA 2 Attributes

The Open Systems Adapter 2 does not support all the facilities available with OSA 1.

Table 13. OSA Supported Facilities

Facility	OSA 1	OSA 2
SNA/APPN	X	X
TCP/IP Passthrough	X	X
TCP/IP Offload for MVS	X	
LFS/ESA-NFS for MVS	X	
LANRES/MVS	X	
OSA/SF for MVS	X	X

3.1.1.10 ADSTAR Distributed Storage Manager (ADSM)

ADSTAR Distributed Storage Manager is a family of client/server storage management products that provides administrator controlled, highly automated and centrally scheduled network-based backup and archive functions for workstations and LAN file servers.

- Provides an enterprise-wide backup and archive facility for a wide variety of LAN file servers and individual workstations such as OS/2, NetWare, Windows, WinNT, DOS, Macintosh, UNIX
- Allows data to be easily shared between users, even those using different operating systems
- Consolidated user file storage for faster retrieval
- Supports interoperability between popular UNIX-based workstations and file servers as client environments: HP-UX, SPARC/Solaris, AIX, SCO UNIX 386, AT&T Global Information Solutions UNIX, Siemens Nixdorf SINIX, DEC ULTRIX, NEC UNIX, Silicon Graphics IRIX

3.1.1.11 Summary

The Open Systems Adapter along with MVS and VM Open Edition, LANRES, LFS, VTAM, TCP/IP and ADSM position the S/390 as a server. OSA supports the following System/390 initiatives:

- Open: OSA helps change the perception of the mainframe as a closed platform. Servers talk directly with LANs, so the S/390 as a server does also. OSA provides customers with a complementary addition to the family of IBM

Networking products that allow S/390 to participate effectively in the network-centric information environment.

- Client/Server: OSA in conjunction with program products such as LANRES, LFS, and ADSM brings the strengths of the mainframe to the distributed environment. Some of these strengths are security, availability, enterprise-wide access to data, and systems management. The goal is to provide customers with system solutions that utilize the strengths of both the distributed and centralized environments in a complementary manner.
- Cost of Computing: From a network perspective, the cost of connectivity can be reduced by attaching S/390 to clients using high volume industry-standard adapters in the workstation.

By providing an integrated industry-standard connectivity to System/390, OSA helps bring the strengths of the mainframe to today′s network-centric environment.

3.2 3172 Gateway

The IBM 3172 Interconnect Controller is a microprocessor-based intelligent controller that enables LAN-attached workstations as well as WAN-attached workstations to access IBM host processor resources and applications.

Three models of the 3172 have been shipped, starting with the Model 1 introduced in 1989. The latest entry into the 3172 family is the Model 3 which was announced and delivered in 1992. The Model 3 is a significant advancement in price/performance over the Model 1 and Model 2, which have now been withdrawn.

The 3172 Model 3 features a pluggable processor card, expandable memory, up to two channel adapters, up to four LAN/WAN adapters and multiple software solutions. The 3172 is connected to a host using either a Parallel channel or an ESCON channel.

3.2.1 3172 LAN-To-Host Mode

The 3172 provides connectivity to the IBM SYSTEM/390 hosts as a LAN gateway, supporting SNA, TCP/IP and IPX data flows. The 3172 supports LAN connections for the following LANs with respective adapters:

- IBM Token-Ring 16/4 Mbps LAN (IEEE 802.5)
- Ethernet LAN (IEEE 802.3 CSMA/CD) and Ethernet Version 2 (only for TCP/IP)
- FDDI (ISO 9314)
- ATM

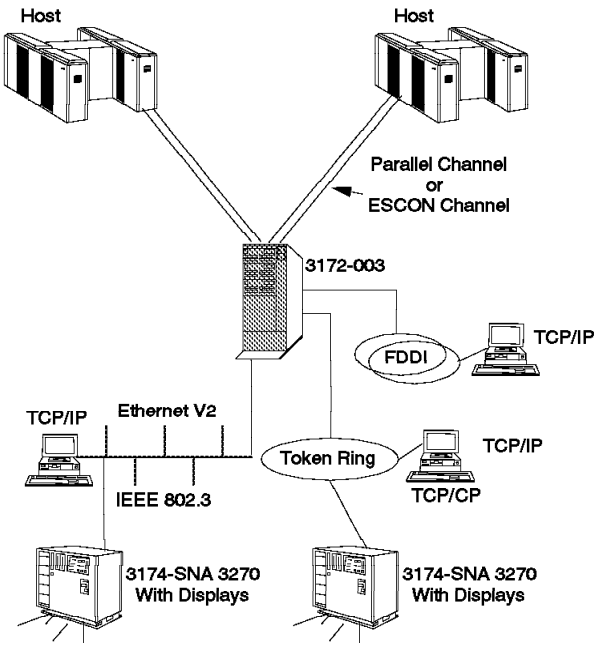

Figure 64. LAN-to-Host Mode

The 3172 Model 3 may run under the Interconnect Control Program (currently Version 3.4) or using OS/2 Version 2.1. When using OS/2, the 3172 can work with Offload, LFS, LANRES, SNA Comm or IP Channel Communications Program.

3.2.1.1 ICP Mode

The Interconnect Controller Program was specifically created to provide transparent access for multiple protocols. ICP, in conjunction with the 3172 Model 3, enables TCP/IP and VTAM-SNA communications between LAN-attached workstations and the host. ICP is normally used when the power of the 3172 for high-speed access to the host is required. ICP code can support two Parallel channel adapters or a single ESCON; EMIF is also supported. All drivers for the channels and LAN adapters are in the ICP code. ICP code *cannot* run concurrently with any other software.

The 3172 configuration, in ICP mode, basically consists of associating each LAN adapter to a corresponding subchannel address pair. In operation, the 3172 acts purely as the physical interface to the LAN, taking data from a subchannel, placing it on a LAN, and vice versa.

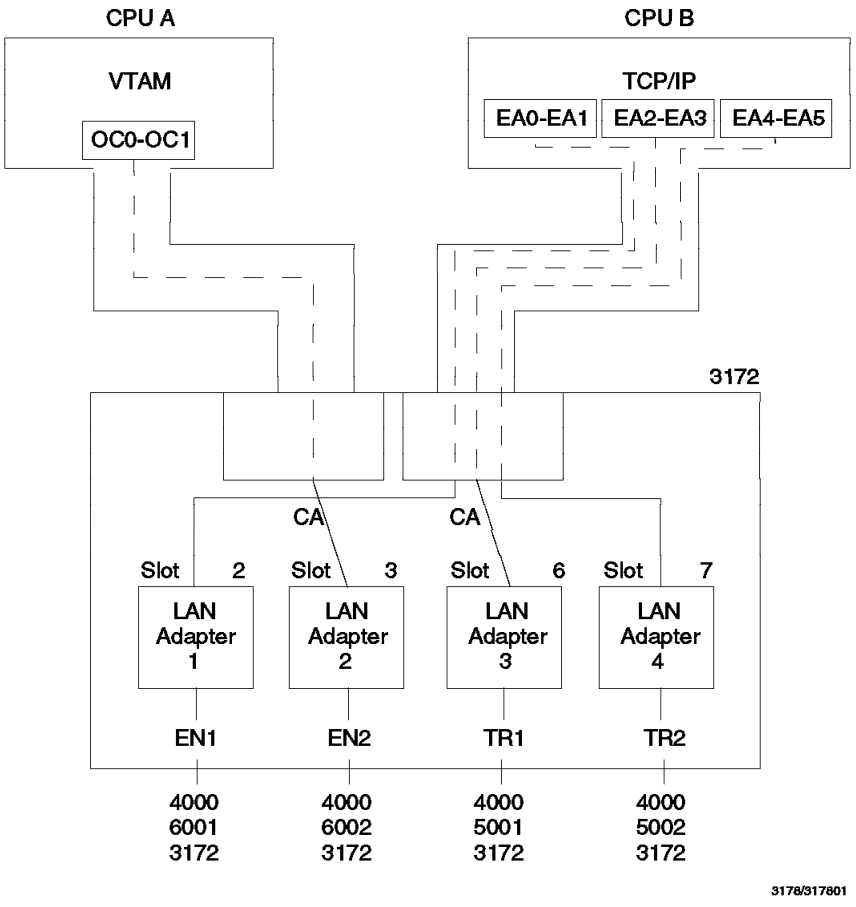

Figure 65. 3172 ICP Mode

SNA Environment

The 3172 Model 3 can support up to a maximum of four LANs. On each LAN, it can support 255 workstations defined as a physical unit (PU). A total maximum of 1020 workstations can be in session with a host.

VTAM supports the definition and activation of LANs with duplicate MAC addresses. Although duplicate addresses are not allowed on the same physical LAN, it is possible to define alternate routes across bridged LAN segments back to the same VTAM. This alternate route and active adapter support increases the level of availability.

Host's View: To VTAM, the 3172 is transparent, which in SNA terms means that it is not addressable. It is not a physical unit (PU) such as a 3745 or a 3174. Figure 66 shows how the host establishes a link to the 3172.

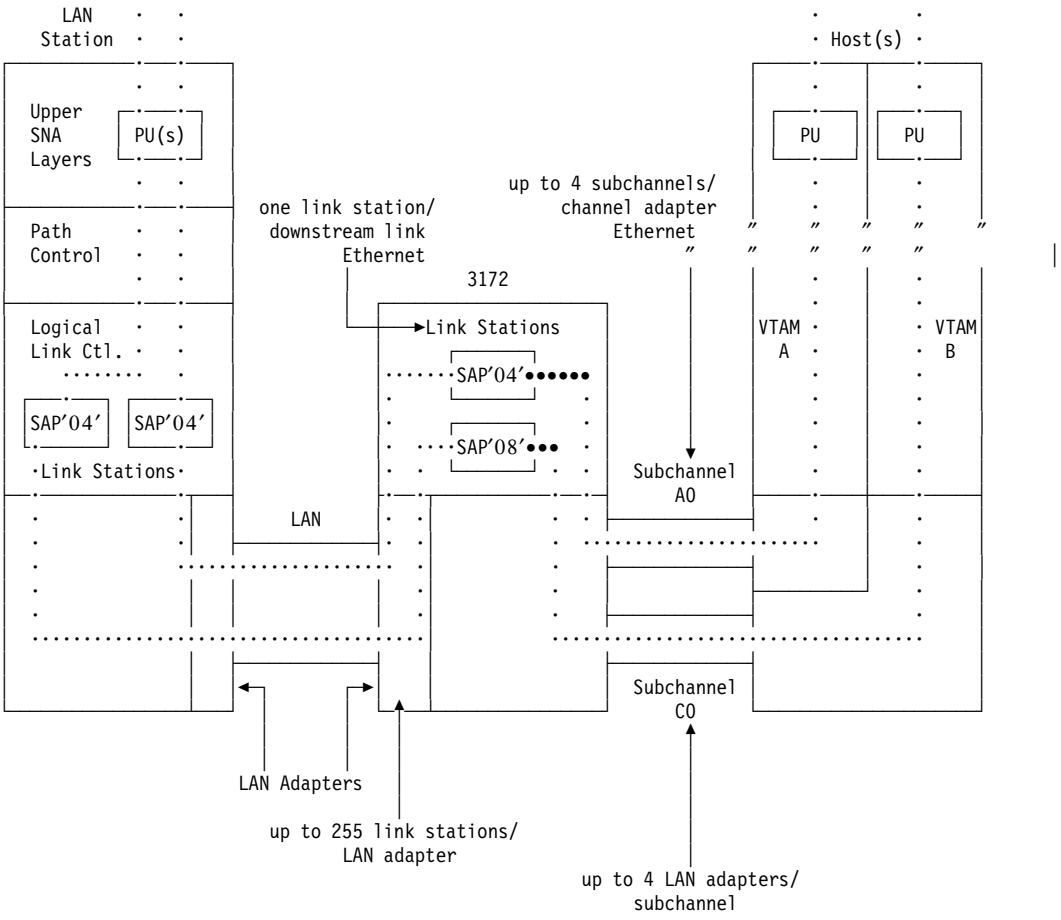

Figure 66. 3172 SNA Support - Logical View

VTAM identifies a LAN through a subchannel. The 3172 is defined to the host control program (MVS or VM) as a channel-to-channel machine, such as a 3088. This provides a physical connection for the LAN. Over this path, VTAM and LAN workstations establish switched-like connections. One path is used for a specific SAP, so that two VTAMs using different SAPs can share a LAN adapter.

LAN Workstation's View: Since this is a logically switched environment, a LAN workstation has to have some means of locating the 3172 gateway. As with an OSA, it must be told the MAC address of the LAN adapter on the 3172 and the

SAP that will take it to the desired VTAM. The PU and LU definitions need to correspond to those in VTAM (either in the Switched or Model Major Node).

Gateway′s View: From the gateway′s point of view, VTAM is a subchannel and a LAN workstation is a link station (of which a LAN adapter can have a maximum of 255). It is up to VTAM and the LAN workstation to make sure that the correct SAP is used. Please refer to the example given in Figure 66 on page 161.

TCP/IP Environment As the 3172 provides LAN-to-host channel connectivity, users are able to exploit the services of applications such as Telnet, X-Windows, File Transfer Protocol (FTP), Simple Network Management Protocol (SNMP), Network File System (NFS), Remote Procedure Call (RPC) and other TCP/IP facilities. LAN workstations are aware of the IP and MAC level address of different TCP/IP nodes via the Address Resolution Protocol (ARP). With ICP support, routing through the 3172 to the TCP/IP host is done via the association of a LAN adapter with a specific subchannel over which TCP/IP communicates to the 3172. Since there is no routing in the 3172 (ICP) itself, the TCP/IP in the host performs the routing function.

Host′s View: TCP/IP for VM or MVS at the host sees each pair of subchannel addresses as a link to reach a LAN segment. Any TCP/IP LAN workstation is seen by a TCP/IP host as an IP address.

LAN Workstation′s View: There is no difference in how a LAN workstation views a host. It knows the IP address it wants to access, and to which MAC address it is associated.

Gateway′s View: From the previous subsections, we can see that apart from configuring the 3172, the only definition necessary to complete the link is to establish the relationship between the subchannel and the LAN adapter. With the use of the relative adapter number, it is possible for a subchannel pair to be associated with more than one LAN adapter. It would appear to TCP/IP as if it is attached to multiple networks and has an IP address for each of them. However, it is not possible for one LAN adapter to be associated with multiple subchannels, and therefore multiple TCP/IPs.

3.2.1.2 Offload Mode (OS/2 Platform)

The power of the 3172 Model 3 is employed to relieve the mainframe from TCP/IP network activity (for example, fragmentation/reassembly and IP routing). Thus, more mainframe cycles will be available for application processing (for example, File Transfer Protocol or Telnet for remote logon), reducing the overall processor usage. The estimated saving of mainframe cycles is between 30 and 50 percent. When operating in an offload capacity, the 3172 is running a TCP/IP stack and has the following software running in it:

- OS/2 operating system
- TCP/IP for OS/2
- TCP/IP for MVS offload code

With this TCP/IP stack within the 3172, it has the ability to do RIP routing of TCP/IP protocols. The 3172 Model 3 offload gateway can dynamically route from one adapter to another. The 3172 TCP/IP stack receives all datagrams destined for the MVS host, processes them and sends them onward. As a result, TCP/IP applications are not aware of the presence of the 3172.

Note

The TCP/IP Offload for MVS saves considerable MVS cycles on the host. However, actual TCP/IP performance will be better when running ICP than when running TCP/IP Offload.

The 3172 Model 3 Offload is also transparent to the TCP/UDP (User Datagram Protocol) applications as it processes the protocols and forwards them to the MVS host. The host processes the socket level applications. The 3172 Model 3 and MVS host *share* an IP address and users will not know where the processing takes place.

Since the 3172 Model 3 (operating with OS/2 TCP/IP) shares the IP address with the host, the host is also a proxy agent for network management information. The SNMP attributes of the 3172 Model 3 are treated as part of the host's standard Management Information Base (MIB). The host TCP/IP agent forwards requests for 3172 management information to TCP/IP in the 3172, which responds with the data and then relays that data to the host. Figure 67 shows the 3172 Model 3 with the offload feature.

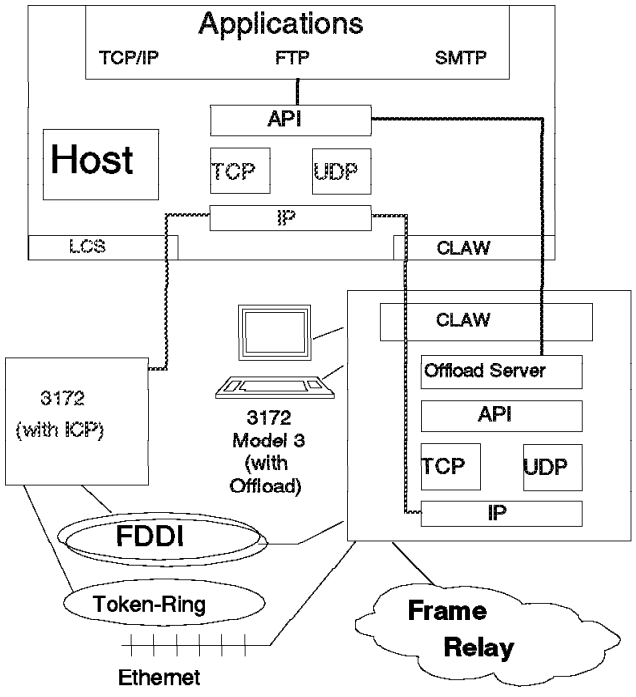

Figure 67. TCP/IP Offload Feature

3.2.1.3 LFS/ESA for MVS and VM (OS/2 Platform)

The IBM 3172 LAN File Services/ESA provides a set of services and functions to LAN servers and clients. The System/390 plus the 3172, appear as a Super OS/2 LAN server in an OS/2 LAN server environment or as a TCP/IP NFS server in a TCP/IP environment, respectively. LFS provides key services to LAN servers and clients through:

- Common Services: A base platform of programming services for server applications across MVS/ESA and VM/ESA operating environments.
- File Services: Storing and sharing of data and applications between clients on different network segments.

The LANRES 3172 configuration, in conjunction with the host, provides services to NetWare servers and clients. LANRES performs the function of extending the MVS or VM environments to NetWare servers and clients. The LANRES gateway provides four key services to NetWare clients and users:

- Disk Serving: Host DASD can be configured to be used by the NetWare server and clients.
- Print Serving: LAN data can be printed on the host system printers, and host data can be printed on NetWare supported LAN printers.
- LAN Administrators: Host systems allow administrators control of the LAN from any authorized host user.
- Data Distribution: Allows authorized host users to manipulate and manage files and directories controlled by NetWare V3.11.

3.2.1.4 3172 IP Channel Communications Program (OS/2 Platform)

The 3172 IP Channel Communications Program Version 1, running under the OS/2 operating system allows network users the ability to access multiple mainframes running TCP/IP applications from LAN-attached (token-ring, Ethernet, FDDI, and ATM) workstations by routing IP traffic between the mainframes (via channels) and workstations (via LANs).

The IP Channel Communications Program is offered either as a factory-installed preloaded feature or is field installable on an existing 3172 Interconnect Controller Model 3.

The 3172 IP Channel Communications Program′s modular design and the connectivity options provide protection against obsolescence. The open-ended design allows for addition of future functions and features.

It can coexist with the 3172 SNA Communications Program and is also similar in terms of installation, subchannel, and LAN adapter configuration, to the 3172 SNA Communications Program.

The 3172 IP Channel Communications Program supports two host connections using the Parallel channel adapter (PCA). Each can provide a channel connection to one host per channel adapter. The 3172 IP Channel Communications Program allows for two Parallel channel adapters to be installed simultaneously to provide for multiple host connections.

ESCON Single Slot Adapter: The new ESCON Single Slot Adapter provides the 3172 Model 3 with direct host attachment through an ESCON channel (10 MBps or 17 MBps), providing multiple host connections. This new ESCON Single Slot Adapter provides the same function as the existing ESCON adapter, but in a single slot configuration at a lower price.

Frame Relay: The 3172 IP Channel Communications Program along with the SNA Communications Program and/or RouteXpander/2 supports access to host applications over frame relay networks via the WAC adapters.

3.2.1.5 SNA Communications Program (OS/2 Platform)

SNA/Comm implements both local and wide area, so it will be discussed in 3.2.2, “3172 WAN-to-Host Mode.”

3.2.2 3172 WAN-to-Host Mode

The 3172 may be configured to operate over a wide area network using public access facilities. Wide area connections are provided through two different platforms.

The first is called the Multiprotocol Networking Software Solution; it is more often referred to by the simpler name of Multiprotocol Extensions (MPE).

The second platform is SNA Communications Program V1.1 (SNA/Comm). MPE and SNA/Comm both have interfaces for a frame relay network.

MPE supports SNA by bridging frame relay to a token-ring, while SNA/Comm supports SNA over frame relay natively and additionally will support SDLC line concentration in a point-to-point or multipoint network.

MPE is wide area only, while SNA/Comm is both local and wide area. Each will connect at a variety of speeds using one of several available interfaces including X.21, V.35, RS422/449 and RS/232. Please note that both platforms vary in their support for these interfaces as well as protocols.

3.2.2.1 Multiprotocol Networking Software Solution

MPE is an OS/2 application that consists of the Offload code originally developed for LAN-to-host environments coupled with IBM RouteXpander/2. This combination, in conjunction with wide-area adapter cards, provides a frame relay interface.

In Figure 68 on page 166, SNA is bridged via frame relay to a token-ring by a 3172 with MPE and then passes through an SNA gateway to a VTAM host. TCP/IP is routed over the same frame relay connection directly through the Offload code and then to MVS or VM TCP/IP.

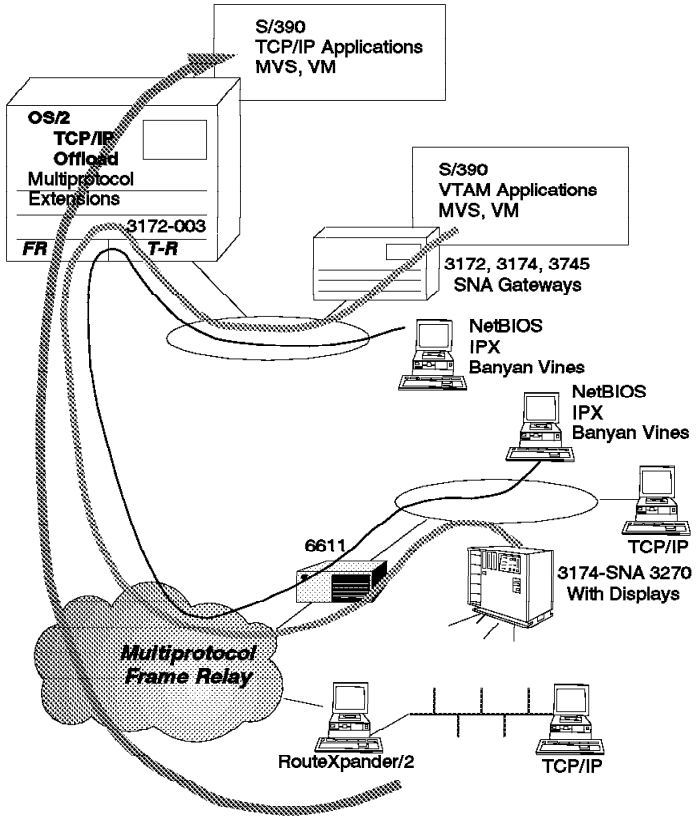

Figure 68. WAN-to-Host Mode Using MPE

In contrast, the method MPE uses to support SNA with SNA/Comm as shown in Figure 69 on page 167. You will note similarities in their support for TCP/IP. In general, MPE is the platform of choice for TCP/IP without SNA requirements, while SNA/Comm is for SNA and SNA with TCP/IP combined. Both packages will support bridging for LAN protocols such as IPX, NetBIOS, and VINES.

3.2.2.2 SNA Communications Program

The second way that the 3172 may operate over a wide area network is with the SNA Communications Program Version 1 Release 1. This product gives the enterprise the ability to concentrate SNA devices over leased lines or through a token-ring, Ethernet or FDDI LAN connection.

It contains support for the ESCON Multiple Image Facility (EMIF) that enables connection to multiple hosts (VTAM only). TCP/IP Offload for MVS or TCP/IP Offload for VM may run concurrently with SNA/Comm over a frame relay connection, thus providing a multiprotocol gateway to a System/390 host.

SNA dial support is provided to stations that implement ISO standard 3309.1. Personal computers that have modems and are running with Communications Manager/2 or Personal Communications/3270 implement this asynchronous SDLC support.

If X.25 connectivity is a requirement for SNA traffic, this is supported with RouteXpander/2 Version 2 using the X.25 Support/2 feature and SNA/Comm. This solution requires that the network nodes use RouteXpander/2 with the X.25

Support/2 feature if LANs are to be connected, or X.25 Xpander/2 for individual nodes. Please note that this solution is based on RFC1356 and does not provide for QLLC (qualified link level control).

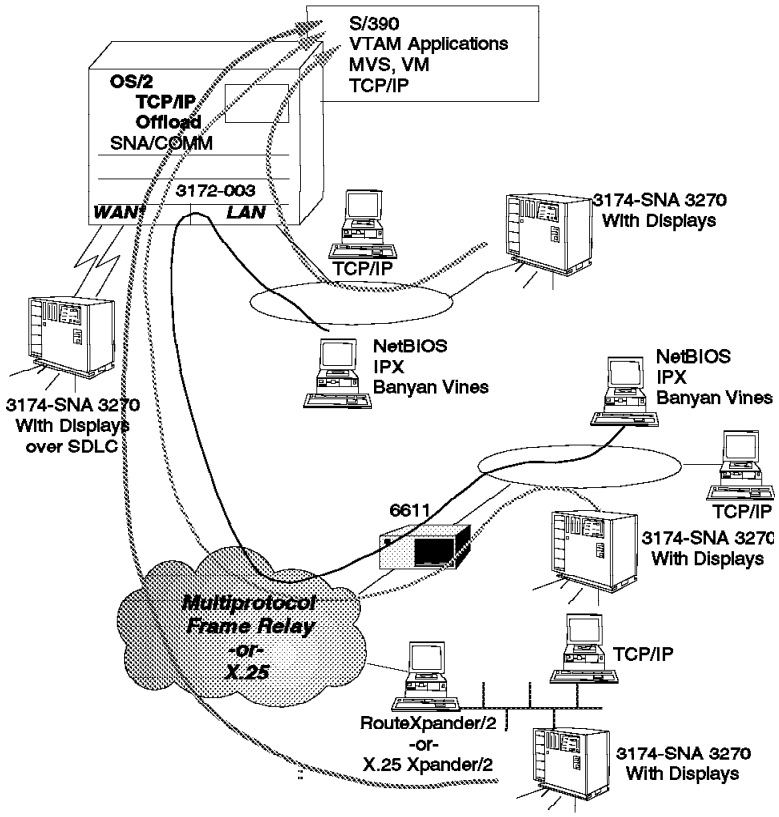

Figure 69. WAN-to-Host Mode Using SNA/Comm

3.3 3174 Gateway

The IBM 3174 family of establishment controllers provides a wide range of establishment gateway options. It has been the gateway of choice for users migrating from the traditional 3270 terminals to the intelligent workstations on LANs. It provides both local and remote gateway functions between the LAN workstations and the host network as shown in Figure 70.

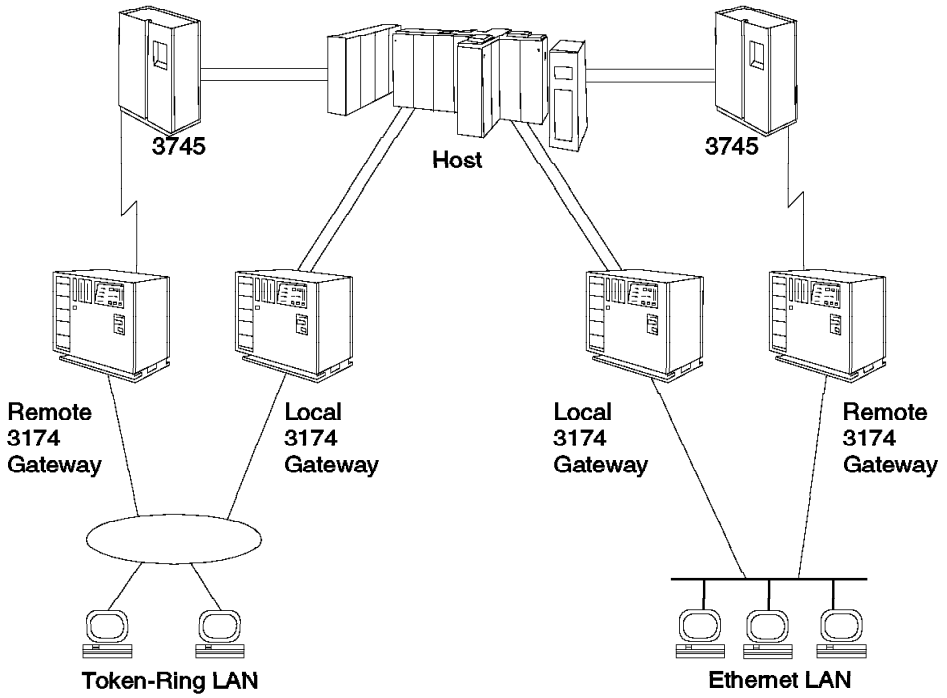

Figure 70. Remote and Local 3174 Gateways

3.3.1 SNA Environment

The 3174 provides an efficient gateway to SNA hosts from a LAN environment. End users on traditional dependent displays attached to a 3174 on the LAN have access to applications on the SNA host. End users on intelligent workstations on the LAN can also access the SNA host using 3270 emulation products.

The 3174 provides host access for PU Type 2 and Type 2.1 downstream devices, attached to the token-ring or Ethernet LAN, in both the local and remote models. The minimum software levels to support each environment are as follows:

	PU Type 2	**PU Type 2.1**
Local	VTAM V2R1 - MVS VTAM V2R1 - VSE VTAM V3R1 - VM	VTAM V3R4 - MVS/ESA VTAM V3R3 - VM
Remote	Basic NCP multipoint support	NCP V5R2.1/V4R3.1 VTAM V3R2 - MVS/VM/VSE

3.3.1.1 Enhancements

Gateway function on the 3174 was announced in 1986 and has been enhanced in the years since. Some performance enhancements include Group Poll and Duplex Multipoint. Functional enhancements include Single Link Multihost and Multihost LAN Gateway, discussed below.

Single Link Multihost Support: The Single Link Multihost support is a function that is implemented entirely in Licensed Internal Code that allows terminals attached to 3174 Models 03R, 13R, 53R, 23R, 63R (token-ring-attached controllers) and 24R, 14R, 64R (Ethernet-attached controllers) to concurrently access up to a total of eight IBM hosts that are attached to the same LAN network as the 3174. The token-ring function is provided as a customization option in 3174 Configuration Support-B Release 1 or higher Licensed Internal Code. The Ethernet feature is in configuration support C release 5 or higher. The attachment to the token-ring for these models can be either a Type 3 Token-Ring (4 Mbps) Adapter or a Type 3A Dual Speed (16/4 Mbps) Communication Adapter.

The Single Link Multihost provides multiple SNA physical unit Type 2 (PU 2.0) appearances, so that terminals attached to the same physical 3174 can communicate with multiple IBM hosts that are also connected to the same LAN network. IBM hosts view each of the multiple PU 2.0 appearances as an independent 3174 Control Unit.

The Single Link Multihost capability supports up to five concurrent sessions for each attached terminal. The five sessions can be distributed among the eight S/370 application hosts in any manner, and the terminal operator can switch among them without having to logoff/logon. The available sessions are allocated during 3174 customization.

Multihost LAN Gateway: Multihost LAN Gateway allows a LAN-attached device to take greater advantage of multiple host connectivity options available in the 3174 Establishment Controller to access hosts through the primary host link and/or the Concurrent Communication Adapter. The user can have sessions with up to three different SDLC, or eight ESCON, or eight frame relay-attached hosts. A maximum of 50 downstream physical units (DSPU) may access the SNA host through each Concurrent Communications Adapter.

LAN-attached devices that require access to multiple hosts via the Multihost LAN Gateway will define the path through the 3174 gateway to the specific host link via the SNA Service Access Points (SAPs) portion of the gateway address. This requires definition during the customization of both the LAN-attached device and the 3174 gateway.

Alerts and problem determination statistics from the LAN will continue to flow through the 3174 gateway to the SNA host over the primary link. Link events, specific to the sessions in progress via the Concurrent Communication Adapter links, will flow to the appropriate host via that link.

Multihost Gateway requires Configuration Support-B Release 2. Additional memory requirements are based on the number of DSPUs to be supported. The frame relay support is available on configuration support C5 or higher.

3.3.1.2 3174 APPN

On March 5, 1991, IBM announced the APPN network node (NN) function on the IBM 3174 Establishment Controller.

The APPN LIC (Licensed Internal Code) feature provides the functions of an APPN network node on the 3174, and allows application programs that use Advanced Program-to-Program Communication (APPC) running on either a workstation or host to communicate with a partner program anywhere else in the APPN network.

The 3174 NN feature provides topology database, directory services, and dynamic routing for the end nodes it supports. End nodes can be connected to the 3174 on the token-ring, via coax, Ethernet, X.25, frame relay or SDLC.

The 3174 NN connects to the VTAM subarea or VTAM APPN, via ESCON or Parallel channel, SDLC, X.25 or frame relay.

The 3174 NN feature coexists with 3174 functions such as LAN Gateway, Peer Communications (LAN-on-Coax), traditional 3270 traffic, and the Asynchronous Emulation Adapter. These features will continue to function as they do today, requiring no changes to the application or to the current network definition or structure.

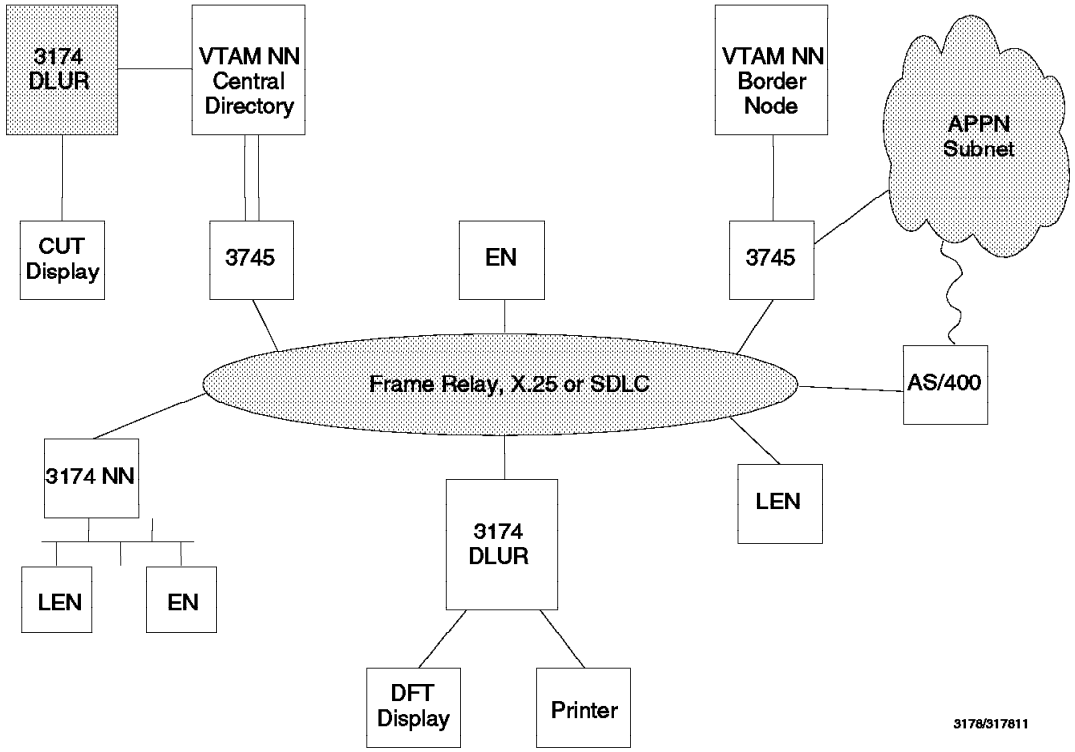

Figure 71. 3174 in APPN Environment

Link Connectivity: The 3174 continues to enhance its APPN support by adding new link types of X.25, frame relay, Ethernet and ESCON.

Dependent LUs over APPN: The DLUR function is used with the Dependent LU Server (DLUS) function in ACF VTAM V4R2 to provide APPN support for 3270 applications. DLUR/DLUS support uses an LU 6.2 session pipe to encapsulate System Services Control Point (SSCP) control data so that equivalent subarea functions are provided. This support allows the SSCP function to remain in VTAM; however, any APPN NN is able to calculate the best route for data.

The session data route is calculated directly between the LU session partners, and the session data is routed natively, not encapsulated, through the APPN network. The route may or may not traverse the DLUS node. The DLUS node and the host applications can be located anywhere in an APPN network as long as the DLUR node has a path to the DLUS node and the DLUS node has a path to the host applications.

3.3.1.3 Host′s View

The following discussion is based on functions that are present in the Configuration Support-C, which includes the APPN feature.

Local Gateway (3174 Models xxL): To VTAM, all dependent LUs are associated with their controlling PUs, which appear as if they are themselves channel-attached PU Type 2s. A logical view of this is shown in Figure 72 on page 172. Each such PU is defined in a local SNA major node with a unique channel address.

Any independent LUs that use the services of the 3174 as an APPN network node (with Configuration Support-C) are associated with the PU Type 2.1 node in the gateway. So, if there are only independent LUs, VTAM only sees one physical unit (the PU Type 2.1 node). With VTAM for MVS/ESA (V3R4 or later), the local SNA major node can be created dynamically. Please refer to the *VTAM V4R1 Implementation Guide* for more details.

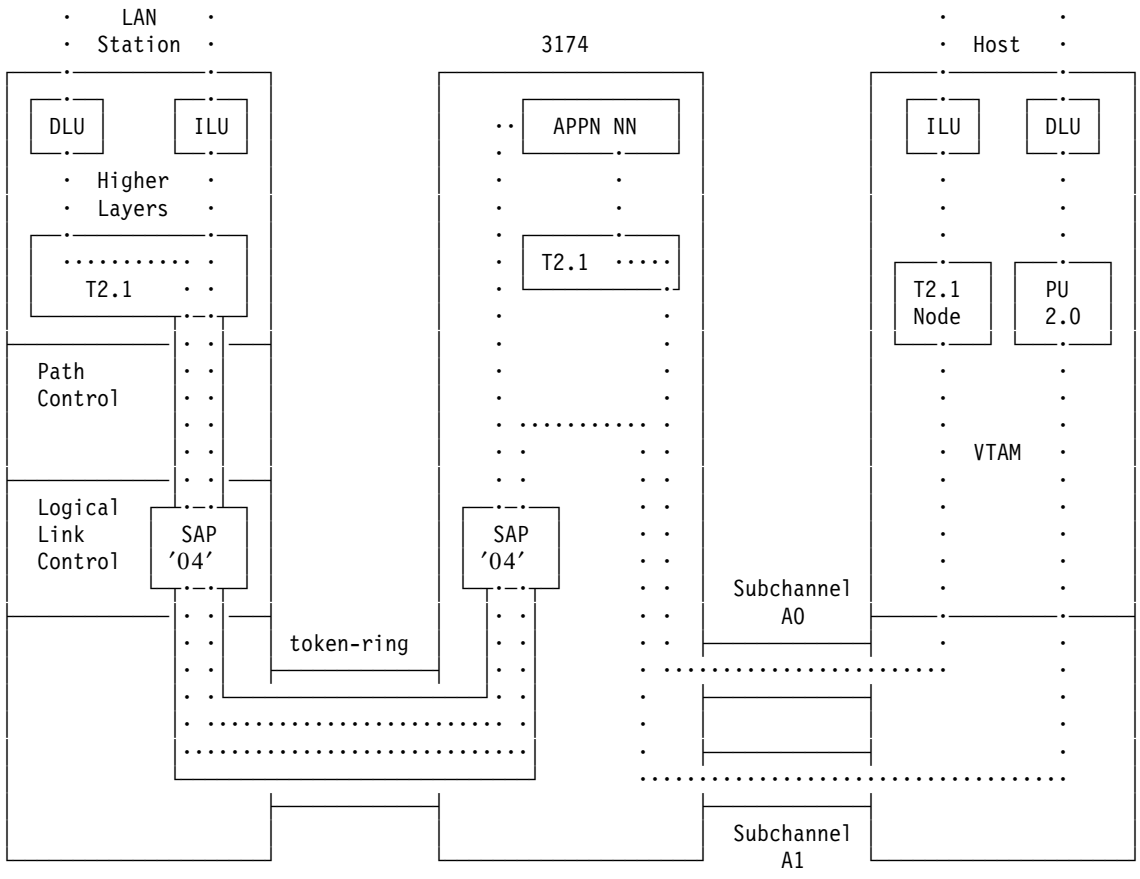

Notes:

DLU Dependent LU

ILU Independent LU

Figure 72. 3174 Local SNA Support - Logical View

Before the availability of APPN support in the 3174, it was possible to support T2.1 downstream PUs through an RPQ - 8Q0800. VTAM saw this environment slightly differently; each T2.1 node was supported individually and independent LUs could be associated with each node. Please see the *3174 APPN Implementation Guide* for more information on RPQ 8Q0800.

Remote Gateway (3174 Models x1R, x2R & 90R): To VTAM and NCP, all dependent LUs are associated with their controlling PUs, which appear as if they are themselves link-attached PU Type 2s. A logical view of this is shown in Figure 73 on page 173. Each such PU is defined as a station on a multipoint line, with a unique polling address.

Any independent LUs that use the services of the 3174 as an APPN network node are associated with the PU Type 2.1 node in the gateway. So, if there are only independent LUs, VTAM only sees one physical unit (the PU Type 2.1 node). Since NCP V5R2.1, a new facility called Group Poll was introduced to improve the performance of remote 3174 gateways. Please refer to the *NCP Planning and Implementation Guide* for more details.

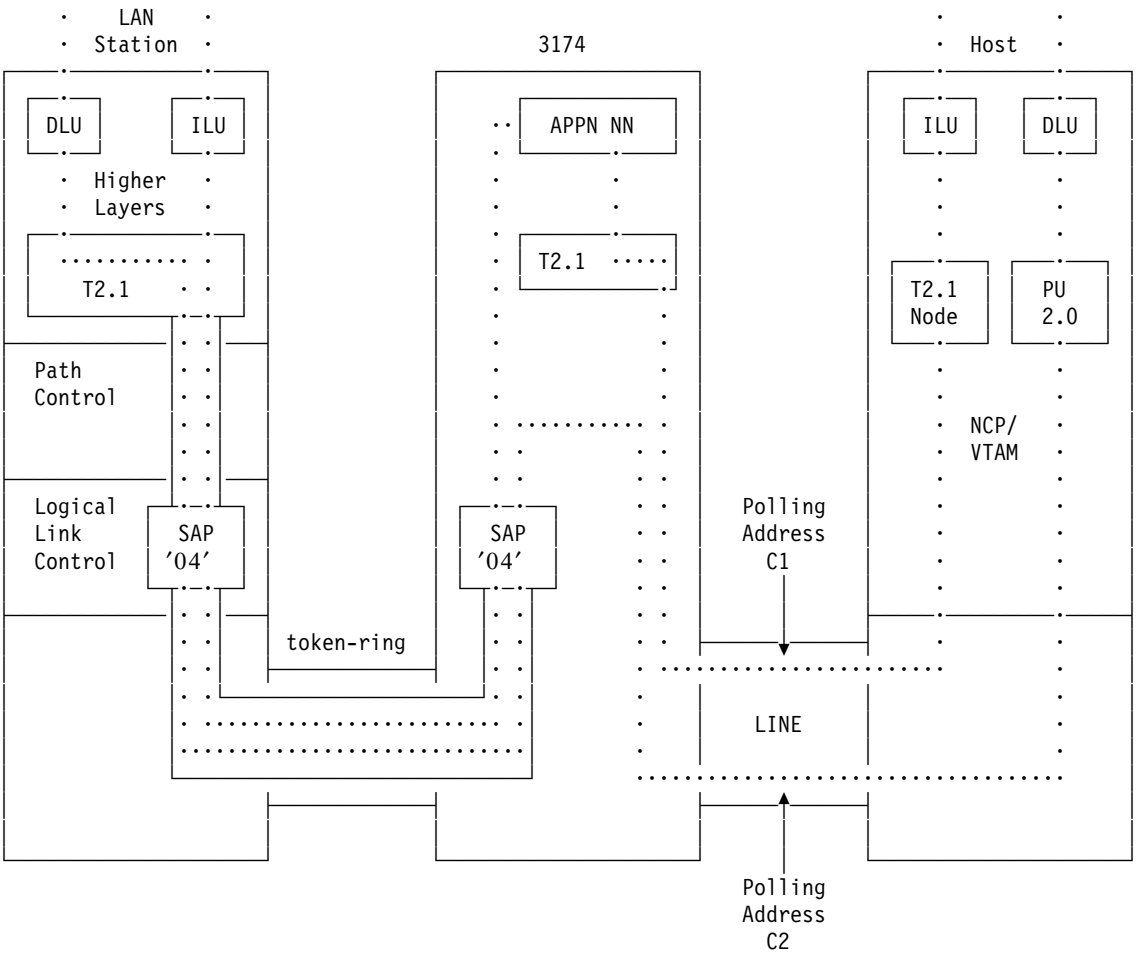

Notes:

DLU Dependent LU

ILU Independent LU

Figure 73. 3174 Remote SNA Support - Logical View

3.3.1.4 LAN Workstation's View

Regardless of whether the 3174 gateway is local or remote, the LAN workstation views it (and the host) in the same way:

- First, if there are dependent LUs (such as 3270s), the LAN workstation will have an SSCP-PU session with the host, and consequently be under its control.
- Otherwise, independent LUs merely see a direction in which they can send BINDs and subsequently exchange session data.
- If the LAN workstation were an APPN node (either network node or end node), then there would be additional APPN functions which would be supported by the 3174. Please refer to the *3174 APPN Implementation Guide*.
- Supporting the LUs is an appropriate PU - T2.1 for independent support, PU 2.0 if only dependent.

- All that the LAN workstation sees of the physical network is that it has a link to the LAN adapter defined by the MAC address in the definitions.

3.3.1.5 Gateway's View

The 3174 sees the LAN workstation as one of the following:

- A Type 2.1 node with no dependent LUs
- A Type 2.1 node with a mixture of dependent and independent LUs
- A PU 2.0 with only dependent LUs

This is determined by coding the LAN workstation's MAC address in, respectively:

- The Network Resources *only*
- The Network Resources *and* the Ring Address Assignment
- The Ring Address Assignment *only*

The last case is only what would have been coded in any non-APPN gateway. In each case, the 3174 presents the following to the host:

- A single T2.1 node with all the independent LUs from all LAN stations, plus independent LUs of coax-attached workstations using peer communications
- A single T2.1 node with all the independent LUs from all LAN stations, plus independent LUs of coax-attached workstations using peer communications; also, a PU 2.0 for each LAN workstation that has dependent LUs
- A PU 2.0 for itself and for each LAN workstation

If it is a remote gateway, these PUs are presented as individual stations on a multipoint line; if local, the PUs are presented as individual channel-attached controllers.

3.3.1.6 Establishing Connection

Unlike the 3172 SNA gateways, the LAN workstations under the 3174 gateway are not seen as switched devices. In particular, VTAM drives the establishment of connections, down to a single path (the line for remotes, subchannel for locals).

VTAM attempts to connect to the DSPUs defined to it as it would to any other PUs. Where the 3174 is mimicking these DSPUs, it will reassure the host (so that, at least for a while, VTAM will not think the effort has been a failure) and then attempt to contact them itself, over the token-ring. When the LAN workstation is ready (and waiting), it will establish an LLC Type 2 connection to the 3174 and VTAM will get a CONTACTED message. At this point, PU and dependent LU activation takes place as normal.

This, of course, only happens for LAN workstations with dependent LUs. T2.1 nodes do not require a presence like VTAM in order to *come to life*, so links between the 3174 and such nodes will be set up only between the two.

Session activation for dependent LUs happens normally, with INITSELF or unformatted logons, directly to VTAM; the 3174 is totally unaware of dependent LUs. An independent LU can initiate its own session by issuing a BIND. This flows to the 3174 and the APPN function then examines the BIND to find the destination LU. Assuming this destination LU is a host application, then the BIND will travel up to VTAM and the session will be established.

If a part of the path between the host and a LAN workstation with dependent LUs is broken, in most cases, an alternate path will require a separate set of definitions in VTAM to be activated.

Any independent LU must be associated with a particular network name. In addition to this, only one 3174 within such a network can provide a host link (through a wildcard entry). So, there are no straightforward alternate routing approaches available, except through the substitution of components.

3.3.1.7 Flow Control and Network Management

A remote 3174 gateway has the full range of SNA flow control measures available (from Virtual Route (VR) pacing, VPACING, PACING, SDLC windows and so on), which are all controllable through VTAM and NCP.

The local gateway does not have the same range of flow control measures; it is hardly a problem because there are no such relatively slow facilities as a remote link.

The 3174 is a full participant in SNA network management:

- It supports generic alerts, central site change management and control facility, and asset management on its own behalf.
- It provides a gateway for LAN management and also implements ring error monitor on the local segment.
- DSPUs can send their own alerts directly to NetView.
- APPN end nodes can send alerts to the network node in the 3174, which will forward them to NetView on the 3174′s SSCP-PU session.

3.3.2 TCP/IP Environment

The 3174, traditionally a cluster controller for 3270 host devices, was enhanced with the TCP/IP Telnet client capability to allow 3270 displays operating in CUT mode, and ASCII displays attached to the Asynchronous Emulation Adapter (AEA), to access TCP/IP Telnet servers in TCP/IP networks. This capability was offered, in March 1992, as a no-charge Telnet RPQ 8Q0935, based on Configuration Support-C Release 2 LIC.

Configuration Support-C Release 3: In May 1993, CS-C Release 3 was announced. Its base microcode includes the base functions of previous releases of 3174 Licensed Internal Code. In addition, the functions provided by the 3174 TCP/IP Telnet RPQ 8Q0935 are now integrated in the base functions of CS-C Release 3. CS-C Release 3 became available in June 1993.

RPQ 8Q1041: In May 1993, Configuration Support-c Release 4 was also announced. Included in this announcement was the 3174 TCP/IP Enhancements RPQ 8Q1041, which provides TN3270 support, TCP/IP-dependent host printer support, and SNMP MIB-II support. This RPQ, available as of April 1994, combines the token-ring support of CS-C Release 3 with the Ethernet support of CS-C Release 4.

Configuration Support-C Release 5: Configuration Support-C Release 5 and frame relay support expands the 3174 connectivity for TCP/IP support. Prior to CS-C Release 5, all TCP/IP access to and from the 3174 assumed LAN (token-ring or Ethernet). With Configuration Support-C Release 5, you are able to

Telnet to a TCP/IP host in the network via wide area network (WAN) using frame relay links.

3174 IP Forwarding RPQ 8Q1289: 3174 IP Forwarding RPQ 8Q1289 enables intelligent workstations, which are not directly attached to the 3174 using Peer Communications, access to TCP/IP hosts via the 3174 the frame relay link(s). The 3174 actually, provides static IP routing for LAN (token-ring or Ethernet) attached intelligent workstations running as TCP/IP hosts.

Configuration Support-C Release 6: The RPQ 8Q1041 was included in the base microcode of Configuration Support-C Release 6. This release made available the 3174 remote source-route bridge function, which allows users to interconnect token-ring LANs attached to the 3174 across a frame relay network.

3.3.2.1 3174 TCP/IP Enhancements RPQ 8Q1041

With the availability of the 3174 TCP/IP Enhancements RPQ 8Q1041, in April 1994 new TCP/IP capabilities were added to the 3174.

The RPQ is based on CS-C Release 4, with token-ring and Ethernet adapter support enabled.

The RPQ 8Q1041 contains the following TCP/IP enhancements:

- **TCP/IP TN3270**

  TN3270 support makes it possible for client terminals to use the TCP/IP protocol to access 3270 applications in full-screen mode. While products such as the RS/6000 do a good job of supporting ASCII terminals in full-screen mode over TCP/IP Telnet, using an ASCII terminal data stream, VM and MVS typically do not. Normal VM/MVS support for ASCII terminal communications is via line-by-line mode.

  TN3270 makes it possible to use full-screen 3270 data stream communication between VM/MVS and a client terminal, instead of the line-by-line ASCII terminal data stream. This makes it feasible for client terminals to access mainframe 3270 applications via TCP/IP, as well as traditional SNA. Customers who prefer to avoid routing SNA over a TCP/IP network can now use the 3174′s TN3270 support to build a pure TCP/IP network with IBM host access.

- **LPD (Line Printer Daemon)**

  TCP/IP dependent host printer support allows TCP/IP hosts to send ASCII print output to the 3174 for printing on the attached printers. Thus, printers can be either coax-attached or AEA-attached.

  Multiple 3174-attached printers provide print distribution for multiple TCP/IP hosts. A pool of printers can be defined as an LPD *queue*. The access to a queue can be open or restricted to a single host.

  The addition of TCP/IP host printing services through the 3174 allows offloading of the TCP/IP printing from the existing TCP/IP printing facilities. Multiple 3174-attached TCP/IP hosts will further offload TCP/IP host printing facilities.

  LPD provides only real time printing. Spooling of the print jobs is not available. If the attached host does not support sending the control file before the data, MLT storage space is used to hold the data until the control file arrives.

- **SNMP MIB-II**

  SNMP MIB-II support enhances the level of network management support offered by the 3174. SNMP now allows access to the network management parameters defined in *RFC 1213 - Managing Information Base (MIB) for Network Management of TCP/IP-based Internets*.

3.3.2.2 Configurations

You can use the TCP/IP Enhancements RPQ on all models of the 3174 that connect to a token-ring or an Ethernet, and that have sufficient memory.

In Figure 74, all CUT-mode terminals can access both TCP/IP hosts and the traditional IBM 3270 host. Using the 3174 Multiple Logical Terminal (MLT) support, the users can *hot-key* among all these sessions. The printer that is attached to the 3174 can receive print jobs from both of the TCP/IP hosts. In this example, the channel-attached 3174 is also performing the 3174's SNA gateway function, enabling the LAN-attached 3174 to access the traditional 3270 host.

TCP/IP support in the 3174 can be used concurrently with all other 3174 functions, including APPN. For example, the two 3174s in Figure 74 might be acting as APPN network nodes at the same time that they are participating in TCP/IP.

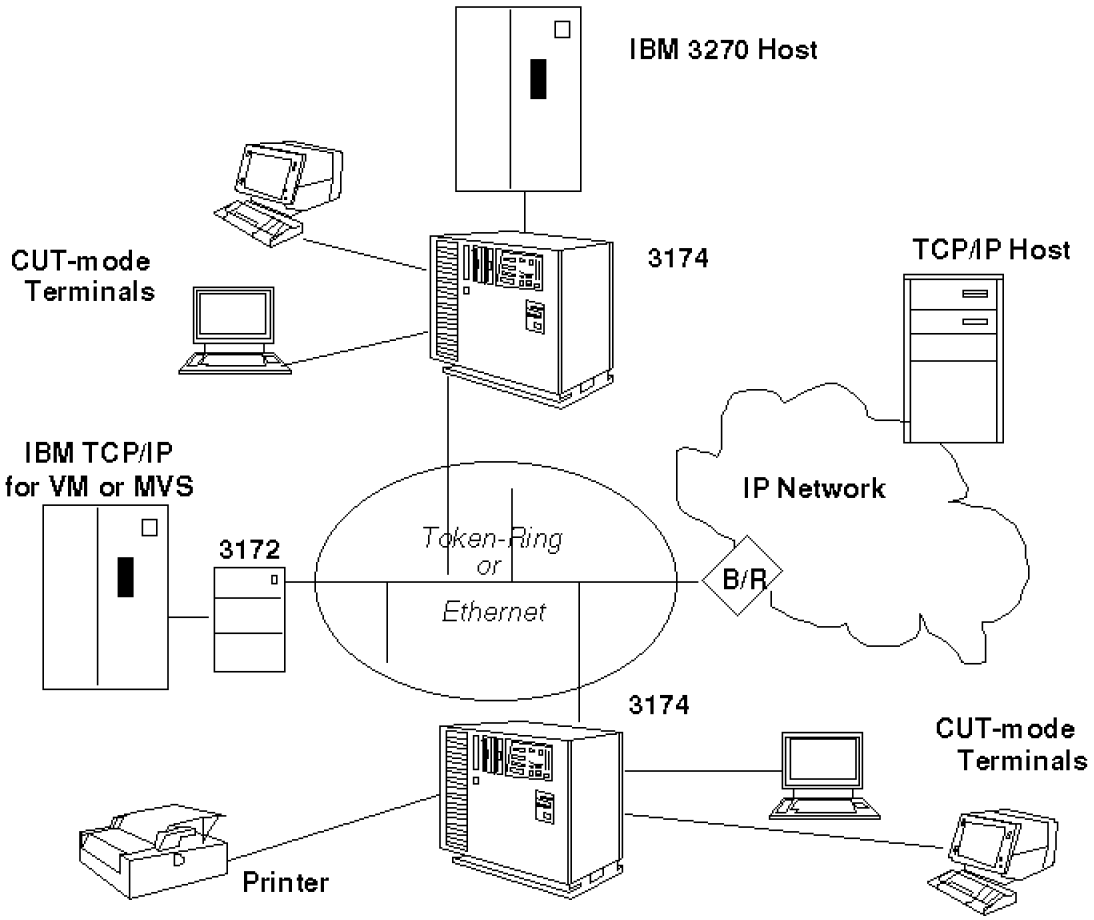

Figure 74. 3174 TCP/IP in a LAN

3.3.2.3 Devices

The TCP/IP Enhancements RPQ allows all displays that attach to the controller, except DFTs, to operate in ASCII Telnet or TN3270 mode:

- Coax displays operating in CUT mode (CUTs, or the CUT side of a DFT-E)
  - ASCII Telnet

    These are supported as DEC VT100, DEC VT200, IBM 3101 or Dasher D210 devices, using the ASCII emulation function of the 3174. (An AEA adapter is not required.)

    The TCP/IP Enhancements RPQ supports these devices in 24X80 mode only; the 132-characters-per-line support for DEC VT220 made available in CS-C Release 2 is not available for the TCP/IP sessions.
  - TN3270

    Full 3270 function is provided for these devices, including color support, models 2, 3, 4 and 5 screen sizes and Write Structured Fields. Local functions, such as copy session-to-session and split screen, are fully enabled while in a TN3270 session.
- ASCII displays that are attached via an AEA
  - ASCII Telnet

    The ASCII data stream is passed to the terminal with no manipulation by the 3174.
  - TN3270

    The 3174 performs 3270 emulation.

All printers that are directly attached to the 3174, either coax-attached or AEA-attached, can be used as LPD printers. HAP printers (printers attached to a port on a display) cannot be used for LPD. A printer that is being used for LPD cannot be used for 3270 host print sessions, nor can it be used for local copy printing.

3.3.2.4 ASCII Telnet vs TN3270

With the RPQ 8Q1041, the 3174 now provides two ways to access TCP/IP Telnet servers: ASCII Telnet and TN3270. When ASCII Telnet is used, an ASCII terminal data stream is carried on the TCP session; the server exchanges ASCII data with the 3174, usually one character at a time. When the 3174 device is a coax terminal, the 3174 performs ASCII emulation, and converts the ASCII terminal data stream so that it can be displayed on a 3270 coax device. When the 3174 device is an ASCII terminal, the 3174 merely passes the data stream on to the terminal, as it is already in ASCII characters.

When TN3270 is used, an EBCDIC 3270 data stream is carried in the TCP session; the server and the 3174 exchange messages, or blocks of data, rather than characters. When the 3174 device is a coax-attached terminal, the 3174 processes the data stream for display on the device. When the 3174 device is an ASCII terminal, the 3174 performs 3270 emulation, and converts the 3270 terminal data stream to the appropriate ASCII terminal data stream.

Figure 75 illustrates the difference between these two types of connections.

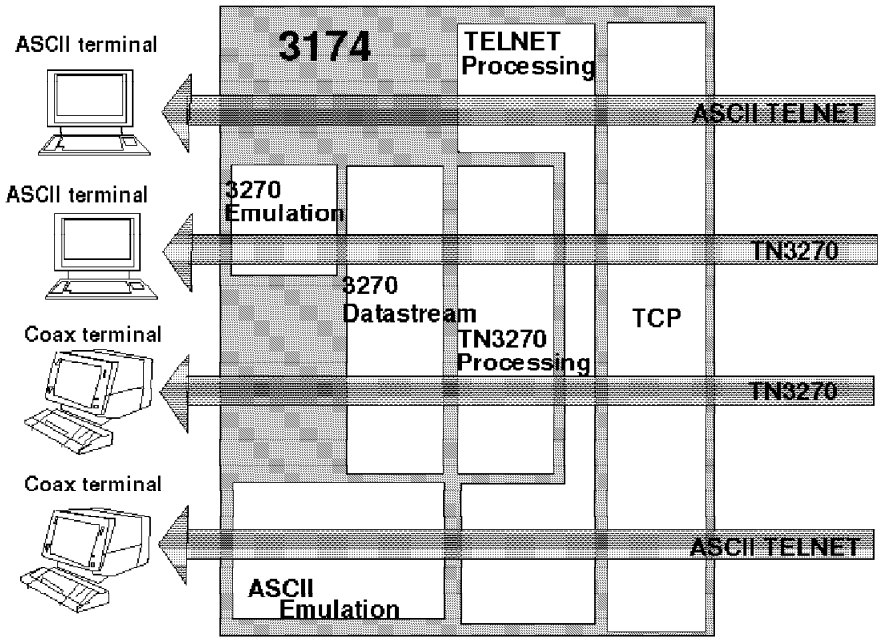

Figure 75. 3174 Processing for ASCII Telnet and TN3270

3.3.2.5 Line Printer Daemon (LPD)

The TCP/IP Enhancements RPQ expands the TCP/IP functions of the 3174 by allowing TCP/IP hosts to send print output to 3174 printers. With an LPD (Line Printer Daemon) server, the 3174 can accept print jobs for either coax or ASCII printers that are attached to the 3174. This function is *not* provided for printers that are attached to displays (HAPs).

The 3174 is not a spooling device; jobs are printed as they are received.

The 3174 assumes that the print data is ASCII. When the output printer is a 3270 printer, the ASCII emulation function of the 3174 is used; when the output printer is an ASCII printer (AEA-attached), the data is passed without conversion to the ASCII device. Figure 76 illustrates this processing.

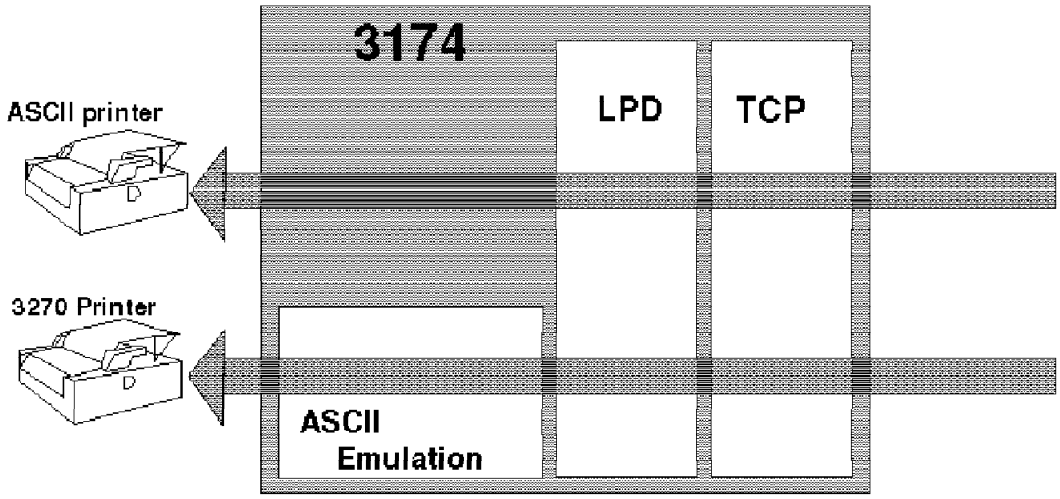

Figure 76. 3174 LPD Support

3.4 3745/3746-9x0 Gateway

The 3745/3746-9x0 family of communications controllers provide a versatile set of network gateway facilities.

They can act as a local gateway (channel attached) or remote gateway (remote linked) for token-ring or Ethernet LAN downstream workstations as shown in Figure 77.

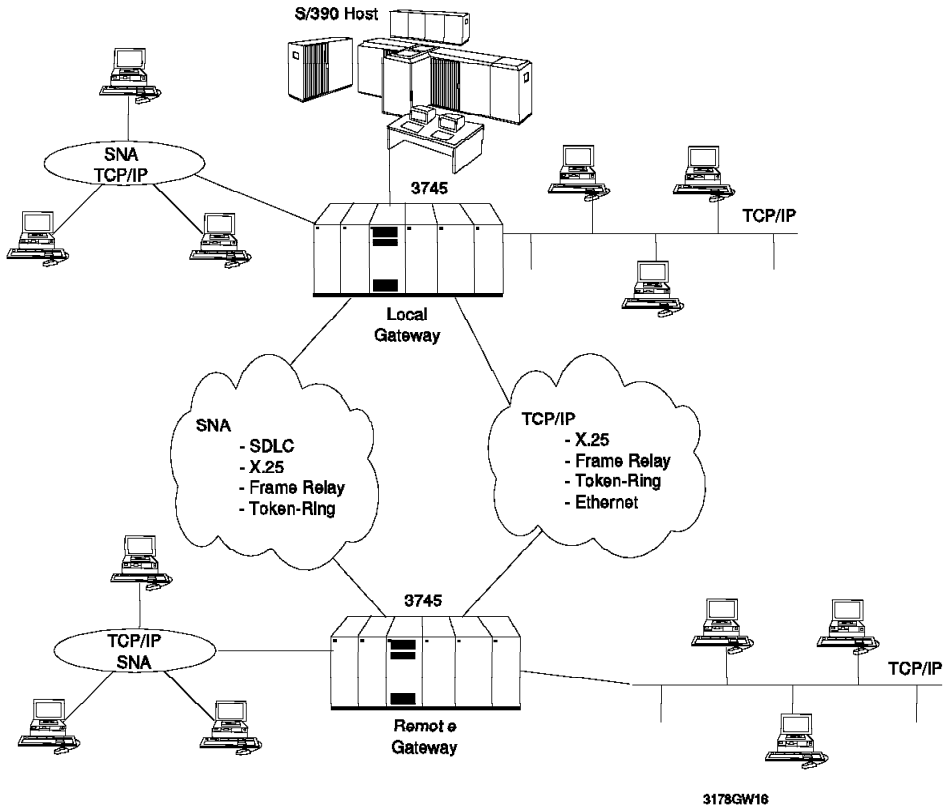

Figure 77. 3745 Gateways

3.4.1 3745/3746-9x0 Models

The IBM 3745 family of Communications Controllers provides a full range of models, answering customer requirements for a wide variety of connectivity, options, performance, and availability features.

The 3745 is available in configurations able to handle almost any possible need and size of network. They can serve equally well as data processing center nodes, concentrator access nodes, and intermediate routing nodes.

The IBM 3745 is available in two basic packages. Models 130, 150, and 170 are housed in a single compact cabinet that holds all the expansion capabilities, while the larger models, 210 and above, can be expanded by the addition of external cabinets.

The compact models have been configured for three practical applications. Model 130 is a special configuration for users needing remote host-to-host connection via high-speed lines. It also supports token-ring and Ethernet LANs, and can be used as a local or remote LAN gateway. Model 150 is particularly suitable for remote concentration, supporting up to 32 low to medium-speed lines, with Ethernet and token-ring LANs, and high-speed lines. Model 170 is an excellent communication controller for small to medium networks, supporting all the options needed in this environment.

Models 210 and above are medium to large communication controllers. Model 210 contains a single Central Control Unit (CCU), and Model 410 contains dual CCUs to provide flexible backup facilities and enhanced performance. Models 310 and 610 are single and dual versions of the 3745 featuring a higher performance CCU, which increases the throughput of the 210 and 410. These models can be expanded to support many token-ring and Ethernet LANs, high-speed lines (up to 2 Mbps), both Parallel and ESCON channels, as well as hundreds of low- to medium-speed lines (up to 256 Kbps). The X1A models offer new facilities for local and remote management of the 3745 as well as increased storage of up to 16 MB (31A and 61A).

The MOSS (Maintenance and Operator Subsystem) is the companion to NetView. It plays a role in overall network management by sending hardware alarms and alerts to NetView, which complement the NCP NMVT (Network Management Victor Transport). In addition, the MOSS allows you to run concurrent diagnostics, swap resources in case of failure, etc. It also contains all the files describing the machine′s configuration and history.

3.4.1.1 IBM 3746-900 Expansion Unit

The 3746-900 expansion unit is supported by 3745 21A, 31A, 41A, 61A and 17A. It offers the following capabilities:

- Increase in the number of serial ports
- Increase in the number of token-ring ports
- Introduces ESCON support
- Performance improvements
- SNA Subarea
- APPN Composite Network Node (CNN)
- Frame Relay DTE and DCE

The 3746-900 makes the 3745 a better offering for interactive traffic and an outstanding performer for the client/server traffic.

When supporting subarea and APPN Composite Network Node (CNN), it still requires the NCP presence.

3.4.1.2 IBM 3746-900 Network Node

The IBM 3746-900 Network Node (NN) is a function which provides stand-alone (independent from NCP) APPN capabilities to the 3746-900. This is done by adding a new separate box, LAN-attached, Network Node Processor (NNP). Token-ring and SDLC ports on the 3746-900 are owned by either the 3746-900 NN function or any of the NCPs running in the attached 3745. ESCON ports can be shared.

The IBM 3746-900 NN supports SNA/APPN traffic, without the need for NCP. It continues to support SNA-dependent LU traffic with the DLUR (Dependent Logical Unit Requester) function.

Although the 3746 Model 900 Network Node function is independent of NCP for the APPN NN and DLUR functions, the 3746 Model 900 (the hardware, not the function) remains connected to a 3745-xxA.

3.4.1.3 IBM 3746-950 Nways Multinetwork Controller

The IBM 3746-950 Nways Multinetwork Controller is a stand-alone machine that provides the APPN NN and DLUR functions without the need of a 3745 and NCP.

The APPN NN and DLUR is provided by the IBM Nways Controller in conjunction with the Network Node Processor (NNP). Opposed to 3746-900 NN the NNP owns all the 3746-950 ports.

The IBM 3746-950 does not support the attachment of a 3745-xxA, thus it cannot offer subarea functions.

3.4.2 IBM 3746-9x0 Connectivity

The IBM 3746-950 Nways and the IBM 3746-900 Network Node provide the same connectivity, including:

- The possibility of installing up to 10 adapters
- A choice of three adapter types:
  - TRA (Token-Ring Adapter)
  - CLA (Communication Line Adapter)
  - ESCA (ESCON Adapter)

Communication Line Adapter (CLA): The Communication Line Adapter (CLA) consists of a Communication Line Processor (CLP) connected to up to four Line Interface Couplers (LICs) types 11 and 12 (LIC11 and LIC12). The CLA supports up to 120 lines. The CLA supports communication lines operating in half or full duplex mode. The protocols supported, with NCP V7.3 and NPSI V3.0, are the following:

- SDLC
- X.25
- Frame relay

The CLA supports the following physical interfaces:

- V.24 leased and switched lines (600 bps up to 19.2 Kbps)
- V.25 bis protocol over V.24 switched lines
- V.35 leased lines (56 Kbps up to 2.048 Mbps)
- X.21 leased lines (600 bps up to 2.048 Mbps)

Note: The IBM 3746-900 provides frame relay networking capabilities in conjunction with NCP (starting with NCP V7R2) running in the 3745 side.

The Communication Line Processor (CLP) can simultaneously connect up to 500 active physical units (PUs), such as PS/2 or IBM 3174, carrying APPN or SNA (dependent LU traffic). In the IBM 3746-900 NN, the CLP can connect up to 2000 PUs, of which a maximum of 500 PUs can be controlled by the 3746 Network

Node Processor; the remaining PUs are controlled by the NCP running in the 3745 side.

Token-Ring Adapter (TRA): The Token-Ring Adapter (TRA) consists of a Token-Ring Processor type 1 or 2 (TRP1 or a TRP2) connected to one or two Token-Ring Interface Couplers (TIC3s) type 3 operating at 4 Mbps or 16 Mbps. The TRP2, an enhancement of the TRP of the IBM 3746-900, supports the 3746 Network Node functions.

Token-ring adapters allow a large number of token-ring LAN stations to be attached to the controller. In the IBM Nways Controller, a single Token-Ring Adapter (TRA) with one or two TIC3s can connect up to 500 active PUs, such as PS/2 or IBM 3174, at the same time.

Each TRA in the IBM 3746-900 NN supports up to 2000 active PUs, of which 500 can be controlled by the 3746 Network Node Processor; the remaining ones are controlled by the NCP running in the 3745 side.

ESCON Adapter (ESCA): The ESCON Adapter (ESCA) connects the IBM 3746-950 Nways Controller to a host via ESCON optical fibers. It consists of an ESCON Channel Processor type 2 (ESCP2) connected to one ESCON Channel Coupler type 2 (ESCC2). In the 3746-900 NN there is either a ESCP1 or a ESCP2 connected to one ESCC1 or ESCC2.

The ESCP2, an enhancement of the ESCP1 of the IBM 3746-900, supports the 3746 Network Node functions. When used in the IBM 3746 Model 900 Network Node, the ESCP2 can be shared by an active 3745 (NCP) and the 3746 Network Node Processor.

In the IBM 3746-900 NN, each ESCA2 supports up to 16 logical connections to the host LPAR (Logical Partition) and stations, of which four can carry the 3746 APPN Network Node Processor traffic. The remaining ones will carry the 3746 NCP-controlled traffic.

3.4.3 3745 SNA Environment

In the SNA environment, the 3745 running an NCP is a fully-addressable gateway. That is, the host views it as a PU with attached PUs (DSPUs) and LUs. Each LAN workstation acting as a DSPU has an SNA address which the 3745 associates to the workstations′ MAC addresses. The 3745, via its NCP, using the SNA addressing structures, has the intelligence to route frames within and between SNA networks.

In this section we discuss the token-ring gateway facilities accessed via the 3745′s Token-Ring Adapter (TRA) and the implications that SNA has on the capabilities of the 3745 as a gateway.

3.4.3.1 Host′s View

VTAM views the NCP on 3745 as a node PU Type 4. It views the downstream LAN workstations as switched devices. The attributes of these switched devices are described on a VTAM configuration file called Major Node, using one of the following two statements:

- VBUILD TYPE=SWNET - Defines a Switched Major Node that contains the PU and LU definitions. Using this definition statement, IDNUM and IDBLK must be coded on the PU statement and their values must match those defined on

the LAN device. For a PU Type 2.1, CPNAME may be used in place of the IDNUM and IDBLK.

- VBUILD TYPE=MODEL - Defines a Model Major Node (VTAM V3R4 or later). Using this statement, an undefined LAN workstation dialing into the host, can dynamically be defined and activated by VTAM referring to the Model Major Node definitions. Each entry for a PU or an LU in a Model Major Node has no relationship with any other PU or LU. Each PU and LU is completely independent. This enables an undefined device to pick up an appropriate PU and LU definition as it dials. For more details on the coding of model major nodes, please refer to the *VTAM for MVS/ESA Implementation Guide.*

3.4.3.2 LAN Workstation′s View

Since this is a switched environment, a LAN workstation has to know how to locate the 3745 gateway; thus, it is necessary to define it on the MAC address of the 3745 and the SAP that will take it to the desired VTAM. Some parameters in the LAN workstations, must match with those in VTAM (either in the Switched Major Node or Model Major Node).

3.4.3.3 Gateway′s View

The 3745 sees the LAN workstation as a switched connection with its own LINE and PU, and it knows the MAC address of each one.

The 3745 views the SNA host as an upstream PU Type 5 SNA subarea node. Since it is not transparent to the SNA host, as in the case of the 3172, NCP must be configured as an SNA node itself and must support and execute some logic to maintain its own host sessions. This is in contrast to the 3172 which is transparent to the host and hence, the host (VTAM) does not expect the 3172 to maintain its own host sessions.

3.4.3.4 Establishing Connection

The method by which connections and sessions are established in the 3745 gateway environment is similar to those used in the 3172 environment. The LAN workstation dials in by establishing a link with the NCP using LLC Type 2 protocols. The LAN workstation then issues its station ID (IDBLK/IDNUM and/or CPNAME) in an XID to the NCP. NCP then passes the device station ID to VTAM for verification. VTAM then finds a Switched or Model Major Node definition that matches with the XID.

NCP sets up a link station for this LAN station. At this point PU and LU activation and session establishment can take place as normal. Due to the switched nature of the connection, VTAM is not concerned with the path that the LAN station uses to establish a session. This provides the LAN station with alternative connections if one of the TICs on the 3745 becomes unavailable; it can insert into any of the other 3745 rings and still establish a host session. All it needs to do is to define the MAC address of the TIC that it will go through as its destination address.

For a more detailed description of the session establishment flow, please refer to *Network Control Program Version 6.2 Planning and Implementation Guide*.

3.4.3.5 Flow Control Network Management

NCP in the 3745, together with VTAM, has implemented the SNA flow control mechanisms in their entirety. Flow control has significant effects on network performance. Severe and prolonged congestion in one part of the network affects other parts of the network. Flow control procedures optimize network throughput.

NCP uses the SNA flow control mechanisms such as virtual route pacing and session pacing to control traffic originating from the DSPU. It also has internal flow control mechanisms that implement NCP Buffer Slowdown to regulate the amount of traffic through the NCP when it detects traffic congestion. The threshold at which this occurs is predefined in the NCP SLOWDOWN parameter on the BUILD definition. At this stage, NCP performs RNR (Receive Not Ready) polling for the token-ring links and withholds virtual route and session pacing responses to reduce the number of buffers in use. When the buffer supply is sufficiently replenished, NCP automatically resumes normal operation.

A more detailed discussion of NCP flow control methods can be found in *Network Control Program Version 6.2 Planning and Implementation Guide*.

3.4.4 3745 TCP/IP Environment

In the IP environment, the NCP IP ports are directly addressable, unlike the 3172. That is, each of its LAN adapters will have its own IP address, just like any other IP node in a network. In addition to this, the 3745 will also act as an IP router. TCP/IP support is available in ACF/NCP Version 6.0 or later. NCP V6.0 supports TCP/IP only on Ethernet. NCP Version 7.0 supports on token-ring as well Ethernet V2 or IEEE 802.3 mode. IP routing tables in NCP V7.0 can be dynamic. SNA is supported only over the IBM Token-Ring LAN and not on the Ethernet LAN.

The relationship between the different components in this environment is shown in Figure 78 on page 187.

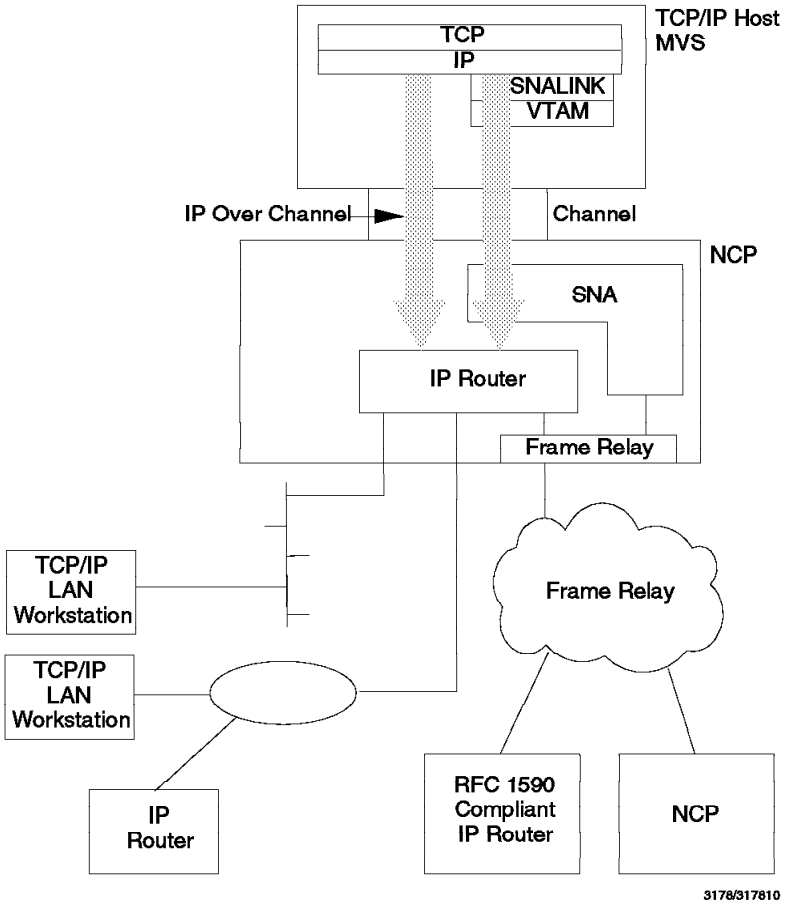

Figure 78. NCP IP Protocol Interfaces

3.4.4.1 Host′s View

A TCP/IP LAN workstation is seen by the TCP/IP host as an IP address. With the 3745 as a gateway, the host TCP/IP knows that the LAN station is on another network that it cannot address directly, and that it must use the 3745 IP router (either explicitly or by default). The host TCP/IP will have a link to this IP router and send all traffic for the target TCP/IP LAN station to that gateway.

Prior to NCP V7.3 the IP datagrams were carried enveloped over an SNA session between the VTAM and NCP. The NCP V7.3 was enhanced to support native IP traffic over the channel as well.

3.4.4.2 LAN Workstation′s View

A LAN workstation has to know that to get access to the TCP/IP host, it will have to go through at least one gateway. The LAN workstation recognizes the host as an IP address through one of the following:

- Static routing definitions in GATEWAY statements
- Dynamic information gained from RIP

It will initiate sending to the 3745 IP router first. Knowing the IP address, the LAN workstation will use ARP to learn the MAC address of the 3745 LAN adapter.

3.4.4.3 Gateway′s View

As shown in Figure 78 on page 187, the 3745 gateway acts as an IP router that provides a link between the LAN workstation and the TCP/IP host. Since this is a router, there is no restriction imposed by the protocol on the number of LAN or host connections you can have. The traffic is supported on an any-to-any basis among all of these connections.

The 3745 does not necessarily know anything about the host applications or LAN workstations (although it is possible to hard code some IP-MAC address relations through the IPHOST generation statement in ACF/NCP V6.0 or later). Thus, when a datagram arrives, addressed for one of the host applications or LAN workstations, the 3745 will use the IPGATE definitions to decide where to send it. If it is to be sent beyond this 3745, to the host, then the datagram is either passed to the appropriate SNA session or to the IP channel.

3.4.5 3746-9x0 SNA/APPN Environment

With the introduction of the 3746 APPN (Network Node) NN and Dependent LU Requester (DLUR) function, the IBM 3746 Model 900 Expansion Unit and the new 3746 Nways Controller (also called 3746 Model 950) can be used for native (not requiring NCP) connections to APPN and non-APPN SNA resources. Currently only ESCON, token-ring, and SDLC attachments are supported. Frame relay and X.25 connectivity has been previewed and will become available in 1996. The remote nodes can be APPN nodes, PU Type 2.1 PU Type 2.0, or PU Type 1.0. PU Type 2.1 are also called Low Entry Networking (LEN) nodes.

Note: All node types mentioned may contain dependent LUs; however, only APPN and LEN nodes support independent LUs. Essential differences between APPN and non-APPN nodes are that only APPN nodes can establish CP-CP sessions.

The 3746 Model 900 (3746-900) and the 3746 Model 950 (3746-950), collectively referred to as 3746-9x0, provide attachments to APPN end and network nodes. They are able to participate as a network node in an APPN network. In addition, exploiting the 3746 NN intermediate session routing and DLUR functions, provide continued support for non-APPN nodes, without having to upgrade these to APPN capable nodes themselves, when migrating from a subarea to an APPN networking environment.

Both 3746 Model 900 (3746-900) and 3746 Model 950 (3746-950) controllers are based on the same hardware and software, and provide identical APPN functions. Conversion of the 3746-900 into a 3746-950 is possible. 3746-9x0 NN and DLUR functions are provided by the 3746-9x0 in conjunction with the (service) LAN-attached network node processor (NNP). The NNP is responsible for session establishment and DLUR functions while microcode running on the 3746-9x0 processors performs intermediate session routing.

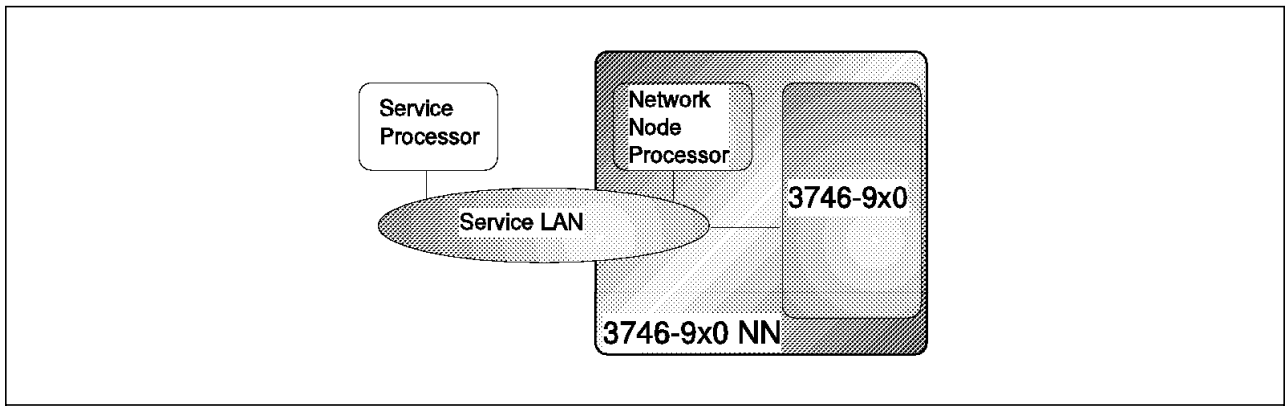

Figure 79. IBM 3746-9X0 NN. 3746-9x0 plus LAN-attached Network Node Processor

The 3746-950 is a truly stand-alone machine currently (IP and frame relay switching, or FRFH, functions have been previewed) providing the APPN NN and DLUR functions only. As can been see in Figure 80, all ports and attached devices are owned by the 3746-950 NN function.

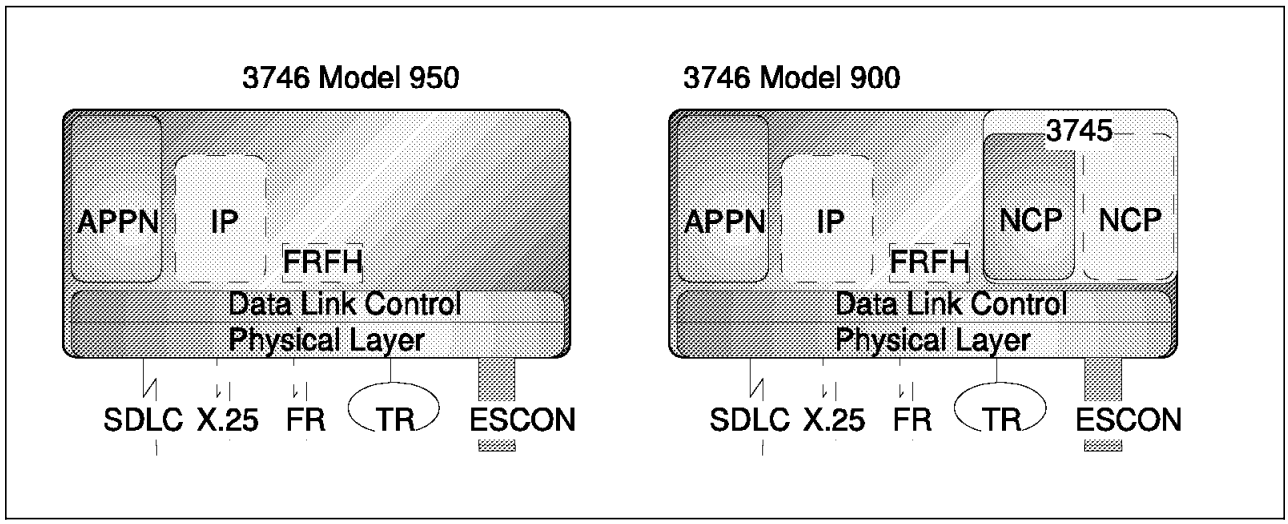

Figure 80. IBM 3746-9X0 Protocol Stacks. Multiple protocol stacks can gain ownership of 3746-9x0 ports; X.25 and frame relay have been previewed.

Figure 80 depicts the protocol stacks that can own 3746-900 ports and attached devices. The 3746-900 is a hybrid machine that continues to be bolted to a 3745, and provides the NN and DLUR functions in conjunction with, and fully independent from, NCP controlled subarea (*and* APPN when part of a VTAM/NCP composite network node) functions. Stand-alone IP and frame relay switching (FRFH) functions have been previewed. Port sharing, to provide simultaneous connectivity for attached devices to both NCP and 3746-900 NN functions, is currently limited to ESCON ports. SDLC and token-ring ports, and attached devices, are owned by either the 3746-900 NN or by one of the NCPs in the attached 3745. Ownership is decided by the order in which a port is activated (first activator gains ownership) and can easily be moved. For example, to move a port from the NCP (subarea/APPN) to the 3746-900 NN (APPN) function, requires inactivation of the port from NCP and activation from the 3746-900 NN.

The 3746-950 and/or 3746-950 NN and DLUR functions make these machines the recommended choice when migrating to an APPN networking environment. Both the 3746-900 and 3746-950 NN functions enable attachment to an APPN network,

while the DLUR function allows non-APPN nodes connected to the 3746-9x0 to participate in APPN networking, but relieves them from being APPN-capable themselves. The sharing of ESCON ports, and easy shift of ownership of token-ring and SDLC ports and devices, on the 3746 Model 900 simplifies the subarea to APPN migration of NCP controlled resources even more.

3.4.6 3746-9x0 TCP/IP Environment

As a statement of direction, the IBM Nways Controller and the IBM 3746 Model 900 Network Node will be provided with a native IP support over token-ring LAN, frame relay links, PPP and ESCON channel, supporting the following dynamic routing protocols:

- RIP (Routing Information Protocol)
- OSPF (Open Shortest Path First)
- EGP (Exterior Gateway Protocol)

The IBM Nways Controller and the IBM 3746 Model 900 Network Node will support SNMP V2 in order to be managed from any SNMP Manager such as NetView from AIX.

The IBM Nways Controller and the IBM 3746 Model 900 Network Node will be able to concurrently support APPN and IP traffic.

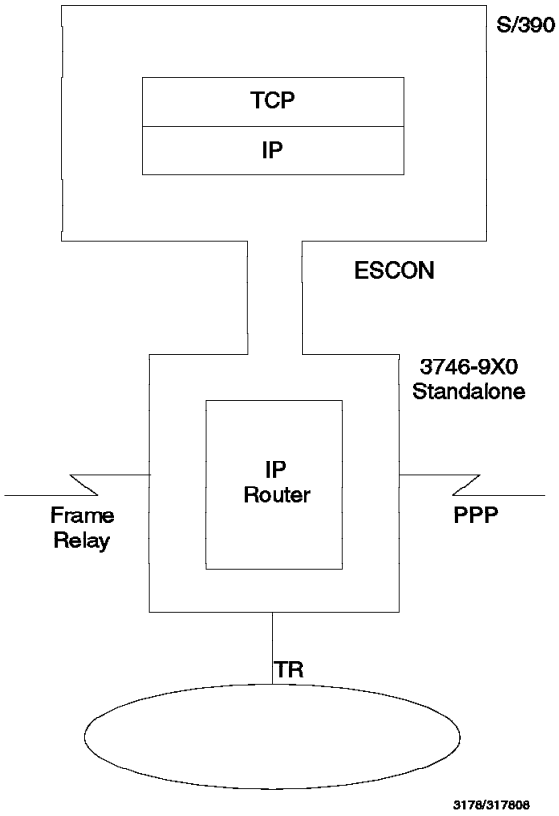

Figure 81. NCP IP Protocol Interfaces

Native IP support over ESCON will no longer need SNALINK in the Host, allowing you to route native IP traffic from any workstation in the network, up to the mainframe.

The IBM 3746-900 and IBM 3746-950 will be able to perform remote IP traffic consolidation while being a high performance, high connectivity IP router.

3.5 RISC System/6000

Several models of the RISC System/6000 have been shipped with a wide range of I/O capacity and CPU power. In general, the I/O capacity is smallest in the 2XX machines (2 slots) and 3XX machines (4 slots), medium in the 5XX machines (8 slots) and high in the 9XX machines. symmetric multiprocessor (SMP) systems are now available also. The RISC comes in desktop (2XX, 3XX), deskside (5XX) and rack-mounted (9XX) packages.

The main advantage of the RISC over the other gateways is that it can be an application platform as well as the communications gateway, providing a complete solution.

3.5.1 AIX TCP/IP

The TCP/IP support for AIX is a feature of the base operating system. Link types supported include token-ring, Ethernet, FDDI, X.25, SLIP (Serial Link Access Protocol: async), and ESCON and Parallel Channels. The TCP/IP support includes socket programming.

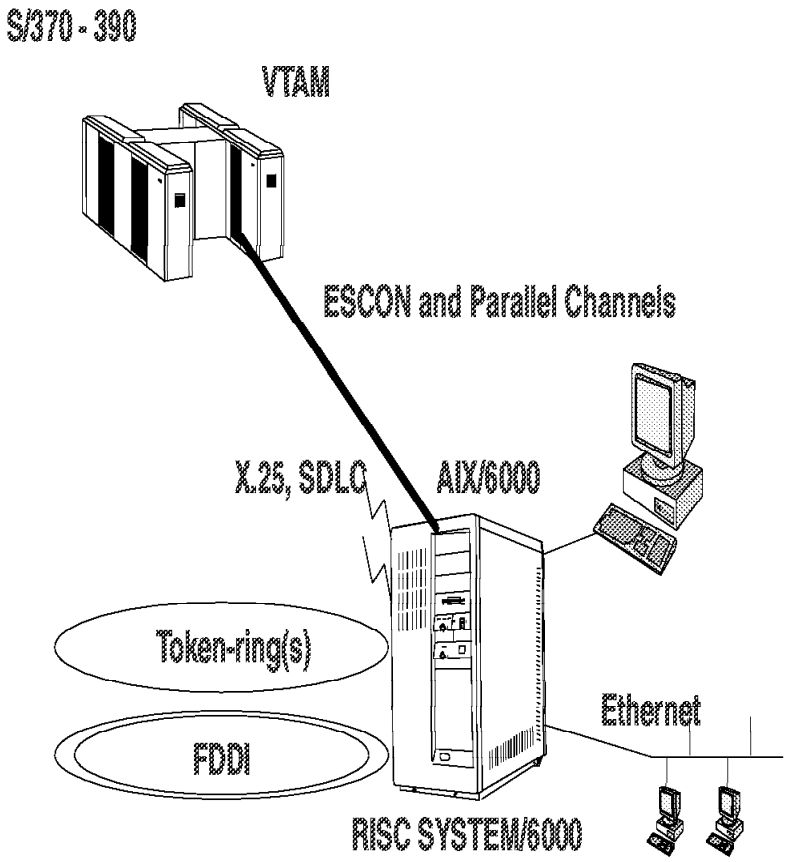

Figure 82. RS/6000 IP Environment

3.5.2 AIX SNA Server/6000

AIX SNA Server/6000 Version 2.1.1 provides end-user configuration functions and application program interfaces that permit user-provided applications to communicate with other applications, residing on other IBM systems, using SNA as the communications protocol.

Network Functions: The following support is provided:

- Subarea networks for dependent communications controlled by a host, using dependent LU types (LU 6.2, LU 3, LU 2, LU 1, and secondary LU 0).
- APPN networks for peer-to-peer communications using independent LU 6.2. This support enables the IBM RISC System/6000 to function as a LEN node, end node or network node.

Application Program Interfaces (APIs): AIX SNA Server/6000 Version 2.1.1 includes APIs that can be used to write Transaction Programs (TPs). These APIs enable the RISC System/6000 to process transactions across an SNA network. The APIs included are:

- Operating System Subroutines for LU Types (1,2,3 and 6.2).
- Library Subroutines for TP Conversations for LU types (1,2.3 and 6.2). These subroutines are also used to send and receive Network Management Vector Transport (NMVT) data between the RS/6000 workstation and the SSCP in an SNA host.
- LU 0 for LU 0 primary and LU 0 secondary.
- Generic SNA. This provides subroutines that interact directly with the services component of AIX SNA Server/6000 Version 2.1.1. These subroutines can be used to program SNA functions that are not available from AIX SNA Server/6000 Version 2.1.1.
- Management Services API enables the RS/6000 workstation, configured as a node in an APPN network, to function as a Management Services (MS) entry point.
- Common Programming Interface for Communications (CPI-C). This API allows program-to-program communications using LU 6.2 and provides maximum portability of applications between various SNA platforms that support IBM′s System Application Architecture (SAA).

LU Types

- LU 6.2 for independent peer-to-peer networks and dependent for subarea networks.
- LU 1, 2, and 3 for dependent communications.
- LU 0 consisting of primary and secondary support. The primary support is for communications with a PU Type 2.0 node such as a Store System Controller or a Terminal Controller. The secondary support is for communications with a host system emulating a PU Type 2.0 node.

Link Connectivity: AIX SNA Server/6000 Version 2.1.1 supports the following link-level communication protocols:

- SDLC
- X.25
- Token-ring
- Standard Ethernet

- IEEE 802.3 Ethernet
- Fiber Distributed Data Interface (FDDI)

AnyNet/6000: The Sockets over SNA feature provides a sockets application programming interface (API) for an SNA environment. It presents a TCP/IP sockets interface to the applications programmer, but uses SNA as the underlying network transport. Sockets over SNA does not require any changes to existing applications written to a sockets interface that uses internetwork addressing; hence, no programmer time is needed to modify sockets applications to run on an SNA network. Also, it does not require an organization to administer two networks, as would be the case if parallel networks or multiprotocol routers were used.

AnyNet/6000 APPC over TCP/IP is compatible with the existing AnyNet family, which includes AnyNet/MVS and AnyNet/2.

With AnyNet/6000 APPC over TCP/IP, customers can use the powerful LU 6.2 Advanced Program-to-Program Communication (APPC) interface to communicate between workstations in a TCP/IP environment. Any APPC or CPI-C application, such as CICS/6000 and DB2/6000, can communicate between workstations or a host to a workstation across a TCP/IP network. AnyNet/6000 APPC over TCP/IP provides this without change to application programs.

AnyNet/6000 APPC over TCP/IP is compatible with the existing AnyNet family, which includes AnyNet/MVS and AnyNet/2.

AnyNet/6000 APPC over TCP/IP provides the customer with the ability to reduce costs by elimination of duplicate networking hardware, software, and communication lines. Additionally, programmer productivity and network management are improved.

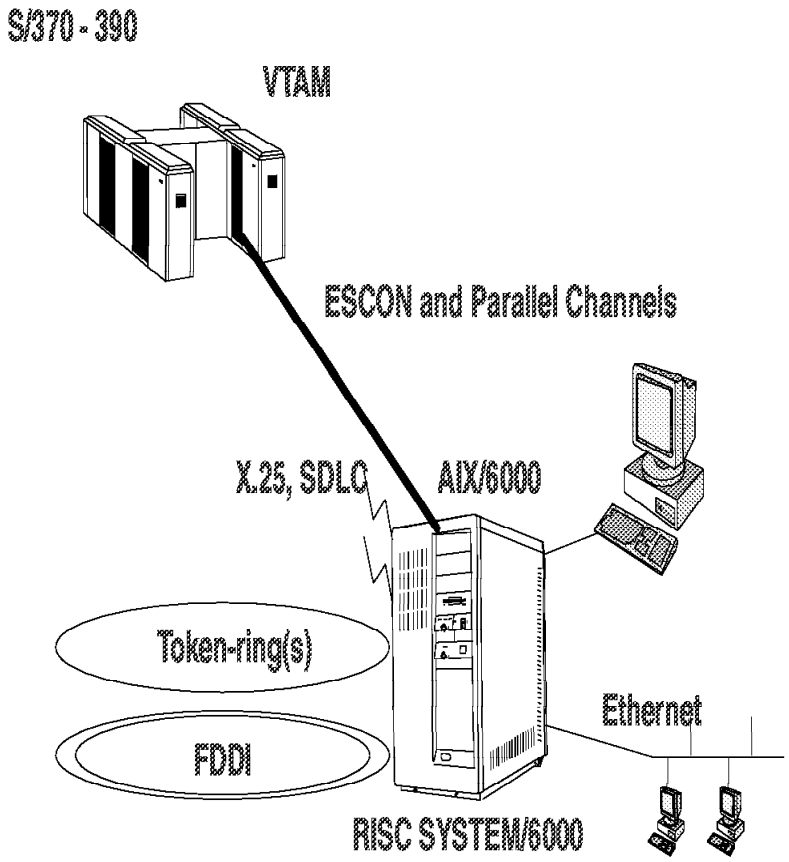

Figure 83. RS/6000 SNA Environment

3.5.3 AIX SNA Gateway/6000

AIX SNA Gateway/6000 Version 2.1.1 is a product designed to be used with AIX SNA Server/6000 Version 2.1.1 to provide network gateway capability for the IBM RISC System/6000 in an SNA network, subarea or APPN.

This program provides a means of distributing and pooling host VTAM resources to downstream PU2.0 and 2.1 nodes, saving host links and storage. The upstream (host) and downstream (workstation) link types can be of different types, effectively making Gateway/6000 a protocol converter. The functional support provided is, in most part, the same as AIX SNA Server/6000 Version 2.1.1.

Link Connectivity:

- SDLC
- Token-ring
- Standard Ethernet
- IEEE 802.3 Ethernet
- Fiber Distributed Data Interface (FDDI)
- X.25

LU Support:

- LU Type 6.2 for independent and dependent communications
- LU Types (1, 2, 3) for dependent communications
- LU 0 for Primary and Secondary communications

The AIX SNA Gateway/6000 product supports communications between multiple downstream (dependent) workstations to one or more host systems, without requiring a direct physical link to the host. The gateway acts as an intermediary between the workstation and the host. From the host perspective, it appears as a PU 2.0. It simulates all of the defined LUs. It takes the session information and passes it to the downstream workstation.

Capabilities/benefits

- Ability for the workstation to have multiple host sessions, to the same or different host, while connected to the gateway via a single link.
- Connectivity options supported by the gateway are X.25, SDLC, token-ring and Ethernet in both the WAN and LAN environments.
- Downstream workstations can be any type that supports IBM′s SNA connections.
- Ability to automatically switch from one host connection to another in the event of a host connection failure.
- Gateway support is provided for LU Types 0, 1, 2, 3 and dependent LU 6.2.
- Support for LU pooling. This is a condition where one or more of the host LUs are grouped together in the gateway in a *pool* and are provided to the dependent workstation upon request.
- Support is provided for up to 254 LUs per host link.

The AIX SNA Gateway/6000 provides the customers with a significantly increased configuration flexibility within their network and allows SNA devices, workstations for example, on a LAN or WAN to communicate with a host.

Growth Enablement: AIX SNA Gateway/6000 expands the communications capabilities of the dependent workstations in the network with minimal changes to the communications link facilities. Downstream workstations are not limited to RS/6000 units but can be an existing or planned customer device that supports SNA connections.

3.5.4 SNA Application Access

SNA Application Access for AIX Version 1.1 (SNA Application Access) brings SNA Host Communications capabilities to the AIX platform and the RISC System/6000. This allows SNA 3270 devices, located anywhere on an SNA network, to gain access to applications residing on the RISC System/6000. The type of applications can be 3270/5250 (both real and emulated), AIX ASCII-based, and Telnet 3270. SNA Application Access extends the networking capabilities for users by participating in both cross domain and local domain support.

SNA Application Access preserves customer investment in 3270 devices by providing access to the full range of existing and emerging AIX/UNIX-based applications. Additionally, since the platform of the existing or migrated applications is transparent to the 3270 user, training and operational impacts are minimized.

Link Connectivity

- SDLC
- IBM Token-Ring

LU Access SNA Application Access supports the following Logical Unit (LUs) types:

- LU Types 0, 1, 2 and 3

Session initiation requests for applications can arrive from three sources:

- Other cross domain manager
- Applications (applications initiate sessions to SNA peripheral devices)
- Direct-connected SNA peripheral devices

In each case, SNA Application Access allocates resources for the pending session and routes the request to the appropriate application.

For all control sessions (SSCP-SSCP, SSCP-PU, and SSCP-LU) and LU-LU application sessions, SNA Application Access enforces the SNA protocols.

SNA Operational Modes SNA Application Access provides two modes of operation: cross domain support and local domain support.

- Cross domain support: SNA Application Access registers as a cross domain manager with the SNA network. Other SNA hosts recognize session requests for applications residing on the SNA Application Access and route the requests accordingly. The sessions between the SNA Application Access-based applications and the SNA peripheral devices controlled by another SNA host are called cross domain sessions.
- Local domain support: SNA Application Access provides services for the directly connected SNA peripheral devices.

To the applications, local domain sessions and cross domain sessions appear the same and SNA Application Access can support both types of sessions simultaneously.

Capabilities/Benefits

- Extends user access to AIX-based applications
- Enables SNA 3270 devices to appear as VT100 terminals or as ASCII dumb terminals to AIX-based applications
- Enables SNA 3270 devices to appear as TN3270 client terminals
- Provides SNA APIs to allow existing, or user-developed, applications to interact with SNA peripheral and SNA 3270 devices
- Extends SNA networking capabilities by providing both cross domain and local domain support

Extends user access to AIX/UNIX-based applications

SNA Application Access provides a consistent platform for allowing SNA devices (3270 terminals, printers, etc.), located anywhere in the SNA network, to communicate with AIX/UNIX applications residing on the RISC System/6000. It includes the appropriate level of function, when combined with the AIX SNA

Server/6000, to provide the SNA communications capabilities of a host system and a communications controller and positions the RISC System/6000 as an SNA applications server. The SNA Applications Access functions include:

- Managing the physical connections
- Providing SNA program emulation
- Servicing Primary Logical Units (PLUs)
- Interfacing with application programs

The SNA Application Access product protects customer investment in the SNA network infrastructure and the installed 3270 devices and enables new business applications to be used by the 3270 environment. Existing host-based applications can be migrated to the RISC System/6000 platform without the knowledge of the end user, which substantially reduces or eliminates the need for any end-user training.

Enables SNA 3270 devices to appear as VT100 terminals SNA Application Access provides extensive 3270-to-VT100 data conversion capabilities that allow 3270 terminals to communicate with the non-3270 AIX-based applications developed for the asynchronous ASCII terminal environment.

The AIX Operating System provides extensive support for ASCII terminals and there is currently a huge array of AIX-based applications software written for these devices. Of these devices, the VT100 is almost universally supported. Block-mode EBCDIC terminals, such as SNA 3270 devices, have been traditionally unable to access these applications. The capabilities of SNA Application Access make the SNA 3270 device appear as an ASCII device to AIX. Both dumb terminal and VT100 emulation support is provided.

SNA Application Access enables new business applications to be used by 3270 SNA terminal users by providing access to previously unavailable applications. These existing applications can run unchanged and user training is reduced or eliminated.

Enables SNA 3270 devices to appear as TN3270 client terminals SNA Application Access converts standard SNA terminal traffic into TN3270 traffic and makes the SNA 3270 device appear as a standard TN3270 client to the AIX/UNIX platform.

SNA Application Programming Interfaces The following Application Programming Interfaces (APIs) are provided for the user:

- **PLU2**

  The PLU2 API provides the user with a platform to build applications that can communicate with SNA LU1, LU2 and LU3 peripheral devices. LU2 peripheral devices are display devices that support the 3270 data stream. LU3 peripheral devices are printers that support the 3270 data stream.

- **PLU0**

  The PLU0 API provides the users with a platform to build AIX applications that can communicate with devices such as terminals and printers on an SNA network. The PLU0 API supports low-level SNA session primitives that allows users to manage the format and flow of data between user programs and devices. An extremely flexible interface, the PLU0 API gives the user access to the entire set of SNA session control messages and allows for free-formatted request units (RUs).

- **FRM**

  The FRM API provides a platform to build forms-based applications that can communicate with 3270 devices over an SNA network.

The SNA APIs enable new business applications by providing the user with increased flexibility to access existing applications and customize new applications.

3.5.5 SNA Client Access for AIX

SNA Client Access for AIX Version 1.1 (SNA Client Access) permits users of programmable workstations (or personal computers), on TCP/IP networks, to gain access to the SNA networking capabilities on the AIX platform and the RISC System/6000. SNA Client Access expands IBM′s offering of multiprotocol interoperability software, which includes the broad line of AnyNet products.

SNA Client Access, in conjunction with AIX SNA Server/6000 V2.1.1 and above, provides the clients with all the SNA networking support that exists for the traditional SNA-attached devices.

The following support is provided:

Link Connectivity All the connectivity capabilities of the AIX SNA Server/6000 is supported.

LU Access SNA Client Access supports the following Logical Units (LU) types:

- LU Types 0, 1, 2, and 3 and independent LU 6.2

The AIX SNA Server/6000 processes the host link activation. IBM′s VTAM activates the SNA Client Access dependent LUs (0, 1, 2 and 3). These dependent LUs rely on the SSCP for activation and session initiation. During activation, the IBM host system establishes an SSCP-LU control session with each activated LU. SNA Client Access maintains this control session. Client programs can access this session through the SNA Client Access LU interface.

SNA Client Access also assists in the establishment of dependent LU-LU sessions. An IBM host system application (primary LU) requests to *bind* with an SNA Client Access client program (secondary LU). If a client program is present at the LU, SNA Client Access establishes an SNA session.

For both the SSCP-LU control session and the LU-LU application session, SNA Client Access supports the lower-level SNA protocols:

- Path Control
- Transmission Control
- Data Flow Control

For independent LU 6.2 sessions, SNA Client Access uses the LU 6.2 to access the APPC and APPN networking capabilities of AIX SNA Server/6000.

Client Support SNA Client Access provides access to the SNA network for the TCP/IP-attached client applications. Support is provided for the following:

- Telnet 3270 clients
- Telnet 5250 clients
- OEM SNA client applications

Capabilities/benefits

- Enables clients, from TCP/IP networks, access to the SNA networking capabilities of the AIX platform and the RISC System/6000
- Allows TN3270 and TN5250 clients to access SNA host systems
- Extends the capabilities of the AIX SNA Server/6000 to support SNA and TCP/IP simultaneously from the same enterprise server device
- Provides greater customer flexibility and use of existing equipment

SNA Access from TCP/IP Networks SNA Client Access, in conjunction with AIX SNA Server/6000, provides a consistent platform for client programs to access an SNA network. It functions as a TCP/IP server providing an SNA network access service to client applications running anywhere in the TCP/IP internetwork. Since the SNA Client Access processes all the lower-level SNA protocols and provides service management, the client programs and can focus on sharing information with the IBM host system′s applications. Client programs can attach to SNA Client Access to gain access to IBM host system′s applications such as JES2, TSO, POWER, IMS, CICS and NetView.

SNA Client Access protects customer investment and provides new business solutions by extending the access to the business-critical data and applications on the IBM host systems.

TN3270/5250 Devices/Emulators Access to SNA SNA Client Access provides Telnet 3270/5250 server functions to the TN3270/5250 clients. It provides the function to convert the Telnet client traffic to the SNA format for communications to the host systems and converts the SNA traffic to the Telnet format for communications to the clients.

This product enables the many AIX/UNIX and PC-based TN3270/5250 devices and emulators to access IBM host systems via SNA.

3.5.6 Summary

These RISC System/6000 solutions provide a wide range of hardware connectivity, CPU performance, and software protocol options. One of the main advantages over other gateway types is that the RISC can also be used as a local application platform in addition to a communications gateway. Products such as DB2/6000 and CICS/6000 provide excellent *right-sizing* and host offloading options. All of the RISC System/6000 communications software listed above can be run concurrently, and can even share hardware links.

3.6 2217 Nways Multiprotocol Concentrator

The 2217 Nways Multiprotocol Concentrator (2217) is an SNA-oriented, mission-critical multiprotocol solution for high-speed local area network (LAN) interconnection across wide area network (WAN) backbones. The initial release supports SNA and non-SNA traffic (TCP/IP, IPX, and NetBIOS) enabling efficient transportation between geographically dispersed LANs.

The strengths of the 2217 come from its ability to do multiprotocol LAN and WAN internetworking across a single WAN network.

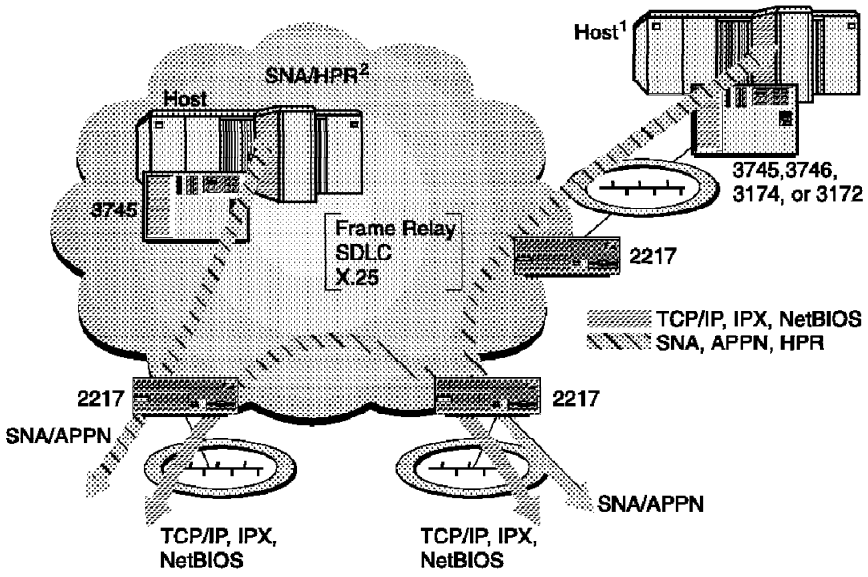

[1] Could also be an AS/400 network.

Figure 84. 2217 Environment

3.6.1 Benefits

There are many benefits in using gateway technology across a WAN:

- Mixing of LAN and WAN data

  Enables transport of LAN and WAN data over an existing SNA backbone without impeding the mission-critical data.

- Network congestion control

  Uses congestion prevention through SNA and APPN flow control.

- Link utilization

  – Enables a higher percentage of link utilization due to flow control.

  – Ensures SNA backbone integrity by preventing LAN broadcasts from traversing the network.

- Performance

  Improves overall network throughput and utilizes higher WAN bandwidth.

- Traffic prioritization

  Enables class of service specification to prioritize different application data.

- Data compression

  Provides higher data rates and improves response times at a lower cost.

- Investment protection

Enables leveraging of SNA investment in skills, network management processes, and operations.

Eliminates the need for parallel networks and extends the scope of the traditional SNA backbone to support high-speed, multiprotocol networks.

3.6.2 Release 2 Enhancements

The Release 2 microcode HPR solution is packaged in the IBM 2217 Nways MpC Model 200 and in the new IBM 2217 Models 202 and 204. The product is shipped with a pre-loaded hard disk and is available as a stand-alone solution. A software configuration package provides a quick and easy-to-use tool to configure the machine to reflect the network topology.

With the Release 2 microcode enhancements for the 2217 Model 2XX family, the IBM Nways MpC enables you to extend the scope of the SNA/APPN backbones to embrace high performance routing (HPR) while enjoying the proven reliability of SNA/HPR network solutions. The 2217 Nways MpC provides cost-effective connectivity for small work groups or remote offices requiring interoperability in any size SNA network.

Some of the new function provided in the Release 2 microcode includes:

- Additional WAC ports and no limits on protocol
- High performance routing (HPR)
- Dependent logical unit requester (DLUR) support
- Downstream SDLC multidrop support
- ISDN Basic Rate support
- Support for ROUTED and RIP over SNA
- Outbound class of service (COS) prioritization

3.6.2.1 High Performance Routing

The 2217 Nways MpC provides support for HPR, the advanced, open technology that quickly and smoothly routes data across a network.

HPR offers the following advantages that are not available with other networking solutions:

- HPR combines the best features of IP, SNA, frame relay and APPN.
- HPR improves network performance and increases network availability.
- HPR implements state-of-the-art congestion control capabilities for reduced bandwidth requirements or more efficient use of existing bandwidth.
- HPR positions customer networks for a smooth transition to ATM.
- HPR provides an excellent transport for the integration of multiprotocol traffic.

3.6.2.2 Dependent LU Requester

The 2217 Network Node (NN) with the Dependent LU Requester (DLUR) function allows customers with remote or LAN link-attached PU 2.0 terminal concentrators, printers and 3270 emulators, to gain access to host applications through the Dependent LU Server (DLUS) function in VTAM V4 R2 with no impact to the existing network.

3.6.3 Hardware Description

The base 2217 Models 202 and 204 hardware platform consists of a system board with three available slots, a 50 MHz 486 DX2 processor, 16 MB memory, a fixed disk with pre-loaded microcode, 3.5-inch diskette drive, an external modem, and a chassis with a power supply and cooling that is packaged in a stand-alone, customer setup unit.

Model 202 comes standard with 1 LAN, 1 WAC and 1 ISDN Basic Rate Interface adapters while the Model 204 comes standard with 1 LAN and 2 WAC adapters.

The 2217 Model 202 supports a ISDN Basic Rate Interface adapter while the Model 204 supports a second Wide Area Connect (WAC) adapter allowing these new models to provide a wide range of connectivity options including HPR, DLUR, frame relay, X.25, SDLC, token-ring, Ethernet and ISDN.

A high-performance 28.8 Kbps external data modem is included with the 2217 Models 200, 202, and 204 to provide configuration support, hard disk maintenance and remote service.

3.6.4 SNA Support

In this section, we discuss the SNA support provided by the 2217. SNA helps ensure standardization of network configurations and accurate transmission of data across networks. An SNA network is organized as a system of nodes and links. Each node is classified according to its capabilities and the extent of control it has over other nodes in the network. The node type is not necessarily associated with a specific type of hardware, and the node′s capabilities can be performed by different devices. For example, a workstation acting as a gateway can perform the same functions as a communications controller. The 2217 provides the following support:

- Multiple Data Link Controls (DLCs)
- LUs (independent and dependent)
- Subsystem Management
- Advanced Peer-to-Peer Networking (APPN)
- SNA Data Compression

3.6.4.1 2217 Support for Multiple DLCs

Data link control (DLC) enables orderly exchanges of data between two nodes through a logical link. The 2217 provides DLC profiles for the following types of network connections:

- Token-ring
- Ethernet
- Frame relay
- SDLC
- X.25

3.6.4.2 2217 Support for LUs

The 2217 can both initiate sessions and respond to session initiation requests. How an LU initiates and responds to requests is determined by the type of LU (independent or dependent).

Independent LU: An independent LU is able to activate an LU-LU session (that is, send a BIND request) without assistance from the SSCP; it does not have an SSCP-LU session. An independent LU is capable of sending as well as receiving BINDs. The BIND sender is referred to as the primary LU (PLU); the BIND receiver is referred to as the secondary LU (SLU).

Only an LU 6.2 can be an independent LU. Independent LUs can have parallel sessions between the same pair of LUs and can have multiple sessions between one LU and several LUs. Their session limits are established on a mode name basis, which can be from 1 to 32,767.

Dependent LU: A dependent LU is an LU that is controlled by an SNA host system. To activate an LU-LU session, a dependent LU requires assistance from an SSCP. It requires an SSCP-LU session to send a BIND. Dependent LU protocols are supported by type 5 subarea nodes using Type 2.0 protocols.

The 2217 supports LU Type 6.2 (APPC). The 2217 provides downstream support for SNA LU Types 0, 1, 2, and 3, which provide communications with host applications that support devices such as:

- LU Type 0 for 3650 and 4700 financial terminals
- LU Type 1 for 3270 printers
- LU Type 2 for 3270 interactive displays
- LU Type 3 for 3270 printers

The SNA Gateway feature supports a wide range of LU Types:

- The SNA gateway enables a S/370 host to support workstations that implement LU 0, 1, 2, 3, or dependent LU 6.2 (APPC). These LUs can be routed to one host at a time. The workstations can be LAN, SDLC, frame relay, or X.25.
- The SNA gateway also supports LU 0, 1, 2, or 3 to AS/400 host computers.

3.6.4.3 The 2217 As an SNA Gateway

The 2217 permits communication between hosts that support PU 2.0 workstations and workstations that use different DLC types. The 2217 as an SNA gateway does the following:

- Acts as a protocol converter between workstations that use DLCs on their links that differ from the DLC on the host link
- Reduces the amount of system definition at the host and workstations
- Reduces host resources and host connections through the use of pooled LUs and automatic logoff of sessions that are unused for a user-specified length of time

Each host views the SNA gateway as an SNA PU 2.0 node, supporting one or more LUs per workstation. As far as the host is concerned, all LUs belong to the SNA gateway physical unit (PU). The SNA gateway can have only one host connection and can direct different workstation sessions to that host. If this host

CP name is equal to the PU name, then this host can act as the focal point. The CP name is appended to all network management vector transports (NMVTs) routed through the gateway.

To the supported workstations, the SNA gateway looks like an SNA PU 4 communications controller and forwards such host requests as BIND and UNBIND. The network LUs are not aware of the SNA gateway. The SNA gateway, however, is aware of all LUs at the workstations. The 2217 enables workstations to be attached on any of the DLCs described in 3.6.4.1, "2217 Support for Multiple DLCs" on page 203.

The SNA gateway is a special type of PU 2.0. As long as a dependent workstation is inactive, the SNA gateway implements the LU functions for the workstation, just as a regular PU 2.0 would. As soon as a workstation is online to the host, the SNA gateway allows the workstation to implement LU functions and merely passes data between workstations and the host.

An SNA gateway enables workstation applications to access remote applications on a subarea network without requiring a separate direct connection to each host in each workstation. From the host's point of view, the host has a single connection to the gateway.

Subsystem Management: The 2217 provides the subsystem management support for sessions, connections, and adapters.

Session Flow Control: To manage the flow of data over a network, the 2217 uses adaptive session-level pacing. The pacing occurs between each pair of adjacent nodes participating in the session route. The pacing between two adjacent nodes is independent of the pacing used between other adjacent nodes in the route.

Adaptive Session-Level Pacing: Adaptive session-level pacing uses a window-based scheme, wherein a sender can send only a limited number, or window, of RUs per explicit grant of permission to proceed. The window size can be changed based on conditions at the receiver. This function permits a node to control the amount of data that is sent and received during normal session operation. The window control allows the receiving node to manage the rate at which it receives data into its session buffers. Adaptive session-level pacing provides a node supporting many sessions a dynamic means to allocate resources to a session that has a burst of activity and to reclaim unused resources from sessions that have no activity. Adaptive session-level pacing allows the receiving node to use its available buffer resources efficiently. Because each session stage between the endpoints is independently paced, both endpoint nodes and intermediate nodes can adapt the pacing for the sessions they handle in accordance with their own local congestion conditions. This action is the basis for global flow control and congestion management in APPN networks.

Adaptive BIND Pacing: BIND traffic can occur in bursts, particularly at node or network startup. Therefore, adaptive BIND pacing exists to control the flow of BINDs between two adjacent nodes. The same window algorithm used for session-level pacing is employed.

Segmenting and Reassembly: To transmit RUs longer than the maximum-size basic transmission unit allowed by a particular link, the 2217 supports data segmentation and reassembly. These segments are reassembled into whole

RUs at the partner node. This action allows the RU size defined for a session to be independent of the link that is used for the route.

3.6.4.4 2217 Support for APPN

APPN adds additional functions to the basic SNA functions supported by the 2217. APPN is a set of functions, formats, and protocols that greatly enhances SNA network management and the usability of APPC applications running in the network. These benefits are realized through reduced configuration requirements, dynamic directory searches, route calculation capabilities, and intermediate session routing.

APPN Node Types: The 2217 can be configured as the following node types and participate in an APPN network:

- Network node
- End node
- Low-entry networking (LEN) node

Each node is distinguished from other nodes in the network by a unique name consisting of two parts: a network ID and a local node name (also known as a control point name). The name identifies each node to all other nodes in the network.

A node can be configured to be an end node or a network node, but when an end node does not have CP-CP sessions to an APPN network node, it acts as an LEN node. An LEN node does not have within itself the APPN functions that network nodes and end nodes have, but it can participate in an APPN network by predefining certain information that APPN functions normally share among APPN network nodes. The node types are described in more detail in the following sections.

The *control point* (CP) is responsible for managing the node and its resources. To obtain APPN network services, the CP in an APPN end node must communicate with the CP in an adjacent network node. Also, to manage the network, the CP in an APPN network node must communicate with the CPs in adjacent network nodes. The CP directs:

- Adapter activation and deactivation
- Link activation and deactivation
- LU assistance in session initiation and termination

3.6.4.5 2217 Support for SNA Data Compression

SNA data compression at the session level increases throughput for large amounts of data across communication links, resulting in:

- Enhanced data throughput on low-speed lines
- Reduced costs on high-cost lines
- Faster response times, resulting in productivity improvements

SNA data compression:

- Is compatible with the S/370 and AS/400 implementations
- Can be used with all LU Types
- Requires VTAM V3R4.1 or higher

3.6.5 LAN Functionality

This section briefly describes the LAN functions that the Nways Multiprotocol Concentrator (2217) supports. The 2217 can transport the following LAN protocols over an SNA network:

- NetBIOS
- IPX
- IP

3.6.5.1 NetBIOS over SNA

NetBIOS is a protocol that uses character-based names for endpoints. It finds machines by name, then connects sessions using Logical Layer Control (LLC) flows that provide connection-oriented links. When connecting sessions over wide area networks (these typically do not provide the efficiency and bandwidth of local area networks), special handling of query and datagram frames is needed to provide session connections. The 2217 uses a list of local names that are exchanged with partner 2217s to route NetBIOS.

3.6.5.2 IPX Traffic over SNA

At the LLC layer, the IPX protocol is a connectionless protocol. IPX relies on higher layer protocols to guarantee delivery. In IPX, servers advertise their services in order for requestors to find server addresses. Routers also help by passing the server′s advertisements around the network. These functions in NetWare IPX are called Service Advertising Protocol (SAP) and Routing Information Protocol (RIP). RIP dynamically keeps track of routes and maintains routing table entries. When a requestor wants to access a server, it sends out frames asking where the nearest server is and then how to get to it. Each router on the network must be able to service SAP and RIP information from other network regions to help local stations reach the proper servers.

3.6.5.3 IP Traffic over SNA

The 2217 Sockets over SNA function enables two socket applications to communicate with each other over an SNA network. If the socket applications are in separate TCP/IP networks connected by an SNA network, two 2217s are needed to route the traffic over the SNA network. The TCP/IP data is routed over SNA using IBM′s multiprotocol transport networking (MPTN) formats.

3.6.5.4 Source-Route Bridging under Frame Relay

The 2217 can support source-route bridging over frame relay. When configured for source-route bridging, the 2217 routes all the protocols it can and bridges those it cannot. For example, DECnet and AppleTalk would be bridged because the 2217 cannot route them. The 2217 uses the protocols′ dedicated routing services, expanding their routing areas to include an entire WAN.

3.6.5.5 LAN Network Manager Agent

LAN Network Manager (LNM) support in the 2217 enables:

- Detection of hard and soft errors on attached token-ring networks
- Detection of configuration changes, such as the addition or subtraction of any locally attached token-ring stations
- Reporting of the above errors and changes, as well as other performance statistics, up to four remotely linked LAN network managers
- Setting of local configuration parameters by a remote LAN network manager

- Forced removal of locally attached token-ring adapters
- Reporting the functional addresses that are enabled for token-ring stations on bridged LAN segments

Note: LNM is not supported over an SNA backbone.

3.6.5.6 Remote Configuration and Operation of the 2217

The 2217 Remote Control Utility (2217 RCU) is an OS/2 APPC application that allows you to remotely configure and operate the 2217. This is the only way to configure a 2217.

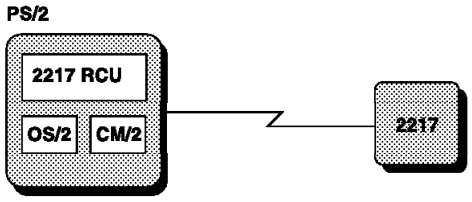

Figure 85. IBM 2217 Remote Control Utility

The 2217 RCU provides an interface that enables you to configure the SNA and multiprotocol functions supported by the 2217. The 2217 RCU configuration panels lead you through specific tasks (for example, configuring the IPX over SNA function) and provides online help.

In addition to configuration panels, the 2217 RCU also provides an interface for you to operate the 2217 and to perform remote diagnostic activities. For example, use the 2217 RCU to activate or deactivate an SNA link on the remote 2217. Also, with one click of the mouse, you can gather all the problem determination files you need to diagnose 2217 errors.

3.6.5.7 Hardware Requirements for the 2217 RCU

The 2217 RCU requires:

- 2 MB hard drive space
- Intel 386 or higher (486 is recommended)
- 12 MB of RAM (recommended)
- Asynchronous modem or LAN connection

3.6.5.8 Software Requirements for the 2217 RCU

The 2217 RCU is shipped on 3.5-inch diskettes. The RCU must be installed on a PC running:

- OS/2 2.1 or higher
- Communications Manager/2 V1.11 or higher

To configure a new 2217, you would use the 2217 RCU to configure connections to the 2217 over a telephone line or over the LAN (token-ring or Ethernet). The

2217 RCU is an APPC application that requires Communications Manager/2 (CM/2) to set up the SNA connectivity.

After the initial connectivity is established between the 2217 RCU and the 2217, you can configure a link from the 2217 RCU to the 2217 over the SNA wide area network for future connections. Connections over the SNA wide area network may provide better performance than connections over a telephone line.

In addition to the functions provided by the 2217 Remote Control Utility, you can monitor and maintain the 2217 using the IBM NetView family of products or other network management software.

3.6.5.9 LAN Adapters

The two LAN adapters supported in the 2217 are:

- Token-ring 16/4 ISA Adapter
- Ethernet ISA Adapter

Only one of these adapters can be installed in the 2217 at a time.

3.6.5.10 WAC Adapter

The 2217 has one WAC ISA adapter, which is a dual-port, high-speed communications adapter. The ports and interfaces supported by the adapter vary according to the feature codes that were ordered for the 2217, and can be the following:

- X.21
- V.35
- EIA 232
- EIA 422/429

The two ports may be used concurrently for any of the three WAN protocols, (SDLC, X.25, and frame relay) with the following limitations:

- Either port 0 or port 1 may be configured for SDLC operation, but not both. The second interface can be configured to support either in the same manner.
- Either port 0 or port 1 may be configured for frame relay operation, but not both.

Both ports may be configured for X.25 operation.

3.6.5.11 Modem

The 2217′s internal modem eliminates the need for an external modem. The modem communicates with other modems at speeds from 300 bps to 14,400 bps. It is Hayes compatible, and can auto-sense and renegotiate speeds to work with most modems. The internal modem has the auto-answer feature that enables communication initialization by incoming calls.

In countries other than the U.S. and Canada, the customer must provide an external 14.4 Kbps modem. This modem attaches to the serial port on the rear of the machine. The modem can be used for RCU and may be used for remote service.

3.7 PC Gateways

A PC gateway is a LAN-attached PC running specialized software that provides gateway services for other devices attached to the LAN to access the host. Its two main connections are:

- Upstream - for communications with the host network
- Downstream - for communications with LAN-attached devices

Two PC communications packages are discussed, outlining their downstream and upstream communications protocols:

- IBM Personal Communications/3270 (PC/3270)
- IBM Communications Server for OS/2 Warp

Each of these products provides basically the same generic SNA gateway functions, except that each has different capabilities in terms of the supported protocols, links, number of LAN workstations and operating systems used.

There are two ways to configure the relationship between the host and the gateway:

Single PU View

- To the host, the PC gateway appears as a single PU Type 2 control unit, with a group of LUs. Therefore, the host is not aware of all the DSPUs (downstream PUs), and thinks that all the downstream LUs belong to the gateway PU.
- To the LAN-attached workstations working as DSPUs, the gateway appears as a PU Type 4.0 communications controller, connected to a system services control point (SSCP) and associated primary partner LUs in the host. Therefore, they have the illusion of being attached directly to the host, and require no special configuration to communicate with the SNA PC gateway.
- The gateway is aware of all the downstream PUs, and the LAN-attached workstations appear to it as PU Type 2, with one or multiple LUs defined. The gateway has to associate the downstream LUs with its own LUs, and executes controls over the DSPUs.

This relationship is illustrated in Figure 86 on page 211.

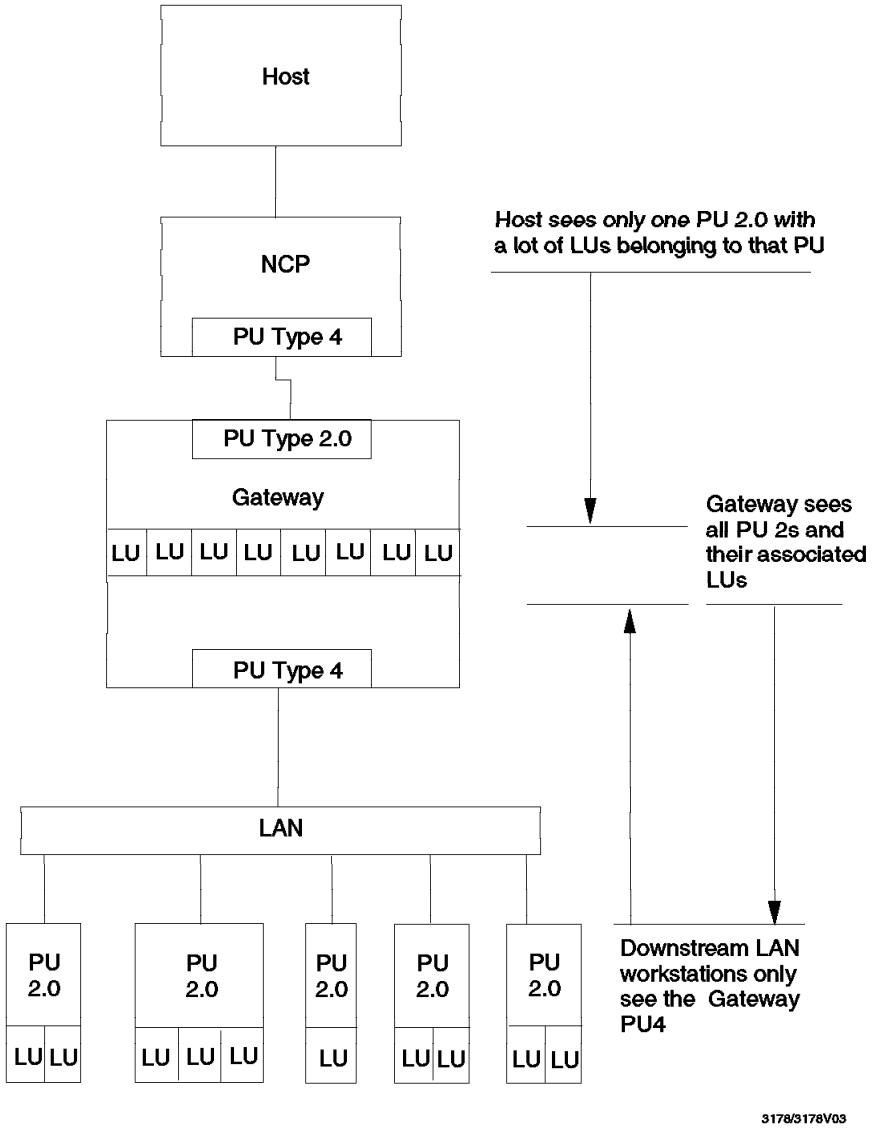

Figure 86. PC Gateway Components View

Downstream PUs View

- To the host, the PC gateway appears as a PU Type 2 with its LUs, and also supports a group of DSPUs. The host is therefore aware of all the downstream PUs, and knows the way to reach them through the gateway PU.
- To the LAN-attached workstations working as DSPUs, the gateway appears as a PU Type 4.0 communications controller, connected to a system services control point (SSCP) and associated primary partner LU in the host. They therefore have the illusion of being attached directly to the host, and require no special configuration to communicate with the SNA PC gateway.
- The gateway is aware of all the downstream PUs, and the LAN-attached workstations appear to it as a PU Type 2, with one or multiple LUs defined.

The gateway has to pass through the SNA command flow and data between the DSPUs and the host.

This relationship is illustrated in Figure 87.

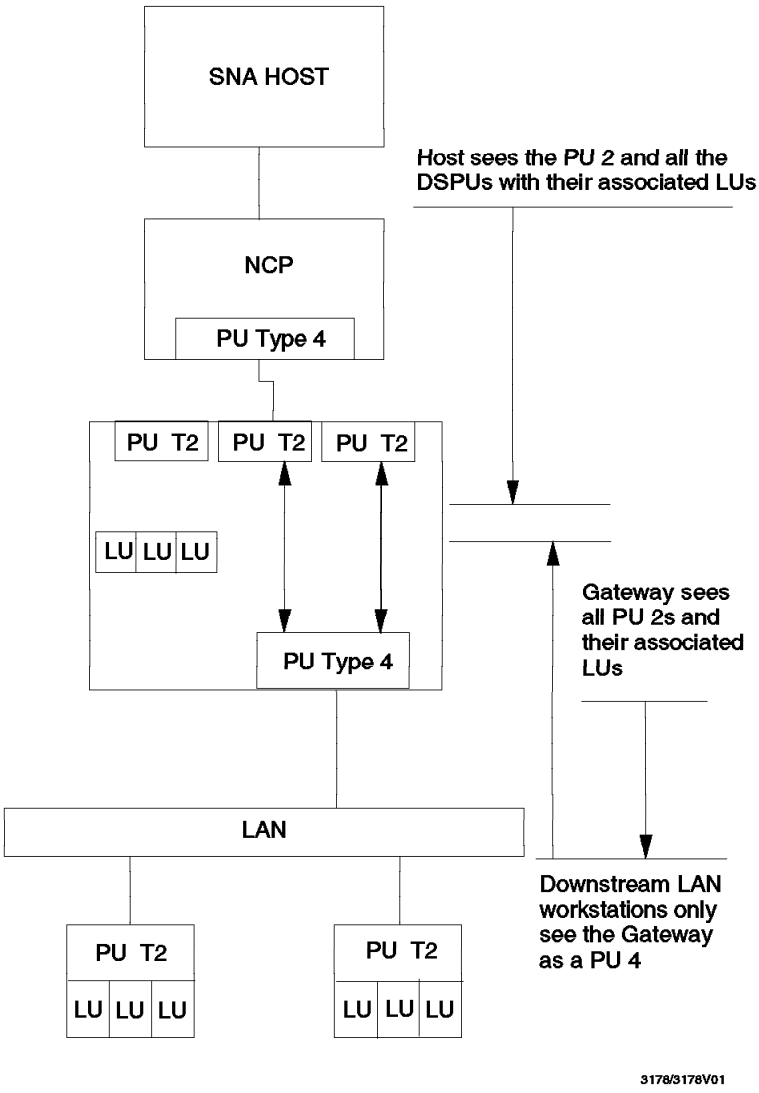

Figure 87. PC Gateway Components View.

When connected via SNA DFT, the gateway appears as a DFT device with up to five sessions associated with it. In this configuration, the IBM 3x74 control unit manages the SSCP-PU session with the host and maps the SSCP-LU session control commands (ACTLU and DACTLU) to the control unit status commands in the adapter buffer. The gateway maintains the SSCP-LU sessions with the host when the LU is not allocated to, and active on, a workstation.

3.8 PC/3270 Gateway

PC/3270 V4.1 for *Full-Function DOS* provides gateway services for workstations running and IBM Disk Operating System (DOS) permits communication between host systems that support PU 2 workstations and workstations that use the same or different DLC types.

The host defines a PU 2 with multiple dependent LUs: LU 1, 2, 3 and 6.2 (for example, an LU 2 for 3270 terminal emulation). Traditionally, there is a communications link connected to a machine, such as a 3174, implementing the PU and sharing the LU function with directly attached devices. The gateway acts as an intermediary between the workstation and the host. From the host's perspective, it still appears as a PU 2. It simulates the presence of all of the defined LUs. It takes the session information, however, and redispatches it to *downstream* workstations. To these downstream machines, the gateway appears as a PU 4, such as a 3745 with an NCP.

PC/3270 as an SNA gateway can do the following:

- Act as a protocol converter between workstations that use DLCs on their links that are different from the DLC that is used on the host link
- Reduce the amount of system definitions at the host and workstations
- Allow you to dynamically change network definitions and add workstations
- Reduce host resources by using pooled LUs, and by automatically logging off unused sessions

The PC/3270 V4.1 for Full-Function DOS gateway supports communication to an IBM System/370, System/390 host using the following protocols:

- Token-ring
- Ethernet
- SDLC
- X.25
- SNA DFT

Up to a maximum of 253 sessions, host display and/or print sessions, are supported when the gateway is connected via IEEE 802.2 (token-ring or Ethernet), SDLC or X.25.

Only five sessions are supported through a gateway attached via SNA DFT and only one host attachment can be made at a time.

The gateway service is provided within PC/3270 V4.1 for Full-Function DOS workstation equipped with a token-ring adapter or an Ethernet adapter and requires the IBM LAN Support Program Version 1.3 or later. Up to two adapters, any combination, are supported simultaneously for downstream workstations.

The gateway provides session management through the use of logical units (LUs). These LUs correspond to the host sessions that the gateway distributes to the workstations on the LAN. The gateway allows workstation LUs to be categorized as one of the following three types: local, pre-allocated or pooled (20 pools).

- Local LUs are assigned to the gateway.

- Pre-allocated LUs are assigned to a specific workstation.
- Pooled LUs allow multiple workstations to share a pool of assigned sessions, thus requiring fewer host sessions to support many casual users with similar or identical host requirements.

When configuring the gateway, the LUs (sessions) may be pre-allocated to individual workstations, or the LUs may be pooled, or a combination of pre-allocated and pooled LUs may be used.

PC/3270 V4.1 may be configured as a dedicated gateway with no local LU sessions or as a gateway with local LU sessions. The local LU sessions are allocated to the gateway in the same manner as LU sessions are allocated for the attached workstations. Each LU allocated to the gateway for a local session reduces the number of LUs available for the attached workstations.

Support is provided for workstations using the 3270 terminal emulation functions, as follows:

- Personal Communications/3270 (all versions)
- Communications Manager/2 V1.0 and later
- Communications Manager of OS/2 Extended Edition V1.3 or Extended Services V1.0
- PC 3270 Emulation Program Version 3.05
- 3270 Workstation Program Version 1.12
- IBM Networking Services/DOS or APPC/PC
- IBM Networking Services/Windows V1.0
- OEM products can connect as long as they are LU 1, 2, 3 and dependent 6.2 types

3.8.1 Gateway Status Utility

This utility provides a link log table to assist the gateway administrator. It displays the link status with the corresponding time of receipt in order to assist in problem determination. Additionally, any new link messages received while this function is active will be displayed in the message area of the screen. The gateway status utility also provides information specific to each session allocated to the gateway. Information is formatted to contain the information applicable to each workstation, followed by the next PUID and so on. Information consists of the PUID, the PU state and the network adapter address, followed by the workstation LU address, the gateway LU address, and current state of each LU entry found in the LU session table.

This utility also provides for the termination of the gateway function and for the return to DOS of all memory previously allocated for the gateway.

The gateway administrator may select between an immediate termination or a *graceful* termination. An immediate termination employs a wait value set during gateway customization and a graceful termination waits for all active sessions to become unbound before termination. With either selection, a warning is shown to the gateway administrator to safeguard against accidental gateway termination.

3.9 IBM Communications Server for OS/2 Warp, Version 4.0 (CommServer)

The CommServer provides PC-to-host and PC-to-PC multiprotocol communications services. OS/2, Windows, and DOS-based PCs can communicate through the CommServer with S/390 and AS/400 hosts and other workstations. The following are some functions provided for the CommServer:

- SNA Gateway support providing full-function, cost-effective support for multiple LAN-attached OS/2, Windows, or DOS workstations, to access System/390 and AS/400 hosts, through one or more physical connections to one or more hosts.
- Remote access to SNA applications over asynchronous, synchronous, Hayes Autosync, digital and cellular connections.
- APPN network node, end node, APPC support and DLUR (Dependent LU Requester) for distributed applications.
- APPN High Performance Routing (HPR) support to improve the availability and throughput of your network. HPR gives you improved network availability through transparent recovery from network failures, improved network performance through enhanced error recovery mechanisms, and state-of-the-art congestion control capabilities for reduced bandwidth requirements. HPR joins all the benefits of SNA and connection-oriented networking, and the connection availability of other connectionless networks.
- Multiprotocol Support
  - Sockets over SNA access node and gateway support, which allow sockets applications to communicate over SNA Networks, or over connected SNA and TCP/IP networks, without changing the applications
  - SNA over TCP/IP access node and gateway support, which allow APPC applications to communicate over TCP/IP networks, or over connected SNA and TCP/IP networks, without changing the applications
- Downstream Protocols
  - AnyNet (SNA over TCP/IP)
  - ATM (LAN Emulation)
  - Ethernet
  - Token-ring
  - SDLC
  - X.25
  - IDLC (synchronous, asynchronous and autosync)
- Upstream Protocols
  - AnyNet (SNA over TCP/IP)
  - ATM (LAN Emulation)
  - Ethernet
  - Token-ring
  - SDLC
  - X.25
  - IDLC (synchronous, asynchronous and autosync)
  - Twinax (for 3270 passthrough)
- Support for a wide range of adapters and modems

3.9.1 SNA Gateway

The gateway function, at the CommServer, is optimized to provide cost-effective host connectivity by sharing communications resources such as adapters and physical connections.

Highlights:

- Host backup links may be defined for automatic reconfiguration should the primary links fail.
- SNA protocols LU 0, 1, 2, 3, and dependent LU 6.2 (APPC). The gateway also supports LU 0, 1, 2, or 3 to an AS/400 host using SNA passthrough.
- Multiple Physical Unit Type 2.0 (PU 2.0) and up to 254 logical units (LUs) per PU. To increase efficiency, LUs may be dedicated to a workstation or pooled among workstations.
- Up to 254 concurrently active workstations per LAN adapter to transparently access one or more System/390 hosts. Implicit workstation definitions simplify gateway configurations and management since downstream workstations are not required to be configured on the CommServer.
- Downstream applications and workstations using standard SNA connectivity protocols for LU 0, 1, 2, 3 and dependent 6.2, and communicating through an SNA Gateway to a host. The downstream workstations include Novell NetWare for SAA and APPLE MAC SNAps Gateway.
- Dedicated and pooled LU definitions. The LUs defined in the gateway can be dedicated to a particular workstation or pooled among multiple workstations. Pooling allows workstations to share common LUs, which increases the efficiency of the LUs and reduces the configuration and startup requirements at the host. You can also define multiple LU pools, each pool associated with a specific application. When a link is defined through the gateway between a workstation and the host, the LU is activated when the session is established and returned to the pool for access by other workstations when the session is ended.
- Session inactivity management. You can also configure an SNA gateway to automatically stop a session after a specified period of inactivity if there are other workstations waiting. This procedure helps increase the number of shared logical units that are available.
- The transmission of network management vector transports (NMVTs) between the gateway and the host. For example, commands coming from the NetView program in the host are received in the gateway and may be passed to and used by another application on the gateway, such as IBM LAN Network Manager.

  Each host views the SNA gateway as an SNA PU 2.0 node, supporting one or more LUs per workstation. As far as the host is concerned, all LUs belong to the SNA gateway PU. The SNA gateway can have multiple host connections simultaneously and can direct different workstation sessions to specified hosts. However, only one host (and it must be on a link with a CP PU) can act as the focal point, and the CP name is appended to all NMVTs routed through the gateway.
- SDLC now supports at least 16 upstream and downstream physical connections, and two-way simultaneous full-duplex communications.
- The WAN line speed has been increased to a minimum of 2 Mbps (T1/E1).

- CommServer can now act as the multipoint primary control for downstream workstations. This support is for workstations and devices connected to the server over SDLC, on a multidrop line over synchronous adapters such as the IBM WAC, MPA, and ARTIC.
- Dependent LU support across a TCP/IP network with SNA subarea and APPN support upstream from the configured gateway.
- Dependent LU support across an SNA network with TCP/IP support upstream from the configured gateway.

3.9.2 Advanced Peer-to-Peer Networking (APPN)

CommServer APPN support provides SNA networking facilities to connect distributed computing applications, peer applications and client applications to their servers. These facilities include directories of partners, route selection, and management services.

APPN provides support to:

- Significantly improve the performance of communications between APPC and Common Programming Interface for Communications (CPI-C) 2.0 applications, especially in a LAN environment.
- Use the CPI-C interface for enabling greater application portability across different platforms.
- Add, delete, or move nodes within the network with limited system definition at the affected node, and no other definition at other nodes.
- Use additional debugging tools and sample programs for easily developing and using APPC and CPI-C applications.
- Use APIs to automate configuration changes and add network management capabilities and use defaults for reducing required system definition.

The CommServer APPN function has the following enhancements:

- *High Performance Routing (HPR)*

  CommServer provides support for HPR, the advanced, open technology that quickly and smoothly routes data across a network. HPR offers the following many advantages that are not available with other networking solutions:

  – HPR combines the best features of IP, SNA, frame relay and APPN.
  – HPR increases network availability by non-disruptively rerouting around network failures.
  – HPR implements state-of-the-art congestion control capabilities for reduced bandwidth requirements or more efficient use of existing bandwidth.
  – HPR positions customer networks for a smooth transition to ATM because HPR was specifically designed for high-speed environments.
  – HPR provides an excellent transport for the integration of multiprotocol traffic.
  – HPR improves network performance.

  HPR is an end-to-end, connection-oriented protocol, which can be visualized as a pipe between two nodes. HPR provides a set of functions which optimize communications across the pipe. Together these features:

  – Improve network availability because connections are end-to-end, and HPR calculates alternate routes and resumes transmissions if an intermediate link or node fails.

  - Improve network performance by only performing error recovery at the endpoints of the network and by selectively retransmitting only lost or erroneous packets.
  - Improve SNA's congestion control by constantly monitoring and adjusting the amount of data flowing between the endpoints ensuring they don't overload.

  CommServer support of HPR provides the following capabilities:

  - CommServer, defined as an APPN node (end node or network node), can function as an endpoint for an HPR session (providing RTP).
  - CommServer defined, as a network node, can also function as an HPR intermediate node.
  - CommServer supports HPR routing over token-ring, Ethernet, FDDI, LAN/ISDN, and frame relay connections.

- *Dependent LU requester (DLUR)*

  DLUR, in conjunction with VTAM's dependent LU server (DLUS) function, enables dependent LUs to operate unchanged in an APPN network, without changing applications. The DLUS establishes an LU 6.2 session with, and provides SSCP services to, the DLUR node. Any number of dedicated PUs can be defined on the LU 6.2 sessions. The CommServer gateway can then provide network management access through the dedicated PU to downstream workstations.

  DLUR allows customers to protect their investment in 3270 emulation, and other dependent LU, applications while migrating new applications to LU 6.2 and APPN.

- *Self-Defining Dependent LU (SDDLU)*

  CommServer can now dynamically identify its dependent LUs to VTAM and eliminate the requirement for a static definition of dependent LUs in VTAM (therefore, no VTAM SYSGEN is required).

- *SNA transmission priority*

  CommServer provides the capability to prioritize SNA traffic for LU 0, 1, 2, 3, and 6.2 sessions over LAN and WAN connections.

- *System Services Control Point (SSCP) takeover*

  CommServer provides the capability to keep sessions active and allow CP name changes when a VTAM/NCP composite node takes over SSCP sessions currently owned by another VTAM/NCP node.

- *CP-CP session reactivation*

  Automatic activation of CP-CP sessions or reactivation of CP-CP sessions that have failed prevents network nodes from being detached from the network whenever at least one link is active.

- *APPC full duplex*

  Advanced Program-to-Program Communications (APPC) now supports full duplex logical conversations and non-blocking verbs for LU 6.2 applications. This enhancement simplifies APPC applications logic and greatly improves performance.

3.9.3 Multiprotocol Support

CommServer and the OS/2 Access feature have incorporated the versatile SNA over TCP/IP and sockets over SNA function from the AnyNet product family. This support enables you to extend and simplify your network by allowing SNA applications to communicate across a TCP/IP network, and Sockets applications to communicate across an SNA network, without changes to the applications. CommServer provides the following multiprotocol functions:

- **Sockets over SNA**

  With Sockets over SNA support, sockets applications can communicate over SNA networks, or over connected SNA and TCP/IP networks, without changes to the applications. In addition, Sockets over SNA users on TCP/IP can benefit from advantages built into SNA networks such as cost-effective bandwidth utilization, and predictable response times.

  Sockets over SNA support can be used in a variety of ways:

  – Sockets applications can run on an SNA network.

    A user can use sockets applications, such as FTP, Telnet, and NFS, across an SNA network, without having TCP/IP connectivity on the network.

  – Sockets applications can run across connected TCP/IP and SNA networks.

    SNMP agents on a TCP/IP network can send management data to an SNMP manager, like NetView/6000, on an SNA network.

  – Sockets applications can run between TCP/IP networks connected by SNA networks.

    For example, FTP or Telnet applications can communicate between two workstations on separate TCP/IP networks that are joined by an SNA network.

- **SNA over TCP/IP**

  With SNA over TCP/IP support, APPC and dependent LU applications (such as printer or emulator applications), can communicate over TCP/IP networks, or over connected SNA and TCP/IP networks.

  You can use SNA over TCP/IP support in the following ways:

  – APPC applications can run on a TCP/IP network.

    APPC applications, such as DB2 and CICS, can run on TCP/IP, with no changes to the applications (and with no need for SNA to be on the network).

  – APPC applications can run between connected SNA and TCP/IP networks.

    APPC applications, such as DB2, on an SNA network can communicate with DB2/2 applications on a connected TCP/IP network.

  – APPC applications can run between SNA networks connected by TCP/IP networks.

    DB2, CICS and IMS applications can be used between two workstations running on two separate SNA networks, which are joined by a TCP/IP network.

3.9.4 Emulator Support

The terminal emulator previously provided in Communications Manager/2 has been removed. Terminal emulation for administrative use at the server is provided by IBM Personal Communications/3270 for OS/2 Entry Level (PC/3270 Entry).

The PC/3270 Entry emulator supports up to two 3270 display sessions via LUA with a System/390. The functions provided by this emulator are a subset of the functions provided by the IBM Personal Communications/3270 for OS/2 emulator.

Chapter 4. Remote LAN Access

In this chapter, we cover remote LAN access environments and technologies. We also discuss several of the remote LAN access solutions available from IBM.

4.1 Remote LAN Access Environments

Remote LAN access generally refers to accessing a computer device using an external line, which most commonly is a switched telephone line. So, you can dial-in to the LAN or dial-out of the LAN over a Wide Area Network (WAN). There are basically four main environments in remote LAN access:

- The remote-to-remote environment
- The remote-to-LAN environment
- The LAN-to-remote environment
- The LAN-to-LAN environment

4.1.1 Remote-to-Remote

A remote-to-remote environment consists of a direct physical connection established between two or more remote workstations.

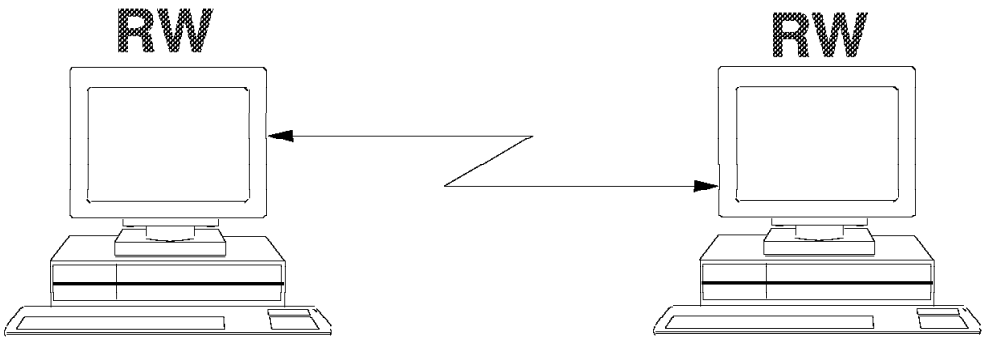

Figure 88. Remote Workstation Dial-In to Remote Workstation

Conferences may be set up between multiple workstations creating an ad hoc LAN over telephone lines. Without LAN adapters and without LAN wiring, remote-to-remote workstations can access each other's LAN resources and LAN-based applications. This environment supports users who need a simple and low-cost WAN connection to support data, resource and program sharing. Another example of a remote-to-remote implementation would be a remote user using the telephone line to run applications on a directly connected LAN server.

4.1.2 Remote-to-LAN

A remote-to-LAN environment, sometimes called dial-in, occurs when a remote workstation initiates a connection to a LAN workstation via some form of WAN/LAN communication server.

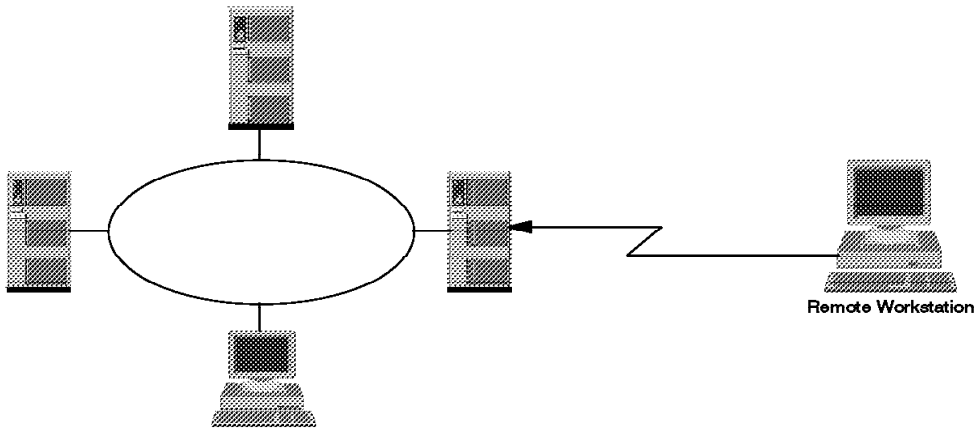

Figure 89. Remote Workstation Dial-In to LAN

Once the WAN connection is established between the remote workstation and the LAN, the remote workstation can directly address any LAN-attached workstation configured to participate within the remote-to-LAN environment. Likewise, because the remote workstation has its own unique address, it can receive information directly from the participating LAN-attached workstations.

4.1.3 LAN-to-Remote

A LAN-to-remote environment, sometimes called dial-out, occurs when a LAN-attached workstation initiates a connection to a remote workstation via a WAN/LAN communication server.

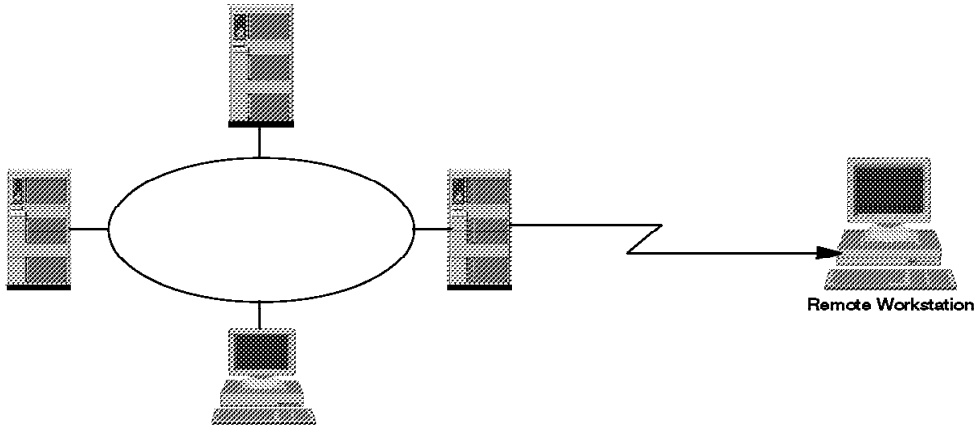

Figure 90. LAN Dial-Out to Remote Workstation

It has the same characteristics and capabilities as the remote-to-LAN environment except that the LAN-attached workstation initiates the connection. An example of LAN-to-remote would be a LAN-attached workstation accessing a remote information server to acquire product pricing data.

4.1.4 LAN-to-LAN

A LAN-to-LAN environment occurs when a LAN-attached workstation connects to another LAN-attached workstation via two WAN/LAN communication servers.

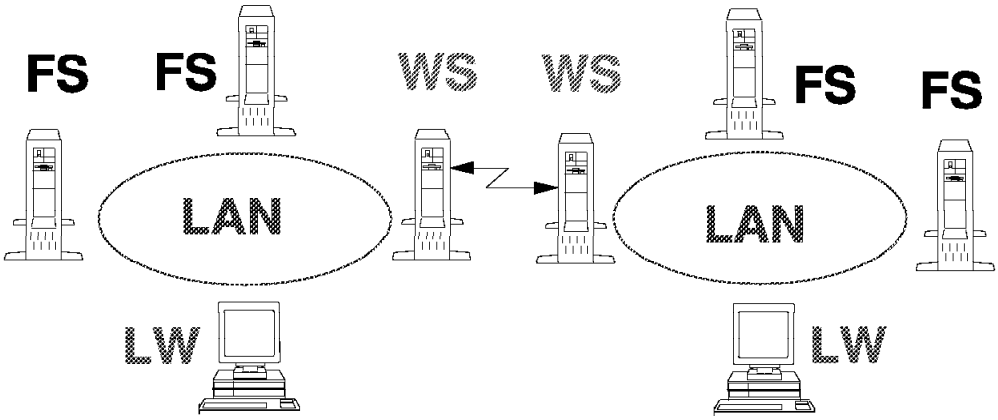

Figure 91. LAN Dial-Out to LAN

It normally combines the functions of the LAN-to-remote and remote-to-LAN environments. The resulting casual bridge allows the customer to utilize switched links rather than leased lines for a more mobile and cost effective solution. The LAN-to-LAN environment provides the capability for LAN-attached machines to access or update information residing in remote locations, and also, to act as a server for other remote workstations connecting to the LAN. Normally, the connections are established on a temporary workstation-to-workstation basis across the WAN.

Note

This LAN-to-LAN environment is different from a split bridge environment. A split bridge establishes a permanent connection between all machines on the two LANs.

The LAN-to-LAN environment is particularly useful for customers with numerous separate LAN networks who have a need to control access on and off the LANs. An example would be banking companies with their many branch offices. It provides an inexpensive mechanism for dynamically connecting the LANs while maintaining control over the origin of traffic flowing between them.

4.2 Remote LAN Access Technologies

There are numerous remote LAN access products available today, which vary widely in cost and functionality. Some use standard hardware devices and are solely software driven, while others may involve special hardware devices.

Products that involve special hardware devices may replace the LAN adapter with a customized WAN adapter in the remote workstation and provide a compatible hardware *tap* on the LAN. This LAN hardware tap varies from a specialized adapter on the LAN file server to a stand-alone multiprocessor box. The implementation of this approach varies widely in sophistication, cost, and performance.

Some products utilize extensions of a remote-to-remote environment to provide remote-to-remote and remote-to-LAN access capabilities, but do not support the LAN-to-remote or LAN-to-LAN environments. Many of the products currently available also do not support graphical interfaces.

Most of the remote LAN access products use one of three known technological approaches:

- The remote control approach
- The remote client approach
- The remote node approach

Each approach provides an inherent level of functionality and limitations.

4.2.1 Remote Control Approach

One of the earliest and most pervasive software approaches is remote control. The remote workstation using this approach dials-in to, and takes control over, a LAN-attached workstation, which executes programs on behalf of the remote workstation over the LAN. Keyboard and screen data from the dedicated LAN-attached system is then routed back to the remote workstation. By routing only keyboard and screen data, this approach minimizes the amount of data that flows across the link, but it requires a dedicated machine on the LAN for each remote workstation dialing-in to the LAN.

Most remote control products transmit keyboard and screen data over the WAN in character mode, although some companies are planning to provide transmission of graphical screen data in the near future. Transmitting graphics images will of course be slower than transmitting characters. However, graphics mode transmission is necessary to support the use of graphics or graphical interfaces, which are gaining significant importance in end-user computing, across the remote link. Lack of graphics support has been a major factor in the loss of popularity for this approach.

One other disadvantage with this approach is security. In addition to the requirement for the LAN-attached workstation to be powered on for remote use, screen data transmitted across the link contains a high percentage of fixed information in a fixed format. Data encrypted in this form is relatively easy to break because the intruder can see the effects of encryption on the fixed information that is transmitted more easily.

Examples of Remote Control Products:

- PC Anywhere

- Carbon Copy
- DCAF
- NetWare Access Server

4.2.2 Remote Client Approach

Gaining popularity today in the remote LAN access market, the remote client approach utilizes a simple mechanism to extend the remote-to-remote environment to service the remote workstation and allow it to share data and applications located on a common WAN/LAN server. This may be accomplished by replacing the LAN device drivers in the remote workstation and LAN-attached server with customized device drivers that will allow them to send and receive LAN frames across a WAN link. This provides LAN application transparency within the remote workstation.

The new device drivers utilize existing protocols to allow remote workstations to connect with each other to form a virtual LAN via the WAN link. In addition, the device drivers provide a mechanism for remote workstations to disconnect from one another upon conclusion of the remote transaction. Since the entire LAN frame is transported between the remote machines over the WAN link, LAN applications running in the remote workstations can support graphical interfaces in the same way as those running on LAN-attached workstations. Also, the LAN frames have much less fixed format information, thus providing a more secure link encryption.

Extending the remote client approach to access information elsewhere on the LAN from a remote workstation requires a LAN-attached server to manage transaction data on the workstation′s behalf. The remote environment is analogous to a standard LAN client/server environment. The remote workstation has addressability only to the network server to which it is connected. Files and programs residing on the common network server can be shared throughout the virtual LAN.

The remote client approach supports small single-server networks, but does not scale well to support large or distributed environments. Bottlenecks in both memory and CPU capacity tend to form in the common network and file server. Thus, most products using this approach are dedicated servers supporting a limited number of remote connections (generally, 1 to 16).

Organizations requiring more connections or greater capacity than can be accommodated by a single network server face potentially complex challenges in duplicating and maintaining data on multiple communication servers. Accessing data and applications that are distributed across multiple servers can be tedious for a remote user in a remote client environment. For instance, a remote user would have to physically disconnect from one server and reconnect to a second server in order to access its resources even though the two servers may be attached to the same LAN.

Examples of Remote Client Products:

- Lotus Notes
- cc:Mail
- Microsoft Windows NT

4.2.3 Remote Node Approach

The remote node approach replaces the device driver within a LAN-attached communication server. The device driver enables the server to take incoming data off a WAN and put it onto the LAN, and also, to take outgoing data off the LAN and put it onto the WAN. In addition to providing the transparency and remote LAN access capabilities of the remote client approach, remote node provides full addressability, allowing the remote workstation to access distributed LAN-attached servers and peer services.

This means that a remote workstation can access information and services wherever they reside on the LAN, rather than the LAN having to be redesigned with a central dedicated server to accommodate access by the remote workstation. It also means that growth in the number of local and remote LAN users can be easily accommodated without duplicating and maintaining data files across numerous servers.

Examples of Remote Node Products:

- LAN Distance (software only)
- 8235 (hardware and software)

4.3 IBM LAN Distance

The IBM LAN Distance family of products provides the capability of extending the ″Office LAN″ to remote users. Using dial connections such as asynchronous communications (async), synchronous communications (Sync), or Integrated-Services Digital Network (ISDN), remote users can access LAN (token-ring or Ethernet) resources, as if they were physically attached to the LAN. In essence, IBM LAN Distance provides users with functionality and features to run LAN applications anywhere, anytime, and provides the systems administrator with effective tools for managing the wide area network (WAN). Figure 92 on page 227 shows the different ways LAN Distance can be connected to the network.

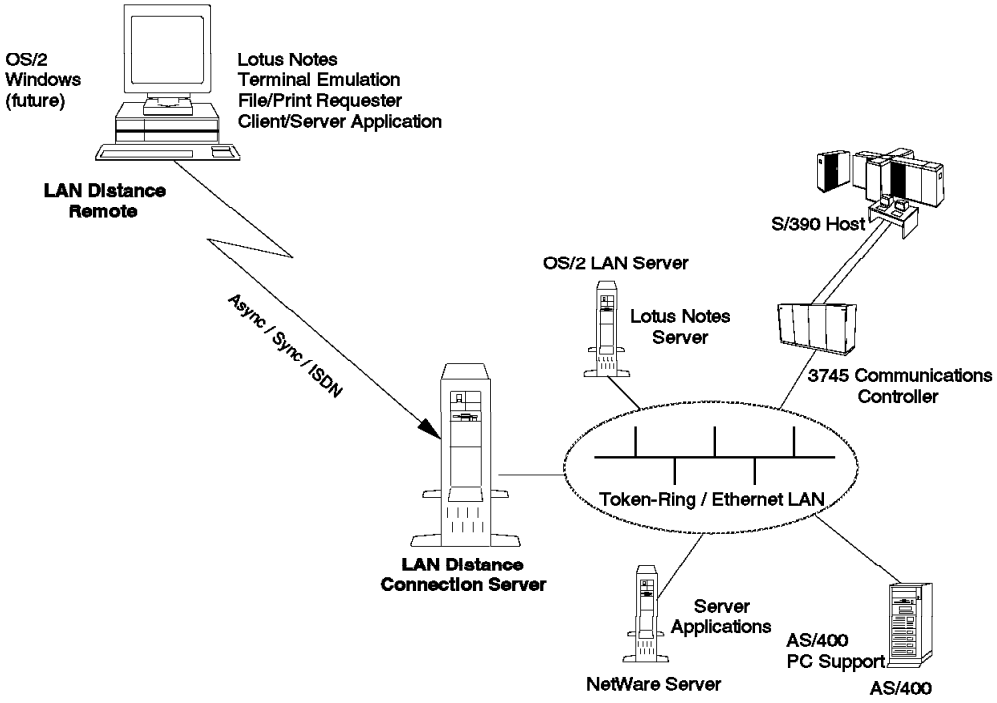

Figure 92. LAN Distance Connectivity Options

4.3.1 Technology

The three basic remote LAN technologies in the marketplace today are discussed in 4.2, “Remote LAN Access Technologies” on page 224. They are the remote control approach, remote client approach, and remote node approach. IBM LAN Distance is a communications package which enables remote users to transparently run their LAN-based applications over switched connections using public switch telephone networks, Private Branch eXchange (PBX), or Computer-controlled Branch eXchange (CBX). It uses the standard hardware devices.

The remote client approach is, in essence, the addition of dial-in capability to a client/server application. This approach is very effective for doing single tasks as the solution can be integrated with the server. However, if you are trying to do a number of different tasks, then solutions based on this approach will be quite cumbersome. For instance, you would dial-in to your Lotus Notes, hang up, dial-in again with the emulation program to access your S/390 host, hang up, and dial-in again to access your NetWare server. Remote client approach is a single purpose solution. With the increase in the number of remote users, the increased dial-in activity would impact the ability of the dial-in system to service both remote users and local LAN users. Thus, most often, this activity has to be offloaded to a separate server. If several servers have to be used, remote users would have to dial in several times to access these servers as the remote client is tied to a server.

The remote control approach is typically thought of as screen mapping. There is a host PC that connects to a LAN and that host PC executes an application. Your remote PC is like a terminal to the host PC. Your remote PC displays exactly what is on the host PCs display and you are able to control that application from your remote keyboard. For primarily text-based applications, the remote control approach works pretty well. However, as you get into graphical-based applications, the cost and overhead of sending bit maps across the WAN starts to degrade performance. This is particularly true with higher density displays such as Super Video Graphics Array (SVGA).

IBM LAN Distance uses the remote node approach. Remote node technology essentially extends the LAN connection across the WAN to a remote PC. Your remote PC is a node on the LAN, just as if it were a node physically in your office. The application support is the same. Your application runs on your remote PC and accesses data from a server that is on your office LAN.

4.3.2 Components and Environments

IBM LAN Distance is a communications package that consists of two components:

- LAN Distance Remote
- LAN Distance Connection Server

They run on ordinary PCs, without any need for specific proprietary hardware. The LAN Distance Connection Server does not need to be a dedicated communications server and can coexist with other software on the same machine.

The LAN Distance Remote can communicate with either another LAN Distance Remote, or a LAN Distance Connection Server. That is, LAN Distance encompasses support for two remote LAN access environments:

- Remote-to-LAN
- Remote-to-Remote

The LAN Distance remote-to-LAN environment is characterized by the remote workstation running LAN applications between itself and one or more LAN-attached workstations via a single WAN connection to the LAN. A separate and direct connection is not required for each LAN-attached workstation with which the remote workstation needs to communicate. The remote workstation can concurrently access multiple LAN-attached workstations without redialing.

There are plans to support the other two remote LAN access environments:

- LAN-to-Remote
- LAN-to-LAN

The LAN-to-remote environment can have the same characteristics and capabilities as the remote-to-LAN environment except that the LAN-attached workstation initiates the connection. As for the LAN-to-LAN environment, it can be a combination of the functions of the LAN-to-remote and remote-to-LAN environments. The connections can be established on a temporary workstation-to-workstation basis across the WAN.

4.3.2.1 LAN Distance Remote

The remote workstation must be running OS/2. Microsoft Windows will be supported at a later date but is currently not supported. The IBM LAN Distance Remote product provides the support for the remote workstation to dial either the Connection Server or another remote workstation and establish a LAN session. It provides a graphical user interface for completing the connection and the necessary interfaces to let your LAN-based applications run remotely. Connecting to the LAN Distance Connection Server allows the remote workstation to access LAN resources as though it were physically connected to the LAN. Through using a peer networking product such as IBM OS/2 LAN Server V3.0 peer services, a LAN Distance remote workstation can connect to another *remote* workstation to share files and printers. The two workstations set up a virtual LAN. The LAN Distance Remote product requires OS/2 Version 2.0 or higher.

4.3.2.2 LAN Distance Connection Server

The Connection Server is an OS/2 application. The IBM LAN Distance Connection Server product provides the ability to *bridge* your LAN to the WAN to allow access from remote workstations running the LAN Distance Remote product. It also enforces security on and off the LAN and manages the *outer* network, which is the WAN. The WAN is inherently slower than the LAN and does not normally have much data integrity and has a tendency to lose connections. The Connection Server manages all of these differences. By using filters, it determines what LAN frames ought to go out on the WAN in order to minimize the WAN traffic. It also provides audit trails for security and administration purposes.

The Connection Server requires OS/2 Version 2.0 (or higher) and can coexist on the same PC with other applications such as IBM OS/2 LAN Server, IBM Communications Manager/2 or IBM Database Manager 2/2. It allows multiple remote PCs to dial-in and concurrently access LAN resources. The IBM LAN Distance Connection Server product is available in two versions:

- Eight-port package

  This is an entry level solution with the same features and functions as the regular package, except that it only supports up to a maximum of eight simultaneous connections. A Connection Server upgrade package is available to upgrade the eight-port Connection Server to a full function Connection Server.

- Standard package

  In this package, the number of simultaneous connections supported is determined by the processing power of the Connection Server PC and the type of WAN communications adapter used. For example, using a 25 MHz i386 DX processor and four IBM Real Time Interface Co-Processor Portmaster cards (each of which supports up to eight asynchronous lines), 32 remote PCs can be supported at one time.

4.3.2.3 Remote-to-LAN Environment

When the remote workstation is configured to run in the remote workstation-to-LAN-attached workstation environment, the remote workstation can dial-in and establish a connection with a LAN-attached workstation that is running the LAN Distance Connection Server product. The LAN Distance Connection Server provides the WAN/LAN connection and supports the same types of WAN connectivity as the remote workstation. In this environment, the

remote workstation has full addressability on the LAN just as if it were physically attached to that LAN. Access to servers on the LAN is restricted to those servers to which authorization has been given.

The LAN traffic received off the WAN at the Connection Server workstation is routed to the target workstation on the LAN. This is accomplished using the bridging capability of the LAN Distance Connection Server product, which provides the capability of transporting LAN frames across the WAN, in a way that is transparent to the applications. The bridging capability is enabled through *LAN device driver replacement technology*.

LAN device driver replacement technology enables the server workstation on the LAN to take data off the WAN and relay it onto the LAN using the correct LAN protocol. This is possible because the data on the WAN link consists of LAN frames sent by the remote workstation. This technology is implemented by replacing the token-ring and Ethernet device drivers in each workstation with device drivers that can send and receive LAN frames across a WAN connection. For example, the LAN MAC driver is replaced with a LAN Distance asynchronous, synchronous, or ISDN MAC driver, to transport the LAN protocol frames such as IEEE 802.5, across the WAN connection to the target workstation.

Since the entire LAN frame is transported between the workstations over the WAN connection, LAN-based applications running in the remote workstation are fully enabled to utilize all graphical interfaces that any LAN-attached workstation can use. The frames are received off the WAN link and routed up to the correct protocol stack and received by the application using the same device driver replacement technology as the sending workstation. The fact that the data arrived across an asynchronous, synchronous, or ISDN link is completely transparent to the protocol stacks and applications which reside above them.

4.3.2.4 Remote-to-Remote Environment

When the remote workstation is configured to run in the remote workstation-to-remote workstation environment, the remote workstation can dial-out (or receive incoming calls) and establish a connection with another remote workstation. Both workstations must be running the LAN Distance Remote product. In this environment, either workstation may be physically on a LAN. However, this LAN connection will not be used, since the connection is established directly between workstations, over a WAN link and will not utilize the physical token-ring or Ethernet connection, even if one is present.

Using this configuration, a *virtual* LAN is created with the stand-alone remote workstations. This enables LAN applications to run over the WAN through existing telephone lines, without the added requirement of a LAN adapter. This virtual capability extends to all LAN Distance Remote workstations that are connected. For example, if multiple LAN Distance Remote workstations are connected to one LAN Distance Remote workstation, they will all be able to communicate with one other. This is accomplished via a discovery/learning algorithm in the LAN Distance Remote product where LAN frames are forwarded in order to discover the remote workstation′s destination location. This provides a low-cost means to allow two LAN applications to communicate without requiring a physical LAN.

4.3.3 Connections, Interfaces and Applications

LAN Distance Remote is connected to the LAN Distance Connection Server through a WAN using an appropriate pair of WAN communications adapters. The Connection Server which sits in the token-ring or Ethernet LAN will have the appropriate LAN communications adapter installed. LAN Distance products support most of the commonly used LAN protocol Application Programming Interfaces (APIs) and LAN-based applications.

4.3.3.1 Supported Connections

IBM LAN Distance supports both token-ring LAN and Ethernet LAN. The appropriate LAN adapter is installed in the Connection Server so it can be attached to the LAN. The same software, LAN Distance Connection Server, supports both token-ring and Ethernet in the same manner.

As for the WAN connection, LAN Distance supports the asynchronous COM port on the PC, as well as a wide variety of asynchronous adapters for both the Connection Server and the Remote. In addition, it supports synchronous connections, ISDN, PBX and CBX. As LAN Distance is a PC application, it can flexibly pick up the support of many different types of adapters and modems that are available for PCs. IBM has externalized the APIs so that adapter vendors can write device drivers to enable them to work with LAN Distance.

4.3.3.2 Communication Interfaces

The LAN protocols provided by the LAN Distance Remote product are NetBIOS and IEEE 802.2. These protocols are written to the Network Driver Interface Specification (NDIS). Protocols which exploit the Open Device Interface (ODI) are also supported through an ODI/NDIS protocol converter.

The following protocols are supported by LAN Distance:

- NetBIOS
- IEEE 802.2 (LLC)
- TCP/IP (which sits as a protocol over the virtual LAN kernel)
- ODI protocols via ODI2NDI (ODI to NDIS protocol converter to support IPX)

Device driver software provides the communication between a network adapter and the transport protocol. The main function of the device driver software is to support network transmissions. The transport protocol provides the communication between an application and the device driver software. is a standard interface for device driver software and the transport protocols over which applications send data. NDIS separates the protocol handling from the hardware manipulation by defining functions needed by transport protocols and device drivers. NDIS has become an industry standard providing a common interface that enables manufacturers to develop network adapters which communicate with LAN software. By complying with NDIS, multiple LAN protocols are enabled without requiring changes to the protocol software or applications and services utilizing these protocols.

Many client/server applications write to the NetBIOS, IEEE 802.2 and ODI/NDIS APIs. Novell′s INTERNET Packet eXchange (IPX) is written to ODI. SNA is written to IEEE 802.2. OS/2 LAN Server is written to NetBIOS. TCP/IP sits as a protocol over the VLAN kernel. Thus LAN Distance supports all these communications interfaces.

Note

Both token-ring and Ethernet LANs are supported. LAN Distance supports the IP router traffic flowing on and off the LAN.

4.3.3.3 LAN-Based Applications

Most of the LAN-based applications will run remotely using IBM LAN Distance. For example:

- IBM OS/2 LAN Server
- Novell NetWare Server
- Microsoft LAN Manager
- Banyan Vines
- Artisoft LANtastic
- AS/400 PC Support
- Lotus Notes/ccMail
- IBM OS/2 Communications Manager/2
- IBM OS/2 Database Manager 2/2

The WAN connectivity is transparent to the LAN-based applications and requires no changes to existing applications or supported LAN protocols. The multitasking capabilities provided by IBM OS/2 enable the remote workstation to perform other tasks while simultaneously supporting a connection to a remote LAN or another remote workstation.

4.3.4 Key Features

LAN Distance provides security functions to protect against unauthorized WAN access to the LAN, as well as other useful security features. Administration and management functions are also provided to help monitor and manage many simultaneous connections from multiple remote workstations.

4.3.4.1 Security

One of the key considerations in implementing a remote LAN access system is security. In addition to the security features built into the LAN and specific application programs, LAN Distance includes a variety of security features that allow you to select as your business needs dictate. Some of the security options available include:

- **User type:** LAN Distance users can be identified as Users, Administrators or Security Administrators, giving them different levels of network authority. The Security Administrator classification allows you to limit the number of people that can access user ID and password information.
- **Passphrase:** LAN Distance supports passphrases which consist of from 4 to 32 case-sensitive characters, including spaces. Passphrases offer improved levels of security, and have proven to be easier for end users to remember.
- **Passphrase encryption:** Standards-based encryption and message authentication codes ensure passphrases cannot be captured and reused.
- **Network address:** Each user account on the connection server can have up to eight PC addresses. Users must call from a valid PC address in order to log on.

- **Valid logon time intervals:** A security administrator can specify days of the week and times of day during which a user is authorized to log on to the network.
- **Configurable logon parameters:** Several security policies that apply to all user accounts are configurable, including minimum and maximum passphrase length, minimum and maximum password age, passphrase duplication frequency and maximum number of unsuccessful logon attempts allowed.
- **Call back:** Users can optionally be identified as call-back users. After user authentication, the line is disconnected and the user is called back at a pre-configured telephone number.
- **Security audit trail:** Both successful and unsuccessful logons can be recorded in the security audit file. In addition, changes to the security database are also logged.

4.3.4.2 Administration and Management

One of the administrative features is the ability for the remote user to go back to the network panel and see the status of the connections. Thus, remote users do not have to worry about whether they are connected to the LAN or not. Errors are logged and you have the ability to look at that log to see what happened. A trace facility is also provided to trace internal events for problem diagnosis. LAN Distance can log users who are connected to the LAN and the durations of their connections for accounting purposes. The administration of multiple Connection Servers can be done from a single local LAN workstation or remote workstation. Network problem determination and resolution can also be done from a remote PC.

System status and management functions are provided on both port usage and call information. These functions show which resources are currently in use and aid in problem determination. Ports may be started and stopped. Bridge management is provided by Connection Server workstations running IBM TCP/IP for OS/2 V1.2. The workstation, which acts as a subagent, responds to the SNMP verbs for GET and NEXT to manage the bridge between the WAN and the LAN.

4.4 IBM 8235 Dial-In Access to LAN Servers

The IBM 8235 Dial-In Access to LAN (DIAL) Servers for Token-Ring and Ethernet is a dedicated multiport, multiprotocol remote access hardware server. This server supports remote personal computer (PC) users dialing-in to applications the same way users access applications from workstations directly attached to a token-ring or Ethernet local area network. With routing and bridging support for the following multiple protocols, a user can remotely access a variety of applications:

- NetBIOS for LAN servers
- IPX for NetWare
- 802.2 LLC for 3270 and SNA
- IP for TCP/IP applications
- AppleTalk Apple Remote Access (ARA) 2.0 (Ethernet Only)

Using standard dial networks, users with PCs and modems who are remote from the LAN can access LAN resources and work with applications as if they were working at locally-attached LAN workstations.

Users in the field, such as agents, sales representatives, and employees who travel or work at home, have the ability to access their applications from any location that has dial-up telephone service. This extends the productivity of the workstation to the remote workplace. Using standard analog modems and dial-up telephone lines, the IBM 8235 and the IBM DIALs Client for OS/2, DOS, and Windows operating in the remote PC allow easy access to resources that users normally access from a workstation connected to a LAN. With support for multiple protocols and with high-performance filtering and compression techniques, excellent performance can be achieved when addressing a variety of applications remotely.

4.4.1 8235 System Components

The 8235 remote access system is made up of three basic components:

1. The Dial-in Access to LAN Client

   A software application that runs on the remote PC providing the dial-in function. The DIALs Client supports DOS, Windows, and OS/2.

2. The 8235 Management Facility

   A Windows application that allows the 8235 to be configured and managed from any LAN-attached workstation running IPX and Windows.

3. The 8235

   A stand-alone hardware device that attaches to either a token-ring or Ethernet LAN and the public switched telephone network. The function of the 8235 hardware and its associated software is to:

   - Provide physical attachment to the LAN and to eight modems.
   - Forward data from the LAN to the remote PCs and from the remote PCs to the LAN using any of the following protocols: IPX, IP, NetBEUI, AppleTalk ARA 2.0 and LLC.
   - Filter and compress data so as to minimize the amount of unnecessary traffic between the LAN and the remote PC.
   - Prevent unauthorized access to the LAN.

4.4.2 Dial-In Access to LAN Servers (DIALs) Client Software

DIALs Client is IBM′s multiprotocol dial-in software for workstations. It allows your modem to fully access resources of remote networks. The DOS and DOS/WINDOWS client requires approximately 850 KB disk and 19 KB RAM.

Note

The DIALs Client is shipped with the 8235 with an unlimited right to copy.

DIALs Client contains the following software:

- OS/2 Drivers (NDIS and ODI):

  These provide support for OS/2-based communication programs. ODI can be provided with LAN adapter and protocol support (LAPS).

- DOS Drivers (NDIS and ODI):

  These provide support for your DOS-based or Windows-based communications programs.

- Connect Application:

  Allows you to create, store, and use connection files to dial-in to remote networks from the OS/2,DOS and windows environments. The connect program:

  – Provides traffic-flow statistics
  – Displays error information
  – Displays the modem status
  – Displays the modem configuration

4.4.3 IBM 8235 New Features

This section describes the new features provided by DIALS Release 2.0 and DIALS Release 4.0.

4.4.3.1 DIALS Release 2.0

1. Dial-In:

   For the dial-in function, 8235 Version 2.0 provides the following features:

   - ARA 2.0 dial-in support for Ethernet 8235s (ARA 1.0 dial-in is not supported). ARA dial-in provides the following features:

     – IP forwarding (MacTCP)
     – Routing or end-node forwarding support for ARA clients
     – AppleTalk device and zone filtering per user, per port, or per 8235

   - Simultaneous PC dial-in over point-to-point protocol (PPP) for the following protocols:

     – NetWare Internet Packet Exchange (IPX support)
     – Transmission Control Protocol/Internet Protocol (TCP/IP)
     – NetBIOS Extended User Interface (NetBEUI)
     – 802.2/Logical Link Control (LLC) (SNA)

   - Support for the Novell Client for DOS/Windows, or Virtual Loadable Modules (VLMs)

   - Windows for Workgroups (WFW) 3.11 support

2. Shared Dial-Out Access

   This is used for access to external asynchronous services such as CompuServe.

3. LAN-to-LAN Support

   - Connections between two networks routing any combination of TCP/IP and IPX over a dial-up link. AppleTalk LAN-to-LAN routing is supported for the Ethernet models of the 8235.

   - Connection features including idle detect, persistence, back-up telephone numbers, dial back, and timed connections.

   - LAN-to-LAN connections established automatically or via the command shell (scripting possible).

   - Leased-line support.

   - AppleTalk device and zone filtering for the Ethernet models of the 8235.

4. Centralized Management
    - All protocols and features are manageable from the 8235 Management Facility for Windows.
    - Management Facility tuning for large IPX networks.
    - BOOTP/TFTP automatic downloading.
    - Command shell via IP Telnet, or dial-in on a PC.
5. Additional Security
    - Security Dynamics ACE/Server (SecurID) support for multiprotocol dial-in.
    - NetWare Bindery authentication for all protocols, including ARA 2.0.
    - 8235 user list.
    - Roaming or fixed dial back.

Note

Release 1.1 and 1.0 DIALs Client for OS/2, DOS, and Windows software is compatible with all 8235 models and previous releases, including Release 2.0. The new DIALs Client software Release 2.0 is shipped with 8235 Release 2.0 and is available in an upgrade kit for previous 8235 models. DIALs Release 2.0 Client software is not compatible with previous models of the 8235, unless the models are upgraded to microcode Release 2.0.

4.4.3.2 DIALS Release 4.0

1. **Dial-In**
    - **Multiprotocol support:** Simultaneous multiprotocol dial-in over PPP: IPX (VLMs and NETX supported) TCP/IP, NetBEUI, 802.2/LLC.
    - **VxD Windows Client Feature Summary:** Client has been re-designed to enable support for:
        - Windows Virtual Device Driver VxD that only uses 2 KB of client conventional DOS memory (verses 34 KB)
        - Multilink PPP protocol (MLP)
        - Channel aggregation (2B)
        - Stac 4.0 compression
        - Port driver for internal ISDN adapters (digital modems, TAs.)
        - Native driver support for IBM WaveRunner digital modem
        - New port driver programming interface (API)
        - Virtual connections
        - New intelligent setup facility
        - Easy client installation scripting
        - Client event logging application
    - **Virtual Connections:** The ability to automatically suspend and resume a physical connection while spoofing network protocols, routing and applications. The physical connection is only brought up on-demand.

- **Spoofing:** When a virtual connection is suspended, the ability for a device to determine what is not *meaningful* traffic. Rather than establishing the connection, the device responds to the source of the traffic with the response that would have been generated by the intended destination device.
- **Dial-in Channel Aggregation:** The ability to use more than one communications channel per connection. By aggregating both 64 Kbps ISDN B-channels users can take advantage of 128 Kbps dial-in connections. Fast 128 Kbps data transfer rates reduce *large file* transfer times.
- **IBM WaveRunner Digital Modem (Internal ISDN terminal adapter):** Provides support for the ISA and PCMCIA versions of the IBM WaveRunner digital modem. The three supported modes are Async V.32 bis modem, ISDN V.120, and Sync Clear Channel.
- **Easy client setup:**
  - An intelligent client setup program that includes a *Connection File Wizard* that walks the user through the installation and modifications to client software.
  - The ability to automatically detect attached communications adapters.
  - Powerful file copy *mastering* capability.
  - Client event logging application provides extensive troubleshooting information. Log information can be displayed to the screen or to a file.
- **Power Switching:** Allows users to switch back and forth between communications adapters. Perfect for employees that use one type of communications adapter when working at home (ISDN) and another adapter (V.34 modem) when traveling.
- **Express Installation:** A new client installation scripting utility that enables network managers to establish defined defaults that make client installation and deployment easier.
- **Third-party client support:** Dial-in access from Windows 95 and Windows NT 3.5, Apple′s ARA, and IBM′s OS/2 DIALS.

  Customers using Windows 95, Windows NT, MAC OS or OS/2 can seamlessly use an IBM 8235 as their dial-in server.
- **Client Event Logging Application:** Events can be displayed on the screen and/or saved in a text file. The logged events include:
  - Buffer allocation/management
  - PPP events and state transitions
  - PPP negotiation options
  - All frames transmitted and received
  - Multilink (MLP)
  - Compression
  - Network protocol decoding (basic IPX, IP and NetBEUI frames)
- **New Port Driver:** The new port driver provides support for internal client ISDN terminal adapters such as the IBM WaveRunner.

- Internal ISDN adapters eliminate the async-to-sync conversion overhead required by external terminal adapters.

2. **New Application Programming Interface (API):** The IBM DIALs 4.0 port driver API enables third parties to independently develop IBM DIALs drivers for their hardware. Many internal ISDN terminal adapters do not present a standard PC 8250/16450/16550 UART interface.
3. **Enhanced Stac 4.0 Compression:** IBM upgraded the Stac compression algorithm from 3.0 to 4.0. Stac 4.0 is faster and more memory efficient. For digital terminal adapters where there is no compression done by the ISDN TA or X.25 PAD, it is essential that the compression algorithm used on the client be as lean and fast as possible.
4. **LAN-to-LAN Features**
   - **Virtual Connections (VC):** The ability to automatically suspend and resume a physical connection while spoofing network protocols, routing and applications. The physical connection is only brought up on-demand.
   - **Spoofing:** When a virtual connection is suspended; the ability for a device to determine what is not meaningful traffic. Rather than establishing the connection, the device responds to the source of the traffic with the response that would have been generated by the intended destination device. Spoofing is done for file server connections (NetWare drive mapping), routing tables (IP RIP and IPX RIP), SAP tables, TCP connections, and SPX connections.
   - **Floating Virtual Connections (FVC):** The ability to resume a suspended virtual connection on a port other than the port on which the original virtual connection was established. It can reduce the need to dedicate ports to specific users.
   - **Juggling Virtual Connections (JVC):** The ability to have more suspended virtual connections than there are ports on the IBM 8235. Customers can have many more suspended users than they have ports. JVC maximizes the utilization of server communications ports.
   - **Persistent Connections (PC):** A IBM 8235 configuration option that allows the server to re-establish the connection in the event of an unexpected line drop.
   - **Timed LAN-to-LAN Connections (TLC):** The ability for network managers to schedule LAN-to-LAN connections. (For example, establish a LAN-to-LAN connection at 10 AM and terminate the connection at 1 PM.)
   - **Piggybacking Updates:** A virtual connection synchronizing mechanism where routing update messages are sent across the link only when the link is open for real data traffic.
   - **Timed Updates:** A virtual connection synchronizing mechanism where at a specified interval the suspended virtual connection is resumed to enable routing update messages to be sent across the link.
   - **Triggered Updates:**
     - A virtual connection synchronizing mechanism where routing update messages are sent across the link only when there is a RIP or SAP database change.
     - Triggered update setup options include additions only, deletions only, or additions and deletions.

- **Channel Aggregation (Multilink PPP, MLP):** The ability to use more than one communication channel per connection. LAN-to-LAN connections can aggregate all IBM 8235 channels (analog or digital) up to the number of ports on the server.
- **Packet Fragmentation:** The ability to configure a default packet size over which packets will be fragmented for more efficient distribution over aggregated communications links.
- **LanConnect Applets:** LanConnect applets for both PC and MAC allow for scripting of on-demand LAN-to-LAN connections.
- **Delta Technology:** Specialized remote adaptive routing protocols for optimizing bandwidth. It prevents unnecessary traffic from being sent over slow WAN connections by only sending the changes (deltas).

5. **Management and Security Features**
   - **PC and MAC server management:** Protocols and features can managed by MAC or Windows versions of IBM NetManager (MAC Appletalk, PC/Windows IPX and IP).
   - **IP Download:** IBM MF will be able to download new code images and configurations when running over either IP or IPX protocol stack.
   - **SNMP Management:** MIB II and others.
   - **Security:** Provides support for agent software from Security Dynamics & Digital Pathways. Centralized authentication via IBM user list, NetWare Bindery, TACACS and most third party hardware security solutions are supported.

4.4.4 What is Virtual Connection?

A Virtual Connection is a standard LAN-to-LAN or PC single-user dial-in connection that is enhanced to detect when no meaningful traffic has been sent over the connection for a period of time, at which time the physical connection is suspended while network protocols (IPX and TCP/IP) are spoofed by devices at either end of the connection. Subsequently, when meaningful traffic is received by either of the devices, the physical connection is automatically resumed and the data is forwarded over the communications link. Virtual connections minimize connect-time costs by physically disconnecting the circuit when there is no meaningful traffic.

Another benefit of a virtual connection is ease-of-use and management. Once the original connection is established, no user or system administrator intervention is required. The physical link is automatically suspended and resumed on demand.

4.4.5 What is Channel Aggregation?

New high performance channel aggregation technology enables dial-in and LAN-to-LAN users to establish more than one communications channel per connection. IBM channel aggregation technology utilizes the industry-standard protocol known as Multilink PPP for maximum client/server device interoperability and investment protection. Packet fragmentation is also available for maximum performance.

4.4.6 Management Facility

The Management Facility program is a Windows application that enables you to configure and manage the 8235s on your network, create user lists, and manage the security of your 8235s. This program is provided with your 8235. The IBM 8235 Management Facility requires a workstation with Windows 3.1 or later, initially attached to the network. All 8235 models operate with the same 8235 Management Facility. You also need to load IPX on the machine running the Management Facility to communicate with the 8235.

In Figure 93 you can see the Management Facility window.

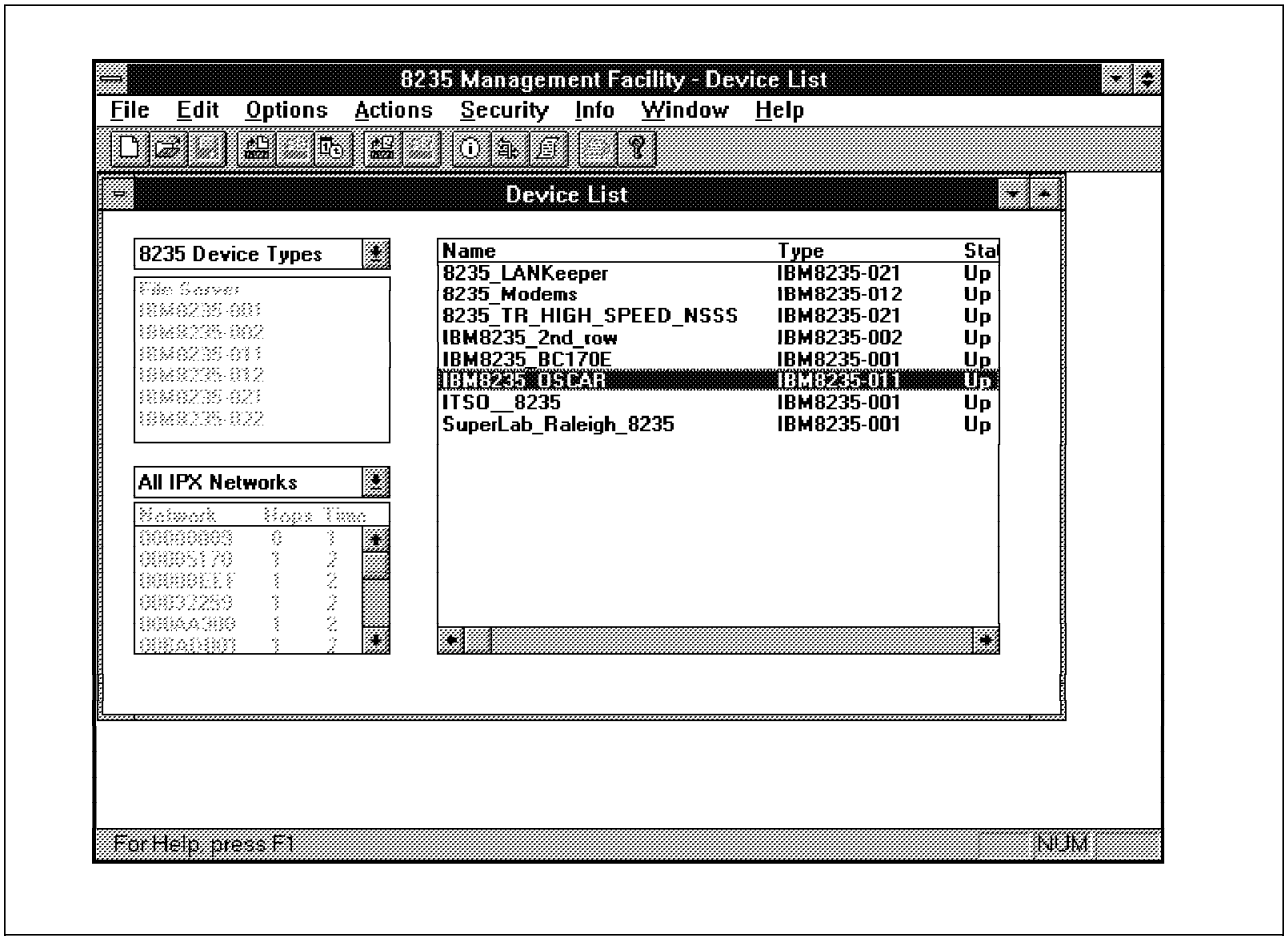

Figure 93. 8235 Management Facility Window

4.4.7 8235 Hardware

Figure 94 on page 241 shows the front panel for all models of the 8235.

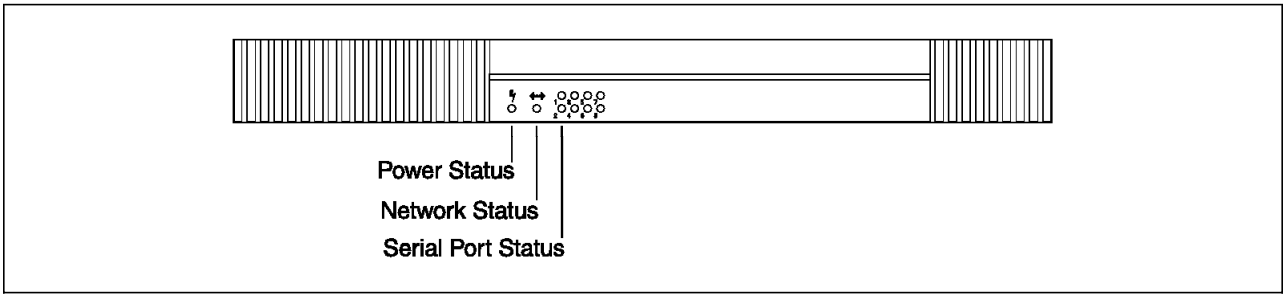

Figure 94. 8235 Front View

The front panel contains LEDs that indicate:

- Power status
- Network status
- Serial port status

Table 14 shows the meanings of the status indicator LEDs on the front panel of the 8235 in various operating modes, and Table 15 shows the meaning of the power LED.

Table 14. Meanings of 8235 Network Status and Port Status LEDs

Status	**Network Status LED**	**Port Status LEDs**
OFF	No power or no network connection	Not in use
Green	Connected to network but idle	User connected
Green flashing (consistent)	Downloading microcode	Downloading microcode
Green flashing (inconsistent)	Connected to the network and transmitting	User connected
Green and Orange flashing	Connected to the network and transmitting with errors	-
Orange flashing (consistent)	Power on self-test	Port configuration errors
Orange flashing (inconsistent)	-	Connected to the modem and transmitting with transmit or receive errors
Orange (solid)	8235 hardware failure	Port or 8235 hardware failure

Table 15. Meaning of 8235 Power Status LED

Status	**Meaning**
ON	Indicates that the 8235 is powered on

4.4.7.1 LAN Connection

As mentioned earlier, the 8235 comes in two models:

- Model 1 contains a token-ring connection port.
- Model 2 has an Ethernet connection port.

The 8235 is also available as a module for the 8250 multiprotocol hub in token-ring and Ethernet models. Figure 95 shows the rear view of the token-ring model 8235-021.

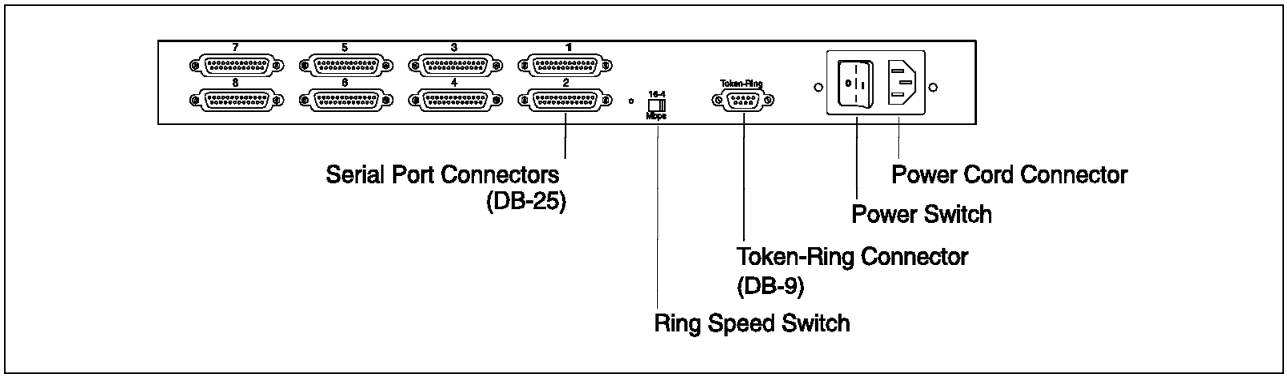

Figure 95. 8235 Model 021 Rear Panel

Figure 96 shows the rear panel of the token-ring model 8235-031.

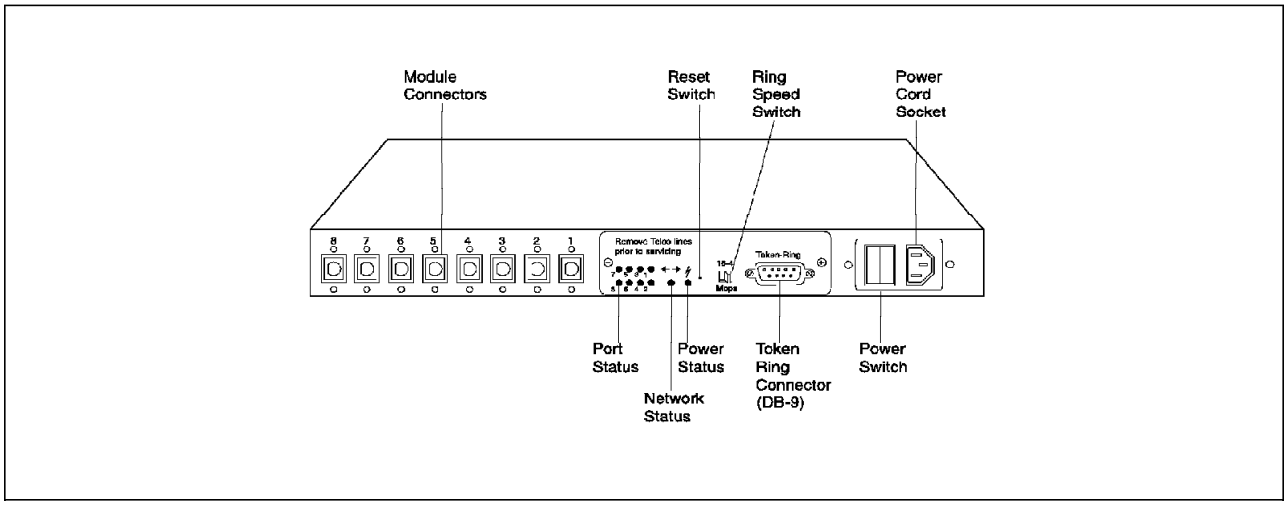

Figure 96. 8235 Model 031 Rear Panel

You make all connections on the 8235 rear panel, so the token-ring model includes one token-ring connector (DB-9) and a ring data rate switch to select the data rate of 4 or 16 Mbps.

Note:

The data rate you set must match the data rate of the token-ring network. Be sure to set the power switch to Off (O) before you set the data rate.

Figure 97 on page 243 shows the rear panel of the 8235 Ethernet model 022.

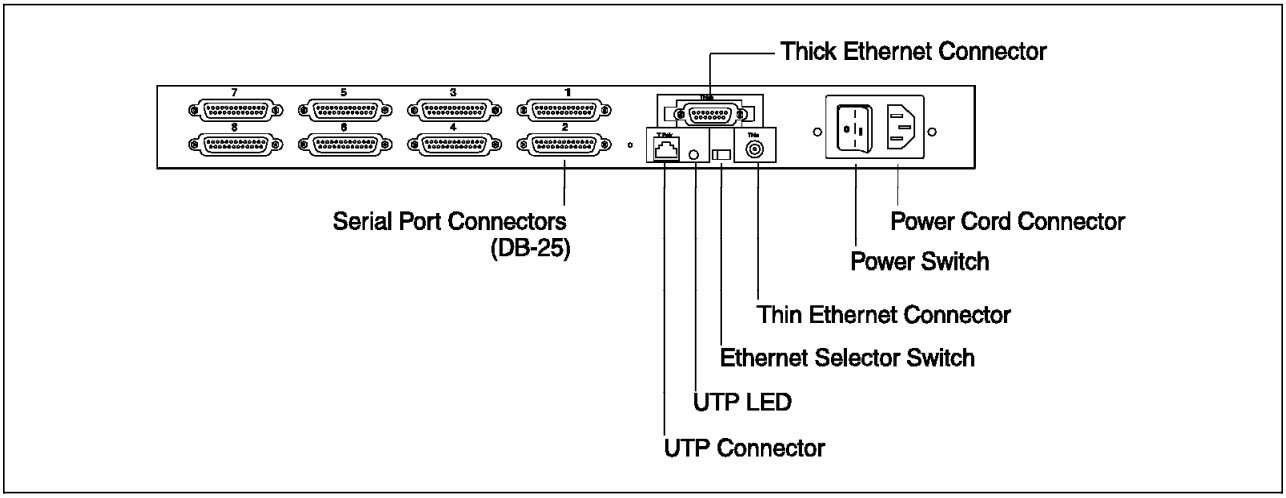

Figure 97. 8235 Model 022 Rear Panel

The 8235 Model 002 (Ethernet) provides three connectors for Ethernet: AUI (Thick Ethernet), BNC (Thin Ethernet) and UTP as shown in Figure 97. You must select the Ethernet connector that you want to use with the switch that is at the back of the 8235.

Three Ethernet wiring schemes are supported:

- Thin (10Base2)
- Thick (10Base5)
- UTP (10Base-T)

When twisted-pair is selected, the LED next to the twisted-pair port on the rear panel of the 8235 Model 2 indicates the network status. Table 16 summarizes what the various flashing patterns mean and what actions, if any, you should take.

Table 16. 8235 LED Error Code Flashing Patterns

LED Pattern	Meaning	Action to Take
On	Normal link is established.	None; normal operation.
Off	10Base-T is not selected	Set the Ethernet connector switch to the 10Base-T (far left) position.
One flash	Link to 10Base-T is down.	Check that the hardware connections are secure. Re-establish the link.
Two flashes	Jabber error (possibly transient). The 10Base-T transceiver has detected a continuous frame transmission of 131 milliseconds or greater by the LAN controller in the 8235 Model 2. Transmission on the network is inhibited.	Wait a few seconds to see whether the problem goes away. If not, restart the 8235 Model 2, or contact IBM Product Support.

4.4.7.2 8235 Code Structure

The software that runs in the 8235 server can be separated into three pieces:

- Boot PROM
- Virtual ROM (VROM)
- The main software image

Boot PROM: The Boot PROM resides in ROM and performs the function of downloading a software image if there is no valid image in the VROM. Otherwise, the VROM performs software downloads. The Boot PROM accomplishes software downloads via Boot Protocol (BOOTP) and trivial file transfer protocol (TFTP) or via SPX. In addition to software downloads, the Boot PROM performs power-on-self test (POST) and switches the device to diagnostic mode if the POST fails.

VROM: The VROM serves to isolate the mainline programs from the hardware by providing the following:

- Device drivers for LAN and serial port I/O
- Buffer and memory management
- Management of non-volatile storage
- LED manipulation
- Message logging
- Acquiring VROM maintained data
- Acquiring hardware configuration information

The VROM also contains a bootstrap application that is capable of acquiring a new download by unattended BOOTP and TFTP or a NetWare SPX download from the Management Facility. The 8235 downloads new images through the LAN port (token-ring or Ethernet).

Main Software Image: The bulk of the run-time function in the 8235 is contained in the main software image. This image consists of the software kernel, frame forwarding support, management, and security.

4.4.7.3 Updating Microcode

The system structure for the 8235 makes it an excellent platform for future enhancements that can be obtained via software updates.

Downloading Modes: The 8235 can be put into several different boot-up sequences under the control of one of the following:

- Management Facility
- Command shell
- Physical interruption (power on and off, pin reset)

The different modes are described in the following paragraphs.

Warm Boot: Under normal circumstances, the 8235 will contain a software image and configuration that has been stored in battery-backed RAM. When the system is rebooted (powered on or restarted due to a configuration change), it goes through a normal cycle. During this cycle, it will temporarily appear to the Management Facility to be in download mode. The device list window will indicate that the device is in DL mode. This condition should last for only a few seconds. If for some reason the 8235 has lost its code image or has been pin reset, it will remain in download mode until a management entity has loaded new code (see “Clear and Download” on page 245 below).

Download Code Only: The 8235 can be instructed to download a new code image only by issuing a Download command from the Management Facility. This means that it will load a new code image, but will maintain its configuration data.

Clear and Download: A Clear and Download command from the Management Facility will put the 8235 into download mode from the Boot Prom on the 8235 and will load both code and VROM, and will cause any configuration data in the 8235 to be lost. It will remain in download mode until a management entity loads a new version of code.

Pin Reset Switch: The 8235 has a tiny pinhole at the back that is not labeled. It is a pin reset which corresponds to an internal switch that performs the hard reset of the 8235 and is often overlooked. It should be used if you lose contact with the Management Facility due to hardware problems or if you lose the administrator's password. It performs the same function as the Clear and Download command. No indication of this pin reset is noted on the hardware itself.

4.4.8 Models Summary

The main difference between all the 8235 models is the communication port that is used.

Table 17. 8235 Models

Model Feature	Token-Ring	Ethernet	HS Serial Port (115.2Kbps)	Internal Modem	Serial Port (57.6 Kbps)
8235-021	X		X		
8235-022		X	X		
8235-031	X		1-8	1-8	1-8
8235-032		X	1-8	1-8	1-8
8250 module	X		X		
8250 module		X	X		

Note:

The models 031 and 032 have empty slots, into which you can install up to eight cards: eight modem cards, or eight serial cards, or a combination of both.

4.4.9 Communication Options

Here is a brief description of the different communication options that the 8235 has:

- Models 021 (token-ring) and 022 (Ethernet)

  The new, high-speed base models, 021 and 022, support serial port speeds up to 115.2 Kbps, enhancing the 8235 model offerings. These new models are shipped with eight RS-232-D (V.24/V.28) ports for attachment of up to eight modems with 115.2 Kbps serial port speed. Excellent performance can be achieved with the high-speed V.34 data compression modems.

- Models 031 (token-ring) and 032 (Ethernet)

  These models do not contain a fixed port configuration. The customer configures the ports to meet their needs with any combination of modems and/or serial cards.

  Model 031 is an unpopulated token-ring base server, and Model 032 is an unpopulated Ethernet base server. Both models provide plug-in slots for V.34 modem cards and serial cards. These models support a total of eight cards (eight modem cards or eight serial cards, or a combination of both cards totaling eight).

  These models can support eight remote users simultaneously with reliable asynchronous transmission speeds up to 115.2 Kbps. With the serial cards, you can configure some or all of the ports to attach external asynchronous terminal adapters for digital services, such as ISDN or Switched 56.

  The Management Facility of 8235 Models 031 and 032 is an extension to the facility provided with the other models of the 8235 and is enhanced to include management of the new V.34 integrated modems and serial cards.

IBM has extended the flexibility of the IBM 8235 Models 031 & 032 remote access server with several new upgrade modules:

IBM 8235-031 & 032 BRI module

- 2B+D with V.110 & V.120 rate adaption.
- S/T and U interface versions are available.
- BRI module can be monitored from IBM MF. Configuration setup, revisions, and troubleshooting can all be managed remotely.

IBM 8235-031 & 032 Sync/Async module

- User can connect synchronous devices (ISDN BRI TAs, CSU/DSUs and modem eliminators) directly to the IBM 8235/MODELS 031 & 032. The direct synchronous connection takes advantage of the faster line speed (128 Kbps vs. 115 Kbps), the elimination of extra timing bits (Async has two extra timing bits per character transmitted), and the overhead of converting a synchronous transmission into asynchronous transmission.
- Supports either synchronous or asynchronous communications channels.

- 8250 Modules

  These modules integrate IBM 8235 remote LAN access server product functions into the 8250 hub.

  There are two kinds of 8235 modules:

  - One for attaching an Ethernet network
  - One for token-ring network attachment

  These modules occupy a single slot in the 8250 hub chassis. The Ethernet module provides one Ethernet attachment switchable to any of the three Ethernet segments on the 8250 backplane. Likewise, the token-ring module provides one token-ring attachment that can operate at either 4 or 16 Mbps. The attachment is switchable to any of the seven token-ring backplane segments.

  Each module has eight serial communication ports. Each port has an RS-232-D (V.24/V.28) interface with a DIN connector for attachment to standard asynchronous modems. Data transfer speed ranges from 2400 bps up to 28.8 Kbps, or even up to 115.2 Kbps when using high-speed data compression modems. The modules come with eight DIN-to-25 pin RS232 patch cables to attach to external modems.

4.4.10 Supported Protocols

The 8235 supports remote clients using any of all the following protocols:

4.4.10.1 NetBIOS and 802.2

The 8235 software filters on LLC service access point (SAPs) and on NetBIOS names based on the filter tables contained in the server. The tables will be set up in the box, but the information can be overridden using the operating system shell. There are no external parameters available to manage filtering as there are for an IBM Token-Ring Bridge or for LAN Distance software. LLC SAP filters allow X′02, X′04, X′05, X′08, X′E0, X′F0 and X′F4 SAPs to be bridged. These are also configurable.

Frame forwarding (that is the process of forwarding data from the client workstation to the LAN and from the LAN to the client) is accomplished

differently depending on the protocol selected during the configuration of the connections.

4.4.10.2 Bridging

The token-ring acts like an IBM token-ring bridge with NetBIOS and 802.2 protocols as shown in Figure 98.

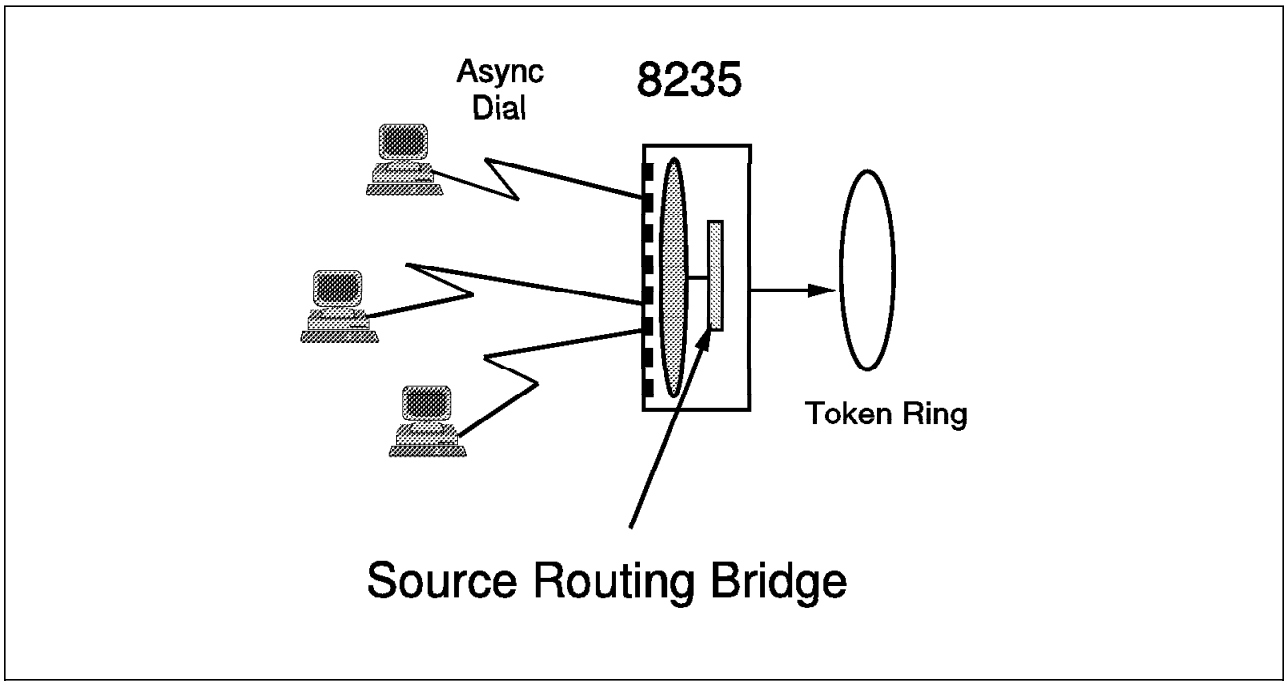

Figure 98. Source-Routing Bridge

The bridged frames appear on the ring as if they came from an adapter. NetBIOS and 802.2 dial-in also supports specialized filtering to protect clients from broadcast traffic on the dial-in links.

The 8235 acts like a transparent bridge for Ethernet as shown in the Figure 99 on page 249.

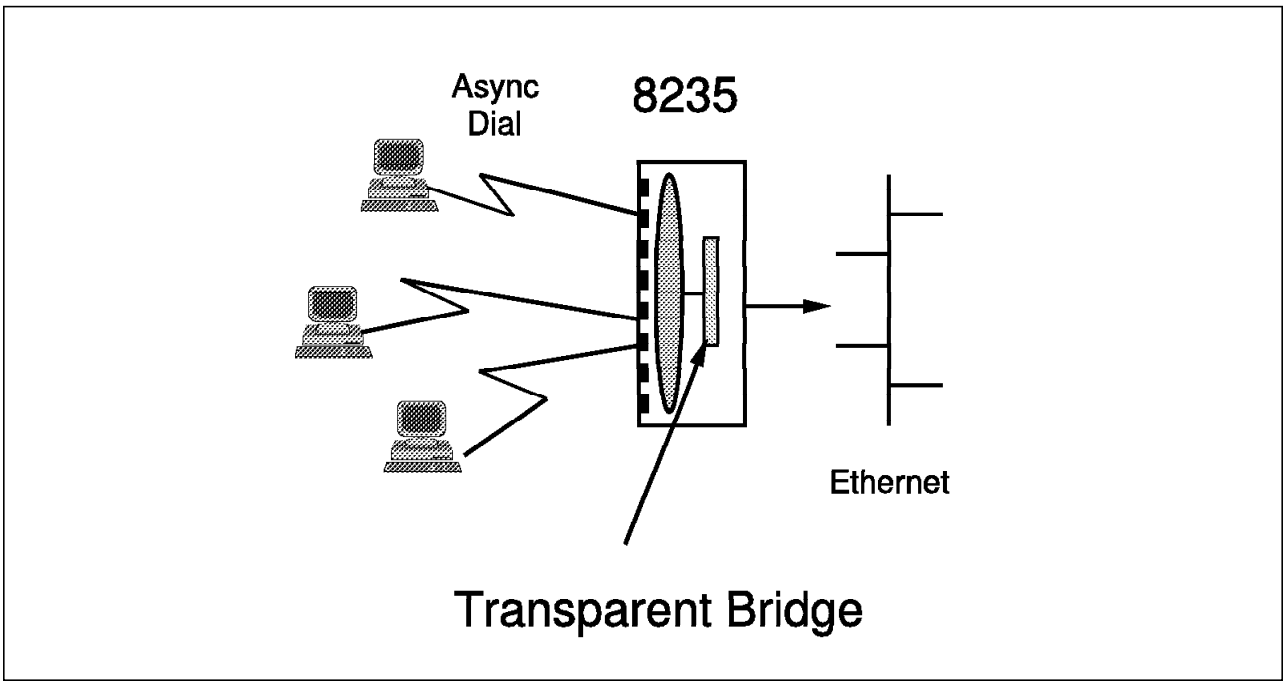

Figure 99. 8235 Acting As a Transparent Bridge

4.4.10.3 Ring Parameter Server

The ring parameter server (RPS) function has been implemented in the case where the 8235 is the only bridge on the ring. Here is an explanation of what the RPS function provides.

The RPS is the target for all request initialization MAC frames that are sent by ring stations during their attachment to the ring segment. The RPS function makes the following parameters available to all ring stations on the ring in response to the request initialization MAC frame:

- Ring number
- Ring station soft error report time value (default of 2 seconds)
- Physical location (not currently implemented)

There can be more than one RPS function active on any given ring segment.

Note

This differs from an IBM source-routing bridge in that LAN reporting mechanism functions are not present in the 8235 which would allow it to report configuration information to LAN Network Manager (LNM) or to accept configuration changes from LNM.

4.4.10.4 IP Traffic

The 8235 will transparently forward IP traffic based on the IP address. The 8235 implements the proxy address resolution protocol (ARP) function to reduce broadcast traffic over the remote lines.

Note

This means that the 8235 will respond to all ARP queries for remote client addresses with its own hardware address instead of having the ARPs go across the WAN. The source stations will then forward packets to the remote clients to the 8235's physical address. The 8235 will then route the packet to the correct client based on the IP address.

An example of how the network would appear is shown in Figure 100.

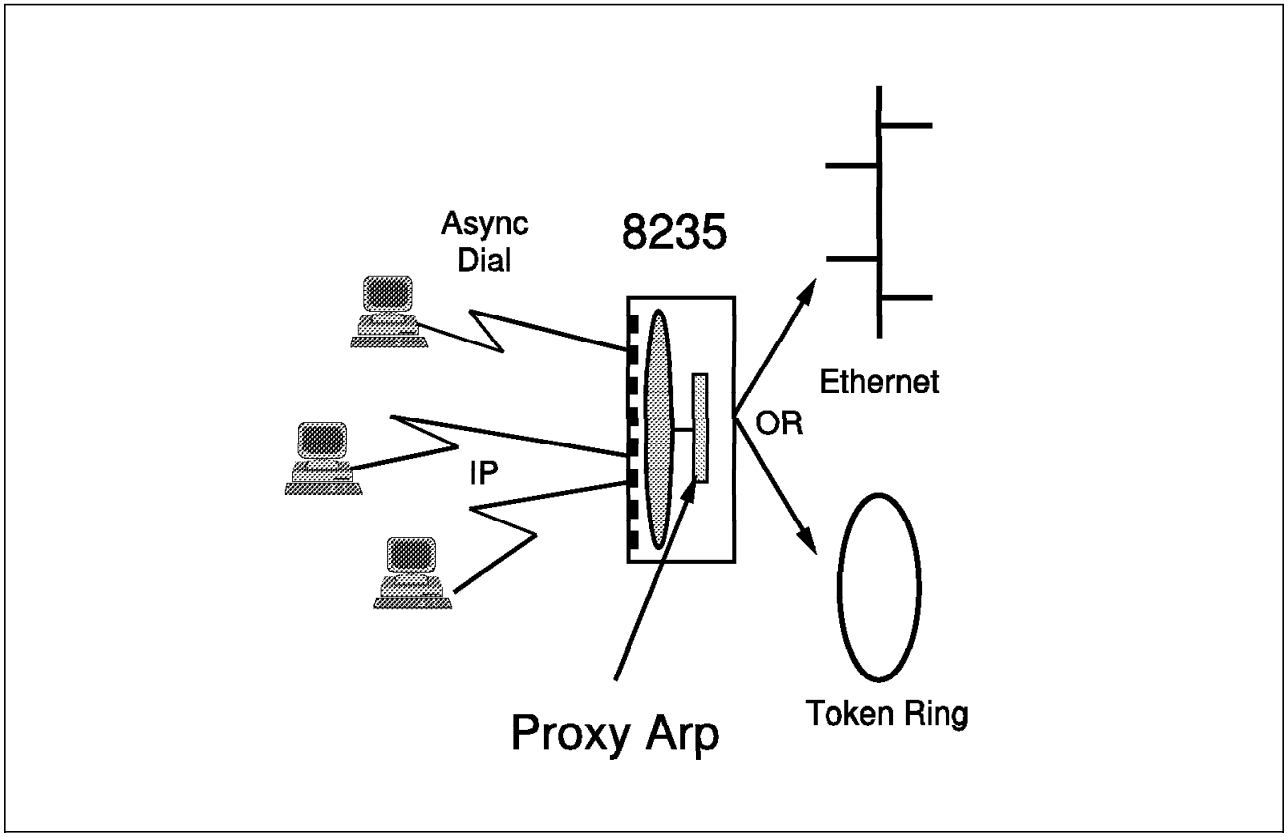

Figure 100. 8235 Proxy ARP

The 8235 will implement the following IP functions:

- IP Address Resolution Protocol (ARP)
- Internet Protocol
- Internet Control Message Protocol (ICMP)
- Transmission Control Protocol (TCP)
- User Datagram Protocol (UDP)
- Trivial File Transfer Protocol (TFTP)
- Boot Protocol (BOOTP)
- Telnet
- Routing Information Protocol (RIP)

For IP traffic, Van Jacobson Header compression is supported. This is transparent to the user, but enhances performance over the telephone network connection.

IP environments pose a unique challenge to dial-in access, as the addresses contain the identification of the network. If the users provide their own IP

address, then they are limited to dialing-in to the network for which they have been preconfigured. There are, however, some environments where the user will be dial-in to the same network all of the time and want to keep the same IP address. Furthermore, because of the nature of IP address discovery (ARP), it is desirable to limit the amount of ARP traffic across the WAN.

Because of this, the 8235 supports address assignment in two ways:

1. Proxy ARP with static client addressing, which has the following properties:
   - Dial-in client has configured IP address, provided to the box by IPCP.
   - A user must dial-in or attach to the same network all of the time.
   - Full end-user TCP/IP application suite support.
   - IP address for each dial-in client is resolved to MAC address of the LAN port (proxy ARP).
   - Packets are routed based on host ID. If the network ID does not match the host ID, the packets will not be forwarded.
   - Remote to remote is a special case (the 8235 recognizes it and forwards the traffic as a special case).
   - Header compression is supported.
2. Proxy ARP with dynamic client addressing, which has the following properties:
   - The 8235 provides unique client IP address through IPCP.
   - Dial-in user can dial into any network that is reachable from the LAN to which the 8235 is connected.
   - The user does not own a well-known IP address. While this may prohibit the use of dial-in clients as servers, it allows the use of most user-oriented software.
   - IP address for each dial-in client is resolved to MAC address of LAN port.
   - Packets are routed based on host ID.
   - Remote-to-remote is a special case (the 8235 recognizes it and forwards the traffic as a special case).
   - Header compression is supported.

Note

The IP address of the 8235 box itself can only be assigned through the Management Facility.

4.4.10.5 IPX Traffic

The 8235 implements an IPX router function as defined by Novell.

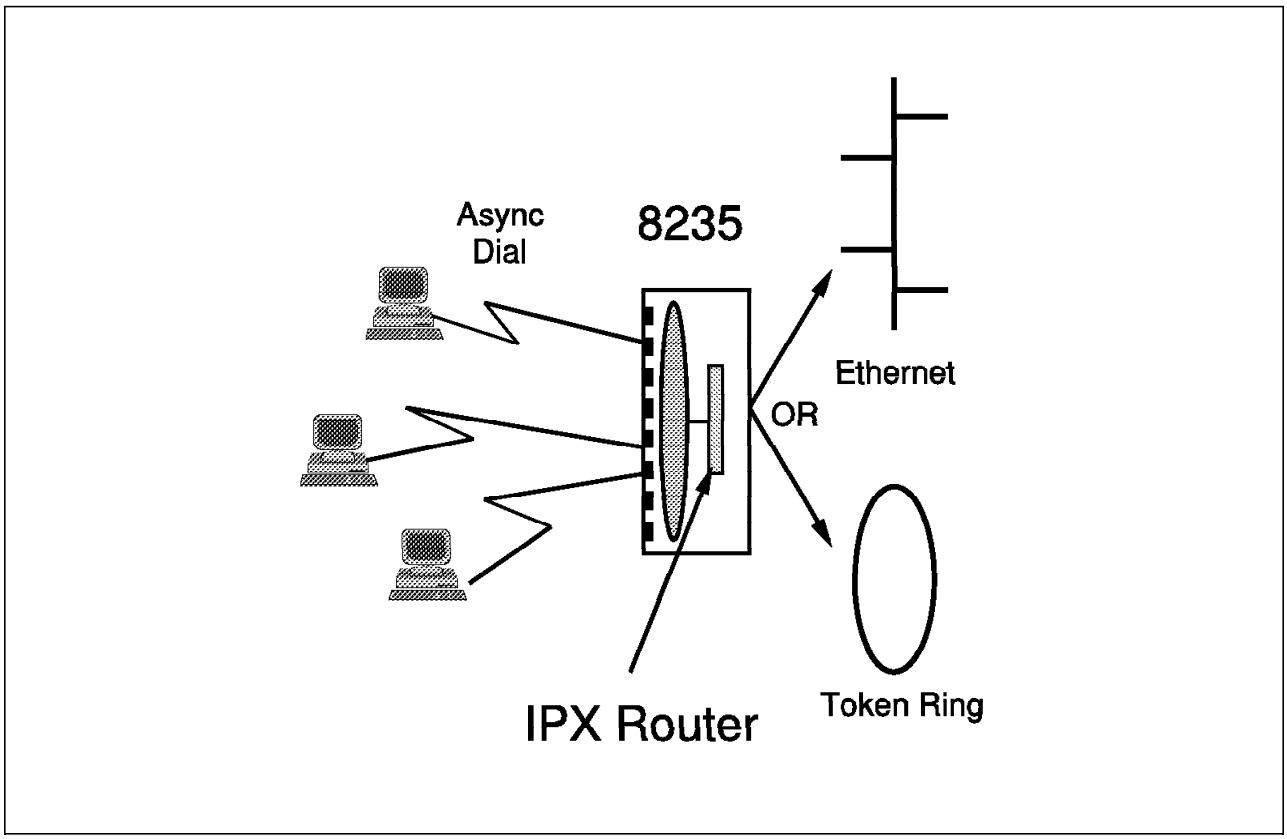

Figure 101. 8235 IPX Router

Basic IPX protocols implemented by the 8235 are:

- Internet packet exchange (IPX) providing the basic network layer transport for NetWare IPX.
- Sequenced Packet eXchange (SPX) for reliable byte stream protocol. This is used for NetWare diagnostics and for downloading code images over IPX.
- Routing information protocol (RIP) which provides a mechanism for IPX routers to exchange network topology information as needed to maintain routing tables. RIP uses a distance vector algorithm to calculate best routes.
- Service advertising protocol (SAP), which provides a mechanism for end systems to locate NetWare services. The 8235 advertises its management via SAP.

The 8235 supports dial-in routing by the remote user for IPX onto the local LAN. The network number of the dial-in port can be assigned by the administrator. If the assigned number is in use on the network when a user dials in, the box can be configured to take on of three actions: use the net number anyway, use a random number, or refuse the connection. If the dial-in client uses a non-zero node address, the server will accept it. If the client uses a zero node address, the server will provide the clients address. The 8235 supports the following IPX frame types:

- Ethernet II (Ethernet)
- 802.3 (Ethernet)
- 802.2 (Ethernet)
- SNAP (Ethernet)
- SNAP (token-ring)

- 802.2 (token-ring)

4.4.10.6 AppleTalk ARA 2.0

You can configure the 8235 as an end node or router and assign it to an AppleTalk zone.

AppleTalk protocols support zones for managing user access to network devices and services. Zones are logical names associated with networks. The network administrator chooses an AppleTalk Phase 2 default zone during the initial setup of the network. The 8235 can be placed in this default zone or in a valid Phase 2 zone in the zone list.

Note: The 8235 supports AppleTalk Phase 2 networks only.

The 8235 may appear as one of the following on the AppleTalk network:

- A node
- A router

End Nodes: Apple Remote Access (ARA) software allows Apple users to connect to an AppleTalk network through a modem/serial link. The ARA remote client calls a locally attached ARA server. The ARA server provides the client with access to LAN resources (electronic mail, file servers, printers, and network applications).

An ARA server operating in end-node mode is responsible for forwarding packets sent to and from the ARA client. The ARA server examines packets sent on the network. If the destination is the ARA server or a remote ARA client, or it is a broadcast packet then the server accepts the packet. If the destination is a remote ARA client the server sends the packet across the serial link to the remote client.

AppleTalk remote access protocol (ARAP) requires the ARA server to prevent broadcast routing table maintenance protocol (RTMP) information from being forwarded to the client over the serial link. The ARA client does not need the RTMP broadcast information.

A packet sent from an ARA client to a user on a different network is forwarded by the ARA server to a router using the *most recent router* method. This method is used because the ARA server operating in end-node mode is not a router and must forward the packet based on the most recent information it has received about the destination. The most recent router method does not ensure the packet is routed to its destination by the fastest available path. The ARA server in end-node mode provides for easy configuration. An end node does not require a new (additional) network number and is less intrusive on large networks because it does not broadcast RTMP packets as a router does.

Advantages using the 8235 in end-node mode

- Easy setup
- Network number not required
- Serial link traffic could be minimized
  - NBP broadcasts not destined for the client are not forwarded.
  - RTMP packets are not forwarded (the 8235 is not a router in this mode).

The end node implementation of ARAP in the 8235 is compatible with Apple′s ARAP implementation. When the 8235 is configured to function as an end node,

the 8235 forwards the data packets to and from the ARA clients in the same way as an ARA server.

With the 8235 functioning as an end node all 8235s on the network can be assigned to one zone in the Phase 2 zone list with the “8235 appears in” option. Network administrators would only need to access one zone to find all the 8235s on the network.

8235 ARA clients can be assigned to a different Phase 2 zone. Assigning ARA users to a different zone can help reduce NBP broadcasts over the serial link if the zone chosen does not receive many NBP broadcasts. This can significantly improve performance over the serial link.

ARA Routers: An ARA server in router mode acts as a router between two networks; the local internetwork on which the server resides and a network into which remote clients are assigned. In contrast to an ARA end-node server, which makes a remote ARA client a node on the network, an ARA server in router mode makes an ARA client a node on a separate dial-in (remote) network. The dial-in network has as many nodes as there are ARA clients connected to the server. This ARA client network can be assigned to any zone on the network including a zone in the Phase 2 zone list, or a newly created zone.

When acting as a router, the ARA server maintains complete zone and routing tables of the internetwork in memory. When a node on the internetwork sends a packet, the router examines the packet header and determines the destination by checking the routing table. If the destination is a remote ARA client the packet is routed to the dial-in network and sent to the node number of the ARA client.

When a packet is sent from an ARA client to the local network over the serial link, the ARA server uses its routing table information to route the packet to its destination by the most efficient path in the routing table.

An ARA server configured as a router can isolate the ARA client from AppleTalk broadcast packets by permitting the client to be located in a dial-in zone. This improves performance over the serial link, as only broadcasts into the dial-in zone are sent over the serial link.

Advantages Using the 8235 in Router Mode: The 8235 can be configured to function as a conforming router or as a seed router. A conforming router obtains routing information from other routers on the network. A seed router provides the routing information to the other routers on the network.

The 8235 operating in router mode provides some advantages:

- AppleTalk broadcast packets sent over the remote link can be limited by placing the remote link into a dial-in zone. Only broadcasts into that zone are sent over the link.
- The 8235 knows the fastest route to all networks and will route client packets by the most efficient path.
- The 8235 can be assigned to a different zone in the Phase 2 zone list. By assigning all 8235s to a particular management zone, network administrators only need to access one zone to find all 8235s on the network.
- The 8235 can isolate ARA clients from the rest of the internet by assigning clients to a dial-in zone. Each client has a different node number in this

zone. The dial-in zone may be a newly-created zone. It does not have to be in the Phase 2 zone list. All dial-in clients can be placed into this dial-in zone. Network administrators can monitor dial-in activity by monitoring this zone.

- Network and zone information is configurable for ARA clients.
- For LAN-to-LAN connections the 8235 must be in router mode.

IP Information: IP forwarding allows the 8235 to provide IP address assignment for dial-in clients. The clients IP address must be part of the Ethernet/IP network. Other IP hosts on the network communicate with the dial-in users through the 8235. The 8235 responds to Address Resolution Protocol (ARP) requests that are destined for a client IP address. This is referred to as *proxy ARP*. When an IP host requests an 8235 client IP address, the 8235 responds to the host with its own Ethernet address, specified on the IP configuration page. The 8235 accepts client packets and forwards the packet to the correct IP client/address.

IP packets are routed across an AppleTalk network by means of encapsulation. The 8235 sends IP packets to Macintosh dial-in clients by encapsulating the IP packet within an AppleTalk packet. The 8235 forwards IP packets from an ARA client to an IP host by de-encapsulating the IP packet.

The 8235 ARA dial-in clients appear as if they are directly connected nodes within the IP network. The IP host and the dial-in client are not affected by the fact that their packets are being routed through the 8235.

The Macintosh dial-in client uses the name binding protocol (NBP) to search for an IPGATEWAY device type in a specified zone. Since the 8235 is the ARA server for the client, the 8235 processes all of the client′s AppleTalk packets and checks its configuration to see if it is configured as an IP gateway for that zone. If it is the 8235 responds to the Macintosh dial-in client that it is an IPGATEWAY.

The dial-in client sends a kinetics internet protocol (KIP) command to the 8235 asking for an IP address. The 8235 responds with the dial-in client′s IP address, subnet mask, broadcast address and the IP address of the name server.

To communicate with an IP host the user must have an IP address. IP addresses are assigned to a Macintosh client as follows:

- Per user: When a dial-in connection is made, the 8235 checks the user list to see if there is a user IP address. If there is a user IP address in the user list, the 8235 assigns this IP address to the client.
- Per port: If there is no IP address in the user list, the 8235 assigns the port IP address to the client.

4.4.11 Security

The 8235 provides several security features. Passwords for both dial-in and LAN-to-LAN connections are automatically encrypted. User lists store user profiles which include user names, passwords, permissions and dial-back. If dial-back is selected in a user profile, the 8235 will hang up after the dial-in or LAN-to-LAN connection is established and then call the user back at a configured (fixed dial-back) number or at a number entered by the user when the connection was established (roaming dial-back). Unauthorized access to the 8235 device configuration or user list can be prevented by assigning the 8235 an

administrator password. This password is stored in the 8235 device configuration information, not in the user list.

The 8235 has a unified security architecture which allows any security server on the LAN to be used to authenticate any user regardless of the protocol being used. This allows a centralized security method to be used for all authentications. 8235 Version 2.0 code or later supports three authentication databases:

- 8235 User List
- NetWare Bindery
- SecurID ACE/Server

The 8235 prompts separately for the user name and password for each method of authentication. Thus, more than one security method can be used simultaneously. SecurID could be used to authenticate an individual user who then logs into a NetWare Bindery group and is granted the access privileges associated with that group. Because the user protocol does not matter, the NetWare Bindery could be used to authenticate an Apple Remote Access (ARA) Version 2.0 dial-in user.

4.4.11.1 8235 User List

Using the 8235 Management Facility a user list can be created, edited, and then saved to a file or loaded into the 8235. The 8235 user list stores the names, passwords, and permissions of users authorized to dial-in to or out of the network or to connect to another network. User lists are stored in battery backed-up RAM in the 8235. Each 8235 can have a different user list or one user list can be downloaded to multiple 8235s. The NetWare Bindery or SecurID is recommended if there are more than 500 users.

4.4.11.2 Using the NetWare Bindery

The NetWare Bindery is a database that resides on a NetWare server. This database contains profiles of network users that define each user′s NetWare name, password, dial-back number, and the permissions to use one or more 8235 functions such as dial-in, dial-out or LAN-to-LAN.

When bindery authentication is enabled, it replaces the 8235 user list authentication.

With bindery security enabled the bindery services utility can be used to create bindery groups for dial-in, dial-out, and LAN-to-LAN users. The group names are 8235_DIALIN, 8235_DIALOUT, and 8235_LAN-to-LAN. The bindery dial-in user groups are used when a user dials into the network using a NetWare name and password. The 8235 logs in to the NetWare server with this user name and password and then logs out. If the 8235 logon to the server was successful, the 8235 allows the user to access the network through the 8235.

4.4.11.3 Bindery and Apple Remote Access (ARA)

To use the bindery, ARA Version 2.0 users must have the 8235 Security Module in their Macintosh systems Extensions folder in the system folder. This module supplies a security drop-in, which provides 8235 password encryption (thereby allowing bindery security to work with ARA Version 2.0.)

4.4.11.4 Using SecurID

Security Dynamics, Inc. manufactures two security solutions that are compatible with the 8235. The first is a multiport, stand-alone device that can be inserted between the 8235 and the modem. This solution requires no particular configuration of the 8235. The device dialing in must be capable of handling the authentication dialog.

Macintosh users who have the external SecurID client box installed for their 8235 can still use their command control languages (CCL) as before; however, SecurID should not be enabled in the 8235 Management Facility, as this will trigger the 8235 internal SecurID client.

SDI's second security solution is the Security Dynamics ACE/Server, which is a system of server and client software and SecurID cards. Once enabled, SecurID authentication is used for all protocols (IP, IPX, NetBEUI, 802.2 LLC, and ARA).

The 8235 can use SecurID to protect its serial ports from unauthorized dial-in access. SecurID authenticates users and may be used in conjunction with the 8235 user list or the NetWare Bindery. See Figure 102 for the SecurID configuration.

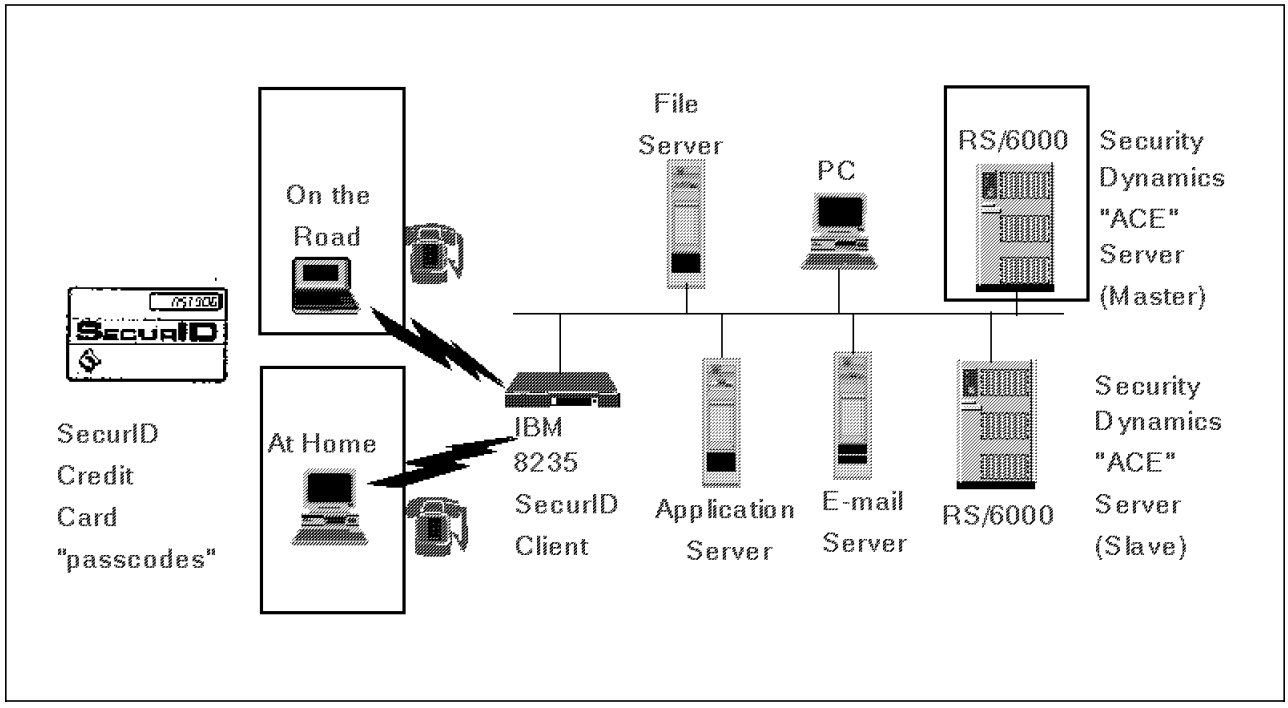

Figure 102. 8235 Security System

SecurID authentication is not required of dial-out users, users managing the 8235 with the command shell, or users managing the 8235 with the 8235 Management Facility. SecurID does not protect the 8235 from dial-out, LAN-to-LAN, or local area network shell access. If the 8235 is using SecurID authentication, incoming LAN-to-LAN connections are not permitted.

The components of a full implementation of SecurID are as follows:

- SecurID server software

This software runs on a UNIX machine. User data protocol (UDP) is used to communicate with the client software running on the 8235. This server software is purchased from Security Dynamics, Inc.

- SecurID client

  This is the component running on the 8235 that communicates with the SecurID server via UDP. It is compatible with SecurID server software Version 1.1 or later.

- SecurID card

  This component is a card that provides the user with a passcode number needed to access the SecurID server.

- Dial-in client software

  This is the standard 8235 Remote Dial-in Client Version 2.0 or later for PC users or Apple Remote Access (ARA) Client Version 2.0 or later for Macintosh users.

4.4.12 The Activity Logger

The activity logger runs under Microsoft Windows and DOS. It provides information about 8235s and their dial-in activity on the network.

The logger carries out the following tasks:

- It records the dial-in activity of the 8235 on the network.
- It notifies the network administrator of 8235 activity according to a set of priorities and classes selected by the administrator.

The 8235 logs its activity to another station using a mechanism of (SNMP) called a trap. Each time the 8235 logs an event, it sends a trap message to its trap host.

The trap host can be one of the following:

- A workstation running the 8235 Activity Logger
- An IP host with an SNMP manager

There can only be one trap host associated with an 8235 at any given time. This trap host is configured in the 8235 Management Facility on the SNMP configuration window. There are two host types to choose from: None and IP.

If you select IP, then you can also specify the IP address of the trap host. This IP host must be an SNMP manager and have some facility for displaying SNMP trap messages if it is to be used as the activity logger. For example, this could be a NetView for AIX management station.

If you select None, then the trap host address cannot be specified via the 8235 Management Facility. Instead, once the 8235 activity logger (which runs on top of IPX) selects an 8235 as a device to be logged to that workstation, the selected 8235 sends all of its trap messages to that workstation. If an 8235 is selected on one activity logger workstation while another Activity Logger workstation is the current trap host, the new workstation becomes the new trap host. This provides flexibility in case a trap host goes down because it is easy to switch over to a backup host.

4.5 Distributed Console Access Facility (DCAF)

IBM Distributed Console Access Facility (DCAF) has been enhanced to support OS/2 Presentation Manager applications, in addition to the already-supported OS/2 or PC/DOS full screen text-mode applications.

Other improvements or new functions include: one-to-many support, mouse support, new security functions, DOS file transfer capabilities and improved performance. In addition, DCAF 1.1 provides APIs for accessing DCAF functions from other applications.

4.5.1 Product Positioning

DCAF 1.1 is the first IBM stand-alone remote console function that supports DOS and OS/2 stations. In addition to full-screen text-mode DOS or OS/2 applications, it also supports OS/2 Presentation Manager applications.

It has also been enhanced to allow encryption and decryption of data, file transfers, user ID and passwords, as well as logging all connections to allow tracking and auditing.

DCAF 1.1 supports connection via switched asynchronous lines as well as SNA LU 6.2 and LU 6.2 to NetBIOS for interactions with stations on the LAN.

DCAF 1.1 allows a variety of ways for central site control and monitoring of workstations distributed across a LAN/WAN or through an SNA network, as well as for performing functions from a remote location, such as a user′s home, to allow the user to take over a workstation in the network.

4.5.2 Highlights

- Enhances management of distributed intelligent workstations through remote console functions.
- Improves network administration/support personnel productivity through central-site control and monitoring capabilities.
- Enables migration and growth from host system environments to systems based on distributed intelligent workstations by providing powerful central Help Desk control and monitoring.
- Supports data encryption and decryption, user ID and passphrase validation by a *third party* authentication server.
- Allows the remote control of workstations installed at locations where physical presence is difficult or inconvenient.

4.5.3 Description

The following items point out the benefits and uses of DCAF.

4.5.3.1 Systems Management

DCAF 1.1 provides the capability for one personal computer workstation to be controlled by another personal computer workstation.

Once a session is established, the target workstation receives all keystrokes entered at the controlling workstation, with the resulting screen images displayed on both the controlling and the target workstation. The target workstation also receives mouse commands from the controlling workstation.

The controlling and the target workstations can switch from this operating mode to a monitoring mode in which the target user controls the keyboard and mouse, but the screen images are echoed on the controlling workstation.

DCAF 1.1 can be used to control most text-mode or PM applications running on an OS/2 station, or full-screen text-mode applications running on a DOS station connected to a LAN. This function is transparent to the application being run.

The two OS/2 workstations communicate via an LU 6.2 (SNA) connection, either directly via telecommunications, across a LAN, or across an SNA backbone. Switched asynchronous links can be used instead of SNA links, or as a backup when the SNA link is not available.

In the IBM Token-Ring environment, the gateway function of the DCAF 1.1 will act as an LU 6.2 catcher, and will communicate with the target station on the ring via NetBIOS protocols. The target station on the ring can either be a DOS-based station or another OS/2-based station that is running DCAF. If connection to a target station on a distributed LAN is required, an alternative to the LU 6.2 connection is a switched asynchronous link from the controlling station to the distributed console gateway function on the LAN. This then provides access to any DOS or OS/2 target station on the logical LAN that has DCAF installed.

New features that provide additional security with the product are:

IBM 4755 Cryptographic adapter support for encryption and decryption for file transfer across the network

User ID and passphrase validation through an authentication server

Logging all connections routed through a DCAF Gateway on the same gateway to allow tracking and auditing

If SNA TIC (token-ring interface connector) attached 37XX control unit is on the ring, direct conversations between stations on the SNA backbone and SNA stations on the ring are supported. The ability to run a workstation remotely from anywhere in the SNA network provides a powerful tool that enhances the system management capability of the distributed intelligent workstation environment.

4.5.3.2 User Productivity

Productivity of administration and support personnel is greatly improved with this remote capability. Use of DCAF from a central site facilitates network management, network administration, and application assistance across an SNA network. Examples of how DCAF 1.1 can be used are:

- Remote control of LAN Manager and LAN Server stations to assist LAN management and LAN administration
- Remote control of most PM or full-screen text-mode applications running on the controlled or target workstation (these could be PS/2-based service consoles, business applications or industrial control applications)
- Monitoring end-user displays and selections for Help Desk environments
- Remote problem determination for trace and dump analysis for remote debugging
- Online education and assistance for newly installed applications

DCAF 1.1 allows the transfer of files between two workstations. This is useful for new code loads, transfer of trace and dump data for problem determination, and distribution of network administration data. Faster problem determination and improved access to user data not only improves system availability, but reduces network support personnel workload. Installing the product documentation online gives instant access to product information for ease of use and efficiency.

4.6 LANHOP/6000

IBM Local Area Network Home Office Program (LANHOP/6000) Version 1.0 provides a connection from a user′s remote computer to a LANHOP/6000 gateway allowing users to access their office LAN workstation from their home or on the road.

The LANHOP/6000 package provides a gateway server program that runs on a RISC System/6000. Also provided is OS/2 Presentation Manager and DOS programs that run on the remote workstation. These programs establish a modem connection with the LANHOP/6000 gateway. Host and/or LAN applications can then be accessed remotely and executed.

This provides users access to their workstations from home or anywhere there is a telephone. The remote workstations can be an OS/2 workstation with Transmission Control Protocol/Internet Protocol (TCP/IP), DOS workstation with TCP/IP, or a DOS/Windows workstation with TCP/IP. A RISC System/6000 running AIX Version 3 with TCP/IP may also be utilized as a remote workstation. Host access is also provided to VM running TCP/IP or MVS running TCP/IP.

4.6.1 Highlights

- Provides the ability to access LAN functions remotely
- Supports TCP/IP, NetBIOS, 802.2
- Runs client/server applications over dial-up asynchronous lines
- Provides double-password scheme
- Provides user-friendly OS/2 PM and DOS programs for a remote computer
- Context sensitive helps
- Requires no callback
- Offers PM and full screen environments
- Allows remote use of job-site network or mainframe
- Provides front-end configuration at the remote computer
- Supports wide range of IBM or OEM modems up to 57.6 Kbps
- Offers gateway/server system based on RISC System/6000
- Allows NetBIOS operations on TCP/IP protocols
- Utilizes TCP/IP File Transfer Utility
- Provides remote execution

4.6.2 Description

Remote LAN Access Connectivity: Finally, a solution to the problem of providing users access to their office workstations from home or on the road is LANHOP/6000.

LANHOP/6000 provides cross platform remote access capability to the following:

- Connection to a PC on the LAN and mirroring of a full screen session (token-ring or Ethernet)
- Connection to a mainframe host session

Remote connecting platforms may be OS/2, DOS, or DOS/Windows.

Growth Enablement - Connectivity Improvements: LANHOP/6000 provides a secure connection from a user′s remote computer to a LANHOP/6000 gateway allowing users to access their office LAN workstation from home or on the road. The programs at the remote end will establish a modem connection with the LANHOP/6000 gateway. The remote programs and the LANHOP/6000 gateway program cooperate to establish a secure link via the SLIP component of TCP/IP. Once the secure link is established, the OS/2 and LANHOP/6000 programs are no longer in the picture and the remote OS/2 or DOS TCP/IP is left connected via SLIP to the network via the LANHOP/6000 gateway.

Growth Enablement - Reduced Environmental Requirements: By providing remote LAN and host access, LANHOP/6000 enables a work-from-home concept which could ultimately reduce space and facilities requirements. This reduction could provide significant cost savings.

End-User Productivity - Improved Worker Productivity: Providing remote access to a user′s office workstation enhances the productivity of a worker who is working remotely. The user may access his/her job-site network and thereby have access to any of the programs provided on that network.

Fencing - Data Access Security: LANHOP/6000 provides the ability to establish group authorization files that will restrict a user′s access to connections (host names). Each user will be authorized to a particular group which in turn will list the connections (host names) available to that group.

Network Management - SNMP Agent: LANHOP/6000 provides management of your gateway network via Simple Network Management Protocol (SNMP). This capability allows you to collect and access statistics on the status of the network interfaces, incoming and outgoing traffic, port utilization, CPU utilization, and error message generation.

Transaction Accounting: LANHOP/6000 provides full collection and mirroring of accounting functions per user session. The following data is provided:

- Start time
- Stop time
- Modem speed
- Gateway name
- Dialing platform (such as DOS, OS/2, and DOS Windows)
- Port number
- User ID

An application programming interface (API) will be provided to access this information.

4.6.3 Technical Description

LANHOP provides cross-platform remote client access to a user's LAN network. Dial-in to a LANHOP gateway server allows access from any remote location with no call-back required. The remote workstations may consist of an OS/2 workstation with TCP/IP, or a DOS or DOS/Windows-based workstation with TCP/IP. Host and/or LAN applications can then be accessed remotely (mainframe access requires TCP/IP for VM or MVS).

The programs on the remote computer will establish a modem connection with the LANHOP gateway server. The remote programs and the LANHOP gateway server cooperate to establish a secure link via the SLIP component of TCP/IP. Once the secure link is established, the LANHOP gateway server program is no longer in the picture and the remote OS/2 or DOS TCP/IP is left connected via Serial Link Internet Protocol (SLIP) to the network.

The remote client programs each have a user-friendly interface with context sensitive helps. Functions provided from the remote client include NetBIOS capability which allows each platform access to a LAN domain for logon as well as mainframe and PC connection, TCP/IP file transfer, remote execution and user configuration.

The LANHOP gateway server provides data access security via fencing. Fencing allows the establishment of group authorization files that will restrict a user's access to connections. Every TCP/IP packet is inspected to verify that user authorization to an intended node is allowed. LANHOP also provides double password protection and access authentication to a gateway server. In addition, full collection and mirroring of accounting functions per user session is provided. Accounting data such as start and stop times, modem speed, gateway name, dialing platform, user ID and port number is collected.

Appendix A. Special Notices

This publication is intended to help customers, systems engineers, services specialists, and marketing specialists understand LANs and IBM LAN solutions and architectures for planning and support purposes. The information in this publication is not intended as the specification of any programming interfaces that are provided by the products mentioned in this book. See the PUBLICATIONS section of the IBM Programming Announcement for IBM LAN products 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.

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.

The following terms are trademarks of the International Business Machines Corporation in the United States and/or other countries:

ACF/VTAM	ADSTAR
Advanced Peer-to-Peer Networking	AIX
AIX/6000	AIXwindows
AnyNet	APPN
AS/400	AT
CICS	CICS OS/2
CICS/ESA	CICS/400
CICS/6000	DatagLANce
DB2	DB2/2
DB2/400	DB2/6000
ES/9000	ESCON
EtherStreamer	Extended Services for OS/2
Extended Services	FFST/2
IBM	IMS

IMS/ESA	LAN Distance
LANStreamer	MVS/ESA
NetFinity	NetView
Nways	Operating System/2
OS/2	OS/400
Personal System/2	Portmaster
Presentation Manager	Processor Resource/Systems Manager
PS/2	RISC System/6000
RS/6000	RT
RXR/2	S/370
S/390	SAA
SystemView	System/390
System/370	Trouble Ticket
VM/ESA	VTAM
Workplace	400

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc.

PC Direct is a trademark of Ziff Communications Company and is used by IBM Corporation under license.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Microsoft, Windows, and the Windows 95 logo are trademarks or registered trademarks of Microsoft Corporation.

Advantis	Advantis
Apple, AppleTalk, EtherTalk, LocalTalk, Macintosh, TokenTalk	Apple Computer, Incorporated
ARCnet	Datapoint Corporation
AT&T, ESS, SLC	American Telephone and Telegraph Company
Attachmate	Attachmate Corporation
Banyan, VINES	Banyan Systems, Incorporated
Cisco	Cisco Systems, Incorporated
Compaq	Compaq Computer Corporation
CompuServe	CompuServe, Incorporated and H&R Block, Incorporated
Dasher	Data General Corporation
DCE	The Open Software Foundation
DEC, DECnet, Digital, ULTRIX, VT100, VT200	Digital Equipment Corporation
DMS-100	Northern Telecom Limited
EtherLink	3Com Corporation
Hayes	Hayes Microcomputer Products, Incorporated
Hewlett-Packard, HP, OpenView	Hewlett-Packard Company
HYPERchannel	Network Systems Corporation
IDNX	Network Equipment Technologies, Incorporated
Intel, i386, Pentium, 386, 486, 80386	Intel Corporation
IPX, LANalyzer, NetWare, Novell	Novell, Incorporated
LANtastic	Artisoft, Incorporated
Lotus, Lotus Notes	Lotus Development Corporation
MOSS	MOSS Systems, Limited
MS-DOS	Microsoft Corporation
NDIS	3Com Corporation and Microsoft Corporation

NEC	NEC Technologies, Incorporated
Network File System, NFS, Solaris, SunOS	Sun Microsystems, Incorporated
NT	Northern Telecom Limited or Microsoft Corporation
Proteon	Proteon, Incorporated
Qualitas	Qualitas
Quarterdeck	Quarterdeck Corporation
SCO	The Santa Cruz Operation, Incorporated
SCSI	Security Control Systems, Incorporated
Siemens	Siemens Company
Sniffer Network Analyzer	Network General Corporation
SPARCstation	SPARC International, Incorporated
Stac, Stacker	Stac Electronics
STB	STB Systems, Incorporated
SynOptics	SynOptics Communication Incorporated
Toshiba	Toshiba Corporation
Wellfleet	Wellfleet Communications, Incorporated
Win32s	Microsoft Corporation
X-Windows	Massachusetts Institute of Technology
X/Open	X/Open Company Limited
Xerox, Xerox Network Systems, XNS	Xerox Corporation
3Com	3Com Corporation
386MAX	Qualitas, Incorporated

Other trademarks are trademarks of their respective companies.

Appendix B. Related Publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

B.1 International Technical Support Organization Publications

For information on ordering these ITSO publications see “How To Get ITSO Redbooks” on page 271.

- *The IBM 6611 Network Processor as an IP Router*, GG24-4064-00 (available on CD-ROM, SK2T-6022)
- *8260 Multiprotocol Intelligent Switching Hub*, GG24-4370-00
- *IBM 2220 Nways BroadBand Switch: Concepts and Products*, SG24-4307-01

A complete list of International Technical Support Organization publications, known as redbooks, with a brief description of each, may be found in:

International Technical Support Organization Bibliography of Redbooks, GG24-3070-14.

B.2 Other Publications

These publications are also relevant as further information sources:

- *IBM 6611 Network Processor: Introduction and Planning Guide*, GK2T-0334-05
- *IBM Multiprotocol Network Program: User's Guide*, SC30-3559-02
- *Nways MRNS V1R3 Protocol Configuration and Monitoring Reference*, SC30-3680-02
- *IBM 2210 Nways Multiprotocol Router Maintenance Information*, SY27-0345-02
- *IBM 2210 Nways Multiprotocol Router Planning and Setup Guide*, GA27-4068-02
- *Nways MRNS Software User's Guide*, SC30-3681-02
- *Nways MRNS Event Logging System Messages Guide*, SC30-3682-02
- *The Nways MRNS Protocol Configuration and Monitoring Reference*, SC30-3680-02
- *Planning for the System/390 Open Systems Adapter Feature*, GC23-3870-00
- *Internetworking with TCP/IP Volume I: Principles, Protocols, and Architecture*, SC31-6144-00

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 URL http://www.redbooks.ibm.com/redbooks.

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 lists of redbooks:

  ```
  TOOLS SENDTO WTSCPOK TOOLS REDBOOKS GET REDBOOKS CATALOG
  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:

  ```
  TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ITSOREGI 1996
  ```

  For a list of product area specialists in the ITSO:

  ```
  TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ORGCARD PACKAGE
  ```
- **Redbooks Home Page on the World Wide Web**

  http://w3.itso.ibm.com/redbooks/redbooks.html
- **IBM Direct Publications Catalog on the World Wide Web**

  http://www.elink.ibmlink.ibm.com/pbl/pbl

  IBM employees may obtain LIST3820s of redbooks from this page.
- **ITSO4USA category on INEWS**
- **IBM Bookshop** — 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.ibmlink.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.

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:

- **IBMLINK**

  Registered customers have access to PUBORDER to order hardcopy, to REDPRINT to obtain BookManager BOOKs

- **IBM Bookshop** — send orders to:

  usib6fpl@ibmmail.com (United States)
  bookshop@dk.ibm.com (Outside United States)

- **Telephone orders**

1-800-879-2755	Toll free, United States only
(45) 4810-1500	Long-distance charge to Denmark, answered in English
(45) 4810-1200	Long-distance charge to Denmark, answered in French
(45) 4810-1000	Long-distance charge to Denmark, answered in German
(45) 4810-1600	Long-distance charge to Denmark, answered in Italian
(45) 4810-1100	Long-distance charge to Denmark, answered in Spanish

- **Mail Orders** — send orders to:

IBM Publications	IBM Direct Services
P.O. Box 9046	Sortemosevej 21
Boulder, CO 80301-9191	DK-3450 Allerød
USA	Denmark

- **Fax** — send orders to:

1-800-445-9269	Toll-free, United States only
45-4814-2207	Long distance to Denmark

- **1-800-IBM-4FAX (United States only)** — 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

- **Redbooks Home Page on the World Wide Web**

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

- **IBM Direct Publications Catalog on the World Wide Web**

  http://www.elink.ibmlink.ibm.com/pbl/pbl

- **Internet Listserver**

  With an Internet E-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an E-mail note to announce@webster.ibmlink.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.

IBM Redbook Order Form

Please send me the following:

Title	Order Number	Quantity

- **Please put me on the mailing list for updated versions of the IBM Redbook Catalog.**

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.

DO NOT SEND CREDIT CARD INFORMATION OVER THE INTERNET.

Index

Numerics

10Base-T (UTP) 243
10Base2 (Thin Ethernet) 243
10Base5 (Thick Ethernet) 243
2210 Configuration Program
  Configuration Window 93
  hardware requirements 93
  Navigation Window 93
  software requirements 93
2210 Nways Multiprotocol Router 84
  indicators 86
  models 85
  remote access 88
  reset button 87
  supported networks 87
2217 Network Node (NN) 202
2217 Nways Multiprotocol Concentrator 200
2217 Remote Control Utility 208
3088 161
3172 159
  gateway 159
  ICP Mode 160
  IP Channel Communications Program 164
  LAN-To-Host Mode 159
  link station 162
  Offload Mode 162
  relative adapter number 162
  SNA Communications Program 164
  subchannel 162
  WAN-to-Host Mode 165
3174
  alternate path 175
  APPN 170
  APPN network node 171
  configuration support 169
  dependent LU 174
  Gateway 168
  independent LU 174
  local SNA gateway 171
  Multiple Logical Terminal (MLT) support 177
  remote SNA gateway 172
  TCP/IP Enhancements RPQ 8Q1041 176
3745 20
  9x0 gateway 181
  Ethernet LAN 186
  IBM token-ring LAN 186
  IP router 186
  link station 185
  LLC Type 2 185
  MAC address 185
  SNA 184, 186
  SNA connection establishment 185
  TCP/IP 186, 187
3746 181
  gateway 181
  IP router 190
  Network Node (NN) 182
  TCP/IP 190
6611 20
  Model 125 121
  Models 145 and 175 121
6611 Router
  Model 120 120
  Model 125 121
  Model 145 121
  Model 175 121
  new adapters 123
802
  8235 247
802.2
  LAN Distance 231
  LANHOP/6000 261
8229 LAN Bridge 20, 61
  benefits 80
  full operational mode 67
  green status LED 61
  hardware reset 61
  minimal operational mode 67
  numeric display 61
  parameters 67
  power light-emitting diode (LED) 61
  Utility Program 66, 68
  yellow status LED 61
8235 226
  8235 hardware description 240
  activity logger 258
  AppleTalk 253
  code structure 243
  communication options 246
  description 233
  DIALs client software 234
  IP Traffic 249
  IPX Traffic 251
  LAN connection 241
  LED, network and port status 241
  LED, power status 241
  management facility 239, 256, 258
  microcode 244
  models summary 246
  NetBIOS and 802.2 247
  remote node 225
  security 255
  supported protocols 247
  system components 234
  technical description 233
8235 Management Facility 234

8235 Management Facility for Windows 236
8Q0800 172

A

access control system 106
access node 27
activity logger 258
ACTLU command 212
adapter activation 206
adapter deactivation 206
adapter types 183
adaptive BIND pacing 205
Adaptive Rate-Based (ARB) 150
adaptive session-level pacing 205
Adaptive Source-Routing Transparent Bridge (ASRT) 38, 114
additional routing 85
address conversion 74
Address Conversion Chart 75
address resolution protocol (ARP) 78, 188, 249, 250
ADSTAR Distributed Storage Manager (ADSM) 157
Advanced Peer-to-Peer Networking (APPN) 203, 217
Advanced Program-to-Program Communication (APPC) 170
AEA 176
age out timer 142
agent 5, 6
AIX 192
AIX SNA Gateway/6000 195
alerts 175
all-routes broadcast 128
anonymous FTP site 94
ANR (Automatic Network Routing) 149
ANSI 19
AnyNet 21, 194, 215
AnyNet APPC 32
AnyNet gateway 21, 25
AnyNet IPX 32
AnyNet/400 29
AnyNet/6000 28, 194
AnyNet/MVS 30
APPC 8, 28, 29, 194
APPC full duplex 218
APPC/APPN 25
APPLE MAC SNAps gateway 216
AppleTalk 50, 89, 125, 235
AppleTalk ARA 253
AppleTalk broadcast packets 254
AppleTalk remote access protocol (ARAP) 253
Application Programming Interface (API) 193, 231
application sessions 197
APPN 170
  High Performance Routing (HPR) 215
APPN network node 171, 172
ARA 235
ARA routers 254
ARAP (AppleTalk remote access protocol) 253
ARB (Adaptive Rate-Based) 150
archive facility 157
archive format 91
ARP (address resolution protocol) 78, 114, 188, 249, 250, 251
AS boundary routing 102
ASCII
  console 107
  emulator 88
  Telnet 178, 179
  terminal 88, 198
  terminal data stream 176
  terminal emulation 66
ASRT (Adaptive Source-Routing Transparent Bridge) 38, 114
asynchronous communications (async) 226
Asynchronous Emulation Adapter 170
ATM 159, 215
attachment cables 83
attachment module 62, 82
AUI (Thick Ethernet) 243
authentication 107
authentication server, third party 259
Autolink 16
automatic logoff 204
Automatic Network Routing (ANR) 149
availability
  3174 174
  alternate path, 3174 175
  SNA and 3745 185

B

backup switched line 115
BAN (Boundary Access Node) 150
bandwidth 113
bandwidth reservation (BRS) 89, 112
Banyan VINES Control Protocol (BVCP) 97
Banyan Virtual Networking System Protocol (VINES) 50, 97, 125
Benchmark Methodology Working Group (BMWG) 51
benchmarking 51
Berkeley Sockets Distribution (BSD) 28
BGP (Border Gateway Protocol) 126
bibliography 269
BIND 174, 199
bit inversion 75, 76
BMWG (Benchmark Methodology Working Group) 51
BNC (Thin Ethernet) 243
boot configuration database 89
boot files 89
boot processes 89
Boot PROM 244
BOOTP (boot protocol) 90, 99, 106, 250
  forwarder 106
  relay agent 106
  server 106
Bootstrap Protocol 99

Border Gateway Protocol (BGP) 126
bottlenecks 54
Boundary Access Node (BAN) 150
bracket, cable management 122
BRI module 247
bridge 3, 47, 61
  Bridge Manager 5
  configuration parameters 67
  connections 59
  hops 43
  mode 68
  number 68
bridging 18, 229, 230, 248
  LAN-to-LAN 95
  LAN-to-WAN 95
bridging switch 97
broadcast address type 101
broadcast packets 59, 254
broadcast storms 26
BRS (bandwidth reservation) 89, 112
BSD (Berkeley Sockets Distribution) 28
BSD Compress-LZW 96
BVCP (Banyan VINES Control Protocol) 97

C

cable management bracket 122
cache 142
call back 233
capacity 54
Carbon Copy 225
CBX. 231
cc:Mail 225
CCITT 19
CCL (command control languages) 257
CCU (Central Control Unit) 182
Central Control Unit (CCU) 182
central site change management 175
central site control facility 175
channel aggregation 239
channel attached 181
CLA (Communication Line Adapter) 183
class names 114
class of service (COS) 33, 202
client 152, 158
client event logging 237
client/server 104, 152, 158, 225, 261
CLP (Communication Line Processor) 183
CM/2 (Communications Manager/2) 20, 25
command control languages (CCL) 257
Common Programming Interface for Communications (CPI-C) 193, 217
CommServer 215
Communication Line Adapter (CLA) 183
Communication Line Processor (CLP) 183
communications controller 181
Communications Manager/2 18
Communications Manager/2 (CM/2) 25, 209
Communications Server for OS/2 Warp 210, 215
community 107
compatibility mode bridging 130, 139
complex bridges 38
Computer-controlled Branch eXchange (CBX). 227
configuration report server (CRS) 96, 140
configuration support 169
Configuration Window 93
congestion control 60
connect application 235
Connection File Wizard 237
connection-oriented protocols 52
connectionless-oriented protocols 52
connections
  floating virtual (FVC) 238
  juggling virtual (JVC) 238
  persistent 238
  timed LAN-to-LAN (TLC) 238
  virtual (VC) 236, 238
Connectivity
  3174 174
  LAN Distance 226, 230
  SNA and 3745 185
console port 89
control (CONT) 76
control point name 206
control sessions 197
COS (class of service) 33, 202
CP-CP session reactivation 218
CPI-C 24, 29, 32
CPNAME 185
cross domain support 197
CRS (configuration report server) 96, 140
CSMA/CD 70
customized device drivers 225
CUT mode 178

D

DA (destination address) 76, 78
DACTLU command 212
data access security, fencing 262
data compression 33, 35, 95, 201
data distribution 164
data flow control 199
Data Link Connection Identifier (DLCI) 130
Data Link Connector (DLC) 130
data link layer 37
data link switching (DLSw) 99, 109, 117, 119, 140, 147, 149
data service unit (DSU) 73
data stream 176
data terminal element (DTE) 61
Datagram Delivery Protocol (DDP) 49
DCAF 225
  description 259
  protocols 260
  remote console 259
  remote control 224

DCAF *(continued)*
security 260
systems management 259
DCAF (distributed console access facility) 259
DCE (Distributed Computing Environment) 30, 73
DDP (Datagram Delivery Protocol) 49
DECnet 50, 125
DECnet IV 97
DECnet V 97
decryption 259
dedicated LU 216
default gateway 105
default router 105
default subnet gateway 105
default subnetwork gateway 100
Deflate - LZ77 95
delta technology 239
deny mode 132
dependent LU 174, 204
Dependent LU Requester (DLUR) 150, 183, 202, 218
Dependent LU Server (DLUS) 171
designated ring 131
destination address 41
destination address (DA) 76, 78
destination service access point (DSAP) 76
device driver replacement 230
dial back 236
dial support 166
dial-in 263
dial-in access 234
dial-in channel aggregation 237
dial-on-demand 116
dial-out 223, 230
dialing platform 263
DIALs Client 234
discovery frame 43
discovery response frame 43
discovery/learning algorithm 230
Disk Operating System (DOS) 213
disk serving 164
distance connection server 229
Distance-Vector Multicast Routing Protocol (DVMRP) 105
Distributed Computing Environment (DCE) 30, 73
distributed console access facility (DCAF) 259
DLC (Data Link Connector) 130
DLC termination 110
DLCI (Data Link Connection Identifier) 130
DLCI (PVC) 103
DLSw (data link switching) 99, 109, 117, 119, 140, 147, 149
DLSw connections 112
DLSw partners 111
DLUR (Dependent LU Requester) 150, 183, 189, 202
domain support 197
double-password scheme 261
downstream LU 210
downstream protocols 215
downstream PU 210
DRAM 84, 85
DSAP (destination service access point) 76
DSPUs 184
DSU (data service unit) 73
DSU/CSU 7
DTE (data terminal element) 61
dual mode bridging 137
dual port token-ring 63
dual-ring station 154
dump analysis 260
Duplex Multipoint 169
DVMRP 98
DVMRP (Distance-Vector Multicast Routing Protocol) 105

E

EasyStart 95, 150
EGP 105
EGP (Exterior Gateway Protocol) 98, 126, 190
EIA 232 61, 66, 209
EIA 422/429 209
ELS (event logging system) 92
emulator 88, 220
encryption 259
end node 1, 47, 173, 206
end nodes 253
end-to-end timers 34
Enhanced Priority Queueing 150
enterprise specific traps 108
enterprise-wide backup 157
enveloped 48
ESCA (ESCON Adapter) 183, 184
ESCON 164
Adapter (ESCA) 183, 184
channel 159
Channel Processor type 2 (ESCP2) 184
channels 182
ESCON is a serial connection using fiberoptics. 153
Multiple Image Facility 166
establishment gateway options 168
Ethernet 182
attachment module 64
BNC (Thin Ethernet) 243
database 136
Ethernet LAN 186
LANHOP/6000 262
Thick Ethernet (AUI) 243
Thin Ethernet (BNC) 243
Version 1 70
Version 2 69, 70
event logging system (ELS) 92
event number 92
Exchange ID (XID) 110
executable memory 90

express installation 237
Exterior Gateway Protocol (EGP) 98, 126, 190

F

fault code 61
FDDI 159
FDDI single or dual-ring station 154
fencing, data access security 262
File Transfer Protocol (FTP) 162
fill-pattern 100
filter program 67, 68
filter programs 69
filtering 126
  database 41, 133
  rate 82
  zone 235
filters 72, 114
filters, LLC SAP 247
FLASH 66
flash memory 85
floating virtual connections (FVC) 238
flow control 175, 186
format manipulation 69
forwarding table 41
four-port serial adapter 123
FR Boundary Access Node (BAN) 150
frame format 74
frame format conversion 70, 76
frame forwarding 72
frame forwarding rate 82
frame header manipulation 76
Frame Loss Rate 57
frame relay 18, 164
frame relay interface 103
Frame Relay Token-Ring Bridge/DOS 6
Frame Relay Token-Ring Bridge/DOS Version 1.0 4
frame types 252
frames 48
FRFH 189
FRM 199
FTP 20
full operational mode 67
fully meshed 104
FVC (floating virtual connections) 238

G

gateway 151, 187
  3172 gateway 159
  3174 168
  administrator 214
  AnyNet 21
  APPLE MAC SNAps gateway 216
  default 105
  default subnet 105
  LAN 182
  LAN gateway 151
  LANRES gateway 164
gateway *(continued)*
  local 181
  local 3174 171
  Multihost LAN 169
  Novell NetWare 216
  PC gateways 210
  PC/3270 213
  remote 181
  remote 3174 172
  server program 261
  status utility 214
general traps 108
generic alerts 175
green status LED 61
group 92
Group Poll 169, 172
GWCON 91

H

HAP printers 178
hardware bridge 61
hardware reset 61
hardware reset button 72
Help Desk control 259
High Performance Routing (HPR) 149, 202, 215, 217
hop count 131
host identifier 48
hot-key 177
HPR (High Performance Routing) 149, 202

I

IBD (Integrated Boot Device) 89, 109
IBM 3172 Interconnect Controller 159
IBM Remote Token-Ring Bridge/DOS Version 1.0 4
ICMP 98
ICMP (internet control message protocol) 250
ICP Mode 160
IDBLK 184
IDNUM 184
IDP (Internetwork Datagram Protocol) 49
IEEE 802.2 12
IEEE 802.2 (LLC) 231
IEEE 802.3 69, 70
IEFT (Interner Engineering Task Force) 51
IGMP 104
InARP 97
inbound filters 131
independent LU 174, 175, 204
INITSELF 174
Integrated Boot Device (IBD) 89, 109
Integrated-Services Digital Network (ISDN) 226
Intelligent Engine 153
Interconnect Controller Program 160
intermediate nodes 1, 141
internal fault 62
internal IP address 99

internal modem 209
internet control message protocol (ICMP) 250
Internet Engineering Task Force (IEFT) 51
Internet Packet Exchange (IPX) 49, 125, 252
Internet Protocol (IP) 49, 125, 250
Internetwork Datagram Protocol (IDP) 49
internetworking 200
IP 207
IP (Internet Protocol) 49, 98, 114, 119, 125
  8235 and IP 249
  access control 100
  address 187
  address resolution protocol (ARP) 250
  address, internal 99
  addresses 14
  broadcast format 100
  channel 188
  download 239
  encapsulation 109
  filters 105
  qualifier list 12
  router 14
  routing 98, 145, 162
  Traffic over SNA 207
  tunneling 104
IP-MAC address 188
IPGATE 188
IPGATEWAY 255
IPHOST 188
IPX (Internet Package Exchange) 207, 231, 251
  8235 251
  LAN Distance 231
  protocol 15
  protocol broadcasts 35
  router function 251
  traffic over SNA 207
IPX (Internet Packet Exchange) 20, 21, 49, 70, 71, 89, 114, 125, 252
ISDN 85, 231

J

juggling virtual connections (JVC) 238
JVC (juggling virtual connections) 238

K

kinetics internet protocol (KIP) 255
KIP (kinetics internet protocol) 255
knee value 54

L

LAA (locally administered address) 110, 148
LAN activity indicators 62
LAN Bridge Manager/2 4
LAN Bridge Server (LBS) 140
LAN Bridging Protocol 130
LAN communication server 222
LAN Distance 226
  administration 233
  connection server 229
  connectivity 226, 230
  description 226
  packaging 229
  protocols supported 231
  remote client 228
  remote node 225
  remote-to-LAN 229
  remote-to-remote 230
  security 232
LAN Distance Connection Server 228
LAN Distance Remote 228, 229
LAN end nodes 1
LAN File Services
LAN File Services/ESA (LFS/ESA) 152
LAN File Services/MVS (LANRES/MVS) 152
LAN Gateways 151, 182
  3172 gateway 159
  3174 gateway 168
  OSA 152
  PC gateways 210
LAN Management
  DCAF 259
LAN Network Manager (LNM) 4, 7, 68, 81, 207
LAN Network Manager Agent 207
LAN Reporting Mechanism (LRM) 140
LAN Resource Extension 152
LAN segment number 68
LAN-on-Coax (Peer Communications) 170
LAN-to-LAN 223
  bridging 95
  LAN-to-LAN Wide Area Network Program 8
  over WAN (LTLW) 25
  timed LAN-to-LAN connections (TLC) 238
LAN-to-remote 222
LAN-to-WAN bridging 95
LAN/WAN combinations 123
LanConnect applets 239
LANHOP/6000 (Local Area Network Home Office Program) 261
  description 261
  Ethernet 262
  protocols 261
  SLIP 262
  token-ring 262
LANRES gateway 164
LANRES/MVS 156, 164
LANStreamer bridge 5
LANStreamer Token-Ring Bridge/DOS Version 1.0 4
latency 53
LBS (LAN Bridge Server) 140
LCC timeouts 60
LDBRG 66, 72
LED (power light-emitting diode) 61

length field 77, 78
LFS/ESA 156, 163
Library Subroutines 193
Licensed Internal Code (LIC) 169, 170
LICs (Line Interface Couplers) 183
Line Interface Couplers (LICs) 183
line printer daemon (LPD) 176, 180
link 16, 69
  activation 206
  connectivity 193, 195
  deactivation 206
  log table 214
  messages 214
  station 162, 185
  utilization 58, 201
LLC (logical link control) 79
LLC SAP filters 247
LLC Type 1 13
LLC Type 2 13, 68, 111, 174, 185
LM10 interface 27
LMI 150
LNM (LAN network manager) 81
Local Area Network Home Office Program (LANHOP/6000) 261
local bridges 39, 128, 133
local data link switching 111, 142
local domain support 197
local gateway 171, 181
local node name 206
local SNA major node 171
Local Token-Ring Bridge/DOS Version 1.0 4
local-wire 100
locally administered address (LAA) 110, 148
logging level 92
Logical Layer Control (LLC) 207
logical link control (LLC) 79
logical unit (LU) 213
logon parameters 233
logon time intervals 233
Loopback driver 12
Loopback Utility 11
Lotus Notes 225
low-entry networking (LEN) node 206
LPAR Support 156
LPD (line printer daemon) 176
LPD queue 176
LRM (LAN Reporting Mechanism) 140
LTLW 10, 11
LTLW (LAN-to-LAN over WAN) 25
LU 6.2 9, 10, 28
LU pooling 196
LU support 196
LU types 193

M

MAC address 9, 37, 131, 161, 185
MAC Filtering (MCF) 115
MacTCP 235
Maintenance and Operator Subsystem (MOSS) 182
management facility 240
management information base (MIB) 19, 66, 108
Management Services 193
manager 5, 6
mapped addresses 69
mastering 237
Max hops 101
maximum load throughput 57
maximum throughput 57
MCF (MAC Filtering) 115
Media Access Control Bridges (802.1D) 133
metric 49
MIB (management information base) 19, 66, 108
Microsoft Windows NT 225
minimal operational mode 67
MLP (Multilink PPP protocol) 236
Model 120 enhancements 120
Model Major Node 185
model upgrade 123
MONITOR process 92
monitoring mode 260
monitoring system 89
MOSPF (Multicast OSPF) 98, 104
MOSS (Maintenance and Operator Subsystem) 182
most recent router 253
MPTN (Multiprotocol Transport Networking) 21, 23
MRNS (Multiprotocol Routing Network Services) 88
MRNS Configuration Program 90
MRNS user interface 92
Multi-Interface Serial Adapter 125
Multicast OSPF (MOSPF) 98, 104
Multihost LAN Gateway 169
Multilink PPP protocol (MLP) 236
multiple data link controls (DLCs) 203
Multiple Logical Terminal (MLT) 177
Multiple Retrieve Function 150
multipoint line 172, 174
multipoint primary control 217
multiprotocol
  connectivity 125
  extensions 165
  networks 109, 202
  router bridging 18
  support 215
Multiprotocol Networking Software 165
Multiprotocol Routing Network Services (MRNS) 88
Multiprotocol Transport Networking (MPTN) 21, 207
MVS Open Edition (OE) 30

N

name binding protocol (NBP) 255
name qualifier list 12
native mode bridging 129
Navigation Window 93
NB30 interface 27

NBMA (non-broadcast multiaccess) 102
NBP (name binding protocol) 255
NCP 213
NCP SLOWDOWN 186
NCP Version 6.0 186
NCP Version 7.0 186
NDIS 234
NDIS interface 18
NetBEUI 24, 26
NetBIOS 9, 13, 21, 24, 71, 119, 207, 231
  8235 247
  data link switching 149
  DCAF 260
  frame size reduction 149
  LAN Distance 231
  LANHOP/6000 261
  names 142
  over SNA 207
  protocol broadcasts 34
  routability 26
NetView 19
NetWare 50
NetWare Access Server 225
NetWare Access Services
  remote control 224
NetWare Bindery 236, 256
network
  address 232
  congestion control 201
  ID 206
  identifier 48
  layer 59
  layer address 127
  layer protocols 49
  node 170, 173, 182, 206
  status 241
network driver interface specification (NDIS) 231
Network File System (NFS) 162
network management vector transport (NMVT) 182, 205, 216
NETX 236
new code loads 261
new port driver 237
NFS clients 152
NMVT (network management vector transport) 182, 205, 216
non-broadcast 102
non-broadcast multiaccess (NBMA) 102
non-disruptive session rerouting 34
non-translational bridge 138
non-volatile RAM (NVRAM) 90
normal response mode 142
Novell NetWare for SAA gateway 216
numeric display 61
NVR 72
NVRAM (non-volatile RAM) 90
Nways Multiprotocol Routing Network Services (MRNS) 88

O

ODI 231, 234
  LAN Distance 231
OE (Open Edition), MVS 30
offload feature. 163
Offload Mode (OS/2 Platform) 162
OPCON 91
Open Shortest Path First (OSPF) 98, 126, 190
Open Systems Adapter 1 (OSA 1) 152
Open Systems Adapter 2 (OSA 2) 152
operating mode 260
operating system 89
Operating System Subroutines 193
operational
  code 87
  microcode 69
  modes 197
  parameters 108
  software 66, 68
operator panel display 122
optical fibers 184
originating default 105
OSA 152
OSI Reference Model 1, 37
OSPF (Open Shortest Path First) 98, 126, 190
  implementation 102
  interoperability 103
out-of-band 66
outbound filters 131
outer network 229
overhead 52

P

PACING 175
Packet Processing Time (PPT) 58
packets 48
  fragmentation 239
  latency 55
  loss 60
Parallel channels 159, 164, 182
parameters 67
partially meshed 102
partners 145
passphrase 232
  encryption 232
  validation 259
path control 199
PBX 231
PC (persistent connections) 238
PC Anywhere 224
PC gateway 210
PC/3270 Gateway 213
PDU (protocol data unit) 80
Peer Communications (LAN-on-Coax) 170
peer-capable adapters 119
performance 123

performance metrics 55
permit mode 132
persistent connections (PC) 238
Personal Communications/3270 (PC/3270) 210
physical layer 71
physical unit (PU) 161, 204
piggybacking updates 238
pin reset switch 245
PING 98, 99
PLU0 198
PLU2 198
plug and play 60
pluggable processor card 159
point-to-point (PPP) 88, 101, 116, 130, 235
pooled LU 204, 213, 216
port defaults 65
port group 63
port status indicators 62
POST (power-on self-test) 87
power light-emitting diode (LED) 61
power status 241
power switching 237
power-on self-test (POST) 87
PPP (point-to-point) 88, 101, 116, 117, 130, 190, 235
PPT (Packet Processing Time) 58
pre-allocated LU 214
Predictor 96
primary link 115
print serving 164
prioritizing traffic 59
priority classes 113
Private Branch eXchange (PBX) 227
problem determination 261
Processor Resource/Systems Manager (PR/SM) 156
programmable workstation 199
protocol data unit (PDU) 80
protocol-dependent 48
Protocols
  conversion 69
  converter 204
  DCAF 260
  ID 77
  LAN Distance 231
  LANHOP/6000 261
  suite 127
proxy agent 163
PSDN 8
public switched telephone network 234

R

rack mount options 122
Rapid Transport Protocol (RTP) 149
RARP conversion 79
rated throughput 57
README 94
reassembly size 100
Receive Not Ready (RNR) 186
relative adapter number 162
reload 90
REM (ring error monitor) 96, 140, 175
remote
  access 25, 215
  bridge 8, 39, 128, 134
  client 224, 225
  client access 263
  connection 225
  console function 259
  control 224
  data link switching 111, 142, 145
  debugging 260
  gateway 172, 181
  linked 181
  logon 162
  node 224, 226
  workstation 229
Remote LAN Access 221
  8235 233
  DCAF 259
  LAN Distance 226
  LAN-to-LAN 223
  LAN-to-remote 222
  LANHOP/6000 261
  overview 220
  remote client 225
  remote control 224
  remote node 225
  remote-to-LAN 221
  remote-to-remote 221
Remote Token-Ring Bridge Program 73
remote-to-LAN 222
remote-to-remote 221
request unit (RU) 198
restart 90
RFCs 51
RIF (routing information field) 43, 44, 76, 128
Ring Address Assignment 174
ring error monitor (REM) 96, 140, 175
ring parameter server (RPS) 96, 140, 249
ring utilization counter 5
RIP (routing information protocol) 15, 98, 101, 126, 187, 190, 202, 207, 250, 252
  implementation 101
  routing 162
  updates 101
RISC System/6000 192
RNR (Receive Not Ready) 186
route information indicator 43
ROUTED 202
routed packet 52
router cache size 99
router capacity 56
router connections 59
router ID 99, 102
router latency 55

router networks 17
Routers 47
RouteXpander/2 17, 166
routing information field (RIF) 43, 44, 76, 128
routing information protocol (RIP) 15, 98, 101, 126, 187, 190, 202, 207, 250, 252
routing table 49, 126
routing table maintenance protocol (RTMP) 49, 253
routing table size 99
RPQ 8Q0800 172
RPQ 8Q1041 175
RPQ 8Q1289 176
RPS (ring parameter server) 96, 140, 249
RS449 88
RTIC Family Communications Adapters 7
RTMP (routing table maintenance protocol) 49, 253
RTP (Rapid Transport Protocol) 149
RU (request unit) 198

S

S/390 host 20
S/390 I/O Subsystem Interface 153
S1 service port 122
SA (source address) 41, 76, 78
SAP (service advertising protocol) 15, 185, 207, 252
screen mapping 228
SCSI port 122
SDDLU (self-defining dependent LU) 218
SDLC (synchronous data link control) 10
SecurID (Security Dynamics ACE/Server) 236, 256, 257
security 224
Security Dynamics ACE/Server (SecurID) 236, 256, 257
segment number 132
self-defining dependent LU (SDDLU) 218
Sequenced Packet eXchange (SPX) 252
serial port 209
serial port status 241
server 152, 158
service advertising protocol (SAP) 15, 185, 207, 252
service port 88
Services Control Point (SSCP) 171
session flow control 205
session-level security 33, 34
shared dial-out access 235
simple bridges 38
simple network management protocol (SNMP) 66, 81, 99, 107, 239, 258
simple password 102
Single Link Multihost 169
single port token-ring 63
single station 154
single-route broadcast 43, 128
SLOWDOWN 186
SMP (symmetric multiprocessor) 192
SNA 8, 21, 71, 119
  3745 184, 185
SNA *(continued)*
  application access 196
  backbone 201
  client access for AIX 199
  communications program 164, 166
  connection establishment, 3745 185
  data compression 34
  data link switching (DLSw) 89, 147
  DFT 212
  dial support 166
  exchange ID 110
  flow control 186
  gateway 204, 215
  IP traffic 207
  IPX traffic 207
  link station 162
  local 3174 171
  NetBIOS 207
  NetBIOS ratio 148
  operational modes 197
  over TCP/IP 219
  remote 3174 gateway 172
  sockets 207, 215, 219
  subchannel 162
  token-ring interface connector (TIC) 260
  transmission priority 218
SNA Server/6000 193
SNAP (subnetwork access protocol) 76
  header 77
  value 132
SNMP (simple network management protocol) 66, 81, 99, 107, 239, 258
  agent 19, 262
  manager 19
  MIB-II 177
  network management 262
socket level applications 163
socket programming 192
sockets 27, 29
Sockets over SNA 194, 207, 215, 219
source address (SA) 41, 76, 78
source service access point (SSAP) 76, 131
source-route - translational bridging (SR-TB) 40, 45, 88, 95
source-route bridging (SRB) 17, 40, 42, 88, 95, 110, 127, 128, 207
source-route transparent bridging (SRT) 40, 44, 88, 95, 135
source-routing 42, 72, 77
spanning tree algorithm 43, 140
spanning tree bridge (STB) 40
split bridge 71, 223
spoke and hub network 102
spokes 103
spoofing 110, 143, 145, 237, 238
spooling 176
SPX (Sequenced Packet eXchange) 252

SR-TB (source-route - translational bridging) 40, 45, 88, 95
SRB (source-route bridging) 17, 40, 42, 88, 95, 127, 207
SRB (sourve-route bridging) 128
SRT (source-route transparent bridging) 40, 44, 88, 95, 135
SSAP (source service access point) 76, 131
SSCP-LU sessions 212
Stac 4.0 compression 238
Stacker-LZS 96
static addresses 69
static route 105, 117
STB (spanning tree bridge) 40
store-and-foreward 55
STP 63
style 100
subchannel 160, 162
subnetwork 100
subnetwork access protocol (SNAP) 76
subnetwork mask 14, 100
subsystem 92
subsystem management 205
Super Video Graphics Array (SVGA) 228
Switched Major Node 184
symmetric multiprocessor (SMP) 192
Sync/Async module 247
synchronous communications (Sync) 226
synchronous data link control (SDLC) 10
System Services Control Point (SSCP) takeover 218

T

T2.1 node 171, 174
TAGs 115
TB (transparent bridging) 88, 95, 127, 133
TCP (transmission control protocol) 99, 250
TCP transport layer protocol 141
TCP/IP 21, 24, 50, 71, 88, 187, 231, 261
  3745 186, 187
  3746-9x0 190
  for OS/2 12, 18
  LAN Distance 231
  LANHOP/6000 261
  Native channel support 187
  offload 152, 156, 163
  passthrough 156
  SNA as carrier 187
Telnet 99, 107, 162, 250
terminal 88
terminal data stream 176
TFTP (trivial file transfer protocol) 89, 99, 107, 109, 250
  client 109
  server 109
Thick Ethernet (10Base5) 243
Thick Ethernet (AUI) 243
Thin Ethernet (10Base2) 243
Thin Ethernet (BNC) 243
third party authentication server 259
throughput 53
TIC (token-ring interface connector) 185, 260
TIC3s (token-ring interface couplers) 184
timed LAN-to-LAN connections (TLC) 238
timed updates 238
TLC (timed LAN-to-LAN connections) 238
TN3270 176, 178, 179, 198, 200
TN5250 200
token-passing 70
token-ring
  adapter (TRA) 183, 184
  database 136
  interface connector (TIC) 185, 260
  interface couplers (TIC3s) 184
  LANHOP/6000 262
  network bridge program 139
tps (transactions per second) 56
TRA (token-ring adapter) 183, 184
traffic class 114
traffic prioritization 148, 201
transaction accounting 262
Transaction Programs (TPs) 193
transactions per second (tps) 56
translation 45
translational bridging 127, 135
transmission control 199
transmission control protocol (TCP) 99, 250
transmission priority 218
transparent bridging (TB) 40, 41, 88, 95, 110, 127, 133
traps 108
triggered updates 238
trivial file transfer protocol (TFTP) 89, 99, 107, 109, 250
tunnel bridge 40, 46, 118
TUNNELING-IP 114
turn-around time 52
Type 2.1 node 174

U

UDP (user datagram protocol) 92, 99, 250
unlinking 69
updates
  intervals 117
  piggybacking 238
  timed 238
  triggered 238
upstream 210
upstream protocols 215
user datagram protocol (UDP) 92, 99, 250
utilization 54
utilization analysis 55
UTP 63, 243
UTP (10Base-T) 243

V

V.35 209
validation, passphrase 259
variables 108
VC (virtual connection) 236, 238, 239
view 108
VINES (Banyan Virtual Networking System Protocol) 50, 125
virtual connection (VC) 236, 238, 239
virtual LAN 225, 230
Virtual Loadable Module (VLM) 235
Virtual ROM (VROM) 244
virtual segment number 110, 148
VROM (Virtual ROM) 244
VT100 198
VTAM 29, 162
VTAM V3R4 171, 185
VxD 236

W

WAC adapter 7, 209
WAN (wide area network) 61, 71, 196, 221
  attachment module 64, 71
  communication server 222
  Re-Route 96
  Restoral (WRS) 96, 115
WAN-to-Host Mode 165
warm boot 244
WaveRunner 236, 237
Wide Area Connect (WAC) 203
wide area network (WAN) 61, 71, 196, 221
Windows Sockets 32
WRS (WAN Restoral) 96, 115

X

X/Open 23
X.21 209
X.25 89
Xerox network systems (XNS) 50, 125
XMODEM 66
XNS (Xerox network systems) 50, 125

Y

yellow status LED 61

Z

ZModem 90
zones 253

Printed in U.S.A.