SG24-5000-00

International Technical Support Organization

IBM LAN Bridge and Switch Summary

January 1996

IBM

International Technical Support Organization
Raleigh Center

SG24-5000-00

International Technical Support Organization

IBM LAN Bridge and Switch Summary

January 1996

Take Note!

Before using this information and the product it supports, be sure to read the general information under "Special Notices" on page xiii.

First Edition (January 1996)

This edition applies to IBM bridge and switch products used for Local Area Network interconnection that are either announced by IBM or already available as of publication.

Order publications through your IBM representative or the IBM branch office serving your locality. Publications are not stocked at the address given below.

An ITSO Technical Bulletin Evaluation Form for reader′s feedback appears facing Chapter 1. If the form has been removed, 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

Abstract

This booklet is a summary of many IBM bridge and switch products commonly used for LAN interconnection. It contains tables, matrices and layouts of these products, their features, functions and frame types. In addition, it contains summarized descriptions of the products with useful information for comparing and contrasting them. It is a concise reference designed to encompass a variety of information in a clear, organized manner. The information presented is gathered from a wide variety of other sources. As always, one should not rely exclusively on any single source of information including this book.

Our intended audience is IBM customers, IBM network specialists and field personnel engaged in network support activities. Some knowledge of network terminology and awareness of planning issues for LAN interconnection is assumed.

(186 pages)

Contents

Tables

Special Notices

This publication is intended to help IBM customers and IBM network specialists find useful information in a single place pertaining to IBM bridge and switch products and their underlying technologies. The information in this publication is not intended as the specification of any programming interfaces that are provided by other manuals for the products referenced herein. See the PUBLICATIONS section of the IBM Programming Announcement for the individual products referenced 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.

Reference to PTF numbers that have not been released through the normal distribution process does not imply general availability. The purpose of including these reference numbers is to alert IBM customers to specific information relative to the implementation of the PTF when it becomes available to each customer according to the normal IBM PTF distribution process.

The following terms are trademarks of the International Business Machines Corporation in the United States and/or other countries:

AIX	APPN
EtherStreamer	IBM
LANStreamer	NetView
NWAYS	Micro Channel
OS/2	PS/2
TURBOWAYS	

The following terms are trademarks of other companies:

Windows is a trademark of Microsoft Corporation.

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.

C-bus is a trademark of Corollary, Inc.

3Com	3Com Corporation
AppleTalk	Apple Computer, Incorporated
AT&T	American Telephone and Telegraph Company
Banyan	Banyan Systems, Incorporated
Cisco	Cisco Systems, Incorporated
Concord	Concord Data Systems, Incorporated
Data General	Data General Corporation
DEC, DECnet, VAX	Digital Equipment Corporation
Fujitsu	Fujitsu Limited
HP, Hewlett-Packard, APOLLO	Hewlett-Packard Company
Intel	Intel Corporation
LANtastic	Artisoft, Incorporated
Motorola	Motorola, Incorporated
Network General	Network General Corporation
NetWare, Novell	Novell, Incorporated
NeXT	NeXT Computer, Incorporated
Proteon	Proteon, Incorporated

Sun	Sun Microsystems, Incorporated
Tektronix	Tektronix, Incorporated
Ungermann-Bass	Ungermann-Bass Corporation
Wellfleet	Wellfleet Communications, Incorporated
Xerox	Xerox Corporation

Other trademarks are trademarks of their respective companies.

Preface

This document is intended to help the reader understand some of the issues involved in using bridges and switches. We also cover lightly the features and functions of the IBM products that implement the technology. We present the details of frame types and contents to enable field personnel and customers to read, understand and diagnose these in given circumstances.

There are many products from IBM and other vendors on the market that offer varying capabilities and degrees of connecting Local Area Networks. This work is primarily oriented towards IBM products that operate at OSI Layer 2 and below. A common technique in the industry today relies on routing LAN protocols. Coverage of this subject is left to other works.

This document is intended for IBM customers, IBM network specialists and those personnel engaged in network support activities. Because it is a summary, it is not intended to be used as the sole reference for the information it contains, but rather as a reminder. Numerous documentation references point the reader to definitive sources of more complete information.

How This Document Is Organized

The document is organized as follows:

- Chapter 1, "LAN Concepts"

  A summary of the concepts of different LAN protocols and the terminology that accompanies them is presented. Frame formats and some tables with common values will help you understand what data flows over your LAN.

- Chapter 2, "LAN Interconnection Techniques"

  This chapter provides a short review of the concepts and terminology of LAN interconnection. Transparent and source-route bridge spanning-trees are covered along with common network design issues.

- Chapter 3, "IBM Bridge Products"

  A short description of the devices that either are or can act as bridges and their capabilities is detailed. Various tables illustrate important facts about the products, such as operating environments, cable requirements and so forth.

- Chapter 4, “IBM 827X Nways LAN Switch Product Family”

  A short description of the 827X Nways family of switch products is presented. Some small scenarios with useful implementation examples are detailed.

Related Publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this document.

- *ATM Introduction and Planning Guide*, GA27-4089
- *IBM RouteXpander/2 Users Guide and Reference*, SC31-6606
- *IBM 8229 Bridge Manual*, GA27-4025
- *NWays 2217 Multiprotocol Concentrator*, G325-3515
- *2210 NWays Multiprotocol Router and NWays Multiprotocol Routing Services*, G325-3435
- *2220 NWays Broadband Switch*, GA33-0292
- *3172 Interconnect Controller*, G221-3680
- *3174 Network Processor*, G325-3438
- *6611 Network Processor and Multiprotocol Network Program*, G325-3436
- *8250 Product Description*, GA33-0317
- *8250/8260 PSPG*, GA33-0285
- *8260 Multiprotocol Description*, GA33-0315

International Technical Support Organization Publications

- *An Inside Look at IBM Workgroup Hubs and Switches*, GG24-2528
- *Local Area Network Concepts and Products*, GG24-3178
- *New and Improved! IBM Multisegment LAN Design Guidelines* , GG24-3398
- *High-Speed Networking Technology: An Introductory Survey*, GG24-3816
- *Multiprotocol Networking with the 3172 Model 3*, GG24-4252
- *IBM 2220 NWays Broadband Switch: Concepts and Products*, GG24-4307
- *Asynchronous Transfer Mode (Broadband ISDN) Technical Overview*, SG24-4625

- *3174 Network Server in Frame Relay and Multiprotocol Networks*, GG24-4376
- *IBM 2210 NWays Multiprotocol Router Description and Concepts*, GG24-4446

A complete list of International Technical Support Organization publications, known as redbooks, with a brief description of each, may be found in:

Bibliography of International Technical Support Organization Technical Bulletins, GG24-3070.

To get a catalog of ITSO redbooks, VNET users may type:

```
TOOLS SENDTO WTSCPOK TOOLS REDBOOKS GET REDBOOKS CATALOG
```

A listing of all redbooks, sorted by category, may also be found on MKTTOOLS as ITSOCAT TXT. This package is updated monthly.

How to Order ITSO Redbooks

IBM employees in the USA may order ITSO books and CD-ROMs using PUBORDER. Customers in the USA may order by calling 1-800-879-2755 or by faxing 1-800-445-9269. Most major credit cards are accepted. Outside the USA, customers should contact their local IBM office. For guidance on ordering, send a note to BOOKSHOP at DKIBMVM1 or E-mail to bookshop@dk.ibm.com.

Customers may order hardcopy ITSO books individually or in customized sets, called BOFs, which relate to specific functions of interest. IBM employees and customers may also order ITSO books in online format on CD-ROM collections, which contain redbooks on a variety of products.

ITSO Redbooks on the World Wide Web (WWW)

Internet users may find information about redbooks on the ITSO World Wide Web home page. To access the ITSO Web pages, point your Web browser to the following URL:

```
http://www.redbooks.ibm.com/redbooks
```

IBM employees may access LIST3820s of redbooks as well. The internal Redbooks home page may be found at the following URL:

```
http://w3.itsc.pok.ibm.com/redbooks/redbooks.html
```

Acknowledgments

This publication is the result of a project conducted at the International Technical Support Organization, Raleigh Center. The project leader and technical advisor for this project was:

Mark DeCain
International Technical Support Organization, Raleigh Center

The author of this document is:

Luc Dekoster
IBM Belgium

Thanks to the authors of the following works for the invaluable advice and guidance provided in their works, freely re-incorporated here:

Local Area Network Concepts and Products, GG24-3178

New and Improved! IBM Multisegment LAN Design Guidelines, GG24-3398

Thank you to the following people for their suggestions in format and technical content provided in the production of this document:

Steve Breese
Mohammad Shabani
International Technical Support Organization, Raleigh Center

Our grateful acknowldegement for their contribution to this work by the following IBM Networking Hardware Division staff:

Norm Strole
IBM RTP, NC

Special thanks to the following people for their help in editing and preparation of this book:

Martha DeCain   Gray Heffner
Linda Robinson   Gail Wojton

International Technical Support Organization, Raleigh Center

Chapter 1. LAN Concepts

This chapter summarizes the basic principles of the LAN protocols: CSMA/CD or Ethernet, token-ring, FDDI and ATM. It is not intended to be a course on LAN concepts or a complete reference document. A good book that you may read to learn about different LAN functions is *Local Area Network, Concepts and Products*, GG24-3178.

1.1 CSMA/CD, Ethernet and IEEE 802.3

The standard ratified by the Institute of Electrical and Electronics Engineers (IEEE) as IEEE 802.3 and approved by the International Organization for Standardization as ISO 8802-3, is slightly different from the protocol that was earlier brought out by a consortium of DEC, Intel and Xerox as (DIX). Ethernet V2

Both Ethernet V2 (DIX) and IEEE 802.3 are widely implemented across the marketplace. They are not identical and use different frame formats, but they can coexist on the same physical LAN. However, stations using one frame type cannot communicate with stations that use the other frame type, unless they are configured to use both frame types concurrently.

Ethernet and IEEE802.3 use a bus topology and a technique called Carrier Sense Multiple Access with Collision Detection (CSMA/CD) to control the operations of the network.

In a CSMA/CD bus, when a station wants to transmit data on the network bus, it first listens to see if the bus is free (carrier sense). If the bus is available, it starts transmitting data immediately. If the bus is not available, the station waits until the activity on the bus stops before it starts transmitting again. All stations receive the data but only the station that recognizes the destination address as its own will copy the frame into its internal buffers.

If there is a *collision* after transmission (that is, another station starts to transmit at the same time), the stations will stop transmitting data immediately after the collision is detected, but continue to transmit a jamming signal to inform all active stations about the collision. In response to this signal, each transmitting station waits a random time before starting the whole process again, beginning with the process of carrier sensing.

The probability of a collision occurring is proportional to the number of stations, the frequency of transmissions, size of frames and length of the LAN. Therefore, care must be exercised in designing LANs with an excessive number of stations. Also, you must ensure that the length of individual *segments* and total length of the LAN does not exceed a certain length as defined by the 802.3 standards.

To be able to detect collisions, a transmitting station should monitor the network for a period of time during which a collision may occur. This is the maximum delay for a transmission to reach the far end of the network and for a collision to propagate back. A compromise can be made between the maximum length of the network, the bit rate to be used and the minimum frame length.

Table 1. Design Rules for Ethernet
The slot time is equal to 512 bit times in a 10 Mbps LAN, which results in a 64-byte minimum frame length.
The maximum frame length is 1500 bytes.
The longest distance between two stations can be as much as 4300 meters.
Maximum number of stations in a collision domain: 1024
No two stations in a collision domain may be separated by more than four repeaters.
No two stations in a collision domain may be separated by more than three coaxial segments. The other two segments must be link segments.
Link segments can be 10BaseT or fiber optic segments.

The maximum length of the network depends on the wiring that is used. The following rules apply:

Table 2. Ethernet Segment Lengths		
10Base5 (Thicknet, coax, AUI)	Max. 500 m	Max. 100 stations
10Base2 (Thinnet, coax, BNC)	Max. 185 m	Max. 30 stations
10BaseT (Twisted pair, UTP, RJ-45)	Max. 100 m	Max. 2 stations
10BaseFx(Fiber Optics)	Max. 2000 m	Max. 2 stations

These are rules of thumb. Please refer to appropriate documentation when designing Ethernet networks.

1.1.1 Ethernet V2 Frame Format

Table 3. Ethernet V2 Frame Format

PREAMBLE	SYNC	DA	SA	TYPE	DATA	FCS
8 bytes		6 bytes	6 bytes	2 bytes	46-1500 bytes	4 bytes

PREAMBLE 62 bits, and

SYNC 2 bits, for synchronization

DA Destination Address - 48 bits

SA Source Address - 48 bits

Three types of destination addressing are supported:

- **Individual**: The unique address of one node on the network.
- **Multicast**: The first bit of the DA identifies a *group* address.
- **Broadcast**: A DA field of all 1s is a *broadcast* to all stations.

The source address (SA) field always contains the individual address of one node, the source station, on the network.

TYPE Type field - 16 bits, identifies the higher layer protocol which is used. Each registered protocol is given a unique 2-byte *Ethertype* identifier. The value assigned to the *type* field in Ethernet is always higher than the maximum value in the *length* field for the 802.3. This is to ensure that both protocols can coexist on the same network. See Ethertypes in Table 41 on page 171.

DATA Data field - This contains the actual data being transmitted and is 46-1500 bytes in length. Ethernet assumes that the upper layers will ensure that the minimum data field size (46 bytes) is met prior to passing the data to the MAC layer.

FCS 32 bits, the Frame Check Sequence is the result of a cyclic redundancy check algorithm on the frame integrity.

Note

The MAC addresses in Ethernet and 802.3 frames are represented in a canonical bit order with the most significant bit on the right. The usual IBM way of writing binary values in a hex representation is with the most significant bit on the left. See 1.3.7, “Canonical Bit Ordering” on page 12.

1.1.2 IEEE 802.3 Frame Format

Table 4. IEEE 802.3 Frame Format

PREAMBLE	SFD	DA	SA	LEN	DATA	FCS
8 bytes		6 bytes	6 bytes	2 bytes	46-1500 bytes	4 bytes

PREAMBLE 56 bits, and SFD (Start Frame Delimiter) - 8 bits, allows for synchronization.

DA Destination Address, and SA (Source Address) - 48 bits. The same format and note applies as to DA and SA in Ethernet V2.

LENGTH field 16 bits, indicates the number of *data* bytes (excluding the pad) that are in the data field.

DATA and PAD 802.3 permits the *data field* to be less than the 46 bytes provided that the whole packet meets the minimum of 64 bytes. In order to ensure that the minimum packet size requirement is met, 802.3 requires the MAC layer to add pad characters to the LLC data field before sending the data over the network.

FCS 32 bits, the result of a cyclic redundancy check algorithm.

1.2 Token-Ring (IEEE 802.5)

In a token-passing-ring network the stations on the LAN are logically connected in a star-wired ring topology. Access to the ring is controlled by a circulating token. A station with data to transmit waits for a free token to arrive. When the token arrives, the station changes the token into a frame, appends data to it and transmits the frame. Each station repeats the frame to its downstream neighbor, while it performs error checking on the bit stream. If the destination station is active, it will copy the frame into its buffer and set the ″frame copied″ and ″address recognized″ bits in the Frame Status (FS) field. The sending station must strip the frame from the ring and release a new token onto the ring.

A maximum of 260 stations are allowed on one 4 Mbps segment, depending on the media in use. A 16 Mbps token ring may support up to 132 devices depending on the media used and the type of wiring hub.

1.2.1 Early Token Release

When the frame transmission time is shorter than the time that it takes to circle the ring, then the originating station must generate idles until the header is received back before releasing a new token.

On a 4 Mbps token-ring LAN where the length of 1 bit is roughly 50 meters, a short frame of 200 bits would be 10,000 meters long. Therefore at 4 Mbps, the idle time can be extremely small. If, however, we consider a 16 Mbps token-ring LAN, 1 bit is 12.5 meters long. In a larger sized network we may wish to optimize the utilization of the medium by reducing the idle time (waiting for a header).

The architecture provides an option called *early token release.* With this option a transmitting station will release the token immediately after completing transmission of the frame before it receives the header back, thereby eliminating the idle time while waiting for the header to reappear. This allows multiple frames, but still only one token on the LAN.

The early token release option is enabled by default on an IBM 16 Mbps token-ring network. It is an option for each station, and is not required that *all* stations implement the option, but it is recommended. Stations that do not use early token release will operate normally with those stations that do.

1.2.2 Token-Ring Frame Format

Table 5. Token-Ring Frame Format

SD	AC	FC	DA	SA	RI	DATA	FCS	ED	FS
1 byte	1 byte	1 byte	6 bytes	6 bytes	2-18 (30)	variable length	4 bytes	1 byte	1 byte

SD Starting Delimiter - 8 bits

AC Access Control 1 byte, bit format: PPP T M RRR

The free token consists of an Access Control field between a Start Delimiter and an End Delimiter. The T-bit in this field identifies this string as a token.

The M-bit, or monitor bit, is set by the Active Monitor and ensures that no frame perpetually circles the ring. A frame with this bit would be removed by the Active Monitor on the second pass.

The other bits in this field PPP and RRR are the priority and reserve-priority bits.

The token-passing protocol guarantees fair access to all participating stations. This fairness is enhanced by an eight-level priority mechanism, based on priority reservations made in the bits of a circulating frame or token.

FC Frame Control - 1 byte, indicates whether a frame carries information destined to the MAC, or to the LLC sub-layer of the Data Link layer.

DA Destination Address and SA (Source Address), - 48 bits

RI Routing Information - (optional), 2 to 18 (30) bytes. See layout in Table 6 on page 7.

DATA Variable length. LLC protocol data or MAC management data.

FCS Frame Check Sequence, a 4 bytes CRC value.

ED Ending Delimiter - 8 bits.

FS Frame status - 8 bits, bit format: A C rr A C rr. A-bits are address-recognized bits, C-bits are frame-copied bits.

A complete description can be found in the *IBM Token-Ring Network Architecture Reference,* SC30-3374.

1.2.3 Routing Information Field (RI)

Bit 0 (the leftmost) of the source address field indicates that a routing information field follows the source address.

Table 6. Routing Information Field Format

Routing Control	Route Designator (Optional)	Route Designator (Optional)		Route Designator (Optional)
2 bytes	2 bytes	2 bytes	variable length	2 bytes

Routing Control 2 bytes with the following format:

B	B	B	L	L	L	L	L	D	F	F	F	E	E	E	r

B: Broadcast bits:

- 0xx Non-Broadcast
- 10x All-Routes Broadcast
- 11x Single-Route Broadcast

L: Length of this routing field.

D: Direction indicator: If this bit is on, interpret the routing information from right to left.

F: Largest frame size indicator bits. See Table 17 on page 73.

E: Largest frame size extension bits.

r: reserved.

Route Designator fields 2 bytes (4 hexadecimal numbers), the first three digits indicate a ring segment number, the last digit is a bridge number. There can be up to 8 route designators for a maximum of 7 hops. New bridges allow 13 hops. The last designator ends with a bridge number of 0. Example: 001,A--002,B--003,0.

1.2.4 Medium Access Control Sublayer (MAC)

The token-passing protocol provides an extensive set of inherent fault isolation and error recovery functions, for implementation in every attaching device. All these functions are part of the Medium Access Control (MAC) sub-layer of the Data Link Layer functions in the adapter of each station.

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

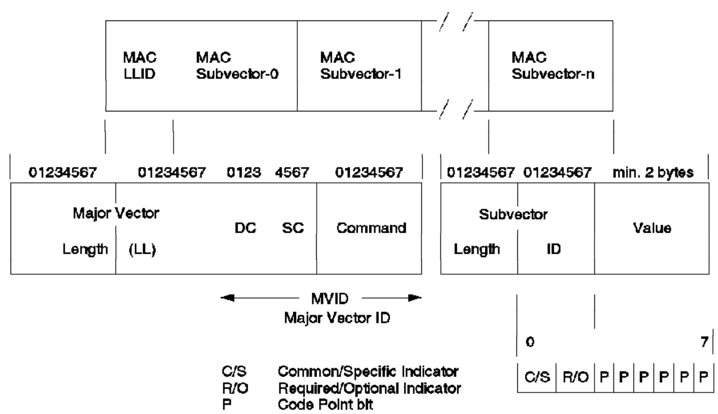

Figure 1. Token-Ring MAC Frame - Data Field Format

MAC frames are processed according to destination (DC) and source (SC) classes, including Ring Station, Ring Error Monitor and RPL Server.

One or more sub-vectors provide additional information depending on the specific major vector identifier.

A complete description can be found in the *IBM Token-Ring Network Architecture Reference,* SC30-3374.

1.3 Common Data, Numbers and Formats

Different LAN protocols have many things in common, such as field layouts and values. Some protocol formats are common, such as Logical Link Control (LLC) for token-ring and IEEE 802.3. This section lists some of these common things.

1.3.1 Destination MAC Address (DA) and Source MAC Address (SA)

Table 7. MAC Address Format

DA:	**I/G**	**U/L**	**reserved**	**FAI**	**address bits**
SA:	**RII**				
	1	1	14 bits	1	31 bits

All ring stations are identified by unique *individual* addresses.

I/G Indicates an Individual (B′0′) or Group (B′1′) address.

RII Since the source of a frame can only be one single station and not a group, the first bit of the SA is now used in the IBM implementation of token-ring to indicate whether or not a Routing Information Field (RIF) follows.

U/L Indicates a Universal (B′0′) or Locally (B′1′) administered address.

The reserved bits are set to B′0′ for locally administered addresses.

Functional Address Indicator
(B′0′) indicates a functional address if I/G = B′1′ (indicating a group address).

For individual locally administered addresses, FAI must be B′0′ by convention.

Universally administered address (UAA): A universally administered or a *burned-in* address is assigned by the IEEE organization.

Some manufacturers have been assigned universal addresses that contain an organizational unique identifier (OUI). For instance, IBM has identifiers X′10005A′ and x′08005A′. All cards that use IBM token-ring chip sets have the first 6 digits of their addresses beginning with these hexadecimal numbers. A few IDs are listed in Table 42 on page 172. Some diagnostic tools can recognize these identifiers and translate them back to brand names.

Locally administered address (LAA): A ring station's individual address can also be *locally administered* by a network administrator.

Group address: A number of destination ring stations can be identified by a *group* address. Defined standard group addresses are listed in Table 43 on page 173.

Functional address: A token-ring LAN also provides a special case of a locally administered group address called *functional* addresses. Each (bit-significant) functional address represents a well-identified server function within the access protocol. Multiple bits in the four low-order bytes indicate multiple functions. They are listed in Table 44 on page 175.

The address X'FFFFFFFFFFFF' identifies all ring stations as destination stations. A frame with X'000000000000' as its destination MAC address is not addressed to any ring station.

1.3.2 Logical Link Control (LLC)

The LLC sublayer of the Data Link Layer is described by the IEEE 802.2 standard. Logical Link Control provides a consistent view of a LAN to the upper layers regardless of the media and protocols being used (IEEE 802.3 or token-ring).

The interface to the upper layers is provided through LLC Service Access Points (LSAPs or SAPs).

A station can have more than one SAP associated with it, just as a station may have more than one session active via one SAP. For example, a LAN station may have an SNA session through SAP X'04' to a token-ring attached 3174-01L, and concurrently be in session with a file server in another LAN station through SAP X'F0'. Logical Link Control provides the capability to manage these independent sessions.

Table 8. LLC Header Format

MAC Header	RI	DSAP	SSAP	CF	I-Field	MAC Trailer
		1 byte	1 byte	1-2 bytes	variable length	

DSAP Destination Service Access Point - the identified SAP for which the LPDU is intended.

SSAP Source Service Access Point - identifies the SAP that originated the LPDU.

CF Control Field - commands, responses, sequence numbers and poll/final bits

I-Field Information Field - variable length.

1.3.3 Sub-Network Access Protocol (SNAP)

A method has been designed by the Internet community to enable the transport of IP-like information over LLC-type networks. It uses SAP AA to identify the SNAP protocol.

Table 9. Using SAP AA to Identify the SNAP Protocol

MAC Header	LLC Header			SNAP Header		Info
	DSAP	SSAP	Cntl	Prot_ID	Type	
	AA	AA	03	000000	0800	IP-frame

1.3.4 LLC over Ethernet V2

LLC-type information can also be carried over Ethernet V2. For that purpose we use an Ethertype X′80D5′ to indicate that an LLC header follows. An Ethernet station is unlikely to understand these frames, but a bridge can recognize them and convert them back to token-ring.

1.3.5 SAP Values

The bit-format of the DSAP is ′DDDDDDUI′.

The bit-format of the SSAP is ′SSSSSSUC′.

I Individual or group SAP

C Command / Response indicator

U User-defined SAP

Since the C bit of the SSAP is not used as a bit within an SAP address, there are in fact only 128 possible SSAP values. The defined LSAPs (LLC-SAP) are listed in Table 39 on page 169.

1.3.6 Full Duplex LANs

On a LAN where stations are directly connected to a switch, a station does not have to contend with other stations on the LAN for permission to transmit. In the situation where transmitting and receiving stations are the only stations on the LAN, no collisions will occur on an Ethernet, and the token-passing protocol becomes unnecessary on a token-ring. This makes full-duplex operation possible; that is, a station can transmit and receive frames simultaneously and virtually double their bandwidth as compared to the nominal bandwidth of the LAN (20 Mbps on a 10 Mbps Ethernet, 32 Mbps on a 16 Mbps token-ring). All idle time can be eliminated because both stations can spontaneously start transmitting whenever they want with minimal regard for normal protocol rules.

1.3.7 Canonical Bit Ordering

In the documentation produced by the standards bodies, a different method of bit ordering is used from the one commonly found in the IBM world. The ″standard″ way of describing the order of bits in a byte is known as canonical bit ordering and is the reverse of the way IBM usually writes binary values. Depending on whose documentation is being studied, and when different sets of documentation are being compared to each other, confusion may result. Ethernet addresses are always in canonical format. Also, when tracing a token-ring, a field containing a station address field may be present within the data. This field may be in canonical bit order, and has to be translated into the IBM order if sense is to be made of it. For example:

The decimal number 71 has a hexadecimal value of 47. IBM would write this value, in binary, as 0100 0111.

In canonical form, this same binary value would be read right to left, becoming 1110 0010, or E2.

If we take a much longer bit stream, such as a station MAC address, then the conversion would be like this:

```
1 0   0 0   5 A   1 7   5 8   6 9

0 8   0 0   5 A   E 8   1 A   9 6
```

The order of the bytes remains the same, but each byte has its binary digits read from right to left when converting from the IBM form to the canonical form.

The hexadecimal conversion table is as follows:

Order

Table 10. Hexadecimal Conversion from IBM Bit Order to Canonical Form Bit															
0	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**A**	**B**	**C**	**D**	**E**	**F**
0	**8**	**4**	**C**	**2**	**A**	**6**	**E**	**1**	**9**	**5**	**D**	**3**	**B**	**7**	**F**

1.3.8 Examples of Frames on Ethernet and Token-Ring

Some examples of frames you are likely to encounter on your LAN are listed below. Many others may be seen, as well as all kinds of encapsulations.

1.3.9 TCP/IP

On Ethernet V2:

Preamble	DA	SA	0800	IP and TCP headers	Data	FCS

On IEEE 802.3 (SNAP):

Preamble	DA	SA	Length	AA	AA	03	0800	IP and TCP headers	Data	FCS

On token-ring (SNAP):

SD	AC	FC	DA	SA	RI	AA	AA	03	000000	0800	IP and TCP headers	Data	FCS	ED	FS

1.3.10 SNA or NetBIOS

On Ethernet V2:

This will not occur

On IEEE 802.3 (802.2 LLC):

Preamble	DA	SA	Length	04 or F0	04 or F0	Cntl	I-field	FCS

On token-ring (802.2 LLC):

SD	AC	FC	DA	SA	RI	04 or F0	04 or F0	Cntl	I-field	FCS	ED	FS

LLC over Ethernet V2:

Preamble	DA	SA	80D5	Length	04 or F0	04 or F0	Cntl	I-field	FCS

1.3.11 Novell NetWare IPX

On Ethernet V2:

Preamble	DA	SA	8137	FFFF	IPX Length	IPX Data	FCS

On IEEE 802.3:

Preamble	DA	SA	Length	FFFF	IPX Length	IPX Data	FCS

On IEEE 802.3 (802.2 LLC):

Preamble	DA	SA	Length	E0	E0	03	FFFF	IPX Length	IPX Data	FCS

IEEE 802.3 (SNAP):

Preamble	DA	SA	Length	AA	AA	03	000000	8137	FFFF	IPX Len	IPX Data	FCS

On token-ring (802.2 LLC):

SD	AC	FC	DA	SA	RI	E0	E0	03	FFFF	IPX Len	IPX Data	FCS	ED	FS

On token-ring (SNAP):

SD	AC	FC	DA	SA	RI	AA	AA	03	000000	8137	FFFF	IPX Len	IPX Data	FCS	ED	FS

1.4 FDDI (SDDI and CDDI)

In many ways, FDDI is similar to the IEEE 802.5 token-ring, although there are some differences. FDDI uses a token-passing protocol in which each station has the chance to transmit data when a free token is captured. Using an algorithm which permits bandwidth allocation, a station can help you decide how many frames it will transmit without releasing the token.

FDDI operates at a bit rate of 100 Mbps.

FDDI uses two rings:

- The *primary ring*, which is similar to the main ring path in token-ring.
- The *secondary ring*, which is similar to the backup ring path of a token-ring.

From a wiring point of view, FDDI is similar to a fiber optic token-ring network; however, a device can be attached to either or both of the primary and secondary rings.

- A *Class A* device attaches to both of the rings directly. If it is a station, it is called a dual attachment station (DAS). If it is a concentrator, it is called a dual attachment concentrator (DAC).

- A *Class B* device attaches to only one of the rings directly or through a concentrator. If it is a station, it is called a single attachment station (SAS). If it is a concentrator, it is called a single attachment concentrator (SAC).

During normal ring operation, the primary ring is active while the secondary ring is idle. In the wake of a failure on the primary ring, the secondary ring will become active when a Class A device wraps the primary ring to the secondary ring establishing a single ring. This functionality is mandatory to maintain the reliability of the LAN.

1.4.1 FDDI over Copper (SDDI and CDDI)

An alternative to FDDI is Shielded twisted pair Distributed Data Interface (SDDI). This proposal is for transmitting FDDI directly on copper wires without converting the electrical pulse stream to optical signals. The data stream remains at the rate of 100 Mbps.

There is also a proposal to the ANSI FDDI TP-PMD workgroup for a copper solution running FDDI over UTP Category 5 Media (CDDI).

1.4.2 Port Types

The standard specifies four port types for FDDI ports:

- A-Type: For dual attachment stations: primary ring-in, secondary ring-out
- B-Type: For dual attachment stations: secondary ring-in, primary ring-out
- M-Type: On a concentrator, to attach to a single attachment station
- S-Type: On a single attachment station, to attach to a concentrator

The connection rules for the different port types are the following:

- A-to-B and B-to-A are peer-to-peer trunk connections
- M-to-S is a master-to-slave connection
- M-to-A and B provides dual homing
- S-to-S is a point-to-point connection

Table 11 shows the connection rules for single and dual attachment stations.

Table 11. Connection Rules for SAS and DAS

Port Type	Port Type	Rule
A	A	Undesirable peer connection that creates twisted primary and secondary rings.
A	B	Normal trunk ring peer connection.
A	M	Tree connection with possible redundancy. Port B shall have precedence for connecting to port M in a single MAC node.
A	S	Undesirable peer connection that creates a wrapped ring.
B	B	Undesirable peer connection that creates twisted primary and secondary rings.
B	M	Tree connection with possible redundancy.
B	S	Undesirable peer connection that creates a wrapped ring.
M	M	Invalid configuration.
M	S	Normal tree connection.
S	S	Connection that creates a single ring of two slave stations.

1.4.3 Dual Homing

A concentrator that is not part of the main ring may be dual-attached via one or two other concentrators to provide greater availability. When connected in this manner, a concentrator is described as a Dual Homing Concentrator (DHC).

Similarly, a Dual Attachment Station can be connected to one or two concentrators using both A and B ports to provide high availability. A station connected in this manner is considered a Dual Homing Station (DHS).

In both of these cases, only port B is active, and the connection to port A remains in standby mode. Should the connection to port B fail, port A would become active without any impact on the users of the Dual Homing Station or concentrator.

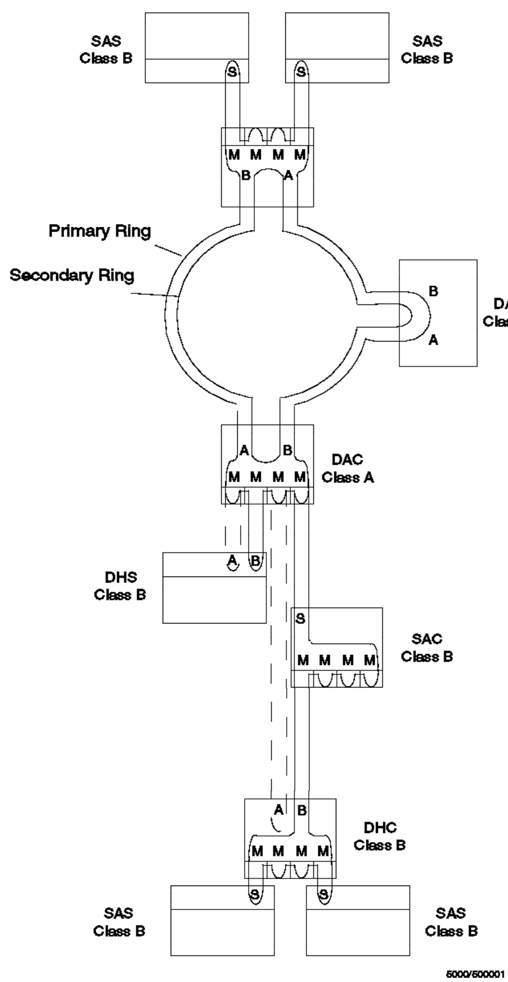

Figure 2. FDDI Attachment Examples

1.5 100 Mbps Standards

Drafts for both Ethernet and token-ring full-duplex standards have been submitted to the respective IEEE groups and are under evaluation. Some products have already reached the market and are forming the basis for these proposed standards.

- **100VG-AnyLAN**

  This is a technology for 100 Mbps operation over Category 3 UTP wire (STP cable can also be used), which is designed to support both Ethernet and token-ring LANs. 100VG stands for 100 Mbps over voice grade. Station access to the media is controlled by a sequential polling mechanism.

  The draft for the 100VG standard is under study by the IEEE (802.12).

- **100BASE-T**

  This is a technology for a 100 Mbps LAN based on the standard Ethernet 802.3 protocols. The drafts are being evaluated by the IEEE and include the following variations:

  - 100BASE-T4 for 4-pair category 3 wiring
  - 100BASE-TX for 2-pair category 5 UTP wiring
  - 100BASE-FX for optical fiber
  - 100BASE-T2 for 2-pair category 3 UTP wiring

1.6 Asynchronous Transfer Mode (ATM)

ATM is a multimedia networking technology that has the ability to provide bandwidth on demand. ATM is protocol independent.

For a more complete discussion of ATM please refer to *Asynchronous Transfer Mode (Broadband ISDN) Technical Overview*, SG24-4625. ATM uses cell relay technology, and is designed to relay data, voice and video in fixed sized cells of 53-byte lengths which consist of a five byte header and a 48-byte ″payload″.

Various types of traffic can be mixed without concern for delay due to this small fixed-length cell size.

Cell relay combines the high throughput and bandwidth optimization of frame relay with the predictability of time division multiplexing. This makes it suitable for both asynchronous data traffic and isochronous voice and video traffic. With its proposed interfaces ranging from 1.544 Mbps to 622 Mbps of full-duplex, dedicated bandwidth, ATM offers throughput capability far beyond all other techniques.

This technology will be used by broadband ISDN and metropolitan networks. ATM is eligible to become the world standard for all high throughput communication in both WAN and LAN.

1.6.1 ATM Structure

The ATM structure consists of a service layer, adaptation layer, ATM layer and physical layer. Figure 3 on page 20 shows the different layers implemented for ATM services.

1.6.2 ATM Adaptation Layer

The ATM Adaptation Layer (AAL) provides ATM service to applications of the same type. There are four defined AALs:

AAL-1 Circuit emulation, which is a constant bit rate service (CBR) suitable for services such as uncompressed video.

AAL-2 Audio/video is a variable bit rate service (VBR), which will be suitable for compressed video. AAL-2 is the service required for multimedia and it is not yet defined.

AAL-3/4 This may be used for connection-oriented data and connectionless data. These are for SMDS compatibility (Switched Multi Megabit Data Service). They create significant cell overhead when compared with AAL-5.

AAL-5 A simplified version of AAL-3/4 without the overhead oriented towards data.

It is the Adaptation Layer that segments and reassembles data into 48-byte payloads and hands these across to the ATM layer.

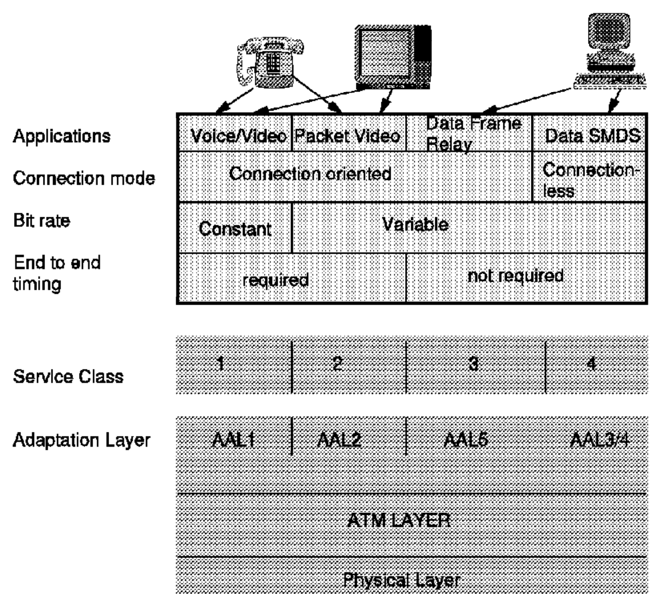

Figure 3. ATM Services

1.6.3 ATM Layer

ATM is like a telephone; it is connection oriented. That is, a "pipe" is first set up, then data is sent and at some point the "pipe" is closed.

Conceptually, on a physical link, there are "thin pipes" and "thick pipes." The thick pipes are Virtual Paths (VPs) and the thin pipes are Virtual Channels (VCs).

The ATM layer adds or removes the header to the 48 byte payload. The 5-byte header contains the Virtual Channel Identifier (thin pipe) and the Virtual Path Identifier (thick pipe), among other things. These identifiers are used by an ATM switch to actually switch. The ATM switch uses a label swapping technique to forward the cell to the next switch or the target device.

1.6.4 LAN Emulation

There are two generic alternatives for building an ATM network within a local area:

"Native" ATM Network

Workstations will be ATM endpoints and will communicate with one another through a central ATM switch, using ATM to perform the networking function.

You could not easily reuse existing LAN software in this environment but you could make use of almost any of the existing wide area networking systems (such as TCP/IP or SNA/APPN). IBM′s Networking BroadBand Services (NBBS) architecture is intended to provide a modern, effective, higher-layer networking structure for this purpose.

This approach is certainly the most efficient but it becomes very difficult to use existing LAN application software or to interface to existing LANs.

ATM Virtual (Emulated) LAN

LAN emulation is not a part of ATM itself. Nevertheless, it is a critical function which uses ATM and which will be needed immediately by most ATM users.

The concept of an ATM virtual LAN is to construct the system such that the workstation application software "thinks" that it is a member of a real shared-medium LAN.

This would enable the reuse of the maximum amount of existing LAN software and significantly reduce the cost of migration to ATM.

The situation is helped by the fact that existing LAN software has usually been built to interface to many different types of LAN. For example, most IBM LAN software will interface with token-ring, Ethernet, token bus, PC-LAN and FDDI networks, all without changes to the application software.

A possibility is open to construct multiple *unrelated* virtual LANs over the same ATM network.

IBM′s LAN Emulation operates in the following way:

- Each ATM end system contains software that provides the appearance of a LAN to the existing LAN software in the workstation. This software, together with the necessary ATM interfacing hardware, is an ATM end point as far as the ATM network is concerned.
- There is a LAN emulation server somewhere in the network which provides address resolution, administration, broadcast forwarding and some regular data transfer facilities to the end-stations.
- End systems (actually, the LAN emulation layer software within the end systems) are connected through the ATM network to the LAN emulation server.
- Broadcasts are performed on behalf of end systems by the LAN emulation server.
- Data transfer takes place on direct VCCs between the end systems. These are set up by the end systems using supporting functions from the LAN emulation server.

This structure maintains full LAN function and can support most higher-layer LAN protocols.

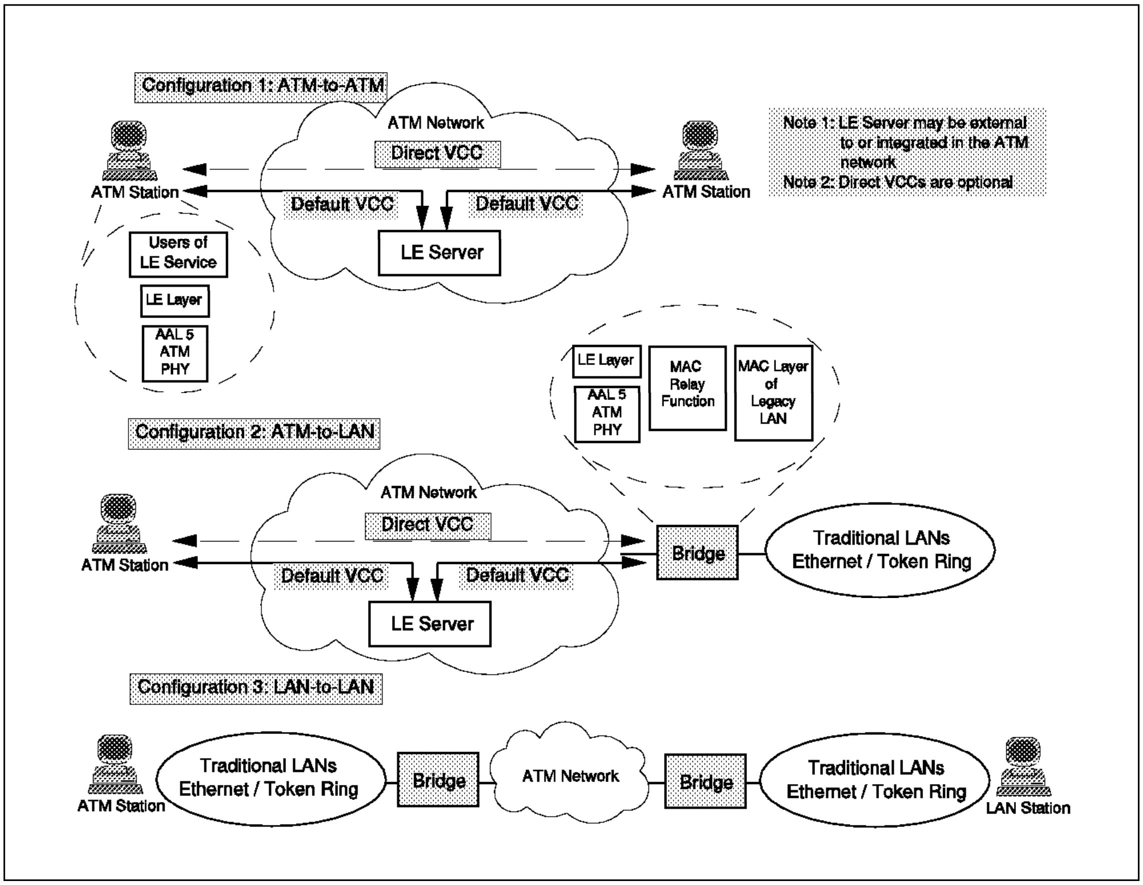

Figure 4. LAN Emulation Configurations

Chapter 2. LAN Interconnection Techniques

This chapter provides an overview of the various techniques available to interconnect LAN segments.

In local area networks, connectivity requirements often exceed the capabilities of a single ring or bus. This leads to segmentation of LANs, with the segments being interconnected with bridges, LAN switches or routers. There are many reasons to segment a LAN, including:

- Maximum number of stations has been reached

  For example, token-ring architecture allows a maximum of 260 stations on a 4 Mbps ring while Ethernet architecture allows a maximum of 1024 stations in a single collision domain.

- Physical network size

  There may be a need to extend local area network capability beyond the cabling guidelines for a single segment.

- Volume of data

  The available bandwidth of a LAN must be shared by all stations. The more stations that are attempting to transmit concurrently, the smaller the share of bandwidth for each station.

2.1 Internetworking and the OSI Reference Model

The computer industry and the standards bodies have chosen to give distinct names (*bridges, switches, routers and gateways*), to network interconnect devices depending on where they operate within the OSI seven layer reference model.

Gateways A gateway operates above layer 3 and supports protocol conversion between unlike protocol stacks. As an example, the communication between OSI and SNA would be handled by a gateway.

Routers All routers make their primary filter/forward decision at Layer 3 of the OSI model. They interconnect and route across many physical media types, including LANs and WANs.

Bridges Normally bridges make their primary filter/forward decision based on the header of a LAN frame and usually at Layer 2 of the OSI model. They can be used to connect homogeneous and heterogeneous LANs. For example, a bridge can be used to connect two token-ring LANs or connect a token-ring LAN to an Ethernet LAN.

Switches A switch is more a cousin of a bridge than a separate LAN interconnect technology, as they both share a common foundation. In a LAN context, they are bridges by another name, except:

- They are generally less expensive than multiport bridges.
- They have higher performance than bridges.
- They are even less protocol aware than bridges.

An ATM switch forwards ATM cells from one port to another; it *switches* cells.

A switching hub would be a LAN hub providing dedicated, non-shared bandwidth to attached devices.

Repeaters A repeater operates at the physical layer and is used to extend the physical characteristics of a LAN. Repeaters do not reduce the amount of traffic on a LAN segment.

In a token-ring environment, repeaters are used to extend the distance between wiring closets, while in an Ethernet environment they are used to increase both the physical span and the number of devices supported within a single collision domain.

The relationship of these devices with respect to the seven-layer OSI reference model is shown in Figure 5 on page 27.

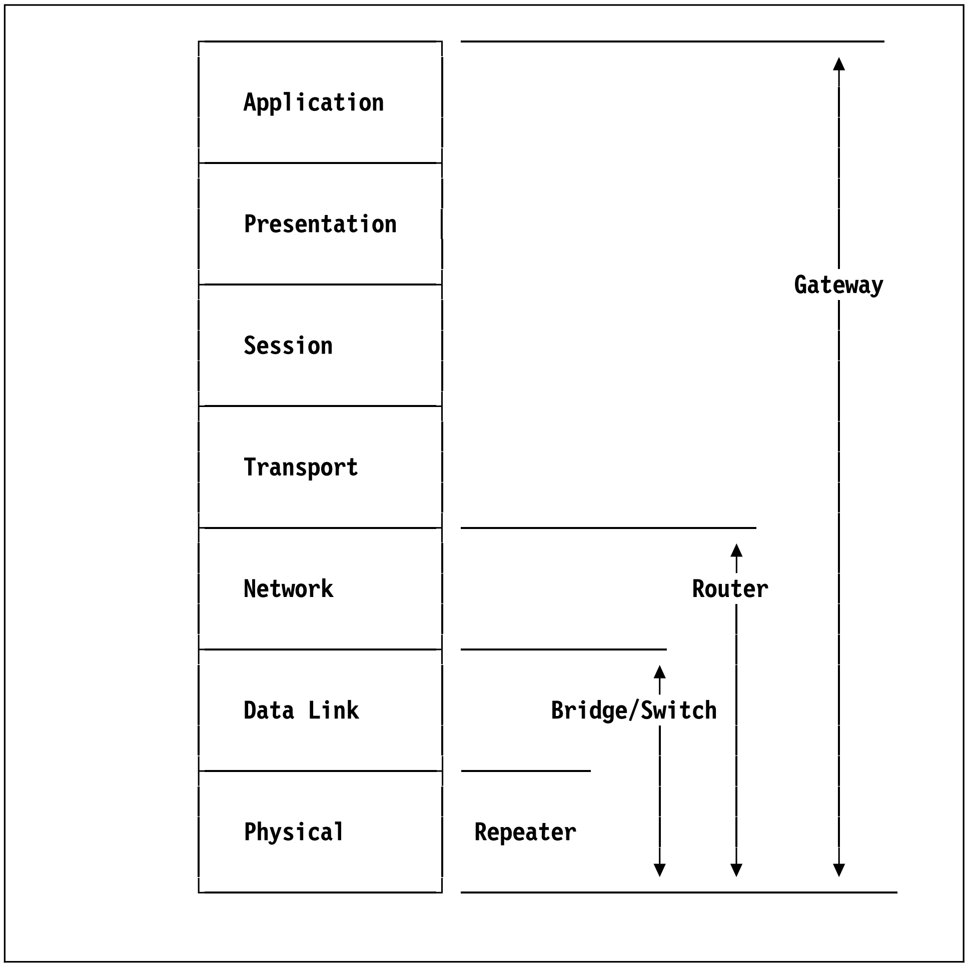

Figure 5. OSI Reference Model and Internetworking Devices

Unfortunately, the same words can mean different things in different contexts, for example:

Segment

In Ethernet LANs, *segment* refers to a contiguous medium between repeaters.

Gateway

In TCP/IP the term *gateway* is synonymous with an IP router.

IBM often uses the term *gateway* to mean an SNA LAN interconnect device. An IBM 3745 LAN gateway is an SNA router.

2.2 Bridging Methods

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

The following types of bridges exist today:

Transparent Bridges
These bridges forward frames based on the destination MAC address of a LAN frame. If the destination address of a frame is known to be on the same LAN as the source address, then no forwarding will take place. If the destination address is not known on the same LAN as the source address, then the bridge will forward the frame. The IEEE standard 802.1d defines the operation of transparent bridges. These bridges are predominantly used with Ethernet LANs. However, they may be used with other LAN types such as token-ring.

Source Routing Bridges
These bridges are used to provide interconnection between 802.5 LANs. They forward frames based on a *Routing Information Field* (RIF) which is part of the MAC frame header. The RIF defines a route which a frame will take to get from its source to its destination.

Translational Bridges
These bridges connect LAN segments of different MAC types. Their operation has facets usually associated with routers, their operation is more protocol specific than the previous bridge types.

A good example of a translational bridge is the Source Route to Transparent Bridge. The SR-TB translates between Source Routing (SR) and Transparent Bridging (TB). These bridges cater for differences in operation at the MAC layer and some MAC specific facets of higher layer protocols. An example of an SR-TB is the IBM 8229 Model 2 Token-Ring to Ethernet Bridge.

There are no standards for SR-TB bridge operation.

SRT Bridges
The Source Route Transparent bridge is able to perform source-route bridging and transparent bridging simultaneously. The IEEE 802.1d base document together with its Appendix C defines the standard for SRT bridge operation.

Details of these bridging methods are discussed later in this chapter.

2.3 Bridge Operation

Bridges operate at the Media Access Control (MAC) layer, which is the lower sub-layer of the Data Link Control (DLC) layer.

A bridge consists of two (or more) physical and MAC layer connections (one for each LAN segment they interconnect). The MAC layer functions contained in the bridge are interconnected by a relay function which passes frames received from one MAC to the other MAC if certain conditions are satisfied.

Figure 6 shows that a bridge implements the physical and data link layers of the OSI Reference Model.

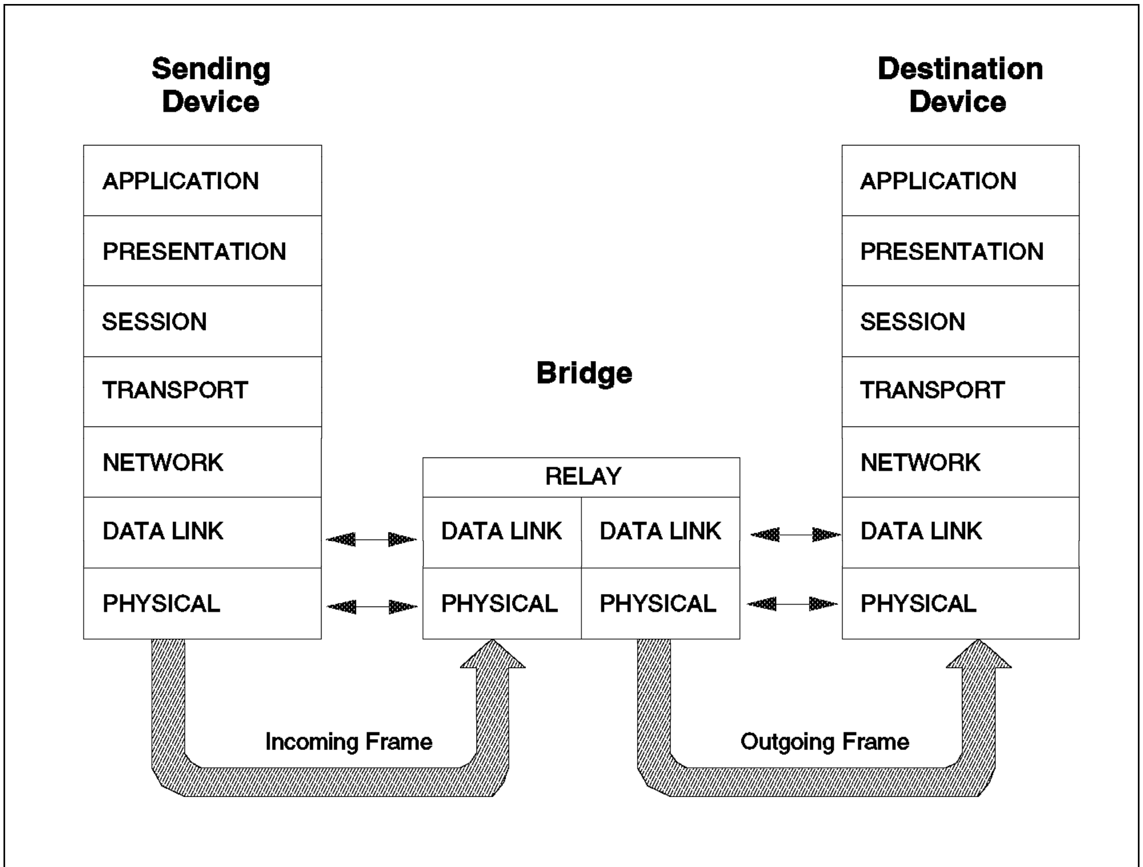

Figure 6. Bridge Implementation

MAC layer bridges are generally transparent to the users of the Data Link Control (DLC) layer. Due to this transparency to higher layers, MAC layer bridges may serve multiple higher layer protocols such as SNA, NetBIOS and TCP/IP concurrently.

Because a bridge is generally *invisible* to higher layer protocols, there are certain considerations that should be taken into account when designing a multisegment LAN. These considerations include:

- Do not inadvertently duplicate MAC addresses in interconnected LANs.
- Bridges forward all layer 2 broadcasts, with a potential danger of broadcast storms with certain protocols.
- Protocols that are connection-oriented at layer 2 must receive end-to-end acknowledgments within a time window. Using multiple bridges and/or remote bridges may result in these timers expiring, which will terminate the LLC-2 session.

2.4 Transparent Bridging

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

There are two key points about transparent bridges:

1. They need a filtering database for the decision to forward/discard a frame.
2. They allow only one active path between a pair of interconnected LANs. This is necessary to prevent frames from looping in the network and is implemented by the so-called *spanning tree algorithm*.

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

The filtering database also contains static entries. These static entries are divided into required and optional support.

Required support:

- IEEE reserved addresses, which the bridge must discard. This includes the *bridge group address* (X′800143000000′) which is used in the *spanning tree Hello BPDU*. See "Filtering Database Update" on page 42.

Optional support:

- Can be set by the user to provide customized forwarding and/or blocking of a frame or group of frames. For example, if a frame is received on port 1, forward on ports 2 and 4 but not on port 3.

Canonical Address Representation

The bridge group address has been written in the IBM format. Normally in Ethernet LANs, the canonical representation X′0180C2000000′ is used.

X′800143000000′ and X′0180C2000000′ represent the *same* MAC address, only they are written with a different bit sequence.

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

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

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

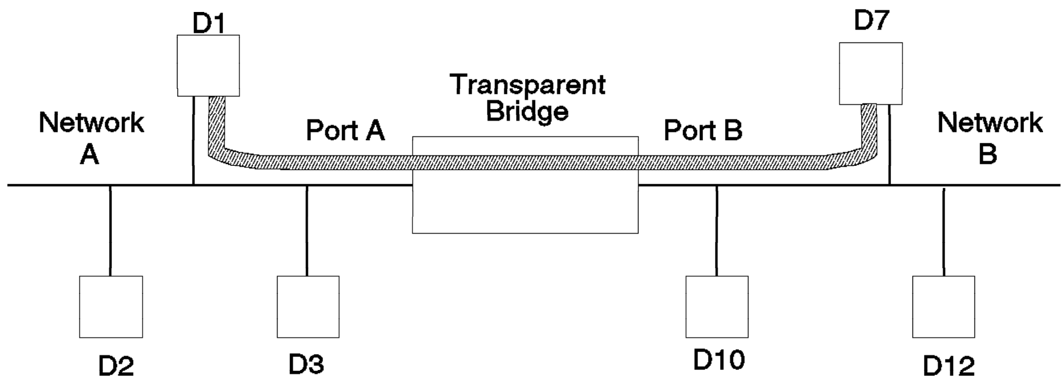

Filtering Database

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

Figure 7. Transparent Bridging Operation

Figure 7 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 can be reached via the LAN on port A. Similarly, if a frame arrives from device D7 on port B it learns that D7 can be reached via the LAN on port B.

Transparent bridges learn from source addresses, and filter or forward on destination addresses.

2.4.1 Transparent Bridge Port States

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

Disabled
: Not participating in spanning tree protocol, not learning, not forwarding frames.

Blocking
: Participating in spanning tree protocol, not learning, not forwarding frames.

Listening
: Participating in spanning tree protocol, not learning, not forwarding frames, in transition from blocking to learning. The purpose of this state is to prevent the bridge from building its filtering database, based on incorrect station information, before the spanning tree becomes stable.

Learning
: Participating in spanning tree protocol, learning, not forwarding frames. The purpose of this stage is to minimize the unnecessary forwarding of frames by ensuring that the bridge has built up its filtering database before beginning to forward frames.

Forwarding
: Participating in spanning tree protocol, learning and forwarding frames.

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

2.5 The Spanning Tree

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

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

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

If there is more than one bridge between any two interconnected LANs, the spanning tree algorithm will ensure that only one of them will be included in the spanning tree. This bridge will be in so-called *forwarding* state and is the only one that will be passing frames between the two interconnected LANs.

All the other bridges that are parallel to the forwarding bridge will be excluded from the spanning tree and will not perform any frame forwarding. These bridges are said to be in *blocking* state.

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

Figure 8 on page 35 shows an example of the spanning tree used to provide the bridging between several interconnected LANs. The upper part of the diagram shows the physical network and the lower part of the diagram shows the spanning tree and the status of the various bridges in the network.

To set up the spanning tree, bridges transmit to each other a special configuration message called *Configuration Bridge Protocol Data Unit* (BPDU). This is also referred to as the *Hello* BPDU and is used by the bridges to:

1. Choose a single bridge on the network to be the *root bridge*. This is the bridge with the lowest *bridge identifier* in the network.

   The bridge identifier is comprised of a two-byte *bridge priority* followed by a six-byte *bridge address*. The bridge priority is a configuration parameter. The bridge address is the MAC address of one of the bridge ports and its purpose is to ensure that bridge identifier is unique.

2. Choose one port on each bridge as the *root port*. This is the port that has the lowest path cost to the root bridge.

3. Elect a *designated bridge* for each LAN in the network, from among the several bridges that may be residing on that LAN. The elected bridge is the one that offers the lowest path cost from the stations residing on that LAN to the root bridge. The port on a designated bridge attached to the LAN for which the bridge is designated becomes the *designated port*. Note that in a multiport bridge, there may be more than one designated port.

Once the bridge and port roles have been calculated:

- The root port of the designated bridge enters the *forwarding* state.
- The designated port(s) on each designated bridge enters the *forwarding* state.
- All other ports remain in *blocking* state.

See “Hello BPDU” on page 38 for details about the contents of the Hello BPDU.

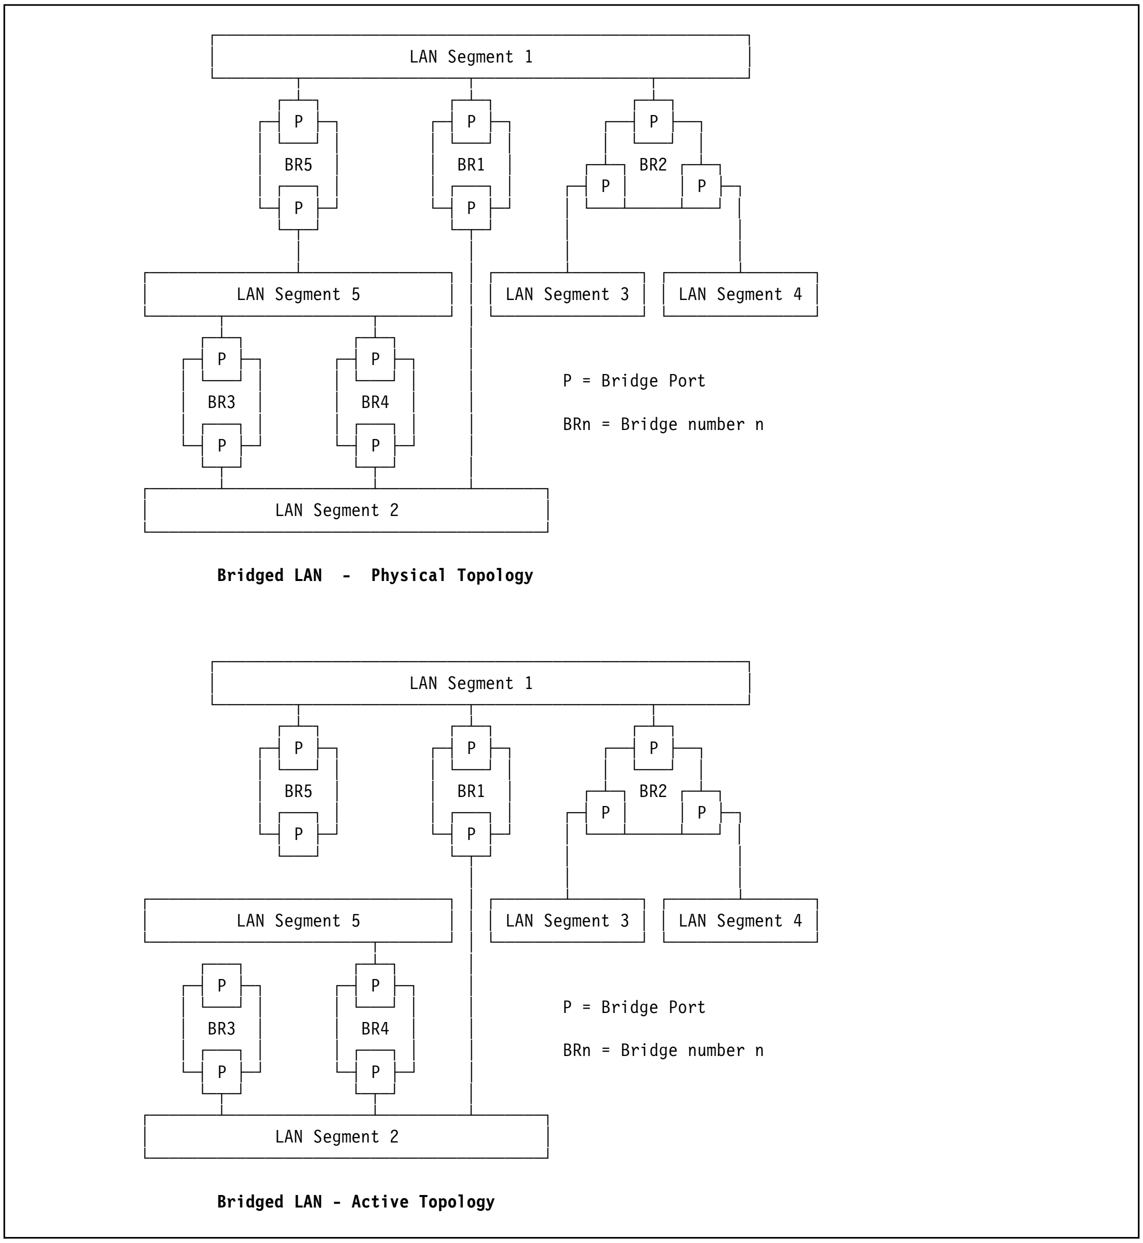

Figure 8. Example of Physical Topology. Also shown is a possible active Spanning Tree topology.

BPDUs are not broadcast frames. They are sent to bridges, and selectively updated and forwarded to other bridges.

2.5.1 The Spanning Tree Protocol

To participate in the spanning tree protocol, each bridge will initially assume it is the root bridge and will transmit a Hello BPDU on each of its ports. This message will be sent every *Hello time*. Hello time is one of the spanning tree configuration parameters that can be specified for each bridge during the bridge configuration. This Hello BPDU will have the following characteristics:

1. The source address will be the address of the transmitting bridge.
2. The destination address will be X′800143000000′.
3. The source and destination SAPs will be X′42′.
4. The Root ID Field will contain the ID of the transmitting bridge.
5. The Bridge ID Field will contain the ID of the transmitting bridge.
6. The Path Cost Field will contain 0.
7. It will be sent out by the bridge on all its ports.

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

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

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

1. The lowest Root ID
2. The lowest Path Cost to the bridge
3. The lowest transmitting Bridge ID
4. The lowest Port ID

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

1. The source address will be the address of this bridge.
2. The destination address will be X′800143000000′.
3. The source and destination SAP will be X′42′.

4. The Root ID Field will contain the Root ID received in the *better* Hello BPDU.
5. The Bridge ID Field will contain the ID of this bridge.
6. The Path Cost Field will be the sum of the path cost received in the *better* Hello BPDU plus the path cost defined for the bridge port on which the *better* Hello BPDU was received.
7. It will be sent out by the bridge on all its ports except the port on which the *better* Hello BPDU was received.

This process will be repeated by all the bridges until:

1. There is one bridge (root bridge) remaining that is still transmitting its original Hello BPDU.

2. One bridge (designated bridge) on each LAN is transmitting the Hello BPDU based on the Hello BPDU received from the root bridge.

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

Note

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

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

Figure 9 on page 38 shows how the spanning tree algorithm is used to determine root bridge, designated bridges and designated ports. In this example, bridge *BR1* becomes the root bridge since it has the lowest bridge identifier. *R* denotes the root ports and *D* denotes the designated ports.

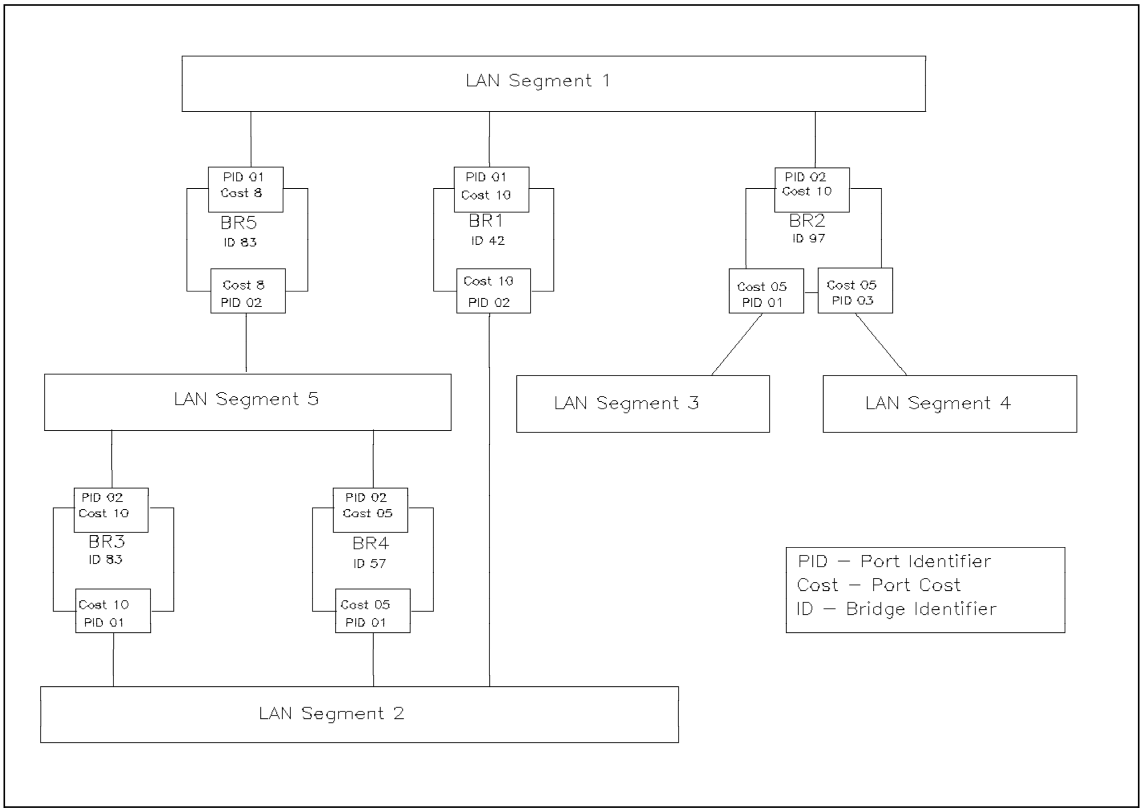

Figure 9. Bridged LAN - Physical Topology

Hello BPDU

The format of the Hello BPDU is shown in Table 12.

Table 12 (Page 1 of 2). Configuration BPDU	
Field	**Size (bytes)**
Protocol Identifier	2
Protocol Version Identifier	1
BPDU Type	1
Flags	1
Root Identifier	8
Root Path Cost	4
Bridge Identifier	8
Port Identifier	2

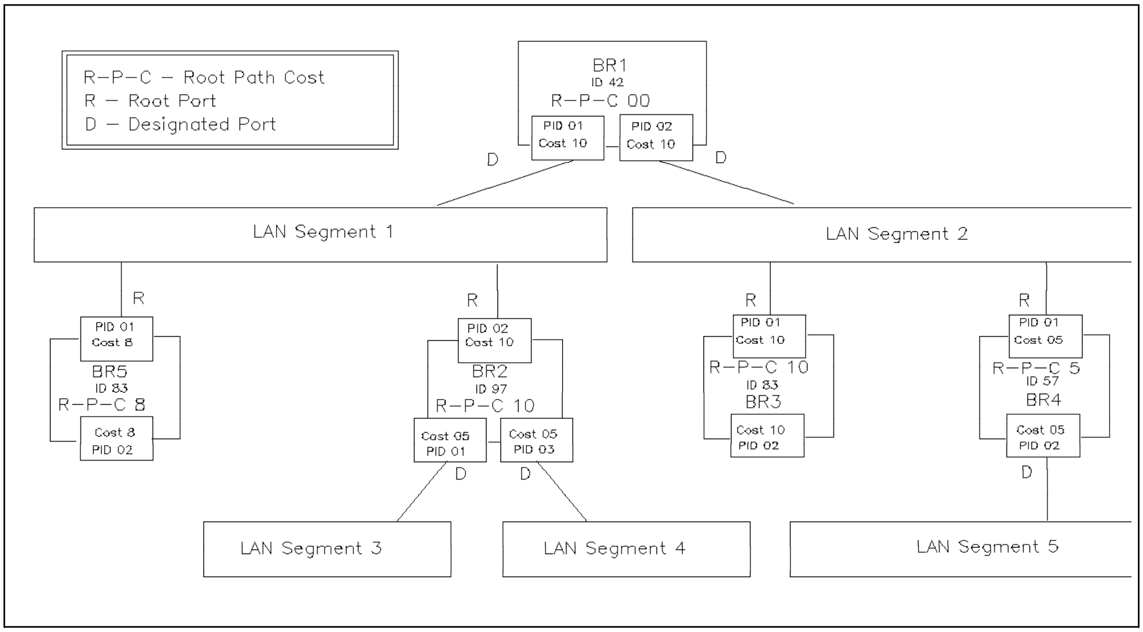

Figure 10. Bridged LAN - Logical Spanning Tree Topology

Table 12 (Page 2 of 2). Configuration BPDU	
Field	**Size (bytes)**
Message Age	2
Max Age	2
Hello Time	2
Forward Delay	2

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

- Protocol Identifier

  Identifies the spanning tree protocol. This field contains 0.

- Protocol Version Identifier

  Identifies the version number of the spanning tree protocol used. This field contains 0. Note that the IEEE 802.1d committee has proposed Version 1 for *remote* bridges only.

- BPDU Type

Denotes the type of BPDU. This field contains 0 for a configuration (Hello) BPDU and 128 for Topology Change Notification (TCN) BPDU. See 2.5.2, "Transparent Bridges and Network Topology Changes" on page 41.

- Flags
  - Topology Change

    This is the least significant bit of the Flag Field, and if set to 1 denotes that the receiving bridge should use the *forward delay timer* rather than the *aging timer* for aging out the entries in the filtering database. See "Topology Change Notification" on page 43 for more details about topology change notification.

  - Topology Change Acknowledgment (TCA)

    This is the most significant bit of the Flag Field and if set to 1, it indicates that the bridge no longer needs to send TCN BPDUs. See "Topology Change Notification" on page 43 for more details about TCA.

- Root ID

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

- Root Path Cost

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

- Bridge ID

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

- Port Identifier

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

- Message Age

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

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

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

- Max Age

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

- Hello Time

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

- Forward Delay Time

  Specifies the length of the time that the bridge should stay in each of the *listening* and *learning* states before moving from blocking to forwarding state. As discussed in “Filtering Database Update” on page 42, this timer may also be used for aging out the entries in the filtering database.

The last three parameters, Max Age, Hello Time and Forward Delay time are global parameters. To avoid confusion, they should be identically set on all bridges. Once the topology stabilizes, the values used throughout the network are those defined in the root bridge.

2.5.2 Transparent Bridges and Network Topology Changes

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

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

The result of any of the above changes is that:

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

Filtering Database Update

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

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

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

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

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

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

1. A longer timer value is to be used in coping with normal changes due to station additions, removals or timeouts. This is a user-configurable parameter and is referred to as *aging time*.

2. A shorter timer value to be used when the bridge is in a state of topology change. See “Topology Change Notification” on page 43. The *forward delay timer* of the root bridge is used for this purpose.

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

Topology Change Notification

Spanning tree topology change is detected by a bridge whenever:

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

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

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

Figure 11 on page 44 shows an example of the result of an active bridge failure or removal. The top of the figure shows a simplification of the spanning tree parameters, *p* denotes the bridge priority, while *c* denotes the path cost assigned to the ports attached to each LAN segment. The second network in the diagram shows the resulting spanning tree. A bridge fails in the third network, and the bottom network in the diagram shows the resulting new spanning tree topology.

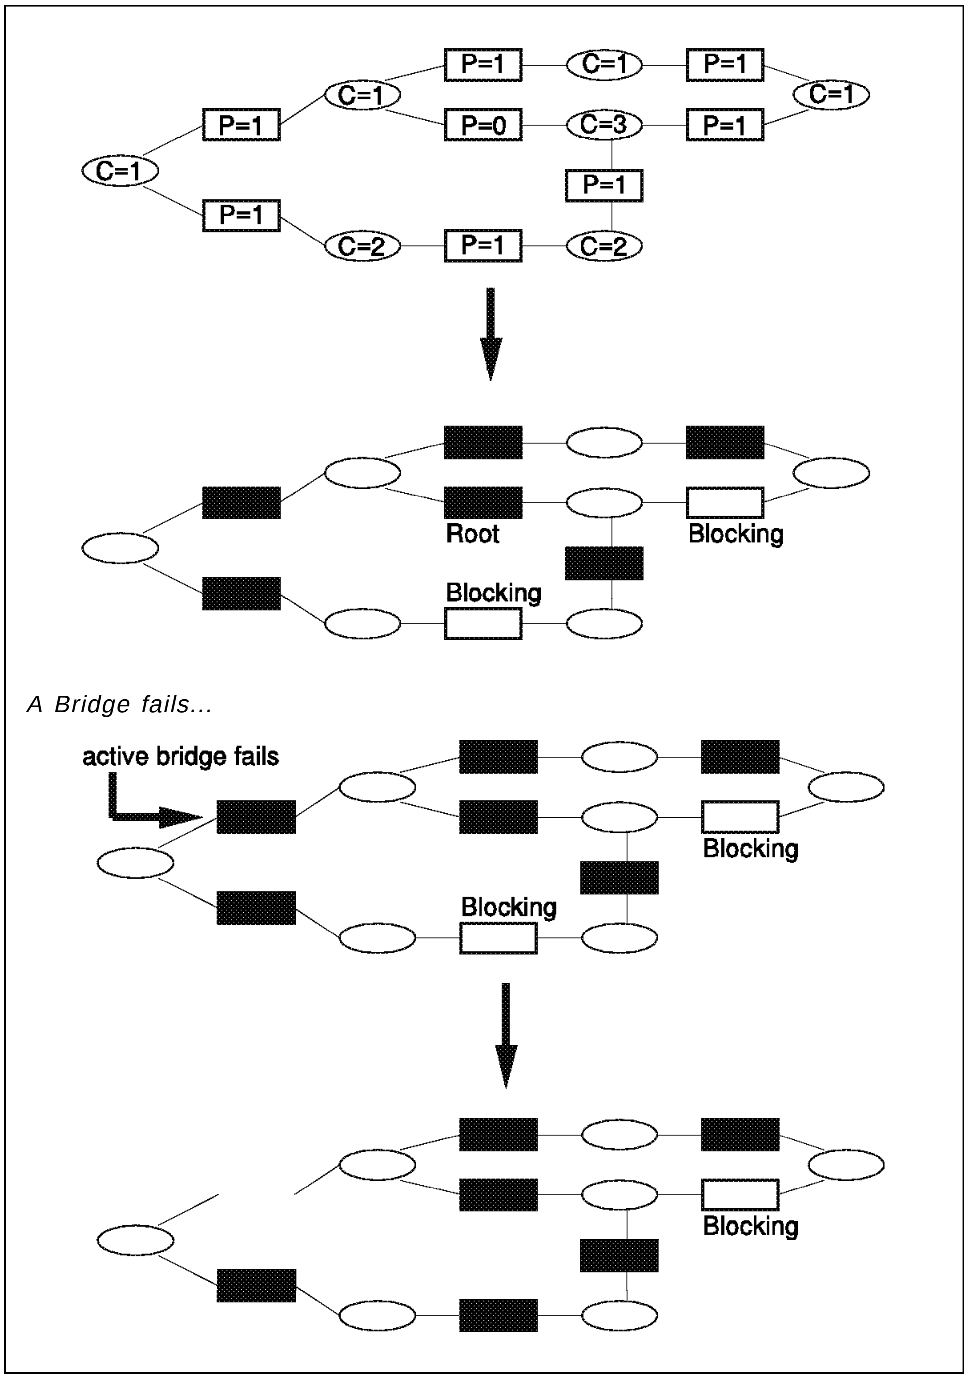

Figure 11. Reconfiguration of a Spanning Tree

2.5.3 Setting the Parameters That Control the Spanning Tree

Table 13 shows the configurable spanning tree parameters which are defined as part of the standard for transparent bridging.

Table 13. Spanning Tree Parameters

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

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

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

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

2.5.4 Summary of the Spanning Tree Algorithm

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

2.6 Source-Route Bridging

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

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

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

If present, the RIF contains at least a two-byte *Routing Control Field* and optionally may contain up to a maximum of eight (or 14) *Route Designator* Fields.

Each Route Designator Field is two bytes in length and consists of a *ring number* (12 bits) and a *bridge number* (4 bits).

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

The route designators map out the route through a multisegment LAN by specifying the sequence of rings and bridges that the frame will traverse.

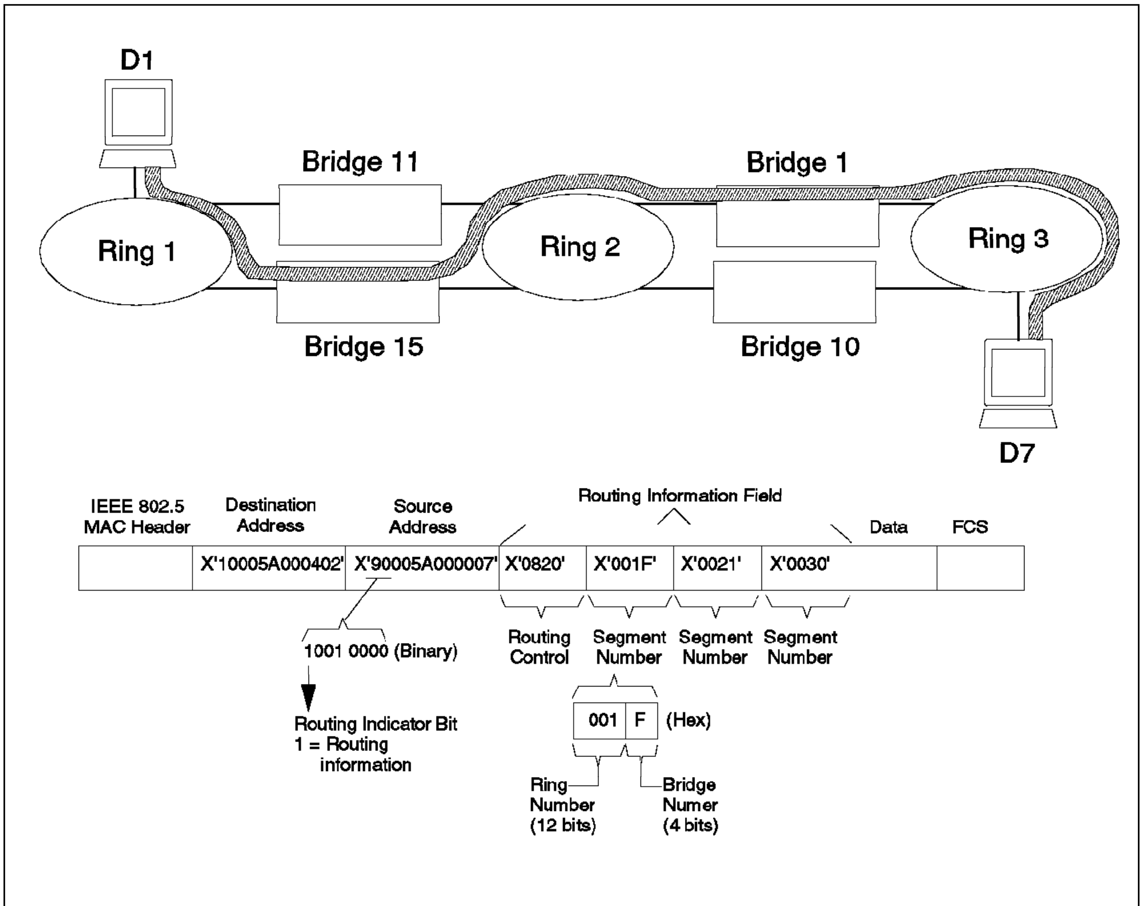

Figure 12. Source-Route Bridging RIF and RII

Figure 12 shows details of the RIF and RII fields carried in a token-ring frame.

The broadcast indicators in the Routing Control Field also control the way a bridge treats the frame. The types of source-routed frames are:

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

The term *Specifically Routed Frame* (SRF) is used by the IEEE to describe these frames.

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

In IEEE terminology, these frames are known as *All Routes Explorers* (ARE).

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

In IEEE terminology, these frames are known as *Spanning Tree Explorers* (STE).

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

2.6.1 Route Determination

Source routing requires that the originating station provides the routing information. This means that when one station wishes to send data to another station, the sending station must first obtain a route to the destination MAC address.

Route determination is often a two-stage process:

- Stage 1: On-segment route determination
- Stage 2: Off-segment route determination

The term *often* is used here because there is no formal method in IBM Token-Ring LANs for route determination. However, the on-segment/off-segment approach described here is relatively common in older token-ring products.

Note

A modification to this technique is required for correct operation in a source-route transparent environment. The modification is shown in 2.8.5, "Route Discovery with SRT Bridges" on page 70.

In the first stage:

- The source station sends a frame, usually a TEST or XID LLC Protocol Data Unit (LPDU) onto the local LAN segment[1]. This frame either has no

Routing Information Field (RII=0) or it is a non-broadcast frame with RII=1 but without any Route Designator Field (RIF=X′0270′). In either case the frame is not forwarded by any of the bridges attached to this LAN.

- The sending station then waits for some time (the time varies depending on the application) and if it does not obtain a response it goes to the second stage of route determination.

For the second stage:

- The sending station resends the TEST or XID LPDU, this time with a stub Routing Information Field with the broadcast bits set. This broadcast may either be an *all-routes* broadcast or a *single-route* broadcast. An example of an all-routes broadcast flow is shown beginning with Figure 13 on page 50 and an example of a single-route broadcast is shown beginning with Figure 14 on page 54.
- If no response is received from the target station within a period defined by the application, it is the application′s responsibility to retry (stage one and/or two) or back out.

[1] This frame is usually sent to SAP 0, a null SAP which is opened automatically by all the adapters and is used to respond to connectionless TEST/XID requests.

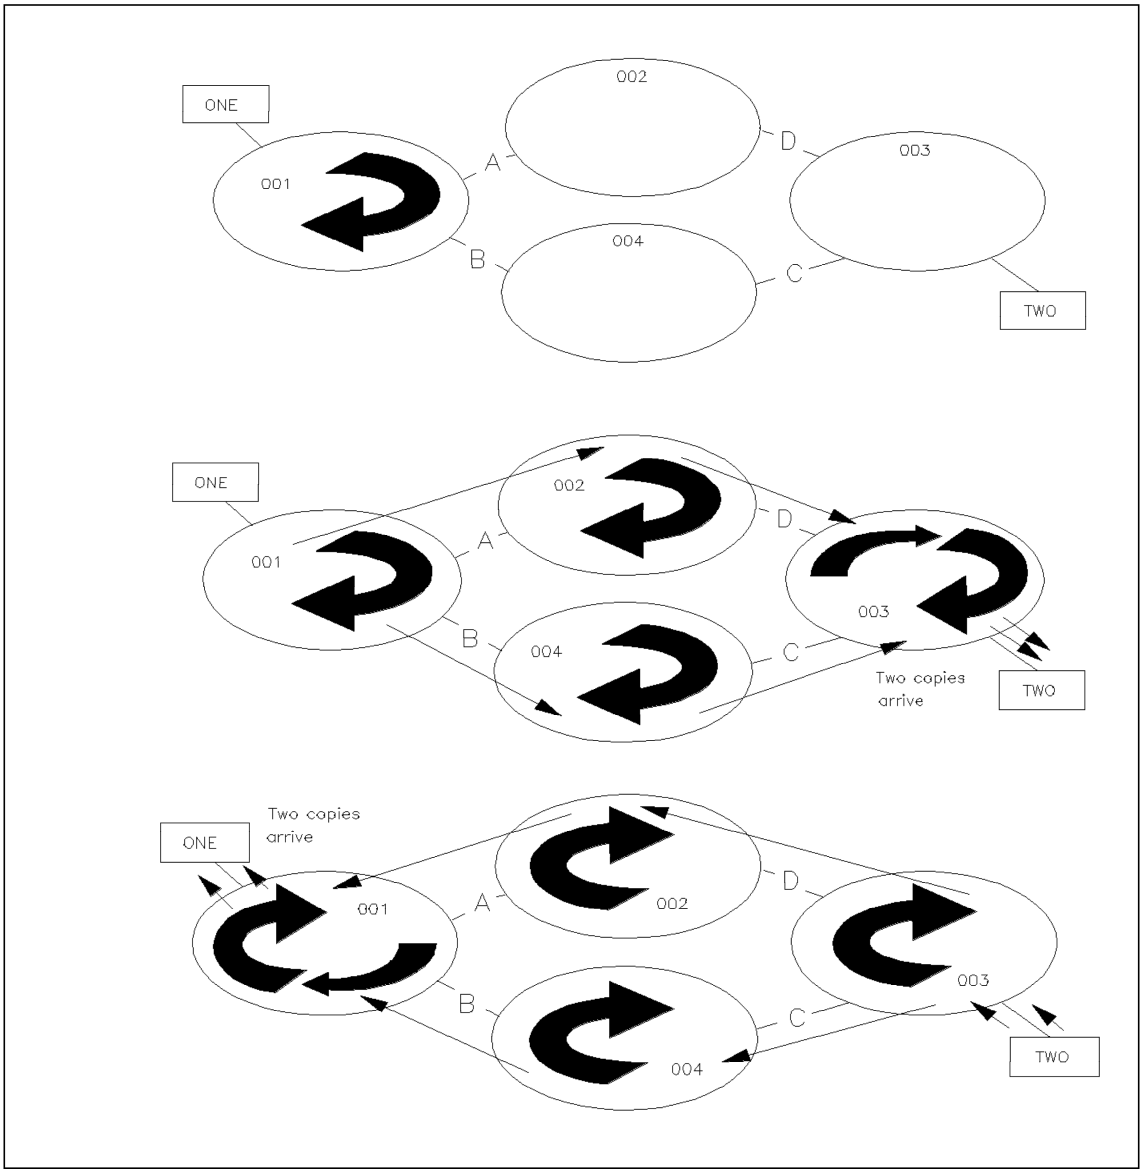

Figure 13. On and Off-Segment ARB Route Determination

On and Off-Segment ARB Route Determination

Top: Station ONE issues a TEST or XID LPDU and waits for a response.

Middle: ARB Off-segment Route Determination

TEST is issued with the all-routes broadcast bits set. Multiple copies reach target station TWO.

Bottom: ARB - The Target Machine Responds

An example of All-Routes Broadcast route determination where Station ONE is setting up communication with station TWO.

The routing information is contained in the specifically-routed frames returned to the source station. When the frames arrive at ONE, the information contained in the RIF is:

1. 001A-002D-0030
2. 001B-004C-0030

With the direction indicator set to 1 - right to left

All-Routes Broadcast Route Determination

Figure 13 on page 50 is an example of a typical route determination scheme using all-routes broadcast frames. After receiving no responses from the on-segment route determination, the sending station issues an all-routes broadcast TEST or XID frame. All the bridges forward this frame unless:

1. The frame has already been through the next segment.
2. Forwarding the frame would exceed the bridge′s all-routes broadcast hop count limit in that direction. IBM limits the number of bridges in a path to seven or thirteen, depending on the product). Each bridge allows the user to further limit how far a frame may travel in a network by setting a hop count limit for that bridge.
3. The bridge filter functions do not allow the frame to be forwarded, if the filter criteria are met.

The RIF is built up as the frame crosses the bridges:

- The sending station provides the stub Routing Control Field.
- The first bridge adds two Route Designator Fields: the first is the starting ring/bridge combination, and the second is the second segment number and a null bridge entry.
- Successive bridges then fill in their bridge number and add another two-byte Route Designator Field containing the next segment number and a null bridge entry.

The routing information of the frame being forwarded through bridge A and bridge D in the example shown in Figure 13, would build up as follows:

After bridge A, the route designators would be 001A 0020.
After bridge D, the route designators would be 001A 002D 0030.

As many frames as there are routes will be received by the target machine. The target machine responds with a non-broadcast frame, for each received frame, flipping the direction bit in the RIF. The response frames then trace back through the network and arrive at the sending station. Usually the route chosen is the route contained in the first reply, although criteria such as the minimum number of hops the frame sizes support would be equally valid. Equally there have been implementations, with programming errors, that have used the last RIF received, resulting in the slowest network path being used.

Single-Route Broadcast Route Determination

Some products implement single-route broadcast route determination for the second stage of route determination. The sending station resends an XID or TEST LPDU[2], with a stub Routing Information Field. It also sets the single-route broadcast fields in the Routing Control Field.

The primary aim of the single-route broadcast function of the IBM Token-Ring Network is to minimize the processing overhead of the target machine (or machines) by only allowing one copy of the broadcast frame on each segment in the LAN. As a result, the target station will receive a single frame.

Note: An all-routes broadcast would cause as many frames as there are possible routes to arrive at the target machine.

The single-route broadcast function is particularly appropriate to LAN server functions, where the number of search requests is significant and where the processor overhead to service multiple broadcast frames could affect response times.

The propagation of single-route broadcast frames is limited by:

- Whether a bridge is part of the single-route broadcast spanning tree
- Filters at the bridge

The single route is derived from one of the following:

2 Generally to the null SAP, SAP 0.

1. A single route derived dynamically by using the spanning tree algorithm

   The spanning tree algorithm for source-route bridges is described in 2.6.2, "The Spanning Tree in Source-Route Bridges" on page 58.

2. A static definition that results from bridge installation parameters

   In this case, the user should specify which bridges are in forwarding state and which bridges are in blocking state. This technique makes the user responsible for ensuring that there is a single route active between any two stations. Manual intervention is required to change the route. This approach is generally not recommended.

On receipt of the single-route broadcast frame, the destination station usually issues an all-routes broadcast frame, directed at the source address. The original sending station receives as many frames as there are routes in the network, from which it usually chooses the route taken by the first frame it receives, for subsequent communication.

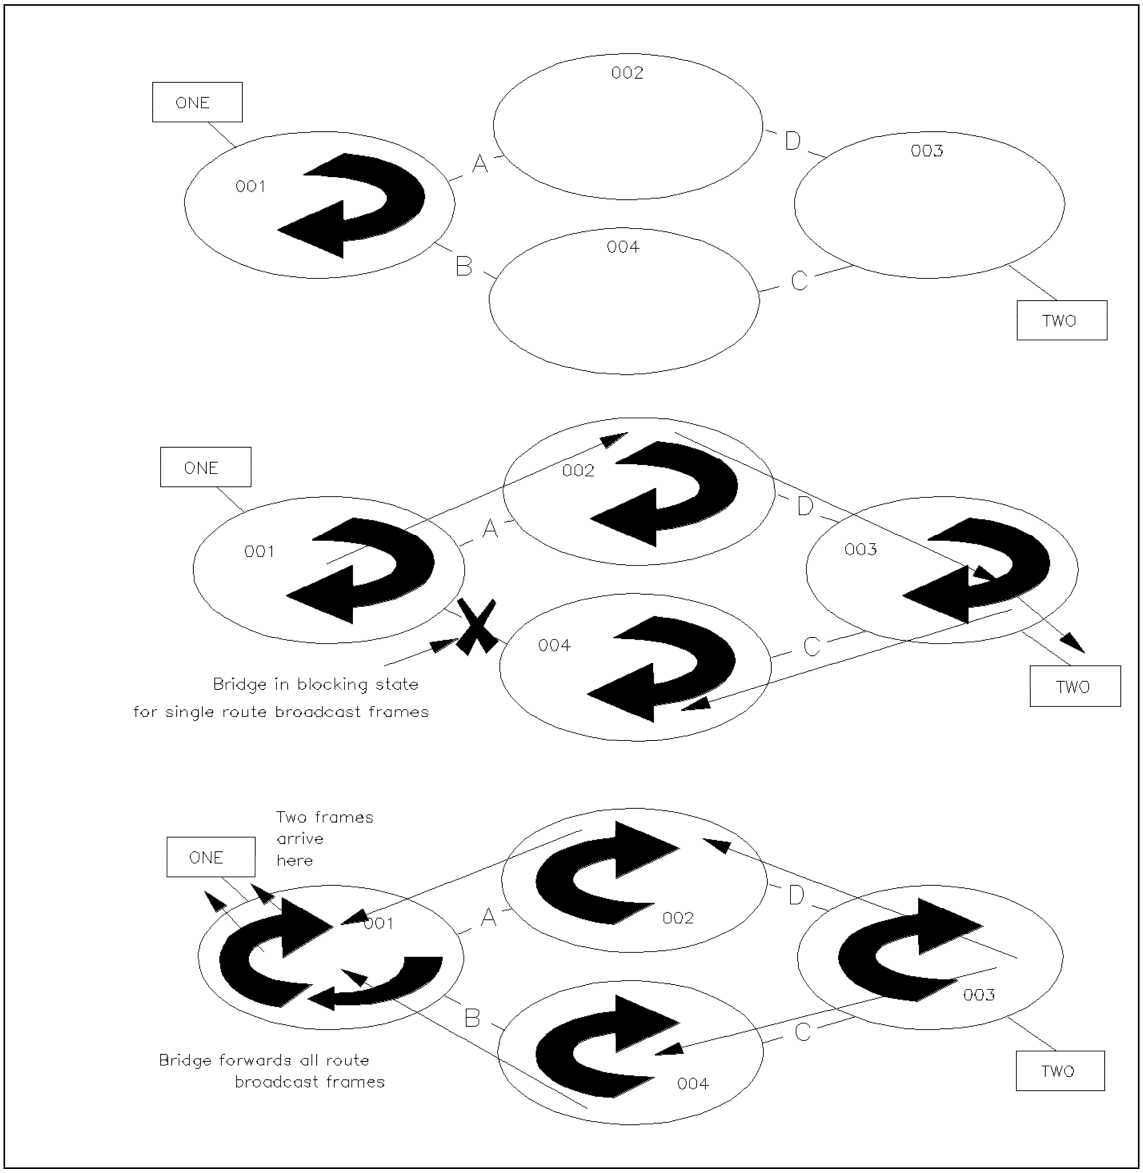

Figure 14. SRB On and Off-Segment Route Determination

SRB On and Off-Segment Route Determination

Top: Station ONE issues a TEST or XID LPDU, and waits for a response.

Middle: A TEST or XID is issued as a single-route broadcast frame. One copy reaches each segment.

Bottom Receiving Station Two Responds

Every frame received is returned using an all-routes broadcast. Two copies are received by the sending station. The Route Information Fields contain these routes:

1. 003D-002A-0010
2. 003C-004B-0010

With the direction bit set to 0 (that is, read left to right).

Other Route Determination Techniques

NetBIOS is another LAN protocol that resolves the MAC address of devices it wishes to communicate with. A typical startup sequence follows:

```
 Requester                                  Server

    SAP F0    ADD_NAME_QUERY             SAP F0    NetBIOS Functional
    ─   ──────────────────────────────────────►        Address
           Single-Route Broadcast                   C00000000080
                    .
                    .     Repeated TC times (Transmit Count)
                    .
                    .
    ─   ──────────────────────────────────────►
                    .
                    .
                    .

    SAP F0       NAME_QUERY              SAP F0
    ─   ──────────────────────────────────────►     C00000000080
           Single-Route Broadcast


    SAP F0      NAME_RECOGNIZED          SAP F0
    ◄   ───────────────────────────────────────    Device Address
             All-Routes Broadcast

    SAP F0           SABME               SAP F0
    ─   ──────────────────────────────────────►     Device Address
           Routed Non-Broadcast
                     .
                     .
```

Figure 15. Common NetBIOS Startup Sequence

Typically, a NetBIOS application first checks to see if its name is unique on the LAN. The NetBIOS ADD-NAME_QUERY frame is used for this, which is sent to the NetBIOS functional address. A NetBIOS parameter (TC) determines how often this frame is repeated.

When a NetBIOS application initiates a session with a partner application, it must first resolve the partner′s name into a MAC address. To do this, a NetBIOS NAME_QUERY containing the partner NetBIOS name is sent to the NetBIOS functional address. The partner responds with an all-routes broadcast frame, so the route and the MAC address are obtained.

This flow is modified when the NetBIOS remote name directory option is used. For more details see the *IBM Local Area Network Technical Reference*.

Another example is provided by TCP/IP. In a token-ring environment, TCP/IP′s Address Resolution Protocol (ARP), typically issues an all-routes broadcast to the all-stations broadcast address (X′FFFFFFFFFFFF′).

RFC1042, *A Standard for the Transmission of IP Datagrams over IEEE 802 Networks*, dated February 1988, provides some advice on implementing IP (Internet Protocol) and ARP (Address Resolution Protocol) on IEEE 802 LANs, including token-ring. Quoting from the RFC:

> The dynamic address discovery procedure is to broadcast an ARP request. To limit the number of all rings broadcasts to a minimum, it is desirable (though not required) that an ARP request first be sent as an all-stations broadcast, without a Routing Information Field (RIF). If the all-stations (local ring) broadcast is not supported or if the all-stations broadcast is unsuccessful after some reasonable time has elapsed, then send the ARP request as an all-routes or single-route broadcast with an empty RIF (no routing designators). An all-routes broadcast is preferable since it yields an amount of fault tolerance. In an environment with multiple redundant bridges, all-routes broadcast allows operation in spite of spanning-tree bridge failures. However, single-route broadcasts may be used if IP and ARP must use the same broadcast method.

TCP/IP RFCs are wonderfully vague at times. TCP/IP does however make extensive use of the all-stations MAC address. In conjunction with token-ring and all-routes broadcast, these all-stations packets can multiply rapidly. All the stations who are not interested or involved, still have to process these packets. Consequently, some devices may be best placed behind a bridge with the broadcast MAC address blocked.

Largest Frame Size Supported by a Bridge

In a multisegment LAN, bridges are only capable of handling certain frame sizes. The largest frame size supported by each bridge is implementation dependent, and may differ from the largest frame supported by end-stations and the largest frame supported by other bridges. Source routing provides a mechanism whereby end-stations can learn the maximum frame size supported by all the bridges in a route.

The Routing Control Field contains an entry for the maximum frame size allowed. This field is initially set by the originator of the broadcast frame to B'111'. During the forwarding process, each bridge examines this value. If it is higher than that supported by that particular bridge, the bridge lowers the value in the field to the maximum it can support. As broadcast frames are forwarded across the bridges in a path, the maximum frame size allowed for the particular route is obtained. Table 14 shows the allowed frame length values[3].

The values shown in this table are published in the *IBM Token-Ring Network Architecture Reference*, SC30-3374. With the advent of source-route transparent bridges, however, they are incorrect; for the correct values, see 2.8.6, "IBM LAN Architecture Statement" on page 71.

Table 14. Maximum Frame Sizes in the Routing Control Field

Largest Frame Field	**Size (Bytes)**	**Comments**
B'000'	516	Minimum for 802.2 LLC Type 1 (connectionless) operation
B'001'	1500	Largest frame size supported by 802.3 LANs
B'010'	2052	Typical 24x80 full screen application
B'011'	4472	Largest frame size supported by FDDI
B'100'	8144	Largest frame size supported by 802.4 LANs
B'101'	11,407	Frame size resulting from this setting
B'110'	17,800	largest frame size supported in 802.5 LANs
B'111'		Used by all-routes broadcast frames

[3] Many applications and devices do take this field into account and adjust the frame sizes they use accordingly. These applications and devices include OS/2 EE, ES/2, CM/2, Personal Communications/3270, 3174 and 3745.

2.6.2 The Spanning Tree in Source-Route Bridges

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

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

1. The Hello BPDU is sent to the bridge functional address X′C000000010 0′.
2. Port ID for a source-routing bridge consists of a *ring identifier* and *bridge number* while the Port ID for a transparent bridge consists of a *port priority* and *port number*.
3. The spanning tree in source-route bridges is used only by single-route broadcast frames. This means that bridges which are in blocking state will only block single-route broadcast frames. They will forward any all-routes broadcast frames, as well as frames which carry the appropriate routing information.

   This means that unlike transparent bridges, source-route bridges support active parallel paths which can be used for load-balancing across bridges as well as providing a backup path in case of bridge failures.
4. As there is no learning process in source-route bridges, they can be in one of three states:
   - Blocking
   - Listening
   - Forwarding
5. Source-route bridges do not support the Topology Change Notification protocol, as it is needed only to update the transparent bridge filtering database.

2.6.3 Token-Ring Spanning Tree and Hop Counts

The objective of the single-route broadcast is that frames arrive once in all LAN segments unless explicitly prevented. In a token-ring network, there is a maximum number of bridges any frame can cross. For the older generation of bridges the limit was *seven* bridges, but for the newer bridges it is *thirteen*.

Consider the token-ring network shown in Figure 16 on page 59, which depicts a complex multi-campus network. Most of the bridges are local bridges, with four wide area links (shown in gray).

The designer of the network has chosen to associate a path cost of 40 with local bridges and a path cost of 800 with remote bridges.

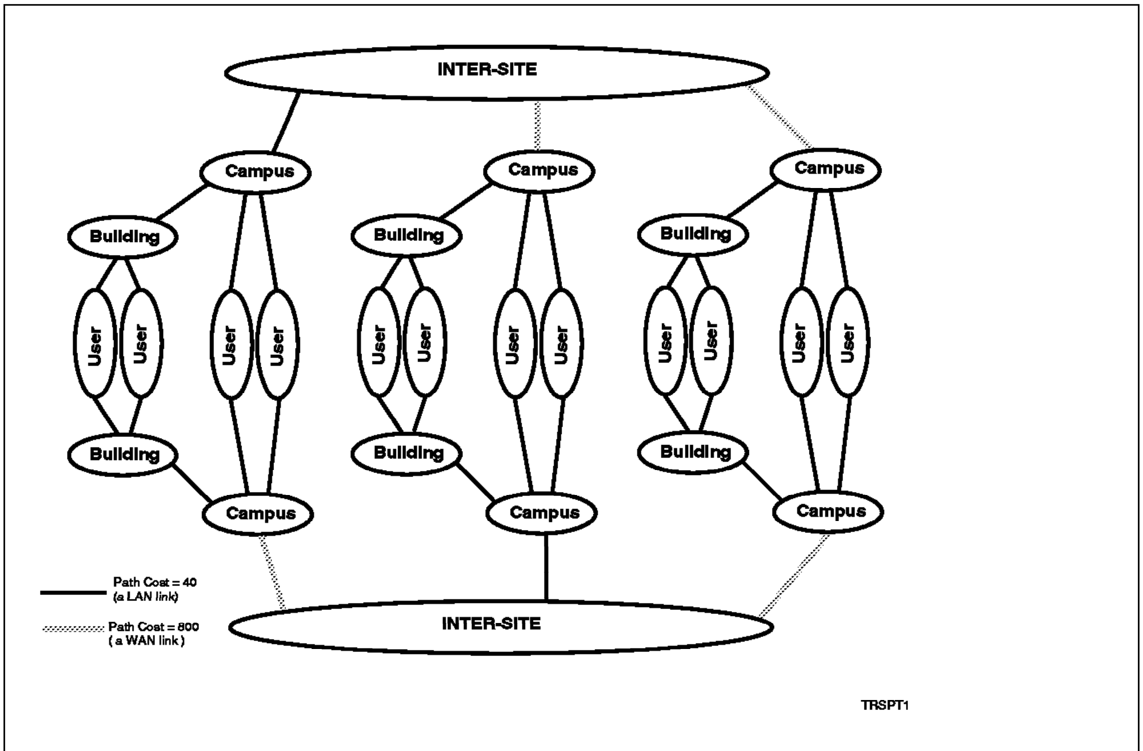

Figure 16. Complex Source-Route Network

The resulting spanning tree is shown in Figure 17 on page 60. The numbers show the path cost from each ring to the root. In this network, there are no more than seven bridges between any two points in the spanning tree. So, unless otherwise prevented, all stations can communicate with each other.

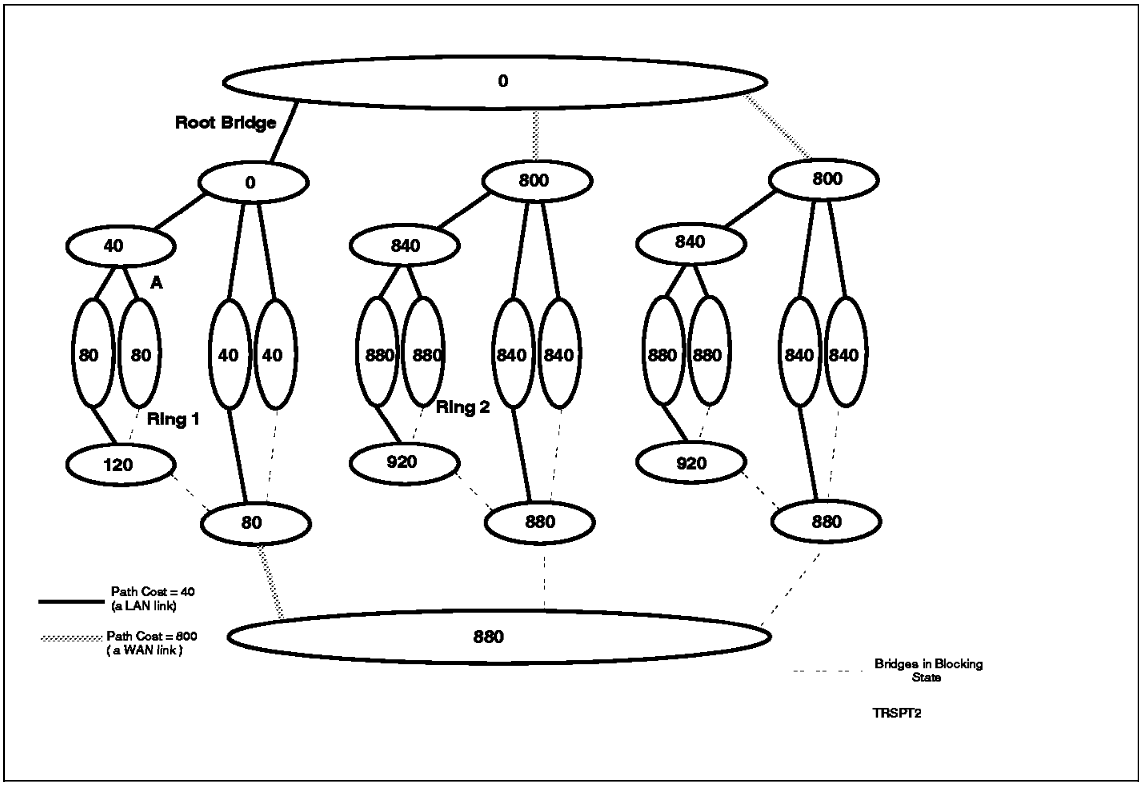

Figure 17. The Resulting Spanning Tree

Suppose however, that bridge A should fail. Users on ring 1 are *eight* bridges from ring 2. If the bridges only support seven hops, and the protocol uses single-route broadcast techniques for session establishment (for example, NetBIOS), then devices on ring 1 cannot communicate with devices on ring 2.

The same issues could be generated in a much smaller network, especially when using some bridge/router products. It is common to associate a *virtual ring* with a wide area link. Consequently, some smaller networks with apparent redundancy may fail to allow any-to-any communication.

2.7 Translational Bridges

Translational bridges interconnect one LAN type to another using bridging techniques common to each network bridged. In an OSI standards-based world this would not cause much difficulty, but unfortunately:

- Not all protocols use IEEE 802.2 LLC headers. One prime example is Ethernet, which often uses Ethernet V2 frame formats rather than IEEE frame formats.
- Some protocols use layer 2 addresses within some of their data frames.
- Some LAN protocols use source routing fields in their MAC headers (token-ring) and some do not (Ethernet).

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

2.7.1 Translational Bridging Summary

Translational bridging occurs when the media types of the connected LANs are different and bridging techniques are used. As is apparent from the previous discussion, translational bridges have many facets normally associated with routers, yet they maintain relative simplicity of setup.

2.8 Source-Route Transparent Bridging

To enable users to employ both source-routing and transparent bridging within a network to satisfy their unique requirements, the IEEE standard 802.1d has also defined Source-Route Transparent (SRT) bridging. SRT is defined in Appendix C of the IEEE 802.1 standard.

Some of the reasons for the coexistence of source routing and transparent bridging within the same network are:

- Connecting two networks which have been developed separately and in which one uses source routing while the other uses transparent bridging
- Requirement to connect both types of networks to a single FDDI network

- Extending multisegment token-ring support for protocols that do not support source routing

SRT provides an international standard for the concept of source- route bridging. The operation of source-route bridges remains a de facto standard.

In practice, SRT bridges are rarely implemented. Ethernet LANs are bridged using transparent bridges, while token-rings are bridged with source-route bridges. Ethernet and token-ring LANs are interconnected with either translational bridges or with routers. Also, FDDI networks are significantly less prevalent than it was expected when the standardization process began.

A bridge with the SRT function performs both the source route and the transparent bridge function simultaneously. To do so, the SRT bridge looks at the Routing Information indicator (RII) in the received frames. If RII=0 the frame will be handled by the transparent bridging logic while the source-routing logic will process the frame if RII=1.

Figure 18 on page 63 shows the operation of an SRT bridge.

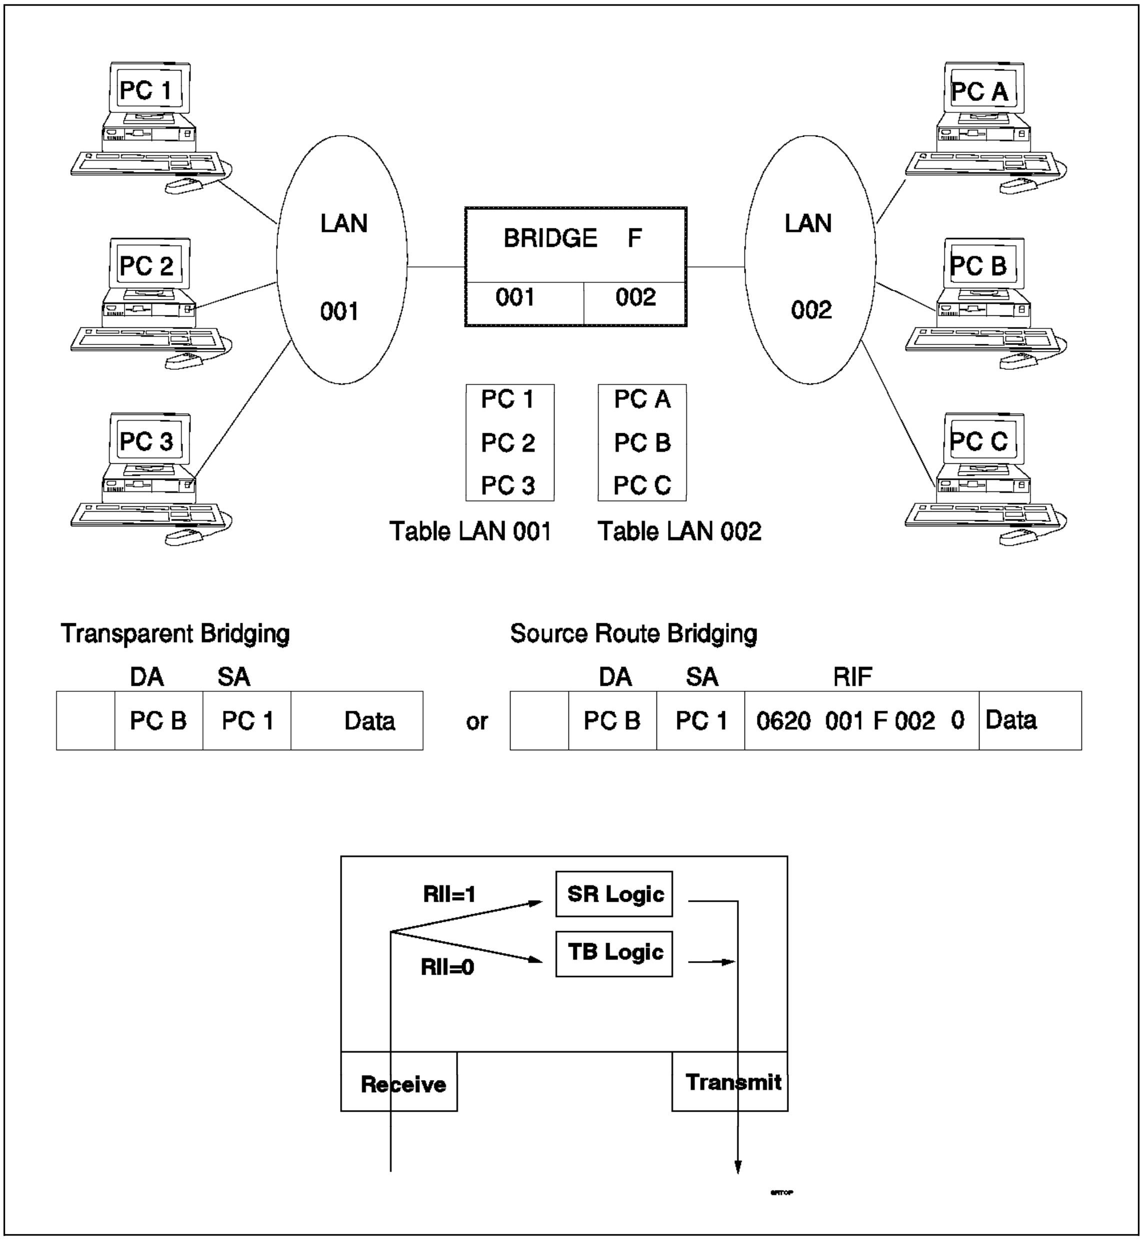

Figure 18. Source-Route Transparent Bridging

The transparent bridging logic of an SRT bridge behaves identically to a normal transparent bridge. This is described in 2.4, "Transparent Bridging" on page 30.

Single-route broadcast frames (spanning tree explorers) are forwarded on the transparent bridge spanning tree. A route designator (descriptor) is added to the frame and the largest frame field is modified if necessary.

All-routes broadcast frames (all route explorers) are forwarded on all ports using the source-route logic described in 2.6, “Source-Route Bridging” on page 46.

Although an SRT bridge acts like an SR bridge for the frames with RII=1, there are a few differences between SR and SRT as described below:

- Hop count limit

  SRT allows a maximum hop count of 13 bridges (14 segments). SRT provides two different settable hop count limits: one for all-routes broadcast and one for single-route broadcast frames.

- Largest Frame (LF) size

  Source Routing uses three bits (bits 1 through 3 in the second byte of the Routing Information Field) for specifying the LF size supported. Table 15 shows the LF sizes for SR bridging published in *IBM Token-Ring Network Architecture Reference*, SC30-3374.

Table 15. LF Size for Source Route Bridging

Code	**LF Size**
000	516
001	1500
010	2052
011	4472
100	8144
101	11407
110	17800
111	Used in all-routes broadcast only

  SRT offers two modes (Base and Extended) for the LF size. The mode used by the bridge is selected by setting an LF mode indicator.

  The base mode uses the same 3 bits as the SR, but some of the values are slightly smaller than those defined in the SR.

  Extended mode uses 6 bits (bits 1 through 6 in the second byte of the Routing Control Field) to specify the LF size.

Table 16 on page 65 shows the Extended mode LF size for SRT bridging.

Table 16. LF Size for SRT								
Base	**Extension**							
	000	001	010	011	100	101	110	111
000	516	635	754	873	993	1112	1231	1350
001	1470	1542	1615	1688	1761	1833	1906	1979
010	2052	2345	2638	2932	3225	3518	3812	4105
011	4399	4865	5331	5798	6264	6730	7197	7663
100	8130	8539	8949	9358	9768	10178	10587	10997
101	11407	12199	12992	13785	14578	15370	16163	16956
110	17749	20730	23711	26693	29674	32655	35637	38618
111	41600	44591	47583	50575	53567	56559	59551	>59551

The two previous tables are different. Consequently a protocol which uses the maximum frame size returned by source-route bridges may not work with SRT bridges.

2.8.1 SRT and Spanning Tree

When a single network consists of TB, SR and SRT bridges, the following considerations apply to the operation of the spanning tree protocol:

- There are two overlapping spanning trees within the network:
  1. TB and SRT bridges form one spanning tree.
  2. SR and SRT bridges form the other.
- SRT bridges participate actively in the TB/SRT spanning tree. This means that they issue and process Hello BPDUs addressed to X′800143000000′.
- SRT bridges participate passively in the SR/SRT spanning tree:
  - This means that they forward BPDUs addressed to X′C00000000100′, but they will neither originate nor process these frames.
  - The SRT network appears to the SR spanning tree as a single segment.
  - The SR network is invisible to the SRT spanning tree.

In summary, the SRT bridge performs both SR and TB bridging, but only one spanning tree algorithm runs in the bridge.

2.8.2 Implications of Using SRT in Existing SR Networks

In general, SRT and SR bridges can coexist in the same network. However, there are configurations that could lead to problems. The following topics deal with these problems.

2.8.3 SRT and IBM Token-Ring Bridge

During their initialization test, the IBM Bridge Program and IBM 8229 token-ring to token-ring bridge send a test frame from one port to the MAC address of the second port with no Routing Information Field (RII=0). This test fails if the frame is received on the second port. If there is an SRT bridge in parallel with the IBM Bridge Program or IBM 8229 token-ring to token-ring bridge, it will always forward this bridge test frame resulting in a test failure. Therefore, the IBM Bridge Program and IBM 8229 token-ring to token-ring bridge cannot be in parallel with an SRT bridge. This means that you should not install an SRT bridge in a loop or in parallel with the IBM Bridge Program or the IBM 8229 token-ring to token-ring bridge.

Note that SRT bridges and IBM SR bridges can be installed within the same network in a serial configuration.

When including SRT bridges in existing SR networks, care should be taken with regard to the correct operation of the spanning tree.

Source-route bridges and transparent bridges use basically the same spanning tree algorithm and protocol. There are differences that are significant to inter-operation:

1. The destination address of the BPDU
   - SR: Bridge functional address X′C00000000100′.
   - TB: Bridge group address X′800143000000′.
   - If a bridge does not recognize the BPDU address, it does *not* process the BPDU.
2. Port ID assignment
   - SR: Ring ID + bridge number.
   - TB: Port priority + port number.
   - Makes a difference to versions of the Token-Ring Network Bridge Program. These versions of the bridge program did not accept BPDUs if they appeared to come from a different LAN segment.

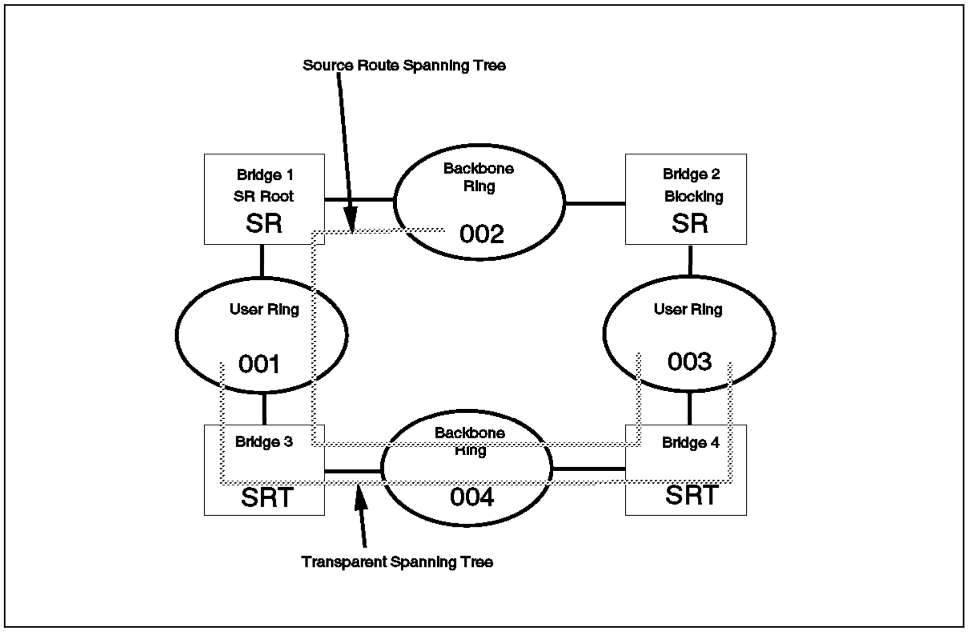

Figure 19. Source-Route and Source-Route Transparent Spanning Trees

Figure 19 shows an example of the two spanning trees in operation. For the source-route spanning tree:

- BPDUs are issued by bridge 1 onto ring 001
  - They are forwarded as TB frames by bridges 3 and 4.
  - They are *not* processed by the SRT spanning tree.
  - They are received by bridge 2 on ring 3.
- Assuming that bridge 1 has a lower bridge ID than bridge 2, then bridge 1 is the root and bridge 2 does not forward single-route broadcast frames.

The SRT network appears to the SR spanning tree as one ring.

For the SRT spanning tree:

- BPDUs are issued from bridge 3 onto ring 001
  - These use the bridge group address.
  - They are ignored (filtered) by bridge 1.
- BPDUs from bridge 4 on ring 003 are filtered by bridge 2.
- No loops appear to exist, so both SRTs are in forwarding state.

The SR network is invisible to the SRT spanning tree. In the example given, there is a sensible overlap in the two spanning trees.

With 8229 translational bridges, the situation becomes more complicated. An example is shown in Figure 20.

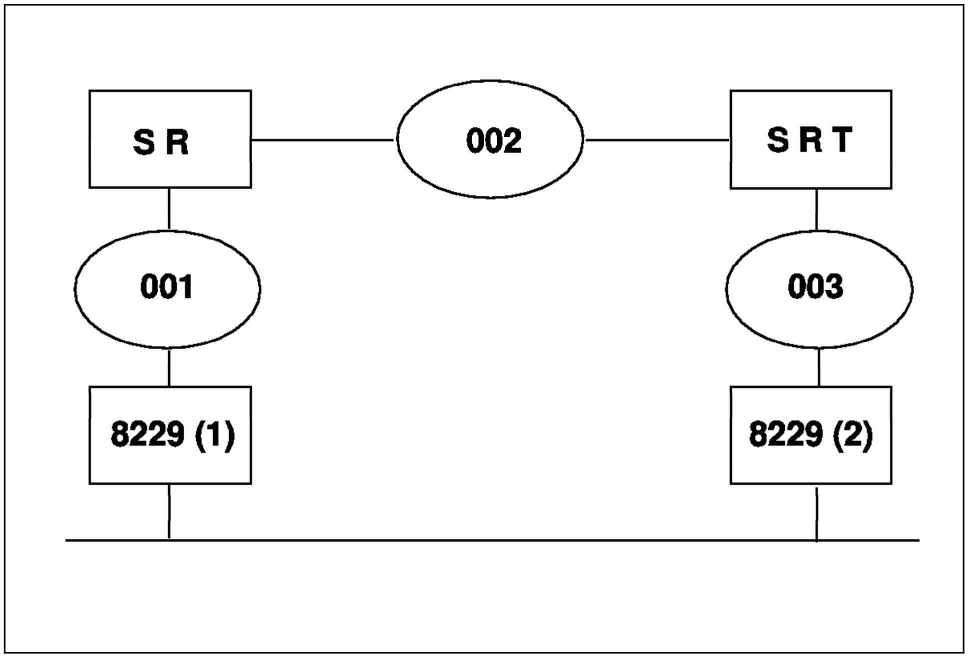

Figure 20. A Network Containing SR, SRT and 8229 Translational Bridges

The 8229 does not participate in the token-ring spanning tree. The 8229 forwards Ethernet BPDUs to the token-ring as single-route broadcast frames addressed to X′800143000000′. Consequently, the token-ring appears to the Ethernet spanning tree as a single segment. In the configuration shown:

- BPDUs issued by 8229 (2) on ring 003 are processed by the SRT bridge.
- BPDUs issued by the SRT bridge on ring 003 are processed by the 8229 (2).
- BPDUs issued by 8229 (1) on ring 001:
  - Are forwarded by the SR bridge to ring 002
  - Are processed by the SRT bridge
- BPDUs issued by the SRT bridge on ring 002 are *filtered* by the SR bridge, and are *not* processed by 8229 (1) as required for correct operation of the transparent spanning tree.

Therefore, this is a configuration that is not recommended. This configuration is characterized by:

- A backup translational 8229
- An SRT in the path between the 8229s
- An SR bridge between the 8229 and the SRT

2.8.4 SRT and High-Availability Design

The high-availability design for token-ring gateway solutions uses a dual backbone ring and duplicate gateway addresses. An SRT bridge is not recommended between rings on which duplicate addresses may appear, because the SRT filtering database associates a bridge port number with an address. This filtering database would be unstable. The value of the port number would depend on which of the duplicate adapters transmitted last.

The recommendations for high-availability design are the following:

- If duplicate addresses are used, make sure there is at least one source-route bridge in each path between the source and target adapters. That way, the duplicate address is guaranteed to appear in a transparent bridge frame on at most one port of any SRT bridge in each path. See Figure 21 on page 70 for an example of high-availability design and SRT. In this example, the roles of the bridges could be reversed.

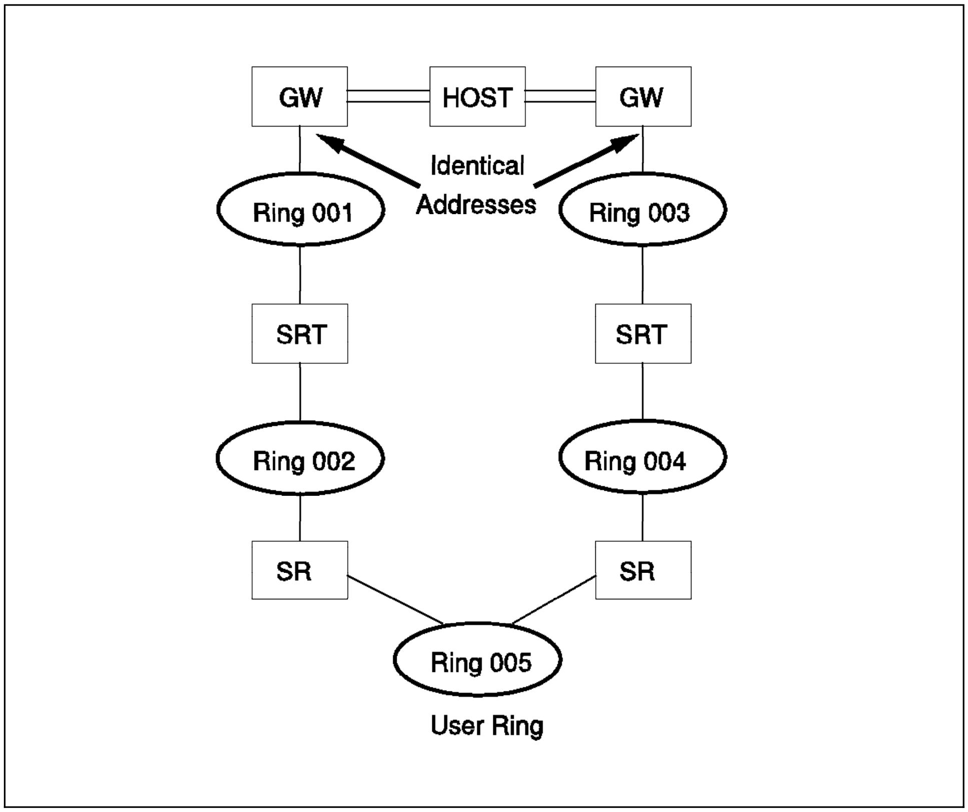

Figure 21. High-Availability Design and SRT. Equally, the SR and SRT bridge positions could be reversed.

2.8.5 Route Discovery with SRT Bridges

A common method for route discovery used by some programs (for example, PCOM/3270) is the following:

- The program sends a local TEST frame that contains no Routing Information Field (RIF) to find its TR gateway.
- An SRT bridge will forward this frame because it does not contain a RIF. This leads to two consequences:
  - The source and the destination station use the spanning tree path of transparent bridging instead of the fastest path of source routing.
  - Timing problems could occur because the source station thinks that the destination station is on the local ring.

Some recent products such as OS/2 Extended Services (ES) use a different route discovery method that anticipates the possible presence of SRT bridges:

- ES sends a local TEST/XID frame as non-broadcast with no *route descriptors*. This means that the Routing Information indicator is set to 1 and the RIF contains X′0270′.

  This frame will not be forwarded by the SRT transparent bridging function as it contains a RIF. It is not forwarded by source routing either because it does not contain a route.

- If no response, ES sends a single-route broadcast TEST/XID frame with RII set to 1. This frame will be forwarded by the source-route bridging.

- If still no response, ES sends a frame without routing information (RII=0). This frame will be forwarded by transparent bridging on the active spanning tree.

2.8.6 IBM LAN Architecture Statement

IBM made the following recommendations to all developers of LAN applications and protocols dated 10 April 1992 to address the differences between SR and SRT bridges:

> SRT bridges are appearing in token-ring, FDDI and other arenas. Changes may be needed in end-station code in the following areas:
>
> 1. Route determination steps (improved route resolution)
>
>   a. Send out the first TEST or XID with RII=1 and RT=000 LEN=2 (instead of RII=0).
>   b. If no response, then send out the second TEST or XID just as you did before (old architecture).
>   c. If no response, then try TEST or XID with RII=0.
>
> 2. Route determination entity - work is underway to include route determination in the standard LLC. The objective is to make source routing "transparent" to end-station applications (or protocol stacks). Future end-stations will use RDE instead of method 1 above.
>
> 3. A & C bits on LLC frames have NO meaning. Do NOT use returned A & C bits for anything (especially not retransmission logic).
>
> 4. LF bits - if you use LF bits to determine max frame sizes then stick to the following:

LF Bits	MAX Send Size	MIN Receive Size
000	516	same
001	1470 *	1500 *
010	2052	same
011	4399 *	4472 *
100	8130 *	8144 *
101	11407	same
110	17749 *	17800 *
111	41600 *	65535 *

* Earlier iterations of the LF bits (over the years) have reached these maximum values. If you are in a network of old bridges and old end stations, you may receive frames that are a little larger than the SRT standard sizes. In order to make sure that the frames you send get through all old and new bridges keep them smaller than the maximum send size of the new SRT bridges.

5. The D bit must be set to 0 on all-routes broadcast frames and limited broadcast frames (otherwise, SRT bridges will discard the frame).

   For specifically routed (non-broadcast) frames the following statement in 2-8 of the *Token-Ring Architectural Reference*, SC30-3374-02, is wrong! "For off-ring non-broadcast frames, the originating station sets the direction bit to B′0′....." This is not the case.

6. End-stations should be able to support the full 30-octet Routing Information Field (maximum 13 hops), even though today′s bridges only provide for seven hops.

2.8.7 SRT Summary

As always, when industry standards follow de facto standards there will be a deviation from what is de facto and what is ″standard″. These deviations will be just enough to trap the unwary.

A good example here are the CSMA/CD access methods: Ethernet V2 and IEEE 802.3. These are very much alike, but different. Similarly, source-route bridging and source-route transparent bridging are very similar, but different. The ramifications of adding SRT bridges to an existing SR-bridged network should be considered very carefully. SRT does, however, offer an international standards basis to the concept of source routing.

2.8.8 Largest Frame Size Supported by a Bridge

In a multisegment LAN, bridges are only capable of handling certain frame sizes. Source-routing provides a mechanism whereby end-stations can learn the maximum frame size supported by all the bridges in a route.

The Routing Control field contains an entry for the maximum frame size allowed. This field is initially set by the originator of the broadcast frame to B‘111’. During the discovery process, each bridge examines this value. If it is higher than that supported by that particular bridge, the bridge lowers the value in the field to the maximum it can support. As broadcast frames are forwarded across the bridges in a path, the maximum frame size allowed for the particular route is obtained. Many applications and devices do take this field into account and adjust the frame sizes they use accordingly. These applications and devices include OS/2 EE, ES/2, CM/2, Personal Communications/3270, 3174 and 3745.

With the advent of source-route transparent bridges, the values shown in *IBM Token-Ring Network Architecture Reference*, SC30-3374, are no longer valid.

SR uses 3 bits (bits 1 through 3 in the second byte of the Routing Information field) for specifying the LF size supported.

SRT offers two modes (Base and Extended) for the LF size. The mode used by the bridge is selected by setting an LF mode indicator.

The base mode uses the same 3 bits as the SR, but some of the values are slightly smaller than those defined in the SR.

Extended mode uses 6 bits (bits 1 through 6 in the second byte of the Routing Control field) to specify the LF size.

Table 17 shows the Extended mode LF size for SRT bridging.

Table 17 (Page 1 of 2). LF Size for SRT

Base	**Extension**							
	000	001	010	011	100	101	110	111
000	516	635	754	873	993	1112	1231	1350
001	1470	1542	1615	1688	1761	1833	1906	1979
010	2052	2345	2638	2932	3225	3518	3812	4105

Table 17 (Page 2 of 2). LF Size for SRT								
Base	**Extension**							
011	4399	4865	5331	5798	6264	6730	7197	7663
100	8130	8539	8949	9358	9768	10178	10587	10997
101	11407	12199	12992	13785	14578	15370	16163	16956
110	17749	20730	23711	26693	29674	32655	35637	38618
111	41600	44591	47583	50575	53567	56559	59551	>

Consequently, a protocol which uses the maximum frame size returned by source-route bridges may not work with SRT bridges. In order to make sure that the frames you send get through all old and new bridges, make sure they are smaller than the maximum size of the new SRT bridges. (Take the column with extension 000.)

2.8.9 Frame Conversion in a Translational Bridge

Translational bridges must convert the frames from token-ring to Ethernet V2 or IEEE 802.2 format. How this is done in the IBM implementation is shown below.

Ethernet V2 and Token-Ring

Table 18. Ethernet V2 to Token-Ring Conversion

Preamble			DA	SA						Typ	Info	FCS		
			xxxxxx		insert					copy				
SD	AC	FC	DA	SA	RI	SAP	SAP	Ctl	P_id	Typ	Info	FCS	ED	FS
						AA	AA	03	0-0					

The frame on Ethernet V2 has an Ethertype field that must be preserved on token-ring. The only way to transport this kind of information on an LLC based frame is in a SNAP type frame. Therefore the bridge inserts (besides routing information from its database) a SNAP header X′AAAA03000000′ (AA is the SAP for SNAP services). In the reverse direction, this header will be removed. Only SNAP type frames can be converted because Ethernet V2 expects an Ethertype to be present.

Table 19. Token-Ring to Ethernet V2 Conversion

SD	AC	FC	DA	SA	RI	SAP	SAP	Ctl	P_id	Typ	Info	FCS	ED	FS
			xxxxxx		delete					copy				
Preamble			DA	SA						Typ	Info	FCS		

SNA or NetBIOS type frames would lose valuable information in this type of conversion.

LLC-over-Ethernet

In order to transport SNA, NetBIOS, or another LLC based protocol over Ethernet V2, a bridge must use a special Ethertype of 80D5 for LLC-over-Ethernet. When this is done, the bridge at the other end will know in which way to convert the frame back to token-ring format.

Table 20. Token-Ring to LLC-over-Ethernet V2 Conversion

SD	AC	FC	DA	SA	Routing Info			SAP	SAP	Ctl	Info	FCS	ED	FS
			xxxxxx		replaced by			copy						
Preamble			DA	SA	Type	Len	Pd	SAP	SAP	Ctl	Info	Pd	FCS	
					80D5									

Bridges require that you specify for which DSAP values this kind of conversion is to be done. The default values are X′00′, ′04′, ′08′, ′F0′ and ′F4′ but you can add or delete any SAP from the list. DSAP X′AA′ is not in the default list: SNAP frames can transfer over native Ethernet V2 in a format that is understandable to all stations. SNAP could travel over Ethernet V2 in the LLC format as well, but stations on Ethernet V2 are unlikely to understand this.

The reverse translation follows:

Table 21. LLC-over-Ethernet V2 to Token-Ring Conversion

Preamble			DA	SA	Type	Len	Pd	SAP	SAP	Ctl	Info	PAD	FCS	
			xxxxxx		replaced by			copy						
SD	AC	FC	DA	SA	Routing Info			SAP	SAP	Ctl	Info	FCS	ED	FS

IEEE 802.3 and Token-ring

Conversion between IEEE 802.3 and token-ring is a lot easier. Both protocols carry LLC information therefore they have comparable formats.

Table 22. Token-Ring to IEEE 802.3 Conversion

SD	AC	FC	DA	SA	Routing Info	SAP	SAP	Ctl	Info	FCS	ED	FS
			xxxxxx		replace by	copy						
Preamble			DA	SA	Length	SAP	SAP	Ctl	Info	Pd	FCS	

And in the reverse direction:

Table 23. IEEE 802.3 to Token-Ring Conversion

Preamble			DA	SA	Length	SAP	SAP	Ctl	Info	Pd	FCS	
			xxxxxx		replace by	copy						
SD	AC	FC	DA	SA	Routing Info	SAP	SAP	Ctl	Info	FCS	ED	FS

2.8.10 To Bridge, Switch or Route? There Is No Easy Answer

General observations:

- Bridges are truly multiprotocol. They forward all packets and protocols regardless of whether the protocols may be routed or not.
- Bridges are generally transparent to end systems.
- Bridges are closer to *plug and play* than routers and require less expertise to install. Once installed, they generally function with minimum attention.
- Switches are faster than bridges in many circumstances.
- Switches are close to plug and play.
- Switches can support individual segments or workstations as desired.
- Switches share many characteristics of bridges.
- Routers may prove more flexible than bridges when interconnection of dissimilar LAN types is required.
- Routers require end systems to be configured with topology information, or to run a routing protocol passively to discover and maintain topology information. Therefore routers are more complex and more difficult to maintain than bridges.

Chapter 3. IBM Bridge Products

This chapter describes IBM bridge and switch products. Its objective is to present the characteristics of these products. For installation instructions and parameter descriptions, refer to the related product documentation.

3.1 IBM Local Token-Ring Bridge/DOS V1.0

The IBM Local Token-Ring Bridge/DOS Version 1.0 is intended to connect two token-ring segments with a data transfer rate of 4 Mbps or 16 Mbps. It runs on a *dedicated* computer that cannot be used for any other operation while performing as a bridge.

Communication across the bridge is transparent to applications written to the IEEE 802.2 standard logical link control interface.

The IBM Local Token-Ring Bridge/DOS Version 1.0 can be configured to automatically configure the network single-route broadcast path. All bridges in the network must participate in the single-route selection. Otherwise single-route broadcast frames may arrive at some segments, and may not arrive at all at other segments.

3.1.1 Maximum Frame Size for Local Bridge/DOS

When the Local Token-Ring Bridge/DOS configuration utility is displayed, a parameter called maximum frame size will be shown. This parameter cannot be modified by the bridge operator. It is set by the bridge program based on the installed adapters.

Adapters

Table 24. Maximum Frame Sizes (Bytes) for Local Bridge with MCA or ISA

		Primary Adapter		
		Adapter II or(/A) (4 Mbps)	**16/4 Adapter (/A) (4 Mbps)**	**16/4 Adapter (/A) (16 Mbps)**
Alternate Adapter	Adapter II or (/A) (4 Mbps)	2052	2052	2052
	16/4 Adapter (/A) (4 Mbps)	2052	4472	4472
	16/4 Adapter (/A) (16 Mbps)	2052	4472	8144

3.2 IBM LANStreamer Token-Ring Bridge/DOS V1.0

The LANStreamer Token-Ring Bridge/DOS is a local bridge program that connects two token-ring segments with data transfer rates of 4 Mbps or 16 Mbps. It is used with the LANStreamer 32-bit adapter, providing high performance throughput, but to take full advantage of the throughput capabilities of the adapters, it must run on a PS/2 PC with a 32-bit MCA bus.

This program runs on a dedicated computer.

The LANStreamer Token-Ring Bridge is based on the IBM Local Token-Ring Bridge and provides the same bridging functions. Major differences are its performance and the hop limit extension from 7 to 13. You must configure the LANStreamer bridge for a hop count maximum of 7, unless every network adapter and its device drivers and applications supports are capable of supporting 13 hops.

The IBM LANStreamer MC32 adapter delivers performance right up to the maximum limits of current token-ring technology. Micro Channel adapters usually transfer data at a minimum rate of about 20 Mbps. The LANStreamer adapter has 40 Mbps maximum streaming data-burst capacity. The design allows frames to be moved directly between the token-ring adapter and system memory without requiring intermediate storage in the adapter. This contributes to a latency reduction (up to 50% lower latency than the IBM Local Bridge/DOS), which results in better overall performance.

The LANStreamer adapter settings must be adjusted, taking the following aspects into consideration:

- The primary adapter is the adapter in the lowest numbered slot.
- The same interrupt level must be used for both adapters.

3.2.1 Maximum Frame Size

The maximum frame size supported by the LANStreamer Token-Ring Bridge is 4472 if any of the adapters are running at 4 Mbps, and 8144 if both are running at 16 Mbps.

3.3 IBM Remote Token-Ring Bridge/DOS V1.0

The IBM Remote Token-Ring Bridge/DOS V1.0 enables communication between devices that are connected to different LAN segments over a wide-area connection. Each LAN segment is connected to a bridge half. *Remote bridges* are also called *split bridges*.

The IBM Remote Token-Ring Bridge/DOS V1.0 maintains an internal protocol over the telecommunications line to transfer the frames between the two bridge halves.

Like the Local Bridge Program/DOS, the Remote Bridge Program/DOS must run on a dedicated computer.

Bridge performance reporting has information that reflects split-bridge specific information. A *frame not forwarded, filtered* counter is maintained to report the result of frame filtering by a user exit. In addition, another counter is used to accumulate the number of frames which are not forwarded because a cyclic redundancy check (CRC) error is detected on the telecommunications link.

Support is provided by the IBM LAN Network Manager Version 1.0 or later.

3.3.1 Line Speeds and Interfaces Supported

As telecommunications network environments vary from country to country, you should refer to the specific announcement letter in that country to verify availability of the following method of network attachment:

- Via synchronous modems, providing the following interfaces at the indicated speeds:
  - EIA RS-232C/CCITT V.24 at 9.6 Kbps to 19.2 Kbps
  - CCITT V.35 at 9.6 Kbps to 2.048 Mbps
  - X.21 bis V.24 at 9.6 Kbps to 19.2 Kbps
  - X.21 bis V.35 at 9.6 Kbps to 1.344 Mbps
  - X.21 (non-switched) at 9.6 Kbps to 64 Kbps
- Via a multiplexor, such as the Integrated Digital Network Exchange (IDNX) Models 20, 70, and 90, through:
  - The USD or HSD-2 communications adapter using CCITT V.35 at 9.6 Kbps to 1.344 Mbps

– The QSD communications adapter using EIA RS-232C/CCITT V.24 at 9.6 Kbps to 19.2 Kbps

– The QSD communications adapter using CCITT V.35 at 9.6 Kbps to 56 Kbps

An additional bridge parameter, *Communications Adapter Electrical Interface*, supports specification of the interface used by the communications adapter to attach to the telecommunications link. Valid options are 1 (RS-232), 2 (V.35) and 3 (X.21).

For those configurations not specifically identified above, the following table indicates the interfaces and speeds with which the communication adapters (in conjunction with IBM Remote Token-Ring Bridge/DOS) can be used.

Table 25. Supported Communications Adapters - Interfaces and Speeds

Interface	**ISA Computer**	**Micro Channel Computer**	
	Realtime Interface Co-Processor	**X.25 Interface Co-Processor/2**	**High Speed Communications Co-Processor/2**
RS 232C/V.24	9.6 to 19.2 Kbps	9.6 to 19.2 Kbps	9.6 to 19.2 Kbps
V.35	9.6 to 64 Kbps	9.6 to 1.344 Mbps	9.6 to 2.048 Mbps
X.21 bis/V.24	--	9.6 to 19.2 Kbps	--
X.21 bis/V.35	--	9.6 to 1.344 Mbps	--
X.21 (leased)	--	9.6 to 1.344 Mbps	9.6 to 2.048 Mbps

Note: The X.21 bis/V.24 electrical characteristics are compatible with EIA RS-232C. The X.21 bis/V.35 interface is equivalent to the CCITT V.35 interface.

The High Speed Communications Co-Processor/2 is not available in all countries.

3.3.2 Dial Support

The IBM Remote Token-Ring Bridge/DOS can be configured to run in switched mode and to use the remote dial feature to dial and connect to its remote partner. This requires:

- Support of a V.25 bis modem, typically operating at 9.6 Kbps, such as the IBM 7855-10 V.32 Modem
- An IBM-provided or user-written application program that will request the “local” half of a remote bridge to make a call

Dial support offers:

- An ability for a bridge to initiate outgoing calls, accept incoming, or both.

- Password protection, so that unauthorized users cannot initiate calls, by indicating the status of the connection attempt.
- Call takedown when:
  - Requested by the operator
  - No frames have been forwarded over the bridge for a user-specified interval
  - Certain line error conditions occur

3.3.3 Maximum Frame Size for Remote Bridge/DOS

The maximum frame size that the split bridges can support is affected by the telecommunication line data rate:

1. Transmitting long frames on slow links takes time, and therefore impacts timer values used between communicating stations.
2. Long frames could mean that the bridge becomes congested more easily.
3. Slow analog links are more prone to suffer errors than higher-speed digital links.

The maximum frame size also depends on the amount of storage allocated within the bridge machine for the communication adapter *transmit buffer size*.

Table 26. Maximum Frame Sizes and Communications Adapter Transmit Buffer Size Default Values

Line Data Rate	**Maximum Frame Size (bytes)**	**Communications Adapter Transmit Buffer Size (bytes)**
9.6 Kbps	516	1100
19.2 Kbps	516	2200
56 Kbps	2052	6417
64 Kbps	2052	7489
512 Kbps	2052	56320
1.344 Mbps	2052	166112
2.048 Mbps	4472	256000

3.3.4 Remote Token-Ring Bridge Partners

Two other IBM products can be one of the halves for a split bridge configuration. They run in remote token-ring bridge compatibility mode. These products are:

- IBM 6611 Network Processor
- IBM 8229 Bridge (must work as a secondary bridge)

3.3.5 Cascading Considerations

Cascading of remote bridges is not restricted; however, several factors must be considered when implementing this configuration. Telecommunications line speed is a critical factor in determining whether applications will work across cascaded remote bridges. The link protocol used (such as LLC Type 2) may have timer values that must be met for the logical link to be maintained. In addition, the application itself may have certain timing constraints. These timing requirements may not be satisfied, particularly at slower line speeds, and must be considered when using the remote bridge in a cascaded configuration.

3.4 The Filtering Facility

The preceding products come with pre-defined filters and a filter program to use in writing custom filters. You can use filters to limit the traffic that can cross a bridge. You can limit any of the following:

- Broadcast traffic
- The number of sessions flowing across a bridge
- The flow of a particular type of traffic, for example, to exclude NetBIOS or TCP/IP
- Traffic to a particular MAC address or group of addresses

The filtering facility is especially useful in a remote bridge environment, where telecommunications links have line data rates that are much lower than their LAN and local bridge transmission rates.

A bridge filter is a TSR (terminate and stay resident) program.

At bridge startup, it is possible to specify up to ten filter program names. The criteria that is used to discard or forward frames is controlled by the user when the filter is started.

The IBM Token-Ring Bridge/DOS programs come with the following filters and filter-related files:

FILTER.ASM A sample filter program

FILTER1.COM The link-limiting filter (supplied only with the Remote Token-Ring Bridge/DOS)

FILTER2.COM The NetBIOS filter

FILTER3.COM The address filter

FILTER3.COM The SAP address filter

FILTER5.COM The SNAP filter

3.4.1 Link-Limiting Filter

The FILTER1.COM file is the link-limiting filter of the Remote Token-Ring Bridge/DOS. The purpose of this filter is to limit the traffic crossing the telecommunications link connecting the two bridge halves by limiting the number of pairs of unique source and destination addresses that can exchange information at one time through a remote bridge.

The filter maintains a table of source and destination address pairs. The size of the table is determined by the LINKS parameter. If the table is full and a new link is desired, the table will be searched for a link that has not had any traffic for the duration specified by the TIME parameter. If there is a table entry that had no traffic within this time period, the new source and destination address pair is put in this entry. If there is no such entry, the new link will not be established (the frame is filtered).

The link-limiting filter applies only to the Remote Token-Ring Bridge/DOS and must only be installed in one bridge half.

3.4.2 NetBIOS Filter

The NetBIOS filter program for both the Local, LANStreamer and Remote Token-Ring Bridge/DOS is named FILTER2.COM. The purpose of this filter is to restrict the proliferation of NetBIOS frames throughout the network.

Filtering NetBIOS traffic is an important performance consideration in large LANs. NetBIOS establishes sessions by broadcasting to the NetBIOS functional address to search for a partner name. All devices with the NetBIOS SAP (and therefore the NetBIOS functional address) opened perform a check to see if the name requested is their own. The use of this filter also allows duplicate NetBIOS names to coexist on the network on opposite sides of the bridge.

3.4.3 Address Filter

The FILTER3.COM file is an address filter. It is shipped with both the Local, LANStreamer, and Remote Token-Ring Bridge/DOS. This filter program forwards or discards (depending on the ACTION specification) all frames that contain source and destination addresses within a range specified on the filter program load command.

The filter can discard an errant all-stations broadcast in order to relieve network traffic. This is particularly useful in a TCP/IP environment, as TCP/IP can use all-stations all-routes frames during its address resolution process.

3.4.4 SAP Filter

The Service Access Point (SAP) filter program supplied on both the IBM Local and Remote Bridge/DOS is FILTER4.COM. You can filter frames that originate from a source SAP (SSAP) range or are sent to a destination SAP (DSAP) range specified on the filter program load command.

3.4.5 SNAP Filter

The Sub-Network Access Protocol (SNAP) filter program supplied on both the IBM Local, LANStreamer, and Remote Bridge/DOS is FILTER5.COM. You can filter frames of Ethernet traffic over token-ring networks based on the SNAP Ethertype range.

The filter program first checks frames for the LSAP header X'AA AA 03'. If found, the Ethertype field is checked for a match with the TYPE value or the range specified in the filter load command.

3.4.6 Filter Combinations

It may be necessary to have multiple filters (a maximum of 10) active concurrently. Filters are executed in the REVERSE order in which they were activated. *Care should be taken in planning and using filters*. It is advantageous to execute filters that discard more frames before filters that discard fewer frames. Execution in this order minimizes the amount of code executed for a frame that will eventually be discarded. A forwarded frame will sequence through and execute all the filters.

3.5 IBM LAN Bridge Manager/2 V1.0

The Bridge Manager provides remote, centralized management of token-ring bridges across multi-segmented 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 Novell Server and the Agent resides on the bridge machine.

The following information can be viewed from your workstation:

- DOS version installed
- Bridge program version installed
- Bridge network address
- Type, microcode level and MAC address of the token-ring adapter

The following are the functions of the Manager component:

- Manage and distribute bridge filters
- Reboot the machine remotely (time delay included)
- Display the bridge hardware and software configuration
- Display statistics obtained from the Agent activity log
- Provide asset inventory information of each bridge machine
- Manage multiple levels of the bridge program
- Mass change the link password
- Provide a bridge installation and setup function for server setup

The following are the functions of the Agent component:

- Simplify the software installation, setup and customization of the bridge machines
- Manage the distribution of all software between the LAN Server and the bridge machines
- Provide information on the machine hardware and software configuration
- Forward alerts to LAN Network Manager

3.6 IBM 8229 Bridge

The IBM 8229 bridge is a modular, rack-mountable plug-and-play PC-based bridge. However, the 8229 has two processors: the 80486SLC processor performs bridging exclusively while an 80186 processor handles management of the 8229. The 8229 is a functional replacement for the IBM 8209 bridge. Three 8229 operation models are available:

- 8229-001 for connections between two local token-ring segments
- 8229-002 for connections between a local token-ring segment and a local Ethernet Version 2 and IEEE 802.3 segment
- 8229-003 for connections between a local token-ring segment and a remote token-ring segment via a wide area network (WAN) up to T1/E1 speeds

The software loaded in the 8229 flash memory determines which attachment modules can be used and what type of network management the 8229 can use: CMOL or SNMP.

The 8229 is a two-port MAC bridge. One of these ports has to be a local token-ring attachment. The second can be either another token-ring, Ethernet, or a WAN adapter.

One or two attachment modules can be installed:

Single Port Token-Ring Module
Dual Port Token-Ring Module
Single Port Ethernet Module
Single Port Wide Area Network Module

The 8229 appears as a source-routing bridge to stations on a token-ring segment and as a transparent bridge to stations on an Ethernet LAN. Additional functions are:

- Filtering
- Novell IPX protocol support
- Frame format conversion

These functions can be enabled, disabled or assigned values by using the 8229 LAN Bridge Utility Program and IBM LAN Network Manager to set advanced configuration parameters.

The 8229 is available as a module in the 8260. With the level of microcode used in this 8260 module, the 8229 is able to act as a SRT (Source-Routing Transparent bridge).

Note: The 8229 participates in spanning tree for SR-bridging. Its ancestor, the 8209, supported the spanning tree protocol but did not participate.

3.6.1 The Filtering Facility

The 8229 supports two types of built-in filters:

Criteria Range Filter

The criteria ranges for this filter consist of an offset and two ranges of allowed values. These 2-byte values compare to a 2-byte field at specified offset into the frame. The filtering field offset 0 begins at the type or length field for Ethernet frames and at the DSAP for token-ring frames.

Range filters may be used to prevent frames of a particular type from traversing the bridge or, conversely, to allow only traffic of a particular type to pass through the bridge.

Address Range Filter

The 8229 Address Range Filter discards all frames that contain source and destination addresses within source and destination address ranges specified in the filter configuration parameters. A single address filter can be set for each adapter on the 8229.

User-Written Filter

You may write your own filter program and have it loaded into the bridge. It will be retained in the 8229 flash memory. The filter program must be compiled for an 80386 processor operating in 32-bit protected mode and be no larger than 32 KB in executable size.

3.6.2 8229 Management Support

The 8229 can be configured to support *one* of two management options at a time: CMOL or SNMP. An example of a management program that uses CMOL is LAN Network Manager or LNM.

IBM LAN Network Manager 2.0 provides token-ring and Ethernet network management and configuration support for the 8229.

3.6.3 Token-Ring to Token-Ring Connectivity (8229-001)

The 8229-001 implements the source-route bridge function between two token-ring segments connected by the 8229. The 8229-001 can be configured with either a dual-port token-ring attachment module or two single token-ring attachment modules.

The 8229 token-ring bridge implements all the token-ring network management functions associated with the IBM Local Token-Ring Bridge/DOS Version 1.0. These include the functional addresses (LAN Bridge Server, Ring Error Monitor, etc.), as well as the capability for single-route broadcast path maintenance.

3.6.4 Maximum Frame Size

The 8229 token-ring to token-ring bridge has a larger maximum allowable frame size than the IBM Local Token-Ring Bridge/DOS. The largest transmit frame size is set by the speed of the ring that the frame is going to be forwarded to:

- For frames forwarded to 4 Mbps rings the maximum frame size is 4472 bytes.
- For frames forwarded to 16 Mbps rings the maximum frame size is 17800 bytes.

The 8229 token-ring to token-ring bridge provides performance improvements over the IBM Local Token-Ring Bridge/DOS program implementation, especially at smaller frame sizes.

3.6.5 Token-Ring to Ethernet/IEEE 802.3 Connectivity (8229-002)

An IBM 8229 with the single-port token-ring network and single-port Ethernet attachment modules installed provides a token-ring to Ethernet V2/IEEE 802.3 bridging function.

Ethernet Version 1 is not supported by the 8229.

The IBM 8229 Token-Ring/Ethernet bridge is designed to perform the necessary conversion of frames. The bridge is:

- Able to interpret the frames of Ethernet Version 2 and convert them into token-ring frames (and vice versa)
- Able to interpret the frames of the IEEE 802.3 protocol and convert them into token-ring frames (and vice versa)

The IBM 8229 also allows an Ethernet backbone to be used for carrying token-ring traffic, though this configuration would mean that frames between token-ring stations would be limited to lengths of less than 1500 bytes, the maximum allowed on an Ethernet. This may be restrictive for applications, given that token-ring LANs can carry much larger frames.

3.6.6 Modes of Operation

The 8229 has two modes of operation when bridging traffic from a token-ring segment to an Ethernet segment. Mode 1 is for token-ring to Ethernet V2 forwarding and Mode 2 is for token-ring to IEEE 802.3 forwarding.

If the automatic mode is enabled, the 8229 determines the format by checking its Ethernet database for that station address. If there is no entry in the database for that station, the format chosen by the mode priority is taken. If no mode priority is yet specified, the 8229 must obtain the default from the selection switch (DIP switch) on the Ethernet attachment module.

3.6.7 LLC over Ethernet

With Mode 1 (Ethernet V2) selected, frames with an LLC header such as SNA or NetBIOS cannot be converted and will be discarded, unless *Forward LLC traffic* is enabled. The SAP values for those frames that you want to forward must be included in the enabling list. By default the SAPs X′00′, X′04′, X′08′, X′F0′, X′F4′, and X′FC′ are enabled. These frames receive on Ethernet a special Ethertype X′805D′ to identify them to the bridge at the other side as LLC-over-Ethernet.

3.6.8 SNAP Support

Native Ethernet V2 to token-ring translation also requires a special format. An LLC header and an SNAP header (Sub Network Access Protocol) of X′AAAA03000000′ are put in front of the Ethertype. The SAP X′AA′ should not be added to the LLC-traffic enable list. The bridge can translate this SNAP format directly into Ethernet V2 again. IEEE 802.3 uses the same SNAP format.

To make it possible for the bridge to choose the right protocol on Ethernet, an adapter must not use both protocols simultaneously. Unfortunately, this is not always the case:

- AIX has the ability to use both formats for TCP/IP with a single adapter.
- DECnet Phase IV uses Ethernet V2, however DECnet/OSI uses IEEE 802.3 and to facilitate migration the device will respond to both formats.

- NetWare allows many different frame formats.
- Some adapters shared by multiple applications will use Ethernet V2 for one application and IEEE 802.3 for another application.

3.6.9 Novell IPX Support

Novell NetWare uses IEEE 802.3 frames without an LLC header. The frames have a value of 'FFFF' just after the length field. These frames are recognized by the 8229 bridge and handled accordingly. Refer to the IPX frame examples.

Table 27. Novell IPX Conversion

Ethernet V2	-->	Token-ring LLC (SAP E0)
IEEE 802.3	-->	Token-ring LLC (SAP E0) **1**
IEEE 802.3 LLC (SAP E0)	-->	Token-ring LLC (SAP E0) **1**
IEEE 802.3 SNAP (SAP AA)	-->	Token-ring SNAP (SAP AA)
Token-ring SNAP (SAP AA)	-->	IEEE 802.3 SNAP (SAP AA)
Token-ring LLC (SAP E0)	-->	IEEE 802.3 LLC (SAP E0) (if mode 2) **2**
Token-ring LLC (SAP E0)	-->	IEEE 802.3 (if mode 2) **2**
Token-ring LLC (SAP E0)	-->	Ethernet V2 (if mode 1)

Notes:

- **1** At the first occurrence of an IPX frame, the 8229 will set an indicator whether to use IEEE 802.3 or IEEE 802.3 LLC. The bridge remembers this indicator until reset.
- **2** The bridge uses prior set indicator to determine if IEEE 802.3 or IEEE 802.3 LLC will be used.
- Just as noted for SAP AA, SAP E0 should not be included in the enabling list for LLC over Ethernet. Doing so would create strange encapsulations.

3.6.10 Token-Ring to WAN Connectivity (8229-003)

The 8229-003 (remote token-ring bridge function) connects a token-ring network operating at 4 Mbps or 16 Mbps to a wide area network (WAN). As in other remote bridge product implementations, each 8229 is half of a split-bridge configuration.

The 8229 WAN attachment module communicates with one of the following bridges at the other end of a WAN:

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

If you are using the Remote Token-Ring Bridge DOS-based programs communicating with the 8229-003, they must be the primary bridge half of the bridge while the 8229 must be the secondary.

3.6.11 WAN Connectivity

The 8229 WAN attachment module supports serial data rates from 9.6 Kbps to 2.048 Mbps with timing supplied by the network interface equipment. Operation is supported for dedicated path connections only.

Table 28. Electrical Interface and Line Data Rates

Electrical Interface	Supporting Interface Cable	Line Data Rate
V.24	EIA-232C Attachment	9.6 Kbps to 19.2 Kbps
V.35	V.35	9.6 Kbps to 2.048 Mbps
X.21	X.21	9.6 Kbps to 2.048 Mbps

3.6.12 Configuration and Operation

In a split bridge configuration, the primary half of the bridge determines the parameters that control the secondary half of the bridge, which is located on the other side of the WAN segment. That is the reason for the following:

- In a token-ring to WAN configuration using two 8229s, the Utility Program must be linked to the 8229 defined as the primary half of the split bridge.
- If a Remote Token-Ring Bridge/DOS program is used with the 8229, you cannot use the 8229 Utility Program. You must designate the workstation with the Remote Token-Ring Bridge/DOS program as the primary half of the bridge and set the parameters on that bridge, using the remote bridge configuration programs or the IBM LAN Network Manager Program. The 8229 accepts the parameters set for the primary half of the bridge.

Each bridge half uses its own filter selections. In this configuration, the 8229 may use any of the supported filters (address, criteria range or user-defined filter) and the remote bridge DOS-based programs may use a user-defined filter.

3.6.13 Maximum Frame Size

As in the case of the Remote Token-Ring Bridge/DOS, the *maximum frame size* that the bridge can process is affected by the line data rate of the telecommunications link connecting the two bridges. When you use the default value for this parameter, the 8229 uses the values shown in Table 29. The maximum frame size is limited to 4472 if either of the token-ring segments to which the 8229 is connected to is running at 4 Mbps. The 8229 discards frames that are larger than the maximum frame size.

Table 29. Maximum Frame Sizes - Default Values

Line Data Rate	**Maximum Frame Size (bytes)**
Up to 9.6 Kbps	516
19.2 Kbps to 38.4 Kbps	1500
38.4 Kbps to 512 Kbps	2052
512 Kbps to 1.544 Mbps	4472
1.544 Mbps to 2.048 Mbps	8144

3.7 IBM Frame Relay Token-Ring Bridge/DOS V1.0

The IBM Frame Relay Token-Ring Bridge/DOS V1.0 is a source-routing bridge that supports:

- Point-to-Point Leased Line mode (PPP) for direct connection to:
  - Another IBM Frame Relay Token-Ring Bridge
  - An IBM RouteXpander/2 Version 1.0
  - An IBM RouteXpander/2 Version 2.0
  - An IBM RouteXpander/2 Version 2.0 with multiport support
  - An IBM 6611 Network Processor.
- Point-to Point Frame Relay mode for indirect connection via a frame relay network to another IBM Frame Relay bridge or equivalent product, such as 6611, 3172, 3174, including OEM equipment, provided it uses the standard RFC 1294 protocol.

The bridge offers support for IBM LAN Network Manager for the ring that is locally attached.

The WAN can transfer data at rates from 9.6 Kbps to 2.048 Mbps using:

<table>
<tr><td colspan="12">Table 30 (Page 1 of 2).</td></tr>
<tr><td colspan="12">IBM Wide Area Connector (WAC) adapter for line speeds from 9.6 Kbps to 2.048 Mbps.</td></tr>
<tr><td colspan="3">CCITT V.35: for speeds up to 2.048 Mbps</td><td colspan="3">CCITT X.21: for speeds up to 2.048 Mbps</td><td colspan="3">EIA RS 422/449: for speeds up to 1.544 Mbps</td><td colspan="3">EIA RS 232D: for speeds up to 19.2 Kbps</td></tr>
<tr><td colspan="12">IBM X.25 Interface Co-Processor/2 and Realtime Interface Co-Processor DOS Support, V3.0 or later.</td></tr>
<tr><td colspan="4">Cable Option V.35: for speeds up to 1.344 Mbps (64 Kbps on ISA-bus)</td><td colspan="4">Cable Option V.24: for speeds up to 19.2 Kbps</td><td colspan="4">Cable Option X.21: for speeds up to 1.344 Mbps (64 Kbps on ISA-bus)</td></tr>
</table>

Table 30 (Page 2 of 2).		
IBM High Speed Communications Co-Processor/2 and Realtime Interface Co-Processor DOS Support, V3.0 or later (MCA-bus only).		
Cable Option V.35: for line speeds up to 2.048 Mbps	Cable Option V.24: for line speeds up to 19.2 Kbps	Cable Option X.21: for line speeds up to 2.048 Mbps
IBM Realtime Interface Co-Processor and Requires Realtime Interface Co-Processor DOS Support, V3.0 or later (ISA-bus only).		
CCITT V.35 Electrical Interface Board: for line speeds up to 64 Kbps	EIA RS-232C/CCITT V.24 Electrical Interface Board: for line speeds up to 19.2 Kbps	

Performance of the IBM Frame Relay Token-Ring Bridge/DOS V1.0 is dependent on which processor and adapters are used.

3.8 RouteXpander/2 V 1.0.1

RouteXpander/2 runs on OS/2 V2.0 or later and provides the following functions:

- Access to frame relay networks
- Source-route bridging and Transparent bridging
- IPX, IP, and SNA/APPN routing

RouteXpander/2 is designed as a low-cost solution for bridging and routing of small LANs within an enterprise. It requires the IBM Network Transport Services/2 (NTS/2) product, or. a similar function such as the newer Multi-Protocol Transport Services (MPTS).

3.8.1 RouteXpander/2 As a LAN Bridge

RouteXpander/2 V 1.0.1 with PTF WR20307 provides the following bridge functions when properly configured:

- Source-routing bridge between two locally attached token-ring LANs or;
- Source-routing bridge between a locally attached token-ring LAN and remote token-ring LANs connected through a frame relay connection or;
- Transparent bridge between two locally attached Ethernet LANs or;
- Transparent bridge between a locally attached Ethernet LAN and remote Ethernet LANs connected through a frame relay connection

The bridge functions provide support for the spanning tree protocol.

Source-routing and transparent bridge functions cannot be active at the same time, nor can RouteXpander/2 perform translational bridging between Ethernet and token-ring LANs, as the IBM 8229 Bridge can do. Also, since RouteXpander/2 V 1.0.1 only supports two adapters, the local and remote bridging functions cannot be active at the same time.

The remote connection uses frame relay protocol over either a point-to-point connection or a frame relay network. The physical connection is via an IBM Wide Area Connector (WAC) adapter.

- CCITT X.21 for up to 2.048 Mbps
- CCITT V.35 for up to 2.048 Mbps
- EIA RS-422/449 for up to 1.544 Mbps

The RouteXpander/2 frame relay device driver can support a maximum of 200 DLCIs over a frame relay connection. Therefore, the traffic from the local LAN can be bridged to LANs connected to a maximum of 200 remote bridges. These bridges can be any combination of:

- 6611 Network Processor
- PS/2 running the frame relay Token-Ring Bridge/DOS program
- PS/2 running RouteXpander/2
- 3172 Model 3 MPE (with RouteXpander/2)
- 3174 with Frame Relay Communication Feature (Config. Support C6)

3.8.2 RouteXpander/2 Bridging Filters

RouteXpander/2 filters bridge traffic based on information that it obtains from an OS/2 file, GSDSRB.INI (for source-route bridging) or GSDTB.INI (for transparent bridging) in the IBMCOM directory. This file is a text file that can be edited using any editor.

For more details about RouteXpander/2 filters, please refer to *RouteXpander/2 User's Guide and Reference*, SC31-6606.

3.8.3 Managing RouteXpander/2

To manage RouteXpander/2 via an SNMP manager, TCP/IP for OS/2 must be installed on the RouteXpander/2 machine.

To manage RouteXpander/2 via NetView or any other SNA management application, Communications Manager/2 (CM/2) must be installed on the RouteXpander/2 machine.

If you want to learn more about the use of RouteXpander/2 as a bridge, its functions and utilization, please refer to *RouteXpander/2 Introduction and Configuration Examples*, GG24-4334.

3.9 RouteXpander/2 V 2.0

RXR/2 V2.0, like V1.0.1, is an OS/2-based LAN/WAN program product. It can function as a multiprotocol router and/or a bridge.

RXR/2 V2.0 does *not* support transparent bridging (this is provided by RXR/2 V1.0.1).

3.9.1 RouteXpander LNM Support/2

By upgrading RXR/2 V2.0 with LNM Support/2, RXR/2 V2.0 configured for source-route bridging is enabled for the following reporting functions:

- Ring Parameter Server
- Configuration Report Server
- Ring Error Monitor

This is equivalent to the LNM support in IBM′s Local and Remote Bridge products. LNM may be used to monitor a RXR/2 bridge with this support installed.

3.9.2 RouteXpander X.25 Support/2 for RXR/2 Version 2.0

With the RXR/2 X.25 Support/2 program, RXR/2 V2.0 can attach to a public or private X.25 Packet Switching Data Network (PSDN) service.

Up to 200 virtual circuits are supported by X.25 Support/2.

Both Wide Area Connector (WAC) adapters and RTIC family adapters are supported.

3.9.3 RouteXpander Multiport Support/2

RouteXpander′s Multiport Support/2 provides source-route bridging for multiple LAN and WAN ports. A total of nine ports are supported:

- A maximum of two LAN ports and seven WAN ports can be bridged.
- A maximum of one LAN port and eight WAN ports can be bridged.

The WAN ports can be a combination of frame relay, X.25 and point-to-point. The following limitations apply:

- A maximum of two frame relay ports
- A maximum of one X.25 port
- A maximum of eight point-to-point connections

3.10 IBM 6611 Network Processor

The 6611 Network Processor is a multiprotocol, multiport router and bridge. Its hardware platform is based on RISC technology, and its software support is provided by the IBM Multiprotocol Network Program, which is preloaded into the 6611 at the time of manufacturing.

In the following sections, we′ll discuss the use of the 6611 as a LAN bridge. For detailed information about other 6611 functions, please refer to *The IBM 6611 Network Processor*, SG24-3870.

There are three models of the 6611: The 6611 Models 120, 140, and 170 have two, four, and seven slots respectively for connection of any of the LAN and WAN adapters. Model 120 comes in eight predefined configurations, each with one WAN and one LAN adapter.

The 6611 supports a variety of LAN and WAN adapters. The following adapters are available:

- Token-ring network adapter
- Ethernet network adapter
- Two-Port EIA 422/449 Serial Adapter
- Two-Port V.35/V.36 Compatible Serial Adapter
- Four-Port SDLC adapter
- X.25 adapter

3.10.1 6611 Functions

The 6611 Network Processor when used with the IBM Multiprotocol Network Program provides three functions:

- Multiprotocol routing
- Bridging (SRB and SR-TB)
- Data Link Switching (DLSw)

As a multiprotocol router, the 6611 provides routing of IP, XNS, IPX, AppleTalk Phase 2, DECnet Phase IV, Banyan VINES, SNA, APPN, and NetBIOS.

3.10.2 The 6611 As a Bridge

The 6611 provides source-route bridging between token-ring networks and translational bridging between a local token-ring and a local Ethernet. The following types of bridging are possible:

- Local bridging
- Remote bridging between 6611 network processors
- Remote bridging between a 6611 and a PS/2 running the RouteXpander/2 Program Version 1.0.1 with PTF WR20307 or Version 2.0
- Remote bridging between a 6611 and a PS/2 running the IBM Remote Token-Ring Bridge Program Version 1.0 with PTF UR37051

The 6611 also provides transparent bridging between Ethernet networks. The following types of bridging are possible:

- Local bridging
- Remote bridging between 6611 network processors
- Remote bridging between a 6611 and a PS/2 running the RouteXpander/2 Program Version 1.0.1

All types of bridging can be combined in a single configuration if required. In addition, the bridging functions can be used concurrently with the other functions provided by the 6611 (routing and data link switching).

The 6611 when functioning as a wide-area bridge uses the two-port EIA 422/449 Serial Adapter or two-port V.35/V.36 Compatible Serial Adapter for speeds between 18.2 Kbps to 2.048 Mbps.

3.10.3 Source-Route Bridging

The 6611 Network Processor provides local source-route bridging between up to seven token-ring segments that are directly attached to a 6611.

Each 6611 that is enabled for the bridging function must be assigned a single bridge number, and each token-ring interface that is enabled for the bridging function must be configured with the segment number for the token-ring segment to which the interface is attached.

3.10.4 Remote SR Bridging between 6611 Network Processors

The 6611 Network Processor provides remote source-route bridging between multiple token-ring segments that are attached to two or more 6611s.

The remote connections between 6611s can use either the PPP or frame relay data link protocols.

Each connection between 6611s can be either of the following:

- 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. This allows a 6611 to establish connections with many other 6611s using a single physical interface.

Additionally, each remote connection between 6611s is assigned a unique token-ring segment number which is referred to as a *link segment number*. The 6611 remote bridging function uses each remote connection (or *link segment*) in the same manner as the local bridging function would use a real token-ring segment.

3.10.5 Remote SR Bridging between 6611 and the Remote Bridge Program/DOS

The 6611 provides remote source-route bridging between multiple token-ring segments that are attached to a 6611 and multiple IBM Personal System/2 PCs running the IBM Remote Token-Ring Bridge Program Version 1.0, each attached to a single token-ring segment.

Remote connections between 6611s and PS/2s use point-to-point links at speeds between 19.2 Kbps and 1.344 Mbps (2.048 Mbps in some countries).

Each PS/2 remote bridge is configured as the secondary half of a remote bridge using the IBM Remote Token-Ring Bridge Program Version 1.0. PTF UR37051 should be installed on the PS/2 remote bridge if management of the remote token-ring segment by IBM LAN Network Manager is required. This PTF should only be installed on a PS/2 remote bridge that is connected to a 6611.

One of the token-ring segments locally attached to the 6611 must be selected to become the *designated segment*. All the PS/2 remote bridges connected to a 6611 are logically bridged to the designated segment.

3.10.6 Remote SR Bridging between 6611 and PS/2 Running RouteXpander/2

The 6611 provides remote source-route bridging between multiple token-ring segments that are attached to a 6611 and multiple IBM Personal System/2s running the IBM RouteXpander/2 Program, each attached to a single token-ring segment.

Remote connections between 6611s and PS/2s utilize frame relay data link protocols.

Each PS/2 is configured as the secondary half of a remote bridge using the IBM RouteXpander/2 Program.

3.10.7 6611 Local Transparent Bridging

The 6611 Network Processor provides local transparent bridging between up to seven Ethernet segments that are directly attached to a 6611.

3.10.8 Remote Transparent Bridging between 6611 Network Processors

The 6611 Network Processor provides transparent bridging between multiple Ethernet segments that are attached to two or more 6611s.

The remote connections between 6611s can use either the PPP or frame relay data link protocols.

Each connection between 6611s can be either:

- A point-to-point communication facility. Such a connection would use PPP data link protocols.
- A DLC (Data Link Connection) across a frame relay service. This allows a 6611 to establish connections with many other 6611s using a single physical interface to a frame relay service.

3.10.9 Remote Transparent Bridging between 6611 and RouteXpander/2

The 6611 provides remote transparent bridging between multiple Ethernet segments that are attached to a 6611 and multiple IBM Personal System/2s running the IBM RouteXpander/2 Program, each attached to a single Ethernet segment.

Remote connections between 6611s and PS/2s utilize the 6611 2-Port Serial Adapter or 6611 two-Port V.35/V.36 Compatible Serial Adapter, using frame relay data link protocols at speeds between 19.2 Kbps and 2.048 Mbps.

Each PS/2 is configured as half of a remote bridge using the IBM RouteXpander Program.

3.10.10 6611 Bridging Filters

When configured for source-route bridging, the 6611 provides the following types of port filters:

- Hop count
- Source SAP
- SNAP value
- RI field
- MAC address
- Sliding window (criteria filter based on value, offset in the frame)

When configured for transparent bridging, the following port filter types are available:

- Source SAP
- Ethernet type
- MAC address
- Sliding window

The order in which the filters are applied can also be defined so that the filters more likely to block unwanted traffic can be applied first, thus improving the performance of the bridge.

3.10.11 6611 Bridge Positioning

The use of 6611s as LAN bridges is recommended when:

- There′s a need for both bridging and routing.
- There′s a need for interconnecting several local token-ring segments.
- There′s a need for interconnecting several local Ethernet segments.
- There′s a need for interconnecting several remote token-ring segments with several local token-ring segments.

- There′s a need for interconnecting several remote Ethernet segments with several local Ethernet segments.

3.11 The IBM 2210 Nways Multiprotocol Router

The IBM 2210 Nways Multiprotocol Router addresses the connectivity requirements of remote branch and regional offices with an extensive choice of models which provide a wide range of connectivity options. This access node has the power and versatility of a full-function, multiprotocol bridge/router, as well as frame relay access device concentration functions. The 2210 offers robust network backup features, including ISDN and V.25bis, to ensure maximum network availability.

The new models of the 2210 offer double the density and performance of the other 2210 models, supporting up to two Local Area Network (LAN) ports plus four Wide Area Network (WAN) serial ports. The LAN ports can be one or two Token-Ring ports, two Ethernet ports, or one of each type. The serial ports support attachment to a variety of wide area networks, or can be used for device concentration. In addition to this base connectivity, the new 2210 models offer an adapter slot which will support over time, optional network interfaces, such as ATM.

The IBM 2210 is available in several models, based on the types of networks you want to support. Discontinued models are Model 123 and Model 124. These models both had one LAN port, two serial connections, 4 MB Flash and 8 MB DRAM. Both of these models have been replaced with the Model 12T and Model 12E. Features are almost identical except for the DRAM size.

Table 31 on page 104 shows the different models and the offerings of the IBM 2210 that are available. Note that the only difference between some of the models is the amount of flash memory. Flash memory is used to store a compressed version of the router′s software. Flash memory is not upgradeable 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 8MB of flash memory for those models.

IBM 2210 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 2210 must support. 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.

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

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

Table 31. IBM 2210 Models

Hardware						Software		
Model	**LAN**	**No. of WANS (See Note)**	**Flash Memory**	**DRAM**	**ISDN**	**Base**	**Additional Routing**	**ISDN**
121	Token-Ring	2	2 MB	4 MB		x	x	
122	Ethernet	2	2 MB	4 MB		x	x	
12T	Token-Ring	2	4 MB	4 MB		x	x	
12E	Ethernet	2	4 MB	4 MB		x	x	
125	Token-Ring	2	2 MB	4 MB	x	x	x	x
126	Ethernet	2	2 MB	4 MB	x	x	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 double density models support an additional service port and an adapter slot which can support ISDN Basic Rate, ISDN Primary Rate and ATM. The availability of these adapter cards are defined in the announcement letter.

3.11.1 Software Package

MRNS is the software that runs on the 2210 and it comes as a base package, plus two separately ordered packages - one containing support for additional routing protocols, and the other containing the ISDN support. The protocols supported by each package are:

- Base suite options
  - IP and bridging (including BAN)
  - IP, bridging (including BAN), and DLS
- Additional routing suite options
  - IP, bridging (including BAN), and IPX
  - IP, bridging (including BAN), DLS and LNM support
  - IP, bridging (including BAN), DLS, LNM support, Banyan VINES, DECnet IV, and DECnet V/OSI

3.11.2 Bridging Methods Supported by the 2210

The supported functionality listed in this sub-section is for the MRNS Program Version 1 Release 1 for the 2210.

The 2210 can act as an adaptive source routing transparent (ASRT) bridge. It is a software collection of the following four bridging options:

- Spanning tree bridge or transparent bridge (STB)
- Source route bridge (SRB)
- Source route transparent bridge (SRT)
- Source route - translational bridge (SR-TB)

The 2210 supports SRB bridging methods over:

- Token-ring
- PPP serial link

It supports STB bridging methods over:

- Ethernet
- PPP serial link
- FR in Version 1 Release 2

It supports SRT bridging methods over:

- Token-ring
- PPP serial link

It supports SR-TB bridging methods over:

- Token-ring
- Ethernet
- PPP serial link

The dial backup or WAN restoral (WRS) feature over a PPP primary link supports any bridging method.

The Dial on Demand feature does not support any bridging methods on serial or on ISDN links.

The 2210 supports bridging over PPP links only. It cannot perform bridging over frame relay, X.25 or SDLC links. If you are planning to transport bridged traffic over frame relay or X.25, you must use tunnel bridge function provided by the 2210. Note that The tunnel bridge is supported over any serial link that supports IP. This includes PPP, frame relay and X.25 serial links.

SRB, STB and SRT bridging methods have to be enabled on the interfaces. While SR-TB and tunnel bridge are enabled at the bridge level.

Table 32 shows the bridging methods supported by each 2210 interface.

With the ASRT bridge, the collection of configuration parameters for the bridge and all its interfaces produces a *bridge personality*. The bridge personality is also called bridge behavior.

In Version 1 Release 2 FR supports all Bridging Methods.

Table 32. Supported Bridging Methods for the 2210

Bridging Methods	PPP	FR	X.25	SDLC	Dial on Demand	Dial backup	Token-Ring	Ethernet
STB	Yes	Yes	No	No	No	Yes	No	Yes
SRB	Yes	Yes	No	No	No	Yes	Yes	No
SRT	Yes	Yes	No	No	No	Yes	Yes	No
SR-TB	Yes	Yes	No	No	No	Yes	Yes	Yes
Tunnel (IP)	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes

The bridge personality may have the following values:

- SRB (pure source routing bridge)
- STB (pure transparent bridge)
- SRT (source route transparent bridge)
- SR-TB (source route - translational bridge)
- ASRT (either SRB, STB, SRT or SR-TB, depending on the source and the destination devices)
- Unknown (When the bridge customization results in no valid bridging)

Table 33 provides a guide to the configuration settings needed for interfaces to produce the desired bridge personality. In this table, port A, B and C could be Ethernet, the token-ring or a serial link on the 2210.

Note that the only bridging methods supported by the Ethernet interface is STB. The token-ring interface supports only SRB, or both STB and SRB. The serial link supports STB, SRB, or both SRB and STB bridging methods.

Table 33 (Page 1 of 2). Bridge Personality Regarding Bridge Configuration

Port A	**Port B**	**Port C**	**SR <-> TB Conversion**	**Bridge personality**
STB	STB	No bridging	Disable	STB
STB	SRB	No bridging	Enable	SR-TB
STB	STB & SRB	No bridging	Enable	ASRT
SRB	SRB	No bridging	Disable	SRB
SRB	STB & SRB	No bridging	Enable	ASRT
STB & SRB	STB & SRB	No bridging	Disable	SRT
STB & SRB	STB & SRB	No bridging	Enable	ASRT
STB	STB	STB	Disable	STB
STB	STB	SRB	Enable	SR-TB
STB	STB	STB & SRB	Enable	ASRT
STB	SRB	SRB	Enable	SR-TB
STB	SRB	STB & SRB	Disable	SRB
STB	SRB	STB & SRB	Enable	ASRT
STB	STB & SRB	STB & SRB	Disable	SRB
STB	STB & SRB	STB & SRB	Enable	ASRT
SRB	SRB	SRB	Disable	SRB
SRB	SRB	STB & SRB	Enable	ASRT
SRB	STB & SRB	STB & SRB	Disable	SRT

Table 33 (Page 2 of 2). Bridge Personality Regarding Bridge Configuration				
Port A	**Port B**	**Port C**	**SR <-> TB Conversion**	**Bridge personality**
SRB	STB & SRB	STB & SRB	Enable	ASRT
STB & SRB	STB & SRB	STB & SRB	Disable	SRT
STB & SRB	STB & SRB	STB & SRB	Enable	ASRT

3.11.3 Bridging Methods Supported between the 2210 and the 6611

The supported functionality listed in this sub-section correspond to the MRNS Program Version 1 Release 1 for the 2210 and to MPNP Program Version 1 Release 3 for the 6611.

A 2210 bridge may communicate with a 6611 bridge over either token-ring, Ethernet or a PPP serial link.

Note that tunnel bridge is not supported between 2210 and 6611.

Table 34 shows the bridging methods supported over each interface of either the 2210 or the 6611.

Table 34. Supported Bridging Methods between the 2210 and the 6611					
Bridging Methods	**Ethernet**	**Token-Ring**	**PPP**	**FR**	**X.25**
STB	Yes	No	Yes	Yes	No
SRB	No	Yes	Yes	Yes	No
SR-TB	Yes	Yes	Yes	Yes	No

Note

When you want to perform source route-translational bridging between the 2210 and the 6611:

1. You may configure either the 2210 or the 6611 to perform the translational function if the source routing (token-ring) side is attached to the 6611,
2. You must configure the 6611 to perform the translational bridging if the transparent (Ethernet) side is attached to the 6611. In this case, if you cannot configure the 2210 to perform the translational bridging.

3.11.4 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 which are sent to the destination IP address. These IP datagrams are routed in the IP network as any 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 datagram 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 as any others bridged frames.

Figure 22 shows you an example of the tunnel bridge implementation.

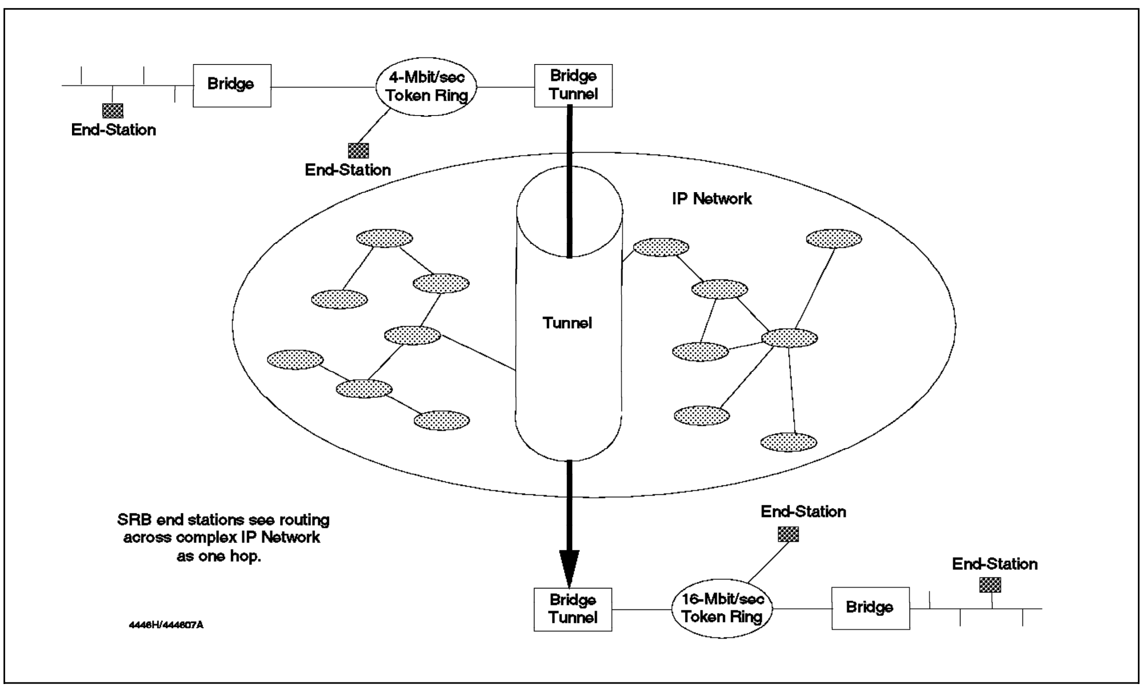

Figure 22. Tunnel Bridging

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 hops count limitation of the source route bridge implementation.

3.12 IBM 3172 Model 3 with Multiprotocol Extension (MPE)

An IBM 3172 Model 3 configured as a Multiprotocol Extension system (MPE) can perform source-route bridging between its own token-ring adapter and Wide Area Connector (WAC) adapter.

3172 MPE implies that the 3172 has, besides other hardware, at least a token-ring and a WAC adapter installed, and that it is loaded with OS/2 V2.1 or later, RouteXpander/2 V1.0.1 with PTF WR20307, and NTS/2.

The WAC adapter appears to the 3172 as a token-ring adapter, and the frame relay network behind this adapter looks like a ring segment. Each DLCI connection to a 6611, RXR/2, or 3174 on the frame relay network, will appear as a bridge on that segment.

The SR bridge participates in the automatic spanning tree protocol.

Bridge filters let you specify:

- Frame type
- Hop count limit
- Destination address
- Source address
- NetBIOS name
- Data and offset in the frame

See 3.8, "RouteXpander/2 V 1.0.1" on page 94 for more details.

3.13 IBM 3174 Communication Controller

Starting with Configuration Support C1, the IBM 3174 Communication Controller, equipped with Type 3A Dual Speed (16/4) Communication Adapter for token-ring and the 3174 Peer Communication LIC Feature, provides source-routing between the token-ring segment and the virtual ring for its coax-attached workstations.

With the 3174 frame relay communication feature and configuration support C6, the 3174 adds source-route remote bridging over the frame relay link.

3.13.1 Multiport Bridge

The 3174 is a multiport bridge that allows connectivity between any station on either:

- The token-ring segment connected to the Type 3A Dual Speed Token-Ring Adapter
- The virtual ring of the Peer Communication (the coax-attached stations)
- The frame relay port

Note: The routing function is part of the Type 3A token-ring adapter microcode. This adapter is mandatory for any source-routing functions in 3174. The 3174 does not support transparent bridging. A 3174 with an Ethernet adapter does not support bridging between the Ethernet-attached workstations and the Peer Communication stations or frame relay port.

The 3174-R43, which is a blade for 8250/8260, also supports source-route bridging. The token-ring connection to the backplane has the bridging support included.

3.13.2 Multitail Bridge

The communication adapter that connects the 3174 to the frame relay network, supports up to 250 DLCIs to remote partners, but only the first eight DLCIs on which a bridge frame is discovered are eligible for bridging. The frame relay bridge port of 3174 is therefore a *multitail bridge* port, with a maximum of eight tails. The selection of the eight DLCIs is done on a first-come-wins-all basis.

The 3174 Frame Relay Communication feature supports RFC 1490 protocol (an RFC 1294 update) therefore the 3174 remote bridging supports other bridging partners such as:

- An 6611 Network controller

- RouteXpander/2 V1.0.1. or V2.0 on PS/2 or 3172
- Frame Relay Bridge/DOS V1.0
- A 3172

The 3174 configures the frame relay network as a single segment. The segment number is entered at configuration. When a frame is to be passed over the same frame relay network, from one remote bridge partner to an other remote bridge partner, then a potential duplicate segment number problem can arise. This is a violation of the bridging standards. To overcome this problem, the 3174 handles a split horizon scenario. The 3174 can discover and accept a different segment number than the one configured by the 3174 for the frame relay network. It assigns a *logical port* to a DLCI with a discovered different segment number.

3.13.3 Target Segment, Virtual Bridge and Extra Hop

The 3174 token-ring adapter can only handle a two-port bridging function. All bridging is done towards a designated target segment. Without frame relay configured, the target segment is the peer segment; with frame relay configured the target segment is always the frame relay segment. Traffic destined to another segment is bridged again to reach its destination.

Consequently, when traffic goes from the token-ring segment to the peer segment, it will first bridge internally to the frame relay segment and then via a virtual bridge to the peer segment, causing an extra hop.

The same principle is used to handle a discovered different segment number on the frame relay network. The frame will be bridged to the frame relay segment and via a virtual bridge to the differently numbered frame relay segment.

3.13.4 Filters

The 3174 allows five different types of filters. The order of execution is important for performance. The default order is as follows:

- Hop count
- Source SAP
- SNAP Ethertype
- Route Designator
- MAC Address
- Frame data and offset

3.13.5 Maximum Frame Size

All frames larger than the maximum frame size are discarded. The current 3174 bridge supports a maximum frame size of 2052, which is hard coded. It can be customized to 516, 1500, or 2052 (default).

Notes:

1. 3174 does not support *automatic* spanning tree.
2. LAN Network Manager V2 for OS/2 is required to support multiport bridging.

3.14 IBM 8250 Multiprotocol Intelligent Hub

The 8250 is designed as a platform to build local area networks meeting the requirements of customers using various types of cabling systems such as shielded twisted pair (STP), unshielded twisted pair (UTP), fiber and coax, and different types of LAN protocols such as token-ring, Ethernet, and FDDI.

The 8250 consists of a chassis that houses the power supplies and the *backplane*, and of modules that are connected to the backplane. There are many different types of modules, such as controller modules, management modules, media modules, bridges, repeaters, transceivers, and switches. Modules are hot-pluggable; that is, they can be installed and removed without powering down the machine.

The 8250 uses an *advanced backplane architecture*, which gives it the capability to run multiple networks using various protocols concurrently. The maximum number of networks supported are as follows:

- Three Ethernet networks
- Seven token-ring networks
- Four FDDI networks

The Models are as follows:

- 8250-06S: six slots, single power
- 8250-6HC: six slots, enhanced architecture
- 8250-6PS: six slots, built-in PS/2
- 8250-06W: six slots, with 3174 blade
- 8250-017: 17 slots
- 8250-17W: 17 slots, with 3174 blade

3.14.1 8250 Token-Ring Bridge Modules

The token-ring bridge modules are two-port, one-slot token-ring to token-ring SR or SRT bridges. One port is provided via the front panel in the form of a DB-9 connector for STP and an RJ-45 connector for UTP cabling. The other port is selectable from among any of the seven internal token-ring segments on the backplane.

3.14.2 8250 Ethernet 6-Port Bridge Module

The 6-port bridge module is a multiport transparent Ethernet bridge. The module occupies two slots on the 8250. Its front panel has three 10Base-T and two AUI connectors for access to external Ethernet networks. One of the six ports (port 6) cannot be accessed from the front panel; it can be connected to any of the three possible Ethernet backplane networks. Table 35 shows the connections allowed for each port.

Table 35. Six-Port Bridge Module Configuration Options

Port	Connectors	Connection Options
1	10Base-T	Front panel 10Base-T connector (RJ-45)
2	10Base-T/backplane	Backplane Ethernet network 1 or Backplane Ethernet network 2 or Front panel 10Base-T connector or (RJ-45) Isolated
3	10Base-T/backplane	Backplane Ethernet network 2 or Backplane Ethernet network 3 or Front panel 10Base-T connector or (RJ-45) Isolated
4	AUI/backplane	Backplane Ethernet network 1 or Backplane Ethernet network 2 or Front panel AUI connector or Isolated
5	AUI/backplane	Backplane Ethernet network 2 or Backplane Ethernet network 3 or Front panel AUI connector or Isolated
6	Backplane	Backplane Ethernet network 1 or Backplane Ethernet network 2 or Backplane Ethernet network 3 or Isolated

8250 Six-Port Bridge Filtering

The six-port bridge module provides destination and source address filtering, broadcast/multicast filtering and an advanced filtering capability that consists of a combination of tests and actions that are user-defined.

3.14.3 8250 Ethernet Two-Port Bridge Module

The two-port bridge module provides transparent bridging between two Ethernet segments. It occupies two slots on the 8250. Its front panel has one AUI connector for access to an external Ethernet network. One of the ports (port 1) can be connected to the front panel AUI connector or to the backplane, while the other can only be connected to the backplane. Table 36 shows the connections allowed for each port.

Table 36. Two-Port Bridge Module Configuration Options

Port	**Connectors**	**Connection Options**
1	AUI/backplane	Backplane Ethernet network 1 or Backplane Ethernet network 2 or Front panel AUI connector or Isolated
2	Backplane	Backplane Ethernet network 1 or Backplane Ethernet network 2 or Backplane Ethernet network 3 or Isolated

3.14.4 8250 Ethernet Six-Port Switch Module

The six-port switch module is physically similar to the six-port bridge module. It provides six ports for connection of Ethernet segments. The connections allowed for each port are the same as those shown on Table 35 on page 114.

The operation of the switch is similar to that of a bridge, except that the switch does not participate in the spanning tree protocol, nor does it provide any of the filtering facilities of the bridge. It is designed as a simple solution for the interconnection of Ethernet segments.

3.14.5 Managing the 8250

Each of the 8250 bridge and switch modules has an EIA-232 serial port to which an ASCII terminal can be connected to manage the module. The initial configuration is done via this connection. The modules can been managed via either:

- This connector
- An ASCII terminal connected to the connector of another module in this 8250 or to an Ethernet managing module
- A remote station running an SNMP management application

3.15 IBM 8260 Multiprotocol Intelligent Switching Hub

The 8260 is a high-end intelligent hub which extends the IBM 8250 LAN capabilities while maintaining full forward compatibility. It has an enhanced passive media backplane that allows for more networks. The 8260 supports a maximum of:

- 8 Ethernet networks
- 17 token-ring networks
- 8 FDDI networks

The 8260 is ATM enabled. Older models are also designed for possible ATM support with a field upgrade, by installing a new ATM backplane.

Models:

- 8260-010: 10 slot chassis
- 8260-A10: 10 slot, with ATM backplane
- 8260-017: 17 slot chassis
- 8260-A17: 17 slot, with ATM backplane

The 8250 bridge and switch modules discussed previously in this chapter, as well as all other 8250 modules, can be installed in the 8260. However, only the 8260 modules can take full advantage of the enhanced backplane capabilities.

8281, 8271[4], 8272[4], 8235, and 3174 are made available as modules for the 8260.

Refer to the 8229 section for information about the 8229 module.

3.15.1 Managing the 8260

The 8260 modules can be managed in the same way as the 8250 modules, by either a local or remote ASCII terminal, or by an SNMP application running on a station connected to the network. The segment management for the bridges can be done with LNM (LAN Network Manager).

3.15.2 8260 Multiprotocol Interconnect Module

The multiprotocol interconnect module allows you to interconnect Ethernet and token-ring networks using bridging and/or routing functions. There are two versions of the module. The one-slot module provides six logical ports for attachment of Ethernet segments to the 8260 backplane. The two-slot module provides the capability to install two additional I/O cards for connection to external token-ring and Ethernet networks.

As a router, the module can route IP, IPX, and DECnet Phase IV traffic. As a bridge, it can perform the following types of bridging:

- Transparent bridging between Ethernet segments
- Transparent bridging between token-ring segments
- Source-route transparent bridging between token-ring segments
- Source-route to transparent bridging between Ethernet and token-ring segments

3.15.3 Filtering

The multiprotocol interconnect module provides advanced filtering functions. These are:

- Destination address filtering
- Source address filtering
- Broadcast/multicast filtering
- Custom filtering

---

4 Statement of Direction

Custom filters are defined in the same way as for the 8250 six-port switch module.

3.16 IBM 8281 ATM LAN Bridge

The 8281 is a stand-alone bridge designed for high performance connection of up to four LANs (token-ring or Ethernet) to a local ATM backbone. It offers a single 100 Mbps ATM port for connecting to an ATM switch. The 8281 is designed to operate at media speeds on the LAN side.

The 8281 provides three basic functions:

1. Locally bridge between its four LANs as a multiport bridge.
2. Bridge between its four LANs and other LANs across a local ATM backbone.
3. Bridge between ATM and LANs directly. ATM workstations and servers must support an ATM LAN emulation protocol, so that they can communicate with LAN workstations using the existing LAN-based applications and communication protocol stacks.

The base unit consists of the following major features:

- Two RJ-45 connectors (and one AUI connector) for two LAN interfaces (of the same type).
- Slot for an optional module with two additional LAN interface.
- Slot for optional ATM interface module. Initially this is a 100 Mbps multimode interface.

When the 8281 provides bridging between token-ring networks across the ATM backbone, then the ATM network is viewed as a single segment for source-routing.

When the 8281 provides bridging between the 802.3/Ethernet V2 networks across the ATM backbone, then the 8281 participates in the 802.3/Ethernet bridge spanning tree protocols, treating the ATM network as a single segment.

The ATM workstations must use different versions of LAN emulation to support source-routing token-ring or transparent routing Ethernet.

The 8281 supports only one single bridge instance, so it either communicates with a “virtual” token-ring or a “virtual” Ethernet on the ATM segment.

The 8281 supports a total of 256 simultaneous VCCs.

3.16.1 Maximum Frame Size

The maximum frame size for token-ring is 4399 bytes (default). Of course, for Ethernet the maximum is 1500 bytes.

3.16.2 Filters

The 8281 supports port filters and allows the filtering order to be specified. The filters apply to inbound frames on the LAN ports.

The following user filters will be supported for source-route bridging ports:

1. Hop count limit for all-routes broadcast
2. Ring number filters
3. MAC addresses
4. Source SAP
5. SNAP header

The following user filters will be supported for transparent bridging ports:

1. Source SAP value
2. Ethernet type
3. MAC addresses

The 8281 provides an SNMP agent for management across a LAN port or using LAN emulation across its ATM port.

The 8281 will become available as a module in 8260.

Chapter 4. IBM 827X Nways LAN Switch Product Family

The IBM 827X Nways LAN Switch family of products provides the benefits of LAN switching to traditional Ethernet and token-ring LAN environments.

The family currently includes the IBM 8271 Nways Ethernet LAN Switch and the IBM 8272 Nways Token-Ring LAN Switch.

The following sections describe these products and the features they offer.

4.1 IBM 8271 Nways Ethernet LAN Switch Family

The IBM 8271 Nways Ethernet LAN Switch product family provides switching for the Ethernet environment. There are currently three models: a basic Model 001 offering eight ports without any expansion (no longer available), a Model 108 offering eight ports with a single expansion slot and a Model 216 offering 16 ports, plus two expansion slots.

Model Highlights 8271-001

- Transports Ethernet frames among up to eight Ethernet LAN segments, each connected to one 10Base-T port on the IBM 8271. A single AUI port is also provided, for use in lieu of one 10Base-T port, for connection using a transceiver to an Ethernet 10Base2, 10Base5, or fiber segment backbone.
- Supports both shared and dedicated LAN segments on any of the ports. Bandwidth to a single LAN station on a dedicated LAN segment can be doubled from 10 Mbps to 20 Mbps by using the full-duplex capabilities of the IBM 8271 and full-duplex Ethernet adapters, such as the IBM EtherStreamer MC 32 Adapter.
- Provides out-of-band console management facilities (accessible using a serial port) as well as an in-band Telnet session, and an SNMP agent that allows control from a central SNMP management station, such as one running IBM NetView for AIX or IBM NetView for Windows.

Model Highlights 8271-108

- Includes all the features described for the Model 001.
- Provides an EtherProbe network monitor port to monitor all switch network traffic.

- Includes a single Universal Feature Slot that will, in the future, support several different optional Universal Feature Cards. These will provide additional connectivity, for example a high-speed ATM uplink, to interconnect IBM 8271s to a backbone ATM network.
- Enhanced filtering capabilities.

Model Highlights 8271-216

- Includes all features already described for the Model 108.
- Provides 16 fixed Ethernet ports as well as two Universal Feature Slots allowing users to support a wider variety of configurations than is possible with the one Universal Feature Slot of the Model 108.

4.2 IBM 8271 Nways Ethernet LAN Switch Model 108

The IBM 8271 Model 108 is a follow-on product to the IBM 8271 Nways Ethernet LAN Switch Model 001 and offers all of the functions of the Model 001, with improvements and additional enhancements.

4.2.1 Basic Operation

The IBM 8271 provides the ability to forward Ethernet frames among up to eight shared Ethernet LAN segments for Ethernet 10Base-T or up to seven shared Ethernet LAN segments for Ethernet 10Base-T and one shared Ethernet LAN segment for attachment unit interface (AUI) connections.

The IBM 8271 employs transparent switching (refer to 4.2.5, "Transparent Switching" on page 124) to forward frames between ports and builds up port forwarding tables as the basis for doing this. It supports a maximum of 1790 active Ethernet LAN station addresses in each port table, and no more than 10,000 total per 8271.

The IBM 8271 is able to forward Ethernet frames at media speeds. With a highly parallel internal design optimized for performance, the 8271 is able to maintain media-speed frame transfer between each of the possible four distinct pairs of ports simultaneously. This feature allows the 8271 to provide an aggregate bandwidth of up to 40 Mbps. Networks with traffic patterns able to take full advantage of the 8271 could sustain throughput equivalent to four 10 Mbps Ethernets.

The eight 10Base-T ports on the IBM 8271 can be connected to a port on an external 10Base-T hub or concentrator (for example, the IBM 8222 6-Port 10Base-T Workgroup Hub, the IBM 8224 Ethernet Stackable Hub, or the IBM

8250 or 8260 Multiprotocol Intelligent Hub). A single AUI port is also provided on the 8271 for use in lieu of one of the 10Base-T connections and can be used to connect the 8271 to an Ethernet 10Base2, 10Base5, or fiber segment. Alternatively, when connecting a 10Base-T port to a dedicated LAN segment, a direct connect cable can be used to connect the port to an Ethernet LAN station without an intervening Ethernet hub.

4.2.2 Full-Duplex

To increase bandwidth available to a server, dedicated LAN segments connected to any or all of the IBM 8271s eight ports can be configured to operate in full-duplex mode.

Full-duplex Ethernet is an emerging extension to the existing IEEE 802.3 standard. It provides for simultaneous, two-way transmission between a device, such as an IBM 8271 and a client station. It provides two independent data paths between the 8271 and a LAN station. Each has a bandwidth of 10 Mbps, for a combined total of 20 Mbps per full-duplex port. In fact, these independent, parallel, full-duplex paths are extended throughout the internal design of the 8271 so that each of the two 10 Mbps paths can be switched to different half-duplex ports.

When all of the ports on the IBM 8271 are configured for full-duplex operation, the IBM 8271 may provide an aggregate bandwidth of up to 80 Mbps.

Full-duplex operation requires a full-duplex capable Ethernet adapter and associated drivers, for example the IBM EtherStreamer MC 32 Adapter or the IBM Dual EtherStreamer MC32 Adapter.

Ports on the IBM 8271 Model 108 are configured for full-duplex Ethernet operation using a configuration utility.

4.2.3 EtherPipe

Configurations larger than eight Ethernet LAN segments can be constructed by interconnecting IBM 8271 ports together. This capability, called EtherPipe, allows two 8271s to communicate by connecting together up to four (user-configurable) full-duplex Ethernet ports on one 8271 to those on another 8271. Each of these full-duplex Ethernet inter-switch links provides up to 20 Mbps of bandwidth, for a maximum of 80 Mbps of bandwidth (using four links) between switches. With multilink EtherPipes, traffic is automatically distributed on each of the links on the basis of destination addresses to achieve load balancing. EtherPipe support provides users the ability to grow and tune network bandwidth as demands dictate,

4.2.4 Virtual Switches

Conversely, it may sometimes be desirable to construct a switch network of less than eight ports, for either security,traffic load management or other reasons. In this case, ports on an IBM 8271 are configured so that no traffic (including broadcast) may flow between ports within a group and ports outside the group. The virtual switch capability of the 8271 allows a single physical switch to be divided into multiple switching domains. Each domain consists of a discrete set of switch ports that may number from two or more aggregate ports.

4.2.5 Transparent Switching

The IBM 8271 provides multiport, transparent connections between ports using an internal 512 Mbps time-slotted bus. Packets are routed to the appropriate destination port based on the destination MAC address of the packet. A packet with an unknown destination address is forwarded to all ports (broadcast) until the location of that address has been learned. Once the location of the destination is learned, direct packet forwarding occurs only between the two ports (source and destination).

Each port has an address filtering and forwarding table that can have up to 1790 active entries. The switch port will learn the 6-byte MAC source address of all stations that are active on its attached LAN. A station may be on the local LAN directly attached to the port or may be several bridge hops removed from the port. Once a station is associated with a port, then any frames entering the switch at another port with a destination for that address will be forwarded directly to that port. Also, any frames that are received at a port with a destination address that is associated with that same port will be discarded.

Figure 23 on page 125 and Figure 24 on page 126 show an example network configuration with a switch and six separate Ethernets, with two port filtering and forwarding tables. Two of the Ethernets (E1 and E2) are attached to separate ports (Port A and Port B) and each has a server, S2 and S3 respectively, attached. Transparent bridges, B1 and B2, connect Ethernets E3 and E4 to the Ethernet, E1 on Port A. Likewise, Ethernets E5 and E6 are bridged to the Ethernet, and this is then switched on Port B. Four client stations (C1-C4) and a server station (S1) are attached to dedicated switch ports. The switch is transparent to the bridges, client stations (C1-C4) and server stations (S1-S3).

Each port learns the MAC address of those stations that can be reached over it. For example, stations C5-C8 and server S2 are associated with Port A, while stations C9-C12 and server S3 are associated with Port B. With transparent switching, the switch uses these MAC addresses to forward frames.

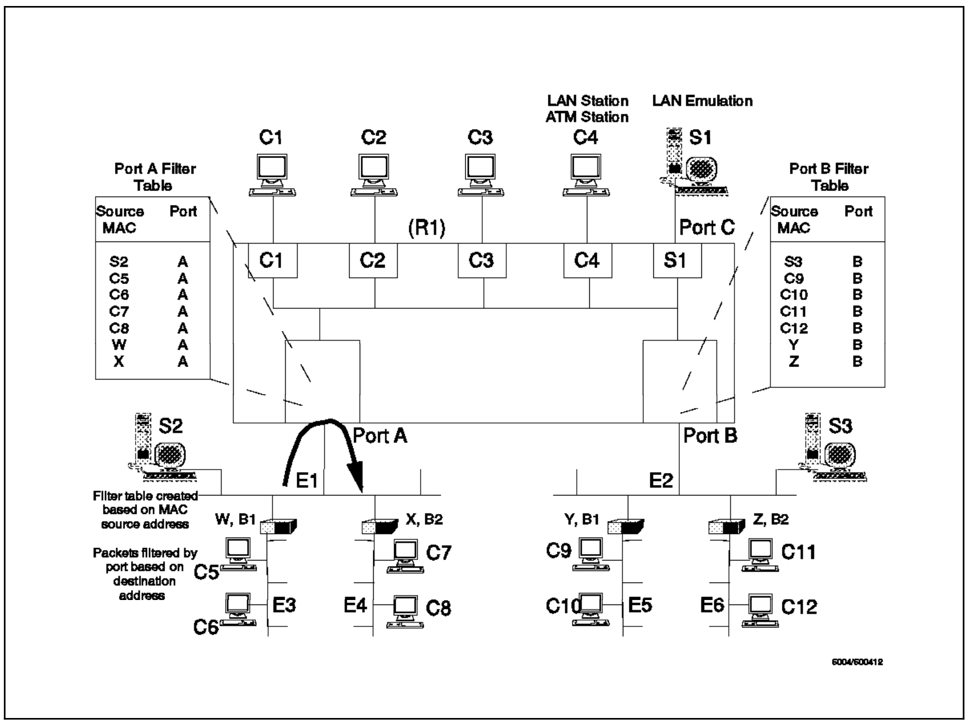

Figure 23. Transparent Switch Learning and Filtering Tables

The 1790 address entries in each port forwarding table are a combination of both local (this port) and remote (another port) stations. A single local station that has active sessions to five remote stations will utilize six address spaces in the local port table (one local and five remote). If multiple local stations are in session to the same remote station, there is only one entry in the table for that remote station (for example, multiple clients to the same server).

Addresses are periodically aged from the port tables, based on the frequency that packets with a particular source address are observed, with the inactive station addresses being removed. The user has control over the aging time.

When a station moves from one port to another, all relevant port tables are updated. Any frames sent to that station prior to adjustment of the port tables will be lost, since they will be forwarded to the last known location of the station. The switch ″learns″ the new location after it has seen the first frame from that station on another port.

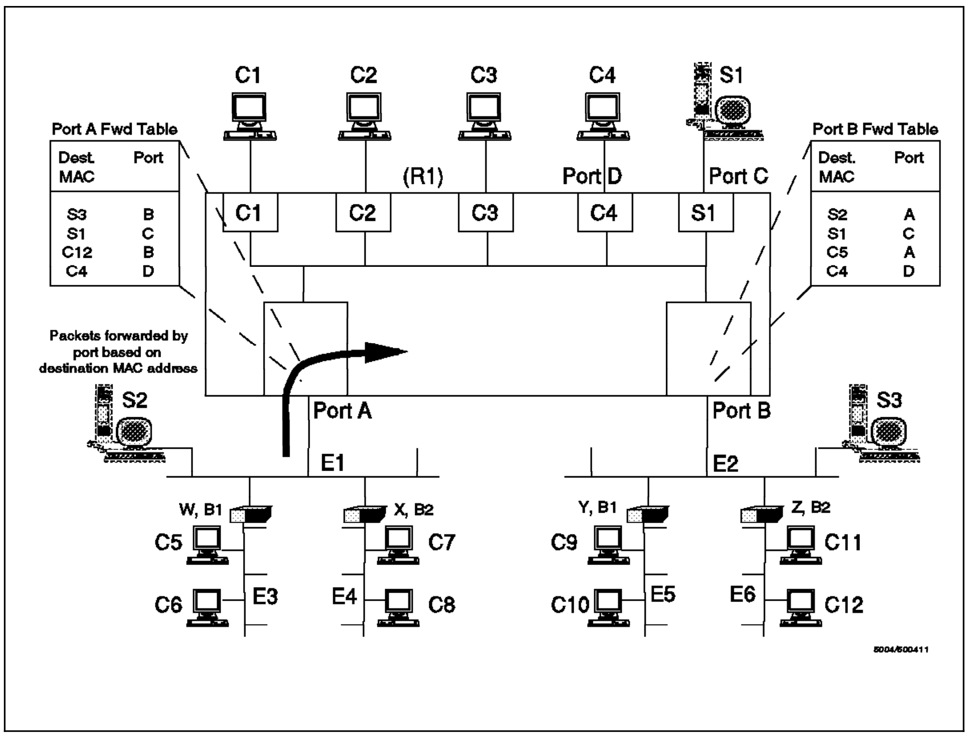

Figure 24. Transparent Switch Forwarding Tables

4.2.6 Operating Modes

The IBM 8271 operates as a conventional transparent LAN switch, as described earlier in this section, but offers enhanced features to optimize performance.

Traditional MAC level bridges and indeed some LAN switches store the entire frame before forwarding it to allow the bridge to verify that the frame is complete and without errors by validating the frame check sequence at the end of the frame. This method, referred to as store-and-forward, adds significant delay to the end-to-end path at each hop.

Cut-through switching, implemented in the IBM 8271, is an alternative scheme that allows frame transmission to begin at the output port prior to end-of-frame reception at the input port. Latency is therefore significantly reduced, with the exposure that a bad frame may occasionally be forwarded. The bad frame, however, would simply be discarded at the next store-and-forward hop or when it reaches the end-station.

The IBM 8271 also provides an adaptive cut-through switching function to reduce the exposure to a large number of bad frames. Frames are monitored for errors while in cut-through mode. If an excessive number of bad frames are detected, the port will automatically change from cut-through to store-and-forward operation so that these frames will be trapped by the port and not forwarded. Frame errors will continue to be monitored, with automatic resumption of cut-through switching when an acceptable error threshold is achieved.

4.2.7 Spanning Tree

To allow the IBM 8271 to participate in complex Ethernet network configurations, the 8271 supports the Spanning Tree Algorithm. The spanning tree implementation in the 8271 is fully compliant with the IEEE 802.1d standard, hence will operate with other network products, with similar capability, to ensure that a single active path only, exists between any two points on the network.

Inter-port connections between two IBM 8271s (including multilink EtherPipes) are viewed as single logical links by the 802.1d spanning tree.

4.2.8 Universal Feature Slot (UFS)

One of the Model 108s most significant enhancements over the Model 001 that it replaces, is the Universal Feature Slot. In the future this slot will support several different, optional UFCs; these will provide additional connectivity to augment the eight 10Base-T Ethernet ports standard in the base product. These Universal Feature Cards will provide high-speed uplink connections, for example, Asynchronous Transfer Mode (ATM). The UFC options are described in 4.8, “IBM 827X Universal Feature Cards” on page 142.

4.2.9 Universal Feature Cards

The Universal Feature Slot of an IBM 8271 will support one optional Universal Feature Card. These include the following (when available):

- 100Base-TX UFC
- 4-Port 10Base-T UFC
- 3-Port 10Base-FL UFC
- 1-Port ATM (155 Mbps SONET) UFC[5]
- One FDDI DAS/Fiber connection (A,B ports)[5]
- One FDDI SAS/Fiber port[5]
- One FDDI SAS/UTP port[5]
- 100Base-Fx UFC[5]
- 100VG AnyLAN UFC[5]

4.2.10 EtherProbe

Another new feature of the IBM 8271 is the EtherProbe network monitoring port. This can be used to view traffic from the segments attached to the 8271 that is forwarded by the switch. The EtherProbe port is a *separate* (in addition to the eight shared or dedicated Ethernet segment ports) AUI port on the front of the Model 108 that can be configured to monitor or ″mirror″ the activity on any single port.

The configuration can be changed dynamically using the console or SNMP. To monitor a full-duplex Ethernet port, the EtherProbe port can be configured to monitor either the transmit or the receive half of the full-duplex Ethernet connection. Using the EtherProbe port, traffic either received or forwarded on any of the eight ports of the Model 108 can be analyzed using a single LAN analyzer.

The EtherProbe port on the Model 108 was designed to be used with a variety of available protocol analyzers, for example the IBM DatagLANce Network Analyzer.

[5] Statement of Direction

4.2.11 Configuration and Management

The IBM 8271 is configured using an ASCII terminal (VT-100 or compatible) connected by using the serial port on the front panel or alternatively using inband signaling with TCP/IP and a Telnet session. When using the serial port, the terminal can be directly connected to the 8271 or can be remotely connected using a modem. Status display and operational control can also be performed similarly.

As an alternative to configuration using the serial connection, the IBM 8271 contains an SNMP (RFC 1213/MIB-II compliant) management agent that will allow a user-supplied SNMP management station, such as IBM NetView for Windows or IBM NetView for AIX, to interrogate and modify management data, to obtain status or to control the operation of the 8271. The definition of the product-specific MIB is included with the product.

Graphical network management applications that support the IBM 8271 are available separately for both IBM NetView for Windows and IBM NetView for AIX. The NetView for Windows program from IBM currently available will be updated to support the IBM 8271 Model 108 (in addition to the Model 001 support it currently provides). This upgrade to the NetView for Windows program will be available with the IBM 8271 Model 108 hardware, or to existing users with PTF UR43605.

Upgrades to the functional code in the IBM 8271 can be effected by a variety of methods. Code may be downloaded from a PC connected through the serial configuration console interface on the 8271, using the Xmodem protocol or by using TFTP from a workstation directly attached to the network.

The IBM 8271 also contains BOOTP/TFTP support that allows it to participate in RFC 951-compliant environments for remote boot. Central BOOTP servers can configure the IP addresses for a network of 8271 switches from a central location.

4.2.12 Filtering

Often, for either security, traffic management or other reasons, it may be desirable to control the traffic that flows through a IBM 8271 switch. Both source and destination MAC address filters are supported by the 8271; together they provide the ability to filter or inhibit frame data flow at the port of entry.

The IBM 8271 can be configured with a filter table consisting of LAN (MAC) addresses and related port numbers. Ethernet frames destined for any of the specified addresses would not be forwarded to any other switch port.

4.2.13 Physical Characteristics

The IBM 8271 can be mounted in a standard 19 inch rack, for which mounting brackets are provided, or can be placed on a flat surface such as a desk or tabletop.

4.2.14 Additional Features

The Model 108 also includes several other enhancements to the Model 001:

- Significant restructuring of the console interface to improve organization and navigation
- Ability to collect management statistics on a per station basis, in addition to the per port and per switch basis available with the Model 001
- Ability to specify address table aging parameters for each switch port in addition to the global aging parameters used in the Model 001
- Ability to filter on source and destination addresses
- SNMP management using IEEE 802.3 frames in addition to Ethernet Version 2 frames (as was implemented on the Model 001)

4.3 IBM 8271 Nways Ethernet LAN Switch Model 216

The Model 216 is identical in function to the Model 108, except that it has sixteen LAN ports and supports two UFCs. At time of publication, the Model 216 has been previewed but has not been announced as generally available.

4.4 IBM 8271 Nways Ethernet LAN Switch Model 001

The IBM 8271 Nways Ethernet LAN Switch Model 001 was the entry level product in the IBM 8271 Nways Ethernet switch product family. It is now withdrawn from marketing. In general, with the exception of the Universal Feature Slot, it shared many of the characteristics of its successor products, the models 108 and 216.

4.4.1 Configuration Scenarios

To an Ethernet network the 8271 appears (for all intent and purposes), identical to a transparent bridge. Since the use of these device-types is well understood and treated in detail elsewhere, we did not include scenarios for their use here. There are many interesting ways an 8271 can be deployed to solve bandwidth contention issues for Ethernet LANs. However, this work is intended to be a summary and in the interest of brevity, detailed scenarios for the 8271 was left to future works.

4.5 IBM 8272 Nways Token-Ring LAN Switch Family

IBM has continued its token-ring innovation leadership with a series of products that bring the new world of switching to this tried and true technology. Over time the announcements surrounding this product will demonstrate a continued commitment to the role of the 8272 in IBM′s Switched Virtual Networking strategy.

4.6 IBM 8272 Nways Token-Ring LAN Switch

The IBM 8272 Nways Token-Ring LAN Switches provide LAN switching technology for a token-ring environment. There are two models either currently or soon-to-be available, a Model 108 offering eight ports with a single expansion slot and a Model 216 offering 16 ports plus two expansion ports.

There will be the following three development and product implementation stages for models of the 8272 and its operational code:

1. Release 1, currently available, supports transparent switching as described in 4.2.5, “Transparent Switching” on page 124.
2. Release 2 will support source-route switching, which is outlined in 4.6.6, “Source-Route Switching” on page 135.
3. Release 3 will support source-route bridging.

The 8272 Nways Token-Ring LAN Switches Models 108 and 216 provide high-speed forwarding of token-ring frames among shared or dedicated token-ring segments attached to any of their ports. They are both *auto-sense* and *auto- configure* devices. That is, they do not require configuration for ring-speed, media-type or connection-type.

The 8272-108 supports a maximum of 1790 active token-ring LAN station addresses per port, and no more than 10,000 in total per unit.

Model Highlights 8272-108

- Transports token-ring frames among up to eight shared or dedicated token-ring LAN segments using twisted pair media (UTP/STP) using RJ-45 connectors, operating at 4 or 16 Mbps.
- Can double the bandwidth available to a single LAN station on a dedicated LAN segment from the usual 16 Mbps to up to 32 Mbps by using native full- duplex capabilities.

- Includes out-of-band console management facilities (accessible over a serial port), as well as in-band using TCP/IP and Telnet. An SNMP agent offers control from a central SNMP management station, such as one running IBM NetView for AIX or IBM NetView for Windows.
- Mounts in a standard 19-inch rack or may be situated on a desk or tabletop.
- Includes one Universal Feature Slot that will, in the future, support optional Universal Feature Cards for additional connectivity.

Model Highlights 8272-216

- Includes all features of the Model 108.
- Provides 16 fixed token-ring ports as well as two Universal Feature Slots, which provides a larger set of configuration options than otherwise possible with the Model 108.

4.6.1 Basic Operation

The IBM 8272-108 provides the ability to forward token-ring frames among up to eight shared token-ring LAN segments over token-ring twisted pair (UTP/STP) media using RJ-45 connectors. Similar in function to a multiport bridge, the IBM 8272 forwards token-ring frames from one of the eight ports to another using either transparent switching, source-route switching or source-route bridging depending on the mode configured.

The 8272-108 has a highly-parallel internal design optimized for performance, hence it is able to maintain media-speeds for simultaneous frame transfers between each of four distinct, possible pairs of ports. This feature allows it to provide an aggregate bandwidth of up to 64 Mbps, when switching among eight half-duplex ports. Networks with traffic patterns able to take full advantage of this ability could sustain throughput equivalent to four 16 Mbps token-rings.

Maximum frame size: The 8272 supports a frame size up to 4540 bytes.

4.6.2 Full-Duplex

The 8272 supports full-duplex (bidirectional) communication with devices on dedicated segments such as other IBM 8272s or workstations that are equipped with both full-duplex token-ring adapters such as the IBM Auto LANStreamer MC 32 Adapter, IBM Dual LANStreamer MC 32 Adapter, or the IBM Auto LANStreamer PCI Adapter, and the appropriate device drivers.

As users segment their LANs to relieve network congestion, they often extend this segmentation to a point where they have placed a single station (such as a high-traffic file server) on a dedicated LAN segment. This may reveal congestion of another sort, congestion at the network access point or adapter within the server.

To increase network bandwidth available to directly connected servers any of the eight ports can be configured to operate in full-duplex mode. Full-duplex token-ring (an emerging extension to the existing IEEE 802.5 standard) provides for simultaneous, two-way transmission between an IBM 8272 and a client. Full-duplex token-ring provides two, independent data paths between an 8272 and a LAN station on a dedicated LAN segment. Each has an available bandwidth for sending and receiving of 16 Mbps, for a combined 32 Mbps per full-duplex token-ring switch port.

In fact, these independent, parallel, full-duplex paths are extended throughout the internal design of the IBM 8272 so that each of the two 16 Mbps paths can be switched to different half-duplex ports. When all of the ports on an IBM 8272 Model 108 are configured for full-duplex operation, the Model 108 may provide an aggregate bandwidth of up to 128 Mbps. Bandwidth on the Model 216 would be commensurately greater.

When an end-station is attached directly to an 8272 port, the capability of the adapter and its associated device drivers will determine the mode that may be used. Full-duplex token-ring requires a full-duplex capable token-ring adapter, such as the IBM Auto LANStreamer MC 32 Adapter, IBM Dual LANStreamer MC 32 Adapter, or the IBM Auto LANStreamer PCI Adapter. If the adapter and its drivers supports full-duplex operation, then the ring activation protocol will indicate to the switch that full-duplex communication is enabled. If the station can only support token-passing mode, then the IBM 8272 port will establish a two-station ring and communicate using normal token-passing protocol rather than full-duplex.

4.6.3 TokenPipe

Configurations of larger than eight token-ring LAN segments can be constructed by connecting IBM 8272 ports together. This feature, known as *TokenPipe*, provides the capability for two IBM 8272s to communicate by connecting together up to four (user-configurable) full-duplex token-ring ports on one 8272 to those of another. Each of these full-duplex token-ring inter-switch links provides up to 32 Mbps of bandwidth between the switches, for a maximum of 128 Mbps of bandwidth (using four links) between switches. With multilink TokenPipes, traffic is automatically distributed across the links within a TokenPipe (using destination addresses) so that the traffic load is balanced. TokenPipe support provides users the ability to grow and tune network bandwidth as demands dictate,

4.6.4 Virtual Switches

Sometimes it is desirable to construct a switched network of less than eight ports, for security, traffic management, or other reasons. In this case, it is possible to cluster ports on a single IBM 8272 so that all traffic (including broadcast) will not flow between the ports of one group and ports outside the group. This *Virtual Switch* capability, as it is known, allows a single physical 8272 to be divided into multiple discrete switching domains, each consisting of a set of two or more switch ports. Since it makes no sense to cluster a single port, these Virtual Switches or domains will usually consist of two or more active ports.

4.6.5 TokenProbe

Another network monitoring feature of the IBM 8272 is its TokenProbe network monitoring port. This can be used to view traffic from the segments attached to the 8272. The TokenProbe port is a *designated* port (in lieu of a shared or dedicated token-ring port) on the front of the Model 108 that can be configured to monitor or ″mirror″ the activity of any single port.

Any of the ports on the 8272 can be designated to be the TokenProbe port. This assignment can be configured dynamically using the console or SNMP facilities of the 8272. To monitor a full-duplex token-ring port, the TokenProbe port can be configured to monitor either the transmit or the receive half of the full-duplex token-ring connection.

Using the TokenProbe port, traffic received or forwarded by the 8272 on any of its ports can be analyzed using a LAN analyzer.

The TokenProbe facility on the IBM 8272 was designed to be used with a variety of available protocol analyzers such as IBM DatagLANce Network Analyzer or the IBM Token-Ring Network 16/4 Trace and Performance Program.

4.6.6 Source-Route Switching

The IBM 8272, with Release 1 software, operates as a transparent switch, forwarding frames between LAN segments based only on MAC addresses. When operating in this mode, all attached rings must have the same ring number and there must not be multiple paths to remote rings. See 4.2.5, “Transparent Switching” on page 124 for more information on this.

The IBM 8272, with Release 2 software, can additionally operate as a source-route switch, forwarding frames between LAN segments based on the routing information field as well as MAC addresses.

This source-route switching technology was developed by IBM to support high-performance switching in token-ring networks. This technology is optimized to take advantage of the benefits of source routing without forcing the user to manually configure ring and bridge numbers for every port on the switch.

With source-route switching, the routing information (RI) field is used for packet forwarding, with the destination address used only by the last switch in the path to identify the specific port associated with it.

Source-route switching enables the use of the RI field in a source routed frame to determine when to forward (or filter) a frame. When source-route bridges are connected to two ports of an IBM 8272 switch, they append their unique bridge and ring numbers to the RI field of all explorer frames that they forward. With source-route switching, the switch port examines the RI field and the destination MAC address for frame forwarding and filtering. Since the switch is transparent to the source-route bridges, they appear as though on the same physical ring. The switch port learns the local ring number, associated with a particular port, by observing the RI field of explorer frames that it receives. When source-routed frames are received at a port, the port locates the local ring ID (that it has learned) within the RI field. Then, based on the direction bit, the next ring ID and bridge ID is located (adjacent to the local ring ID). This information can then be associated with the port for filtering and forwarding of frames. It is maintained in port forwarding tables along with the MAC addresses of frames that are locally accessible through that port. If the destination MAC address is directly accessible on one of the switch ports, then the local ring number will be the last ring number in the RI string. In this case, the MAC address rather than the next hop is used to select the destination port.

This lookup function is performed in hardware at each port and does not delay the frame forwarding process.

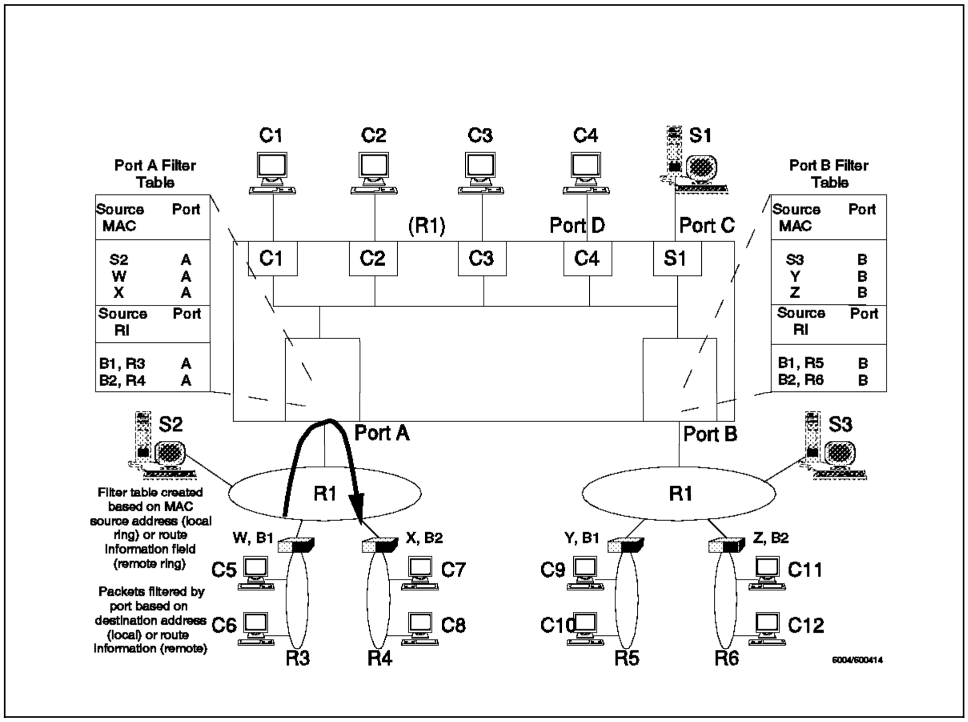

Figure 25. Source-Route Switch Learning and Filtering Tables

A source-route switch configuration example is shown in Figure 25 and Figure 26 on page 138 for the same topology that was used in the transparent switching example. The significant difference is in the contents of the Port A and Port B port forwarding tables. The list of addresses associated with client stations (C5-C12) on the remote rings (R3-R6) are now replaced with a portion of the route information field to identify the remote ring and bridge number. Since any frames that pass through the switch to reach these stations must be specifically routed frames, the switch can utilize only the RI field contents to make the correct forwarding decision.

All stations that are local to the switch on ring R1 remain in the port tables as in the previous example and the switch uses the MAC address to forward frames to these stations. In essence, a portion of the RI field is learned and maintained just as MAC addresses are learned and maintained in transparent switching. If a bridge moves, the RI change will be seen by the switch port and a new route will be learned. Routes will be aged from the table in the same manner as MAC addresses are aged.

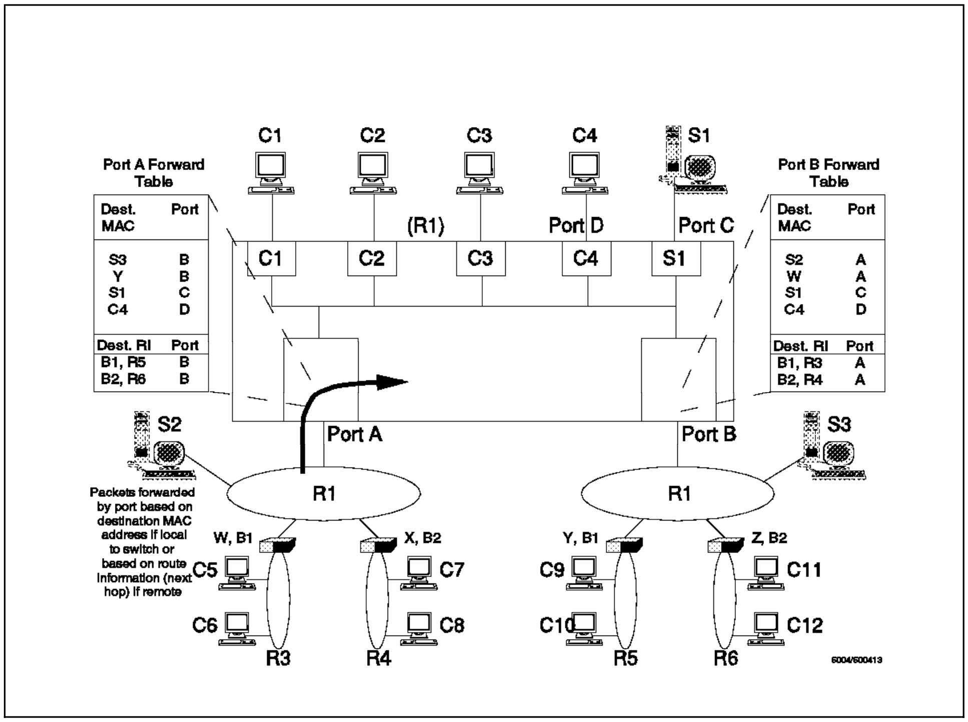

Figure 26. Source-Route Switch Forwarding Tables

4.6.7 Operating Modes

The IBM 8272 can be configured to operate in cut-through, store-and-forward or adaptive cut-through switching mode in the same way as the IBM 8271 already described in Section 4.2.6, "Operating Modes" on page 126. This configuration is independent of whether transparent or source- route switch mode is being used.

There is, however, a potential conflict between cut-through mode of operation and the need to capture the token at the target port on a shared-media link. The 8272 design allows the output port to temporarily buffer the incoming frame as it is being transferred, while at the same time attempting to acquire the token. Frame transmission will begin as soon as the token is seized. If the target ring load is extremely high, cut-through operation will revert to a store-and-forward mode, with the entire frame being stored at the output port buffer while waiting to capture the token.

Note

Cut-through transfer mode is supported only on 16 Mbps ports when the target port is operating at 16 Mbps or 4 Mbps.

4.6.8 Spanning Tree

The IBM 8272 Model 108 also supports the Spanning Tree Algorithm and is fully compliant with IEEE 802.1d. This allows the IBM 8272 to participate in complex network configurations where several IBM 8272s are interconnected by redundant paths.

4.6.9 Universal Feature Slot

Each Universal Feature Slot of an 8272 will support one optional Universal Feature Card. These include the following (when available):

- 4-Port Token-Ring UTP/STP UFC
- 2-Port Token-Ring Fiber UFC
- Ethernet Bridge UFC[6]
- 100 VG AnyLAN UFC[6]
- 1-Port ATM (155 Mbps SONET) UFC[6]
- One FDDI DAS/Fiber connection (A,B ports)[6]
- One FDDI SAS/Fiber port[6]
- One FDDI SAS/UTP port[6]

4.6.10 Universal Feature Card (UFC) Capabilities

The IBM 8272 Model 108 includes one Universal Feature Slot, into which a variety of optional, field-installable UFCs can be inserted to provide additional connectivity. See 4.8, "IBM 827X Universal Feature Cards" on page 142 for details.

6 Statement of Direction

4.6.11 Additional Features

The IBM 8272 also provides several additional features that contribute to reducing the cost of installing, operating, and owning switched token-ring networks.

With adaptive cut-through capabilities described earlier, the IBM 8272 can be self-optimizing when problems arise, which could affect overall network reliability and performance. Additionally, the 8272 will automatically sense which of the following token-ring connection types is being employed on each port:

- A shared media segment on a token-ring concentrator
- A directly connected token-ring LAN station
- Operating in half duplex or full-duplex mode
- Operating at 4 Mbps or at 16 Mbps
- Another 8272

It will also automatically configure each port to operate at the highest level of capability (without operator actions). The auto-sense/auto-configure capability minimizes network administration associated with both initial installation and ongoing network changes and helps assure optimum use of the network without having to manage a large number of configuration variables. The auto-sense/auto-configure capability can, however, be overridden by explicit configuration.

No special cross-over cables are required for token-ring stations on dedicated media segments or for switch-to-switch connections. These connections can use the same cabling as can be used to connect the IBM 8272 port to shared media segments.

4.6.12 Configuration and Management

The IBM 8272 is configured using a terminal (VT-100 or compatible) connected through the serial port on the front panel of the 8272, or by a Telnet session. When using the serial port, this terminal can be directly connected to the 8272 or can be remotely connected with a modem. Status display and operational control can also be performed using this same terminal.

As an alternative to the configuration terminal the IBM 8272 contains an SNMP (RFC 1213/MIB-II compliant) management agent that allows an SNMP management station, such as IBM NetView for Windows or IBM NetView for AIX, to interrogate and modify management data to obtain status or to control the operation of the 8272.

Upgrades to the functional code in the IBM 8272 can be accomplished in a variety of methods. Code may be downloaded from a PC connected to the serial interface on the 8272, using the Xmodem protocol or using TFTP protocol from a workstation directly attached to the network.

The IBM 8272 also contains BOOTP/TFTP support that allows the 8272 to participate in RFC 951-compliant environments for remote boot. Central BOOTP servers can configure the IP addresses for a network of IBM 8272 Nways Token-Ring Switches from a central location.

Remote Boot Capability

Additionally, BOOTP servers can designate (by IP address) the Trivial File Transfer Protocol (TFTP) server that contains an image of the IBM 8272 microcode from which a switch can load its operating environment when starting.

4.6.13 Filtering

Often, for either security, traffic management or other reasons, it may be desirable to control the traffic that flows through the IBM 8272 switch. Both, source as well as destination MAC address filters are supported on the 8272 providing the ability to filter or inhibit frame flow at port of entry. Users may configure a filter table in the 8272 consisting of LAN (MAC) addresses and related port numbers. Frames originating from or destined to any of the specified addresses sent from any station on the associated switch port would not be forwarded to any other switch port.

4.6.14 Physical Characteristics

The IBM 8272 Nways Token-Ring LAN Switch Model 108 can be mounted in a standard 19 inch rack. Mounting brackets are provided.

4.7 IBM 8272 Nways Token-Ring LAN Switch Model 216

The Model 216 is identical in function to the Model 108, except that it has sixteen LAN ports and supports two UFCs. At time of publication, the Model 216 has not been announced as generally available.

4.8 IBM 827X Universal Feature Cards

The IBM 8271 and IBM 8272 Models 108 and 216 can have optional Universal Feature Cards (UFCs) installed to provide additional functions such as extra twisted-pair (UTP/STP) ports, high-speed uplink connections, or other shared-media LANs.

Currently available UFC options are described in the following sections.

4.8.1 IBM 8271 Universal Feature Cards

100Base-TX UFC

The 100Base-TX UFC provides one 100Base-TX MDI-X port with an RJ-45 connector suitable for connection to a shared 100-Mbps Ethernet segment with a compatible 100 Mbps repeater or directly to a compatible 100 Mbps Ethernet adapter, such as IBM 100/10 PCI Ethernet Adapter.

Fully compliant with IEEE 802.3u specifications, the 100Base-TX UFC supports connections over an unshielded twisted-pair, Category 5 (UTP-5) cable at distances up to 100 meters.

The 100Base-TX port can be configured as either half-duplex for connection to a shared 100 Mbps Ethernet segment or full-duplex for connection to a dedicated Ethernet device. These features are shown in Figure 27 on page 143.

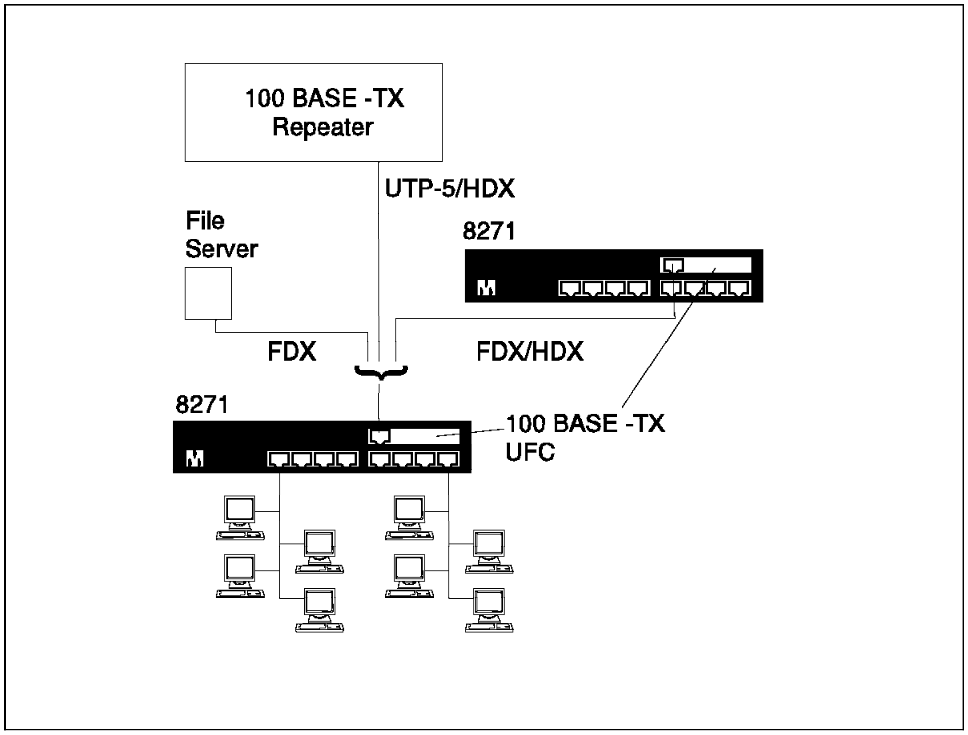

Figure 27. 100Base-TX UFC Interconnections

The 100Base-TX UFC faceplate contains the RJ-45 connector for the 100Base-TX port, and four LEDs that indicate the operational status of the UFC port. The 100Base-TX UFC port is configured in a similar way to the fixed 10Base-T ports on the IBM 8271 Models 108 and 216. Although the 100Base-TX UFC port can be used to connect to another Model 108, Model 216 or dedicated full-duplex 100-Mbps connection, it cannot be mixed with any of the fixed 10Base-T ports to form a multi-link EtherPipe connection. It can, however, be used to interconnect several IBM 8271 Model 108s and Model 216s with a shared 100 Mbps Ethernet backbone.

The EtherProbe port on the Model 108s and 216s cannot be configured to monitor the 100Base-TX UFC port.

4-Port 10Base-T UFC

The 4-Port 10Base-T UFC provides four additional 10Base-T MDI-X ports with RJ-45 connectors. Any of these four UFC ports can be configured in the same way as the fixed 10Base-T ports to provide either shared (half-duplex), 10-Mbps Ethernet connections or dedicated (full-duplex), 20-Mbps connections. This is shown in Figure 28.

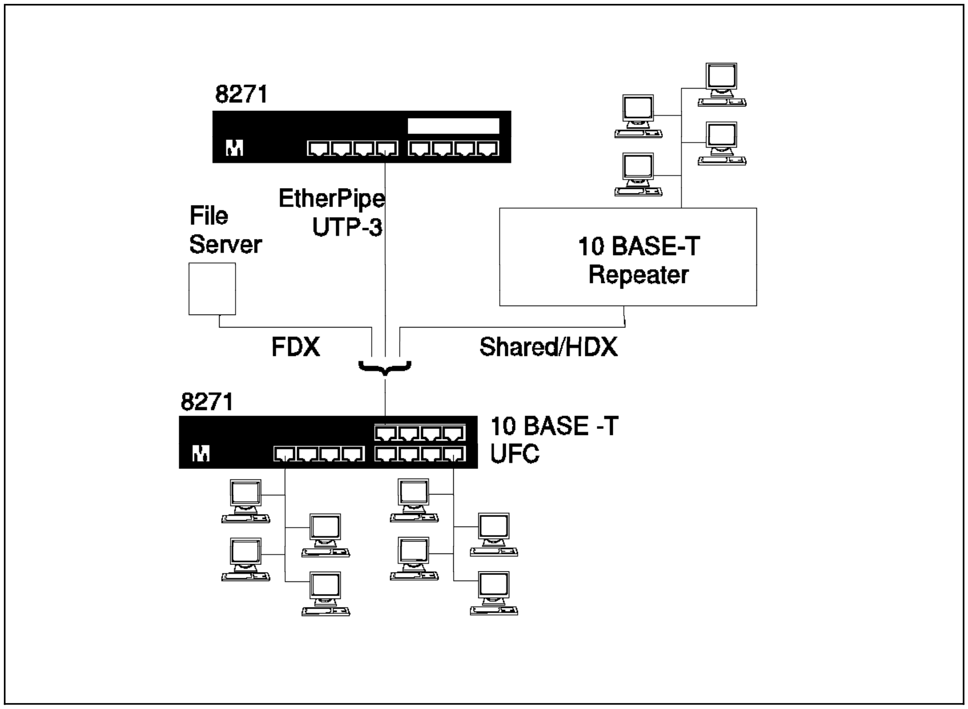

Figure 28. 4-Port 10Base-T UFC Interconnections

Each of the UFC ports can be configured (in combination with any of the fixed ports, or not) in multilink EtherPipe connections and included in virtual switches. The UFC ports support address filters, BOOTP/TFTP, Telnet, or SNMP sessions, and can be monitored by the EtherProbe port.

In addition to the RJ-45 connectors for the 10Base-T ports, there are 13 LEDs (one set of three for each port plus one for the UFC overall) on the UFC faceplate. These LEDs indicate the operational status of the UFC and its ports.

3-Port 10Base-FL UFC

The 3-Port 10Base-FL UFC provides three multimode fiber connections for ST connectors. Any of these three UFC ports can be configured in the same way as the fixed 10Base-T ports to provide either shared (half-duplex), 10 Mbps Ethernet connections or dedicated (full-duplex), 20 Mbps connections at distances up to 2 kilometers (6600 feet).

Shared connections are established by connecting a port on the UFC to a port on a compatible Ethernet repeater or hub such as the IBM 8224 Ethernet Stackable Hub or IBM 8250 or 8260 Multiprotocol Intelligent Hub. Dedicated connections can be established by connecting these ports directly to a compatible Ethernet adapter. This is shown in Figure 29.

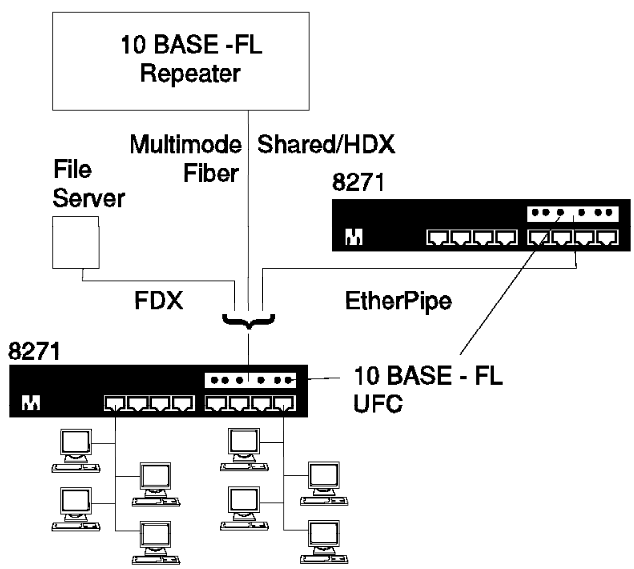

Figure 29. 3-Port 10Base-FL UFC Interconnections

In addition to the ST connectors for the 10Base-FL ports, there are 10 LEDs (one set of three for each port plus one for the UFC overall) on the UFC faceplate. These LEDs indicate the operational status of the UFC and its ports.

1-Port ATM (155 Mbps SONET) UFC

The 1-Port ATM UFC is a fully ATM Forum-compliant SONET interface that will attach to an ATM backbone, such as the IBM 8260 ATM switch, over multimode optical fiber cables of up to 2 Km (6600 feet). See Figure 30.

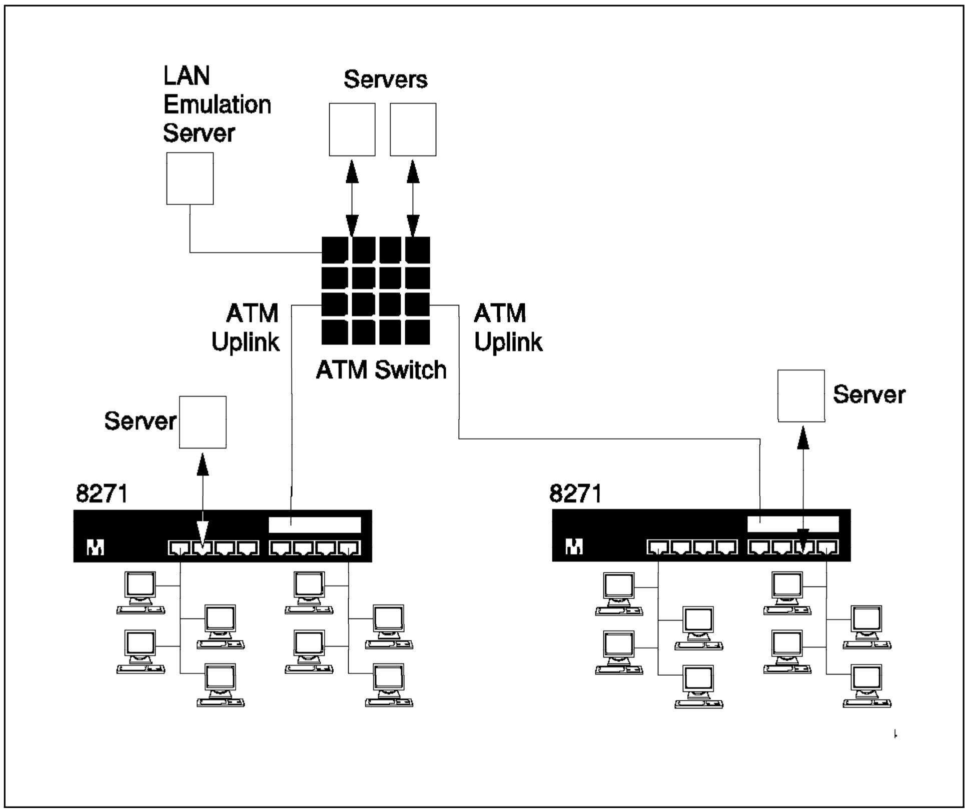

Figure 30. 1-Port ATM UFC Interconnections

The ATM UFC supports ATM Forum-compliant LAN Emulation permitting existing applications to be used without software changes.

FDDI UFCs

IBM plans to offer the following FDDI UFCs:

- Dual Access Station (DAS) multimode fiber connection.
- Single Access Station (SAS) multimode fiber connection.
- Single Access Station (SAS) UTP Category 5 connections

These UFCs provide 8271 interconnection to an FDDI backbone.

The UFCs are compatible with industry and international FDDI standards (ISO 9314/ANSI X3T9.5) and support FDDI station management (SMT 7.3). They bridge frames between Ethernet ports and the FDDI segment and support IP fragmentation (RFC 1188). This is shown in Figure 31 on page 147 and Figure 32 on page 148.

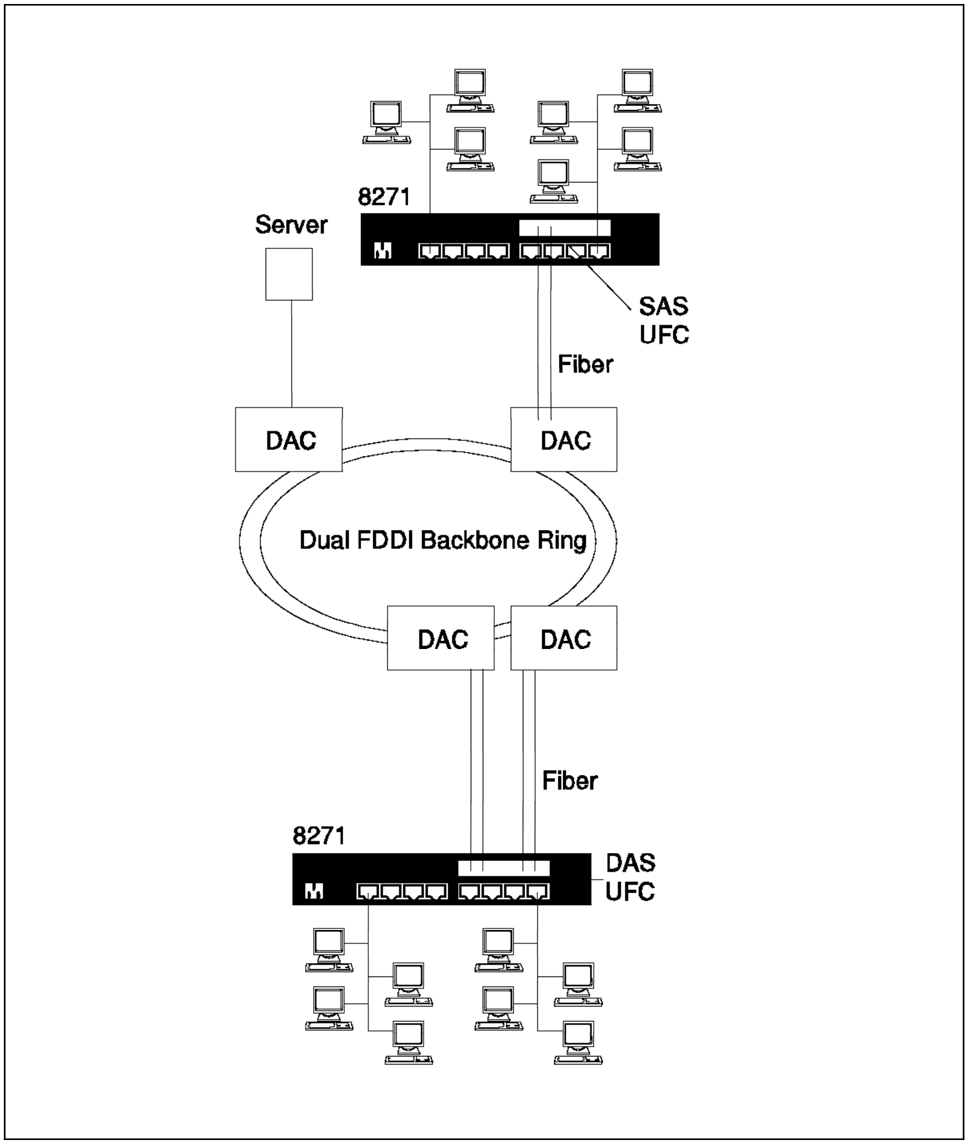

Figure 31. DAS and SAS FDDI Interconnections 1

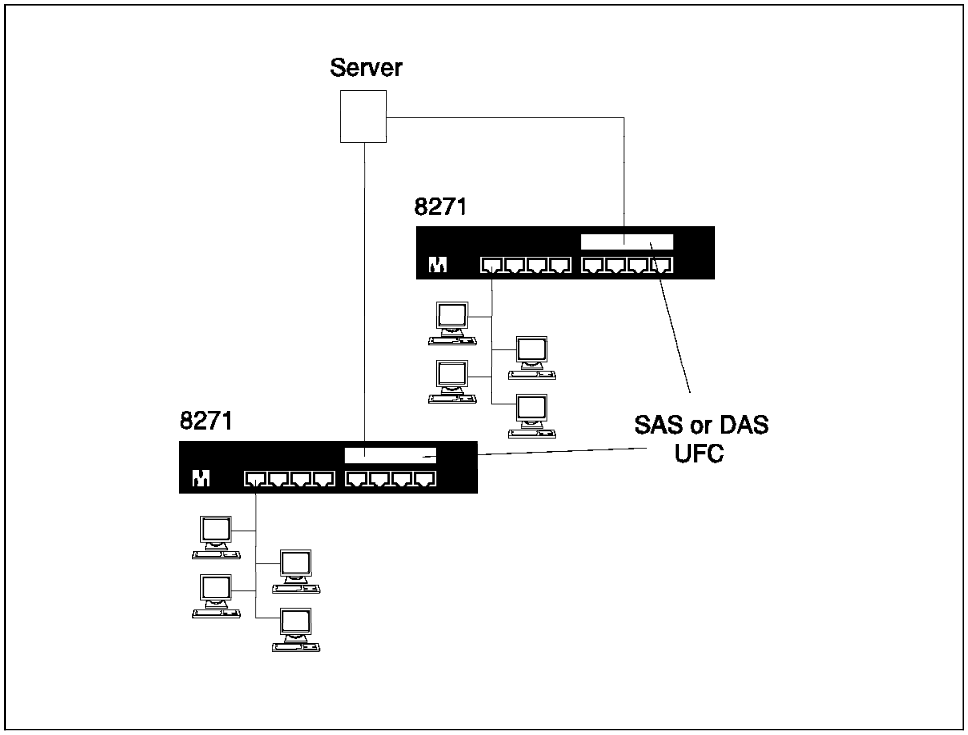

Figure 32. DAS and SAS FDDI Interconnections 2

100Base-Fx UFC

This UFC provides one 100Base-Fx multimode fiber connection for an SC connector. It is similar in function to the 100Base-TX UFC.

100VG AnyLAN

This UFC has a 100VG connection capability for links to stations that implement this high-speed technology. See Figure 33 on page 149.

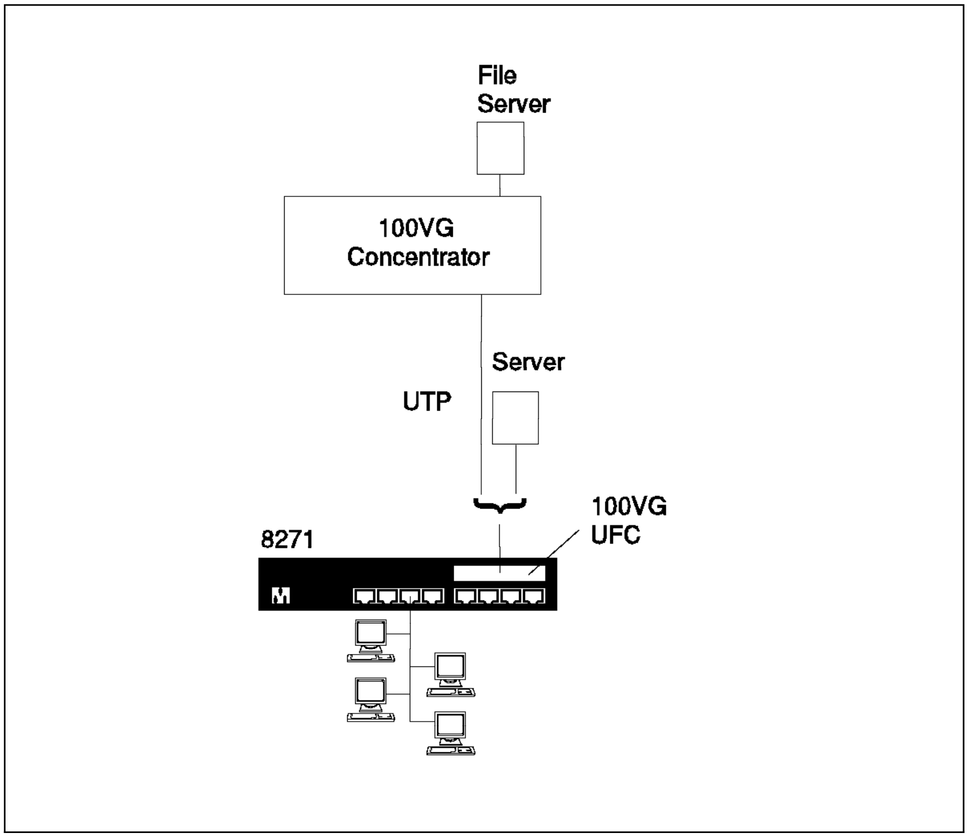

Figure 33. 100VG AnyLAN UFC Interconnections

4.8.2 IBM 8272 Universal Feature Cards

These UFCs provide additional token-ring ports for Models 108 and 216 which can be configured similarly to the fixed ports. Each or all of these UFC ports can be configured (in combination with any of the fixed ports if required) in multilink TokenPipe connections and as members of virtual switches. These UFC ports will support address filters and will support BOOTP/TFTP, Telnet or SNMP sessions.

4-Port Token-Ring UTP/STP UFC

The 4-Port Token-Ring UTP/STP UFC provides four token-ring ports, supporting token-ring twisted-pair media with RJ-45 connectors, which act identically to the remainder of the 8272 ports.

In addition to the port RJ-45 connectors there are 14 LEDs (one set of three for each port plus two for the UFC overall) on the face plate of this UFC. These LEDs indicate the operational status of the UFC and its ports.

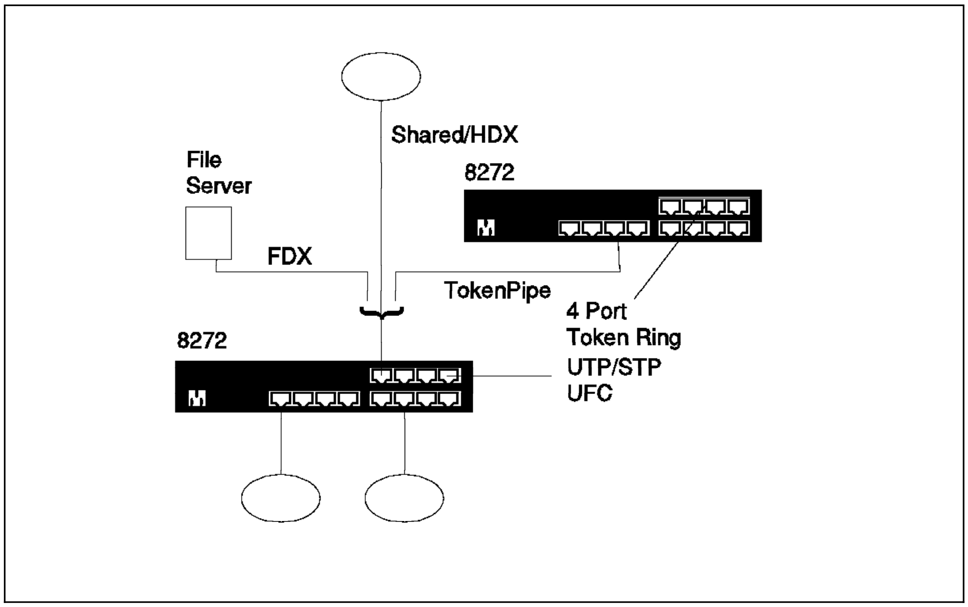

Figure 34. 4-Port Token-Ring UTP/STP UFC Interconnections

2-Port Token-Ring Fiber UFC

The 2-Port Token-Ring Fiber UFC provides two token-ring multimode fiber connections with ST connectors. Similar in capability to the 8272s token-ring UTP/STP ports, each port on this UFC can be connected to either a shared token-ring segment using fiber Ring-In (RI) or Ring-Out (RO) ports on a token-ring concentrator or hub, or a token-ring fiber port on another 8272. This is shown in Figure 35 on page 151.

With the first release of this UFC, token-ring segments between interconnected 8272s will utilize both the backup ring path as well as the main ring path. Consequently, the redundant path is in use and not available for ring recovery.

IBM intends to enhance this UFC to support automatic token-ring recovery, a function available in token-ring concentrators that automatically recovers from ring faults by utilizing the backup ring path of a dual token-ring.

A token-ring fiber connection between two 8272s can be combined with other token-ring connections (UTP/STP or fiber) to form a multilink TokenPipe. When used as TokenPipe connections, these fiber ports can be configured as full-duplex token-ring.

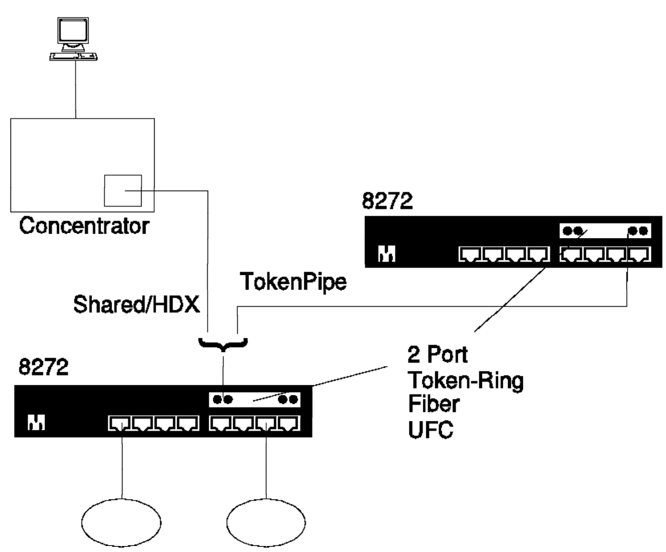

Figure 35. 2-Port Token-Ring Fiber UFC Interconnections

In addition to the ST connectors for the token-ring fiber ports, there are nine LEDs (one set of three for each port plus three for the UFC overall) on the faceplate of this UFC. They indicate the operational status of the UFC and its ports.

Ethernet Bridge

This UFC contains a translational bridge that forwards traffic between any of the switched token-ring ports on the 8272 and one Ethernet 10Base-T port on the UFC. It supports Ethernet V2 or IEEE 802.3 framing, and provides media-speed bridging between token-ring and Ethernet.

The UFC also contains a separately addressable SNMP agent.

1-Port ATM (155 Mbps SONET) UFC

The 1-Port ATM UFC, already described in "1-Port ATM (155 Mbps SONET) UFC" on page 146, is also fully supported by the IBM 8272.

FDDI UFCs

The FDDI UFCs, already described in "FDDI UFCs" on page 146, are also fully supported by the 8272.

4.9 8272 Transparent Mode and Source-Route Switching/Bridging

The following discussion sheds some light on some of the interesting implementation questions that will occur as the 8272 is incorporated within existing networks.

4.9.1 8272 Implementation Guidelines

The initial IBM 8272 product will be a transparent switch. In this microcode release, the 8272 does not need to be configured with a bridge number or the segment numbers to which it is attached. This means that all active ports in a domain must have the same segment number, which is external to the 8272 switch configuration. The 8272 relies on the network designer to ensure that all attached segments are configured properly. Each domain participates uniquely in a spanning-tree as if it were a separate virtual switch. See the description of 8272 virtual switches in 4.6.4, "Virtual Switches" on page 135. A domain may consist of one or more ports. The number of active ports in an 8272 determines the maximum number of domains.

Source-route bridging is distinguished from source-route switching by the bridge hop requirement. A bridge is uniquely identified within a RI field using the ring pair and bridge ID, while a source-route switch does not appear as a hop within the RI field. With source-route switching, the Routing Information Field is used for packet forwarding; the Destination Address is used by the last switch in the path to identify the specific port that is associated with the MAC address.

The first release (R1) of the 8272 software will perform transparent forwarding of frames based on the destination address of the frame. This is the same algorithm that has been implemented in transparent Ethernet bridges for the past decade. In contrast, source-route bridging enables token-ring bridges to recognize and use a unique field (the RI field) of a source-routed frame for determining when to forward (or filter) a frame. Special care should be used to ensure that there is only one physical path to a source-route bridge domain. Parallel paths will be supported with source-route switching. This is discussed further in 4.6.6, “Source-Route Switching” on page 135.

Every token-ring segment will have a unique segment number that is used by all bridges attached to the segment. Bridges will append their unique bridge and ring ID information to the RI field of all explorer frames that they forward. The door is opened to source-routing confusion when one connects unique segments with a transparent switch. The switch will pass frames from one segment to another (including explorer frames). An 8272 and all attached segments and/or stations will appear to be on the same physical ring (since the switch gives the appearance of a logical ring) to a source-route bridge. An 8272 switch is essentially invisible to source-route bridges.

For the first release of microcode for the 8272, an ideal use of the switch is to enhance available bandwidth for file servers. If ring segments are in need of reconfiguration (to minimize latency issues), be sure you understand the above discussion before implementing an 8272 as an alternative to a source-route bridge.

Release 2 (R2) of the 8272 microcode will implement source-route switching. Here, the switch port examines the RI field and/or the destination MAC address for frame forwarding and filtering. The switch port learns the local ring ID that is associated with a particular port by observing the RI field of explorer frames that it receives.

When source-routed frames are received at a port, the port locates the local ring ID (that it has learned) within the RI field. Then, based on the direction bit, the next ring ID and bridge ID is located (adjacent to the local ring ID). Essentially, a portion of the RI field is ″learned″ and maintained just as MAC addresses are learned and maintained in transparent bridging.

If a bridge moves or a new one is added, the RI change will be seen by the switch port and a new route will be learned. Routes will be aged from the table in the same manner that MAC addresses are aged. This information can be associated to a unique port for both filtering and forwarding of frames and is maintained in the port tables along with the MAC addresses of stations that are locally accessible via that port.

If the destination MAC address is directly accessible on one of the switch ports, then the local ring ID will be the last ring ID in the RI string. In this case, the MAC address rather than the next hop is used to select the destination port. This lookup function is performed in hardware at each port and does not delay the frame forwarding process.

Release 3 (R3) functions for the 8272 will provide full SRB capabilities across all ports. This means the 8272 can be configured to act as an SRB on some or all ports in store-and-forward mode, and act as a switch on remaining ports in cut-through or adaptive cut-through mode. Adaptive cut-through mode allows the user to configure thresholds for an 8272 to shift from cut-through mode to store-and-forward mode based on congestion parameters.

4.9.2 Configuration Scenarios

The 8272 initially supports transparent switching; later it will support source-route switching, and may use either MAC addresses or RI field information. Source-route switching is the function needed to operate with SR bridged networks. Some users however, (especially Novell environments) have implemented SRT bridging. The 8272 will operate in an SR and SRT environment. The following scenarios detail some of the issues involved in using an 8272 in typical token-ring environments.

Note

Special attention is drawn to several salient facts about the 8272 that network designers should take into consideration. The first is that the 8272 has a *largest frame size* of 4540 bytes.

The second is referenced in 2.8.6, "IBM LAN Architecture Statement" on page 71. Please note that not all vendors (including IBM) have been consistent with this statement in their treatment of this area. If your application is not in conformity with this statement, there is a strong likelihood that your 8272 implementation will be affected.

Lastly, any and all segments attached to an 8272 domain must have identical segment numbers.

4.9.3 Release 1 Transparent Switching

In this release, the 8272 will act as a transparent switch and will use the 802.1d spanning tree. This release can be used to implement networks consisting of both an 8272 with Release 1 microcode and source-route bridges, with care.

This means that you can have a mixture of transparent switching and source-route bridging in your network. Since the spanning trees in transparent switching and source-route bridging operate differently, (and independently of each other) one must avoid configuring networks in which a loop is created. Loops are normal in a source-route environment but not so in a transparent bridge world. Loops may occur most readily between different segments, where the segments are connected by a mixture of 8272s and source-route bridges.

4.9.4 Spanning Trees

When there is a mixture of 8272 transparent bridges and source-route bridges, there will be overlapping spanning trees (source-route and 802.1d transparent). These will be distinct, separate and autonomous trees with no interaction, but a definite influence on each other.

The spanning tree of the IBM 8272 will only apply to contiguous segments of the token-ring network and 8272s. The spanning tree BPDUs sent by an 8272 *will not* be used by SR bridges (and *will not* be forwarded by them either). Whereas, the spanning tree BPDUs sent by the SR bridges will not be used by the 8272 but will be forwarded by them. This will result in configurations similar to the diagram in Figure 36 to be invalid.

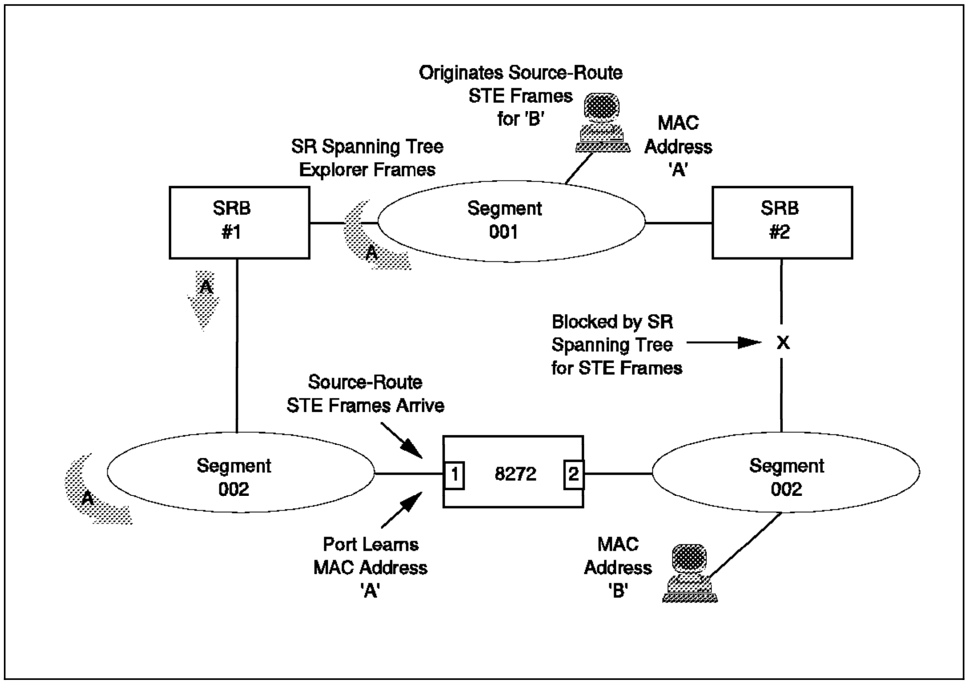

Figure 36. Invalid 8272 Configuration with Source-Route Bridges

This is because a source-route spanning tree is built to send Spanning Tree Explorer (Single Route Broadcast) frames along one network path only.

Since the SR Hello BPDUs are passed by an 8272, when the SR spanning tree is constructed one of the SR bridge ports (perhaps SRB #2 to Segment 002) will be in a *blocking* state to build a single path for STE frames, whereas the TB spanning tree is built to avoid network loops for all traffic. Because the Hello BPDUs of the 8272 are not passed by the SR bridges, the 8272 tree that is built assumes there are no loops in the network, when in fact there are loops. The resulting tree will have both ports of an 8272 (as shown in Figure 36 on page 155) to be in a *forwarding* state.

The end results of this are that if a station (shown as MAC address ′A′) sends an STE frame, the SR spanning tree will correctly forward it along a single path to the destination station. The 8272 will learn the address of the source station (A) on the original port of arrival (1), as shown in Figure 36 on page 155. When any All-Route Explorer (ARE) frames are sent out by stations attached to Segment 001 (shown as MAC Address ′A′ in Figure 37) they will reach both ports of the 8272 (one over Segment 001-SRB #1-Segment 002 and another one over Segment 001-SRB #2-Segment 002). This will result in an unstable filtering database in the 8272 (showing the station to be on port 1 at one point and at another point showing it on port 2). Therefore, avoid configurations where a loop exists between segments connected by both SR bridges and 8272s.

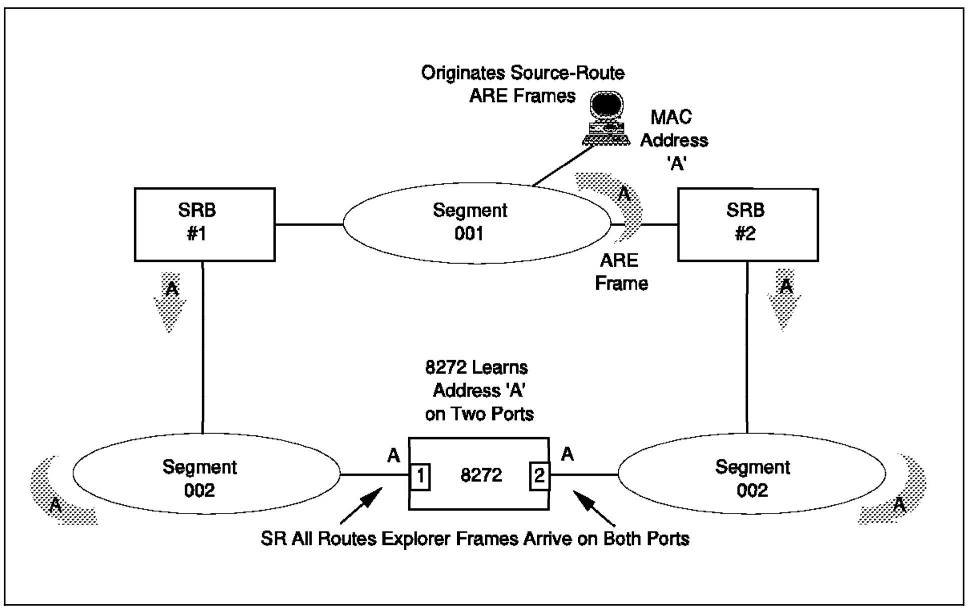

Figure 37. Effect of Duplicate Paths

However, one can prevent this problem by assigning path costs to the ports so that the blocked port belongs to an SR bridge and not to a transparent bridge/switch. This is displayed in Figure 38 on page 157. In this case, since a blocked port does not forward any frame (Specifically-Routed frame, Spanning-Tree Explorer frame, or All-Routes Explorer frame), then the problem that is described above will not arise. However, please note that if the configuration of the spanning tree changes for any reason, then your configuration may become invalid. Especially, if there is a loop and the port which is *blocked* belongs to a transparent bridge/switch.

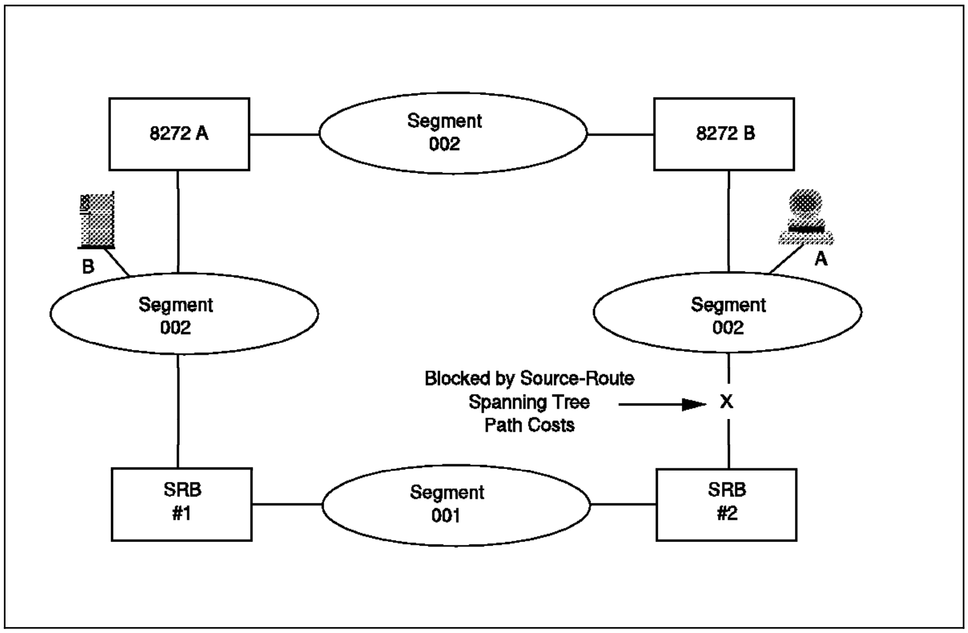

Figure 38. An SRB Network with 8272s

4.9.5 SRT Bridge and 8272 Switch

See 2.8, "Source-Route Transparent Bridging" on page 61 for a discussion of Source-Route Transparent Bridging. When a network has SRT bridges instead of SR bridges (contrast Figure 36 on page 155 with Figure 39 on page 158), one must avoid the same type of looping that is described for SR bridges. This is because although the SRT bridges actively participate in the 802.1d spanning tree protocol with the 8272, (which would result in a port, as shown in the diagram in Figure 39 on page 158 being in a *blocking* state), one is still faced with a problem. SRT bridges examine TR frames to see if routing information is present. If there is a RIF present, they will act on the frame as an SR bridge and forward the frame. Any ARE frames from stations attached to Segment 001 will reach both ports of the 8272, as shown in Figure 40 on page 159. Again, this will result in an unstable address filtering database as described previously.

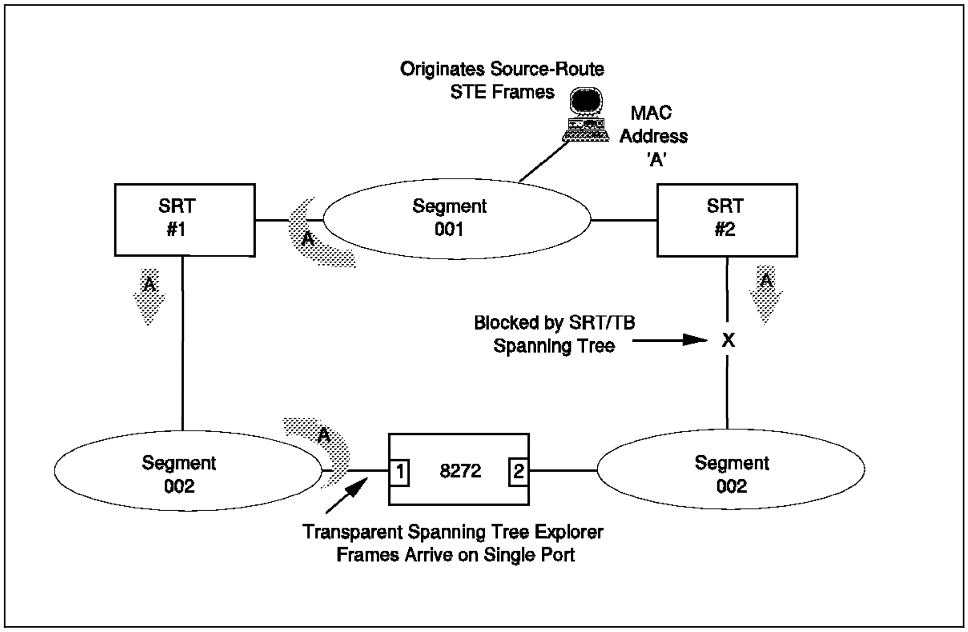

Figure 39. Blocked SRT Spanning Tree

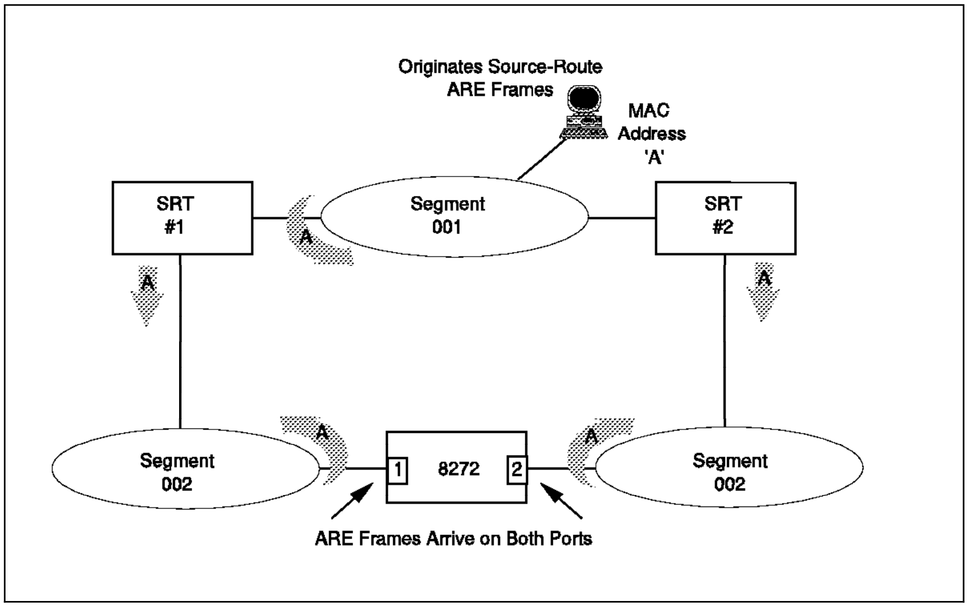

Figure 40. ARE Frames Passed by SRT Bridges

Since SRT bridges take part in the 802.1d spanning tree (as does an 8272), if an SR bridge is replaced with an SRT bridge (for example, Figure 38 on page 157 compared to Figure 41 on page 160) then a situation may develop that results in an undesirable spanning tree. This is shown where the port between 8272 B and Segment 002 is in a *blocking* state in Figure 41 on page 160. In this case, one will have a situation where there are multiple segments with the number 002 assigned to it, and these multiple segments 002 are connected to each other with source-route bridges. It is well known that this is an invalid configuration in source-route networks (segment numbers must be unique). Therefore, station A will not be able to communicate with station B.

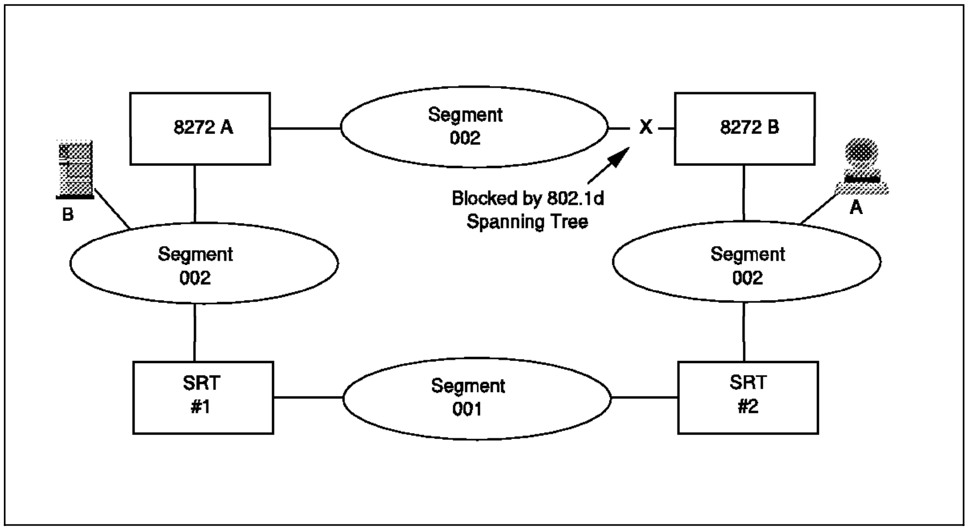

Figure 41. Undesirable 802.1d Pruning Results

If the configuration of the spanning tree changes due to the failure of ports, bridges, or switches, then any configuration may become an invalid one. Especially if one has a loop and the blocked port belongs to an 8272 switch rather than an SRT bridge.

4.9.6 Release 2 Source-Route Switching

In this release the 8272 will act as a source-route switch. This means that it will use source-routing as opposed to transparent bridging. See 4.6.6, "Source-Route Switching" on page 135 for a more complete discussion of how source-route switching will work. The main differences between source-route switching and source-route bridging are:

1. All the segments connected to a source-route switch will have the same segment number.
2. The SR switches use 802.1d spanning tree protocol instead of the source-route spanning tree protocol.

The 8272 acting as a source-route switch can be mixed with source-route bridges. Even if there is a loop as shown in Figure 38 on page 157 or the examples given for Release 1, the configuration will be valid. There are no transparent bridge address table management issues (as described in 4.9.3, "Release 1 Transparent Switching" on page 154).

4.9.7 High-Availability Designs with an 8272

Classic, dual token-ring backbone installations for high-availability require special attention when integrating an 8272 transparent switch. A typical redundant network design is shown in Figure 42. Because the 8272 is a transparent switching device (in its first release of microcode) any and all stations and/or networks attached to it must appear to be on the same segment. To use an 8272 in this network requires the support of source-route switching in the second release (R2) of microcode. To use an 8272 in a dual backbone design requires some advance planning.

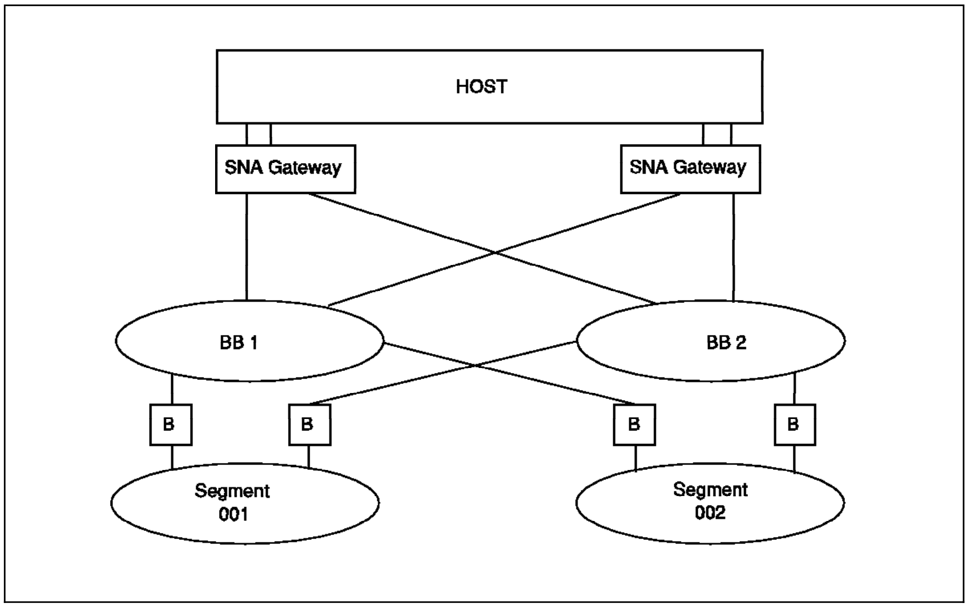

Figure 42. High-Availability Network Design with Source-Route Bridges

Figure 43 on page 162 shows what the design shown in Figure 42 would become if IBM 8272 Nways Switches were added to the network. The overall integrity of the design has not been impacted; users still have a high-availability design. The addition of the 8272 switches will provide a net increase in total bandwidth available to attached devices on the backbone segments.

The diagram in Figure 43 on page 162 shows two connections between the 8272s and their respective backbones. While this is possible, it should be pointed out that doing so requires that the spanning tree option be enabled. When this is done, one of the connections will be *blocked*. This connection will become active upon failure of an *active* port. In some environments the backbone may actually reside inside the 8272, in which case stations will directly attach themselves to the device. In this circumstance, it is not necessary to enable the spanning tree.

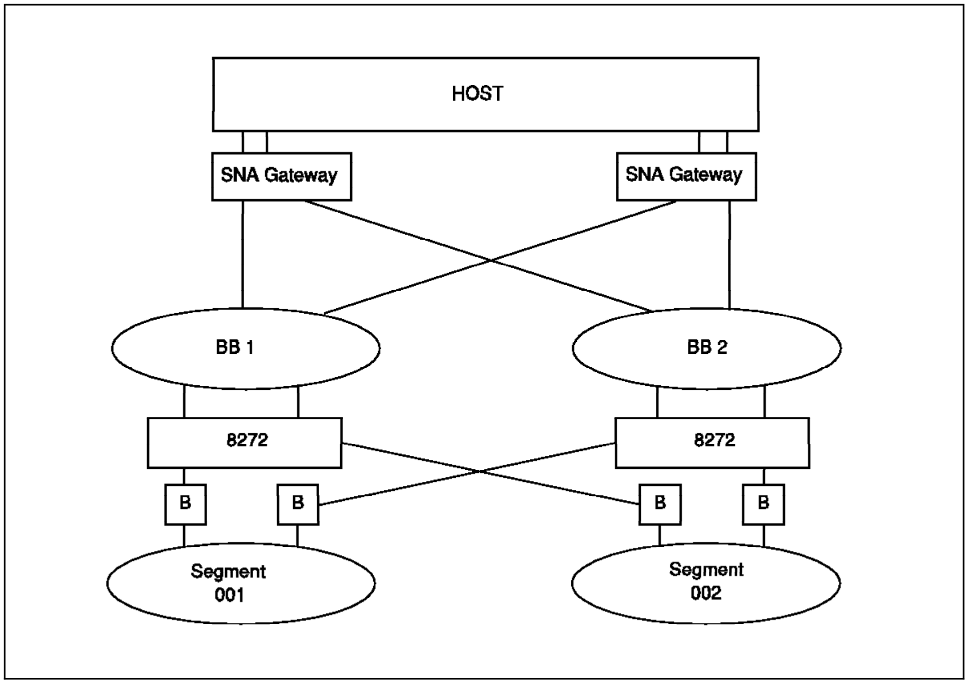

Figure 43. High-Availability Network Design Incorporating 8272s

4.9.8 Release 3 Source-Route Bridging

In this release, the 8272 will provide support for source-route bridging. This means that the 8272 will take part in the token-ring spanning tree; each segment will have its own unique segment number.

The interoperability issues are the same as what currently exists between SR and SRT bridges. However, if one configures a network that consists of 8272 Release 1, 2 and 3 as well SR (or SRT) bridges, then one must take into account the limitations previously described.

Appendix A. Bridge Matrix Table

Table 37. Bridge Product Matrix

	8229 Bridge	8250 T-R Bridge Modules	8250 2-port Ethernet Bridge	8250 6-port Ethernet Bridge	8260 Multiprotocol Interconnect	8281 ATM LAN Bridge	Local T-R Bridge/DOS V 1.0	LANStreamer Bridge/DOS V 1.0	Remote T-R Bridge/DOS V 1.0	Frame Relay T-R Bridge/DOS V 1.0	6611 Model 120	6611 Model 140	6611 Model 170	RouteXpander/2 V 2	RouteXpander/2 V 1.0.1	3172 MPE	3174 Config. Support C6	8271 EtherStreamer Switch	8272 LANStreamer Switch
Local bridging																			
Token-ring - token-ring	Y	Y	-	-	Y	Y	Y	Y	-	-	-	Y	Y	Y	Y	Y	Y		Y
Ethernet - Ethernet	Y	-	Y	Y	Y	Y	-	-	-	-	-	Y	Y	-	Y	-	-	Y	-
Token-ring - Ethernet	Y	-	-	-	Y	-	-	-	-	-	-	Y	Y	-	-	-	-	-	-
ATM - LAN	-	-	-	-	-	Y	-	-	-	-	-	-	-	-	-	-	-	Y	Y
Max. bridged segments	2	2	2	2	8	4	2	2	-	-	-	4	7	2	2	2	2	16	16
Remote bridging																			
Token-ring - token-ring	Y	-	-	-	-	-	-	-	Y	Y	Y	Y	Y	Y	Y	Y	Y	-	-
Ethernet - Ethernet	-	-	-	-	-	-	-	-	-	-	Y	Y	Y	-	Y	-	-	-	-
Max. local segments	1	-	-	-	-	-	-	-	1	1	1	3	6	2	1	1	2	-	-
Frame Relay max. ports	-	-	-	-	-	-	-	-	1	1	2	6	12	-	1	1	1	-	-
X.25 max. ports	-	-	-	-	-	-	-	-	-	-	-	-	-	1	-	-	-	-	-
Point-to-point max. ports	1	-	-	-	-	-	-	-	1	1	2	6	12	1	1	8	1	-	-
Max. total remote ports	1	-	-	-	-	-	-	-	1	1	2	6	12	8	1	1	1	-	-
Max. remote segments	1	-	-	-	-	-	-	-	1	note 6	note 6	note 6	note 6	note 6	note 6	note 6	8	-	-
Notes	-	-	-	-	*1	-	-	-	-	-	*2	*2	*3	*4,5	-	-	*6	-	-

Notes:

1. 8260 has six Ethernet ports. Two additional Ethernet or token-ring ports are optional.
2. 6611 model 120 has one LAN and one WAN adapter, and cannot be used as a local bridge.
3. 6611 model 140/170 can have a maximum of four/seven LAN or WAN adapters respectively.
4. RouteXpander/2 Version 2 base product supports one LAN and one WAN port. Multiport support requires RouteXpander/2 Multiport Support Program, which supports a maximum total of 9 LAN and WAN ports.
5. X.25 support requires RouteXpander X.25 Support/2.
6. Each Frame Relay connection up to 200 DLCIs. 3174 uses only eight for bridging.

A.1 LAN Products Table

The following table lists some common LAN-based products.

Table 38 (Page 1 of 5). LAN Products Table

Type	Feat	Description	TR	Eth.	FDDI	ATM	ICS	RJ-45	AUI	Fiber	Back plane	Comm.
8220	003	Fiber Optic Converter	X				1			1		
8222	001	Ethernet Workgroup Hub		X				6	1			
8222	008	Ethernet Workgroup Hub		X				8	1			
8222	016	Ethernet Workgroup Hub		X				16	1			
8224	001	Ethernet Stackable Hub (unmanaged)		X								
	002	Ethernet Stackable Hub (SNMP)		X				16	1	opt.		
8226	001	MAU RJ-45	X					8				
8228	001	Multiple Access Unit	X				8					
8229	001	TR Local SRT Bridge	X				2					
8229	002	TR to Eth. TR Bridge	X	X			1		1			
8229	003	TR Remote Bridge	X				1		1			V.35
8230	001	Controlled Access Unit (up to 4 LAMs)	80				RI/RO					
	5501	CAU 20-port STP LAM	X				20					
	5502	CAU 20-port UTP LAM	X					20				
	6738	CAU 20-p STP RJ-45 LAM	X					20				
	3048	Optical Repeater	X							RI/RO		
8230	002	CAU (same feat. as 001)	80				RI/RO					
	6748	CAU Active UTP LAM	X					20				
	4125	CAU Active 16/4 UTP LAM	X					20				
8230	003	CAU without Dual Ring (5 ICS, 3 LAMs)	140									
	7737	ICS RI/RO module	X				RI/RO					
	2007	STP RI/RO module	X					RI/RO				
	2010	Optical Fiber RI/RO	X							RI/RO		
	7751	ICS RI, Fiber RO	X				RI			RO		
	7754	Fiber RI, ICS RO	X				RO			RI		
	2011	ICS 2 port LIU	X				2					
	2008	UTP 3 port active LIU	X					3				
	2009	UTP 4 port LIU	X					4				
	8230	RLAM Attachment	X					(2)				
	8710	Remote LAM (RLAM)	X					16				
	8225	Active Remote LAM (RLAM)	X					16				
8230	013	CAU Same as 003, Dual Ring	140									

Table 38 (Page 2 of 5). LAN Products Table												
Type	**Feat**	**Description**	**TR**	**Eth.**	**FDDI**	**ATM**	**ICS**	**RJ-45**	**AUI**	**Fiber**	**Back plane**	**Comm.**
8230	213	CAU Same as 013, Active	140									
8230	04P	CAU (no LAMs)	X					16				
	5450	RI/RO Model 4 RJ-45	X					RI/RO				
	5458	RI/RO Model 4 Optical	X							RI/RO		
8230	04A	Active CAU (no LAMs)	X					16				
8235	001	DIAL-IN Access T-R 57Kbps V.24	X				1					8
8235	002	DIAL-IN Access Eth. 57Kbps V.24		X			1					8
8235	011	DIAL-IN Access T-R RJ-11	X				1					8
8235	012	DIAL-IN Access Eth. RJ-11		X			1					8
8235	021	DIAL-IN Access T-R 115Kbps V.24	X				1					8
8235	022	DIAL-IN Access Eth. 115Kbps V.24		X			1					8
8235	031	DIAL-IN Access T-R 115Kbps V.34	X				1					8
8235	032	DIAL-IN Access Eth. 115Kbps V.34		X			1					8
8238	Px1	T-R Stackable Hub B/S/G	X				16					
8238	Ax1	T-R Active Stackable Hub B/S/G	X				16					
	6240	Passive Expansion	X				16					
	6230	Active Expansion	X				16					
	6201	Fiber RI/RO	X							RI/RO		
	6202	Copper RI/RO	X				RI/RO					
8240	001	FDDI Concentrator (6 opt.slots) (wdfm)			24					A + B		
8244	06F	FDDI Workgroup Hub 4-p Fiber			X					2 + 4		
8244	06S	FDDI Workgroup Hub 4-p STP			X					2 + 4		
8244	06U	FDDI Workgroup Hub 4-p UTP			X					2 + 4		
8244	12F	FDDI Workgroup Hub 10-p Fiber			X					2 + 10		
8244	12S	FDDI Workgroup Hub 10-p STP			X					2 + 10		
8244	06U	FDDI Workgroup Hub 10-p UTP			X					2 + 10		
8250	06S	Intelligent Hub (6 slots, single power)	X	X	X							
8250	6HC	Intelligent Hub (6 slots, adv.arch.)	X	X	X							
8250	6PS	Intelligent Hub (6 slots, PS/2 build-in)	X	X	X							

Type	Feat	Description	TR	Eth.	FDDI	ATM	ICS	RJ-45	AUI	Fiber	Back plane	Comm.
8250	06W	Intelligent Hub (6 slots, 3174 blade)	X	X	X							
8250	017	Intelligent Hub (17 slots)	X	X	X							
8250	17W	Intelligent Hub (17 slots, 3174 blade)	X	X	X							
	3800	8-port 10BaseT STP		X				8			8	
	3801 3802	12-port 10BaseT UTP Telco		X				12Tel			12	
	3822	24-port 10BaseT UTP Telco		X				24Tel			24	
	3804	2-port Repeater		X					2			
	3814 -15-16	4-port 10BaseFL		X						4	4	
	3817	BNC Connector		X					6BNC			
	6773 - 7395	2/4 port 10BaseFL		X						2/4	2/4	
	3820	T-R MAU RI/RO RJ-45	X				2	8			8	
	7386	Token-ring Repeater	X				2					
	3821	20-port RJ-45	X					20			20	
	7400	18-port Active RJ-45	X					18			18	
	3822	Fiber Repeater +RJ-45		X				?		2		
	3883	T-R SR Bridge	X				1	1			1	
	3958	T-R SR/SRT Bridge	X				1	1			1	
	3142	8229 SRT Bridge Module	X	X			(?)	(?)	(?)		(?)	
	3182	8229 SR Bridge Module	X				(?)	(?)			(?)	
	3179	8229 SR/SRT Bridge	X				(?)	(?)			(?)	
	3828	Ethernet Bridge 2-ports		X					1		1-2	
	7384	Ethernet Bridge 6-ports		X				3	2		5	
	6714	Ethernet Switch 6-ports		X				3	2		5	
	6769	Ethernet Router 6-ports		X				3	2		5	
	3825	FDDI 8-port			X					8	8	
	3826	FDDI STP 8-port			X					8STP		
	3174	3174 Module	X				1					y
	7020	Frame Relay for 3174										V.35 .
	3155	8235 Module for T-R	X									V.24 .
	3160	8235 Module for Eth.		X								V.24
8260	010	Intelligent Switching Hub 10-slots	X	X	X							
8260	017	Intelligent Switching Hub 17-slots	X	X	X							
8260	10A	Intelligent Switching Hub 10-slots ATM	X	X	X	X						

Table 38 (Page 3 of 5). LAN Products Table

Table 38 (Page 4 of 5). LAN Products Table

Type	Feat	Description	TR	Eth.	FDDI	ATM	ICS	RJ-45	AUI	Fiber	Back plane	Comm.
8260	17A	Intelligent Switching Hub 17-slots ATM	X	X	X	X						
		100 Mbps ATM 4-port				X				4	4	
	1024	24-port 10BaseT, Telco, PPS		X				24Tel			24	
	1020 -40	20 / 40 port RJ-45, PPS		X				20/40			20/40	
	1036	36-port Telco, Bank Switching		X				36			36	
	11-,12-, 1310	10-port 10BaseFB PPS		X						10	10	
	7106 7206	Interconnect Module (2-slot for I/O module)	(X)	(X)							6 (+2)	
	8202	10BaseT I/O card		X				1				
	8203	10Base2 I/O card		X					1BNC			
	8204	10Base5 I/O card		X					1			
	8205	T-R I/O card	X				1					
	5204	8281 ATM LAN Bridge Module	(X)	(X)		X		4			4	
	sod(*)	8271 EtherStreamer Switch module		X		X		12			12	
	sod(*)	8272 LANStreamer Switch module	X			X		12			12	
	3018	Active T-R 18-ports RJ-45	X					16			18	
	3020	Passive T-R 20-ports RJ-45	X					18			20	
3172	003	Interconnect Controller	X	X	X	X						
	2225	Ethernet Busmaster Adapter		X					1			
	2225	EtherStreamer MC 32		X				1				
	2235	Auto LANStreamer MC 32	X					1				
	2300	FDDI Adapter			X					X		
	2310	TURBOWAYS 100ATM				X				X		
	2900	WAC Adapter X.21										X.21
	2910	WAC Adapter V.35										V.35
	2920	WAC Adapter RS-422/449										RS422
6611	120/125	Network Processor (2 slots)	1	1								1
6611	140/145	Network Processor (4 slots)	4	4								4
6611	170/175	Network Processor (4 slots)	7	7								7
	2025 2125	2-port Multi-interface Serial Adapter										2
3174		Communication Controller Conf. Support C6	X	X			1		1			
		Frame Relay Feature										1

Type	Feat	Description	TR	Eth.	FDDI	ATM	ICS	RJ-45	AUI	Fiber	Back plane	Comm.
		Dual speed Token-ring Adapter type 3A	X									1
		Peer Communication Feature (Coax)	X	(X)								32
8271	001	EtherStreamer Switch		X	X	X		8				
8271	108	EtherStreamer Switch		X	X	X		8+4		1		
8272	108	LANStreamer Switch	X		X	X		8+4		1		
8281	001	ATM LAN Bridge	(X)	(X)		X		4(6)		1		
	6812	100 Mbps ATM Adapter								1		
	6800	2-port Expansion Feature	(X)	(X)				2				
8282	001	ATM Workgroup				X		8-25M		100M		
	001	4-ports 25 Mbps ATM Module				X		4-25M				
----	0268	25 Mbps ATM Adapter ISA				X		1				
----	0269	100 Mbps ATM Adapter MCA				X				1		

Table 38 (Page 5 of 5). LAN Products Table

Note: (*) Statements of directions

Table 39. LSAP Reserved Addresses		
Bits	**HEX**	**Description**
0000 000x	00,01	Null SAP
0000 0010	02	IEEE 802.1B LAN Management
bbbb 001x	x2,x3	Net Management (Standards)
0000 0110	06	Department of Defense Internet Protocol (IP)
nnnn 011x	x6,x7	General Standards
pppp 101x	0000 1110	xA,xB
General Standards	0E	PROWAY Net Management and Initialization
0001 0010	12	4333 Management
0100 0010	42	MAC Bridges BPDUs
0100 1110	4E	Layer messaging service for factory automation
0111 1110	7E	ISO 8208 X.25 packet layer protocol
1000 1110	8E	PROWAY active station list maintenance
1010 1010	AA	Sub-Networking Access Protocol (SNAP)
1111 111x	FE,FF	All Stations Address
qqqq 111x	xE,xF	General Standards
1111 1110	FE	OSI Network Layer Protocol

Table 40 (Page 1 of 2). LSAP User-Defined Individual Addresses		
Bits	**HEX**	**Description**
0000 010x	04	SNA Path Control
0000 110x	0C	SNA
0001 000x	08	SNA
0001 000x	10	Novell NetWare / 4033 Data
1000 000x	80	user-defined - XNS
1000 010x	84	user-defined
1000 100x	88	user-defined
1000 110x	8C	user-defined
1001 000x	90	user-defined
1001 010x	94	user-defined

Table 40 (Page 2 of 2). LSAP User-Defined Individual Addresses		
Bits	**HEX**	**Description**
1001 100x	98	user-defined
1001 110x	9C	user-defined
1010 000x	A0	OEM reserved
1010 010x	A4	OEM reserved
1010 100x	A8	OEM reserved
1010 110x	AC	OEM reserved
1011 000x	B0	OEM reserved
1011 010x	B4	OEM reserved
1011 100x	B8	OEM reserved
1011 110x	BC	OEM reserved VINES IP
1100 000x	C0	IBM reserved
1100 010x	C4	IBM reserved
1100 100x	C8	IBM reserved
1100 110x	CC	IBM reserved
1101 000x	D0	IBM reserved
1101 010x	D4	Resource Manager (8230 Discovery)
1101 100x	D8	IBM reserved
1101 110x	DC	Address Discovery Protocol
1110 000x	E0	Novell NetWare
1110 010x	E4	PC Network Net Management
1110 100x	E8	Terminal Controller Comms for Retail POS
1110 110x	EC	IBM reserved
1111 000x	F0	PC Network NetBIOS
1111 010x	F4	LAN Network Management
1111 100x	F8	Remote Program Load, Initial Microprogram Load (IMPL)
1111 110x	FC	GDS RPL Discovery
Note: It is important to remember that all of the addresses above are user-defined in the industry.		

A.1.1 Ethertype Values

This table lists some of the most common EtherTypes. A more complete list can been found in RFC1700 (Request for Comments) or in *TCP/IP Tutorial and Technical Overview*, GG24-3376.

Table 41. Some Interesting Ethertypes

x′0000 - 05DC′	Length field of an IEEE 802.3 frame
x′0600′	Xerox IDP
x′0800′	IP protocol
x′0801′	X.75 Internet
x′0802′	NBS Internet
x′0803′	ECMA Internet
x′0804′	Chaosnet
x′0805′	X.25 Level 3
x′0806′	ARP
x′0807′	XNS
x′0900′	Ungermann-Bass net debugger
x′0A00-0A01′	Xerox IEEE802.3 PUP
x′0BAD′	Banyan Systems
x′6010-6014′	3Com Corporation
x′6000 - 6009′	DECnet
x′7809′	AppleTalk
x′8005′	HP Probe
x′8019′	Apollo Computers
x′8035′	RARP
x′8038-803F′	DEC
x′8046-8047′	AT&T
x′805D′	LLC over Ethernet V2
x′8069′	AT&T
x′809B′	AppleTalk
x′80F3′	AARP
x′8132-8136′	Bridge Communications
x′8137′	Novell IPX
x′814C′	SNMP
x′81CC-81D5′	Apricot Computers
x′8700-8710′	Motorola
x′9000′	Loopback
x′9001-9003′	3Com (Bridge) System Management

A.1.2 Burned-In Manufacturers IDs (OUI)

A more complete list of burned-in IDs (OUI) in universally administered MAC addresses can been found in RFC1700 (Request for Comments).

Table 42. Some Interesting Burned-In IDs

Canonical bit order	Token-ring bit order	Manufacturer′s name
00000Cxxxxxx	000030xxxxxx	Cisco
00000E	000070	Fujitsu
00000F	0000F0	NeXT
000010	000008	Sytek
00001D	0000B8	Cabletron
00005E	00007A	IANA
000065	0000A6	Network General
000093	0000C9	Proteon
0000A2	000045	Wellfleet
0000A7	0000E5	NCD X-terminals
0000A9	000095	Network Systems
0000AA	000055	Xerox
0000C0	000003	Western Digital
00AA00	005500	Intel
00DD00	007700	Ungermann-Bass
00DD01	007780	Ungermann-Bass
02608C	400631	3Com
080002	100040	3Com
080009	100090	HP
080011	100088	Tektronix, Inc.
08001A	100058	Data General
08001B	1000D8	Data General
080020	100004	Sun
08002B	1000D4	DEC
080041	100082	DCA
08005A	10005A	IBM
10005A	08005A	IBM
800010	010008	AT&T
AA0003	5500C0	DEC Global physical address for some DEC machines
AA0004	550020	DEC Local logical address for systems running

A.1.3 Defined Group Addresses

These are defined group addresses. The list is subject to change.

Table 43 (Page 1 of 3). Standardized Group Addresses

GROUP	Token-ring bit order	Canonical bit order
LAN traffic monitor	X′D500 D400 00C0′	X′AB00 2B00 0003′
VAXELN	X′D500 D400 0040′	X′AB00 2B00 0002′
DNA level 1 routing layer	X′D500 00C0 0000′	X′AB00 0003 0000′
DNA routing layer end nodes	X′D500 0020 0000′	X′AB00 0004 0000′
DEC DNA remote console (MOP)	X′D500 0040 0000′	X′AB00 0002 0000′
DEC DNA Dump/load assistance (MOP)	X′D500 0080 0000′	X′AB00 0001 0000′
Customer use	X′D500 2000 XXXX′	X′AB00 0400 XXXX′
System Communication Architecture	X′D500 2080 XXXX′	X′AB00 0401 XXXX′
LAN	X′FFFF 0002 0080′	X′FFFF 0040 0001′
LAN	X′FFFF 0006 0020′	X′FFFF 0060 0004′
LAN	X′FFFF 8007 0020′	X′FFFF 01E0 0004′
Loopback assistance	X′F300 0000 0000′	X′CF00 0000 0000′
Concord DTQNA	X′0000 9640 XXXX′	X′0000 6902 XXXX′
Shadow for prom 23-365A1-00	X′1000 D4AX XXXX′	X′0800 2BX5 XXXX′
Shadow for prom 23-365A1-00	X′1000 D4B6 XXXX′ through X′1000 D4EX XXXX′	X′0800 2B6D XXXX′ through X′0800 2BX7 XXXX′
Prom 23-365A1-00	X′1000 D4C4 XXXX′ through X′1000 D4CX XXXX′	X′0800 2B23 XXXX′ through X′0800 2BX3 XXXX′
VAXft 3000 fault-tolerant LAN addresses	X′1000 D40F XXXX′	X′0800 2BF0 XXXX′
Prom 23-365A1-00	X′1000 D40X XXXX′	X′0800 2BX0 XXXX′
VAXft 3000 fault-tolerant LAN addresses	X′1000 D407 XXXX′	X′0800 2BE0 XXXX′
Shadow for prom 23-365A1-00	X′1000 D42X XXXX′	X′0800 2BX4 XXXX′
Bridge management	X′1000 D444 0000′	X′0800 2B22 0000′
Prom 23-365A1-00	X′1000 D48X XXXX′	X′0800 2BX1 XXXX′

Table 43 (Page 2 of 3). Standardized Group Addresses		
GROUP	**Token-ring bit order**	**Canonical bit order**
NI20 products	X′5500 C0C0 XXXX′	X′AA00 0303 XXXX′
UNA prototype	X′5500 C000 XXXX′	X′AA00 0300 XXXX′
Miscellaneous	X′5500 C040 XXXX′	X′AA00 0302 XXXX′
H400 - TA Ethernet transceiver tester	X′5500 C040 0000′	X′AA00 0302 0000′
Prom 23-365A1-00	X′5500 C080 XXXX′	X′AA00 0301 XXXX′
DECnet phase IV station	X′5500 2000 XXXX′	X′AA00 0400 XXXX′
ISO 10589 level 2 1 intermediate stations	X′8001 4300 00A8′	X′0180 C200 0051′
All cons snares (ISO 10030)	X′8001 4300 00E8′	X′0180 C200 0017′
Reserved for transparent bridging	X′8001 4300 000x′	X′0180 C200 00x0′
Bridge	X′8001 4300 0000′	X′0180 C200 0000′
All LANs bridge management group address (802.1D)	X′8001 4300 0008′	X′0180 C200 0010′
Bridge management	X′8001 4300 0008′	X′0180 C200 0010′
ISO 10589 level 1 1 intermediate stations	X′8001 4300 0028′	X′0180 C200 0041′
Loadable device	X′8001 4300 0048′	X′0180 C200 0021′
All cons end systems (ISO 10030)	X′8001 4300 0068′	X′0180 C200 0016′
Load server	X′8001 4300 0088′	X′0180 C200 0011′
Reserved for FDDI	X′8001 4300 10X0′	X′0180 C200 080x′
FDDI all root concentrator MACs (ANSI X3T9.5)	X′8001 4300 1004′	X′0180 C200 0820′
FDDI RMT directed beacon	X′8001 4300 8000′	X′0180 C200 0100′
FDDI status report frame	X′8001 4300 8008′	X′0180 C200 0110′
OSI network layer end-stations	X′9000 D400 00A0′	X′0900 2B00 0005′
NetBIOS emulator (PSCG)	X′9000 D400 00E0′	X′0900 2B00 0007′
Local Area Transport (LAT)	X′9000 D400 00F0′	X′0900 2B00 000F′
OSI NL intermediate stations	X′9000 D400 0020′	X′0900 2B00 0004′
CSMA/CD encryption	X′9000 D400 0060′	X′0900 2B00 0006′

Table 43 (Page 3 of 3). Standardized Group Addresses		
GROUP	**Token-ring bit order**	**Canonical bit order**
Local Area System Transport (LAST)	X′9000 D420 XXXX′	X′0900 2B04 XXXX′
DNA level 2 routing layer routers	X′9000 D440 0000′	X′0900 2B02 0000′
FDDI ring purger advertisement	X′9000 D440 80A0′	X′0900 2B02 0105′
LAT directory service solicit	X′9000 D440 80D0′	X′0900 2B02 010B′
DNA naming service advertisement	X′9000 D440 8000′	X′0900 2B02 0100′
LAT directory service solicit	X′9000 D440 8020′	X′0900 2B02 0104′
DNA naming service solicitation	X′9000 D440 8080′	X′0900 2B02 0101′
All bridges	X′9000 D480 0000′	X′0900 2B01 0000′
All local bridges	X′9000 D480 0080′	X′0900 2B01 0001′
AppleTalk support	X′9000 E000 0000′	X′0900 0700 0000′
AppleTalk highest address within range except broadcast	X′9000 E000 003F through X′9000 E0FF FFFF′	X′0900 0700 00FC through X′0900 07FF FFFF′
Vitalink gateway	X′9000 3CA0 0080′	X′0900 3C02 0001′
Vitalink diagnostics	X′9000 3C40 00A0′	X′0900 3C02 0005′
Novell IPX	X′9000 7200 0040′	X′0900 4E00 0002′
Apollo domain	X′9000 7800 0000′	X′0900 1E00 0000′
Hewlett-Packard DTC	X′9000 9000 0020′	X′0900 0900 0004′
Hewlett- Packard probe	X′9000 9000 0080′	X′0900 0900 0001′

A.1.4 Current IEEE Functional Addresses

This is the current list of the IEEE defined functional addresses.

Table 44 (Page 1 of 2). Current IEEE and IBM Functional Addresses		
FUNCTION	**Token-ring bit order**	**Canonical bit order**
Active monitor	X′C000 0000 0001′	X′0300 0000 0080′
Ring Parameter Server	X′C000 0000 0002′	X′0300 0000 0040′
Network Server Heartbeat	X′C000 0000 0004′	X′0300 0000 0020′

Table 44 (Page 2 of 2). Current IEEE and IBM Functional Addresses		
FUNCTION	**Token-ring bit order**	**Canonical bit order**
Ring Error Monitor	X′C000 0000 0008′	X′0300 0000 0010′
Configuration Report Server	X′C000 0000 0010′	X′0300 0000 0008′
Synchronous bandwidth manager	X′C000 0000 0020′	X′0300 0000 0004′
Locate - directory server	X′C000 0000 0040′	X′0300 0000 0002′
NetBIOS	X′C000 0000 0080′	X′0300 0000 0001′
Bridge	X′C000 0000 0100′	X′0300 0000 8000′
IMPL server	X′C000 0000 0200′	X′0300 0000 4000′
Ring authorization server	X′C000 0000 0400′	X′0300 0000 2000′
LAN gateway	X′C000 0000 0800′	X′0300 0000 1000′
Ring wiring concentrator	X′C000 0000 1000′	X′0300 0000 0800′
LAN manager	X′C000 0000 2000′	X′0300 0000 0400′
User-defined	X′C000 0000 8000′ through X′C000 4000 0000′	X′0300 0000 0100′ through X′0300 0200 0000
ISO OSI ALL ES	X′C000 0000 4000′	X′0300 0000 0200′
ISO OSI ALL IS	X′C000 0000 8000′	X′0300 0000 0100′
IBM discovery non-server	X′C000 0001 0000′	X′0300 0080 0000′
IBM resource manager	X′C000 0002 0000′	X′0300 0040 0000′
TCP/IP	X′C000 0004 0000′	X′0300 0020 0000′
6611-DECnet	X′C000 2000 0000′	X′0300 0400 0000′
LAN Network Manager & 6611	X′C000 4000 0000′	X′0300 0200 0000′

List of Abbreviations

ATM	Asynchronous Transfer Mode
BPDU	Bridge Protocol Data Unit
DAC	Dual Attachment Concentrator
DAS	Dual Attachment Station
DHC	Dual Homing Concentrator
DHS	Dual Homing Station
DLCI	Data Link Connection Identifier
FDDI	Fiber Distributed Data Interface
IBM	International Business Machines Corporation
IEEE	Institute of Electrical and Electronics Engineers
ITSO	International Technical Support Organization
LES	LAN Emulation Server
LLC	Logical Link Control
LNM	LAN Network Manager
MAC	Medium Access Control
PPP	Point-to-Point Protocol
RTIC	Real Time Interface Co-processor
SAP	Service Access Point
SAC	Single Attachment Concentrator
SAS	Single Attachment Station
SNAP	Sub-Network Access Protocol
SR	Source-Routing (Bridge)
SRT	Source-Routing Transparent (Bridge)
SR-TB	Source-Routing to Transparent Bridge

Index

Numerics

A

B

G

H

I

L

S

T

U

V

W

X

Printed in U.S.A.

SG24-5000-00