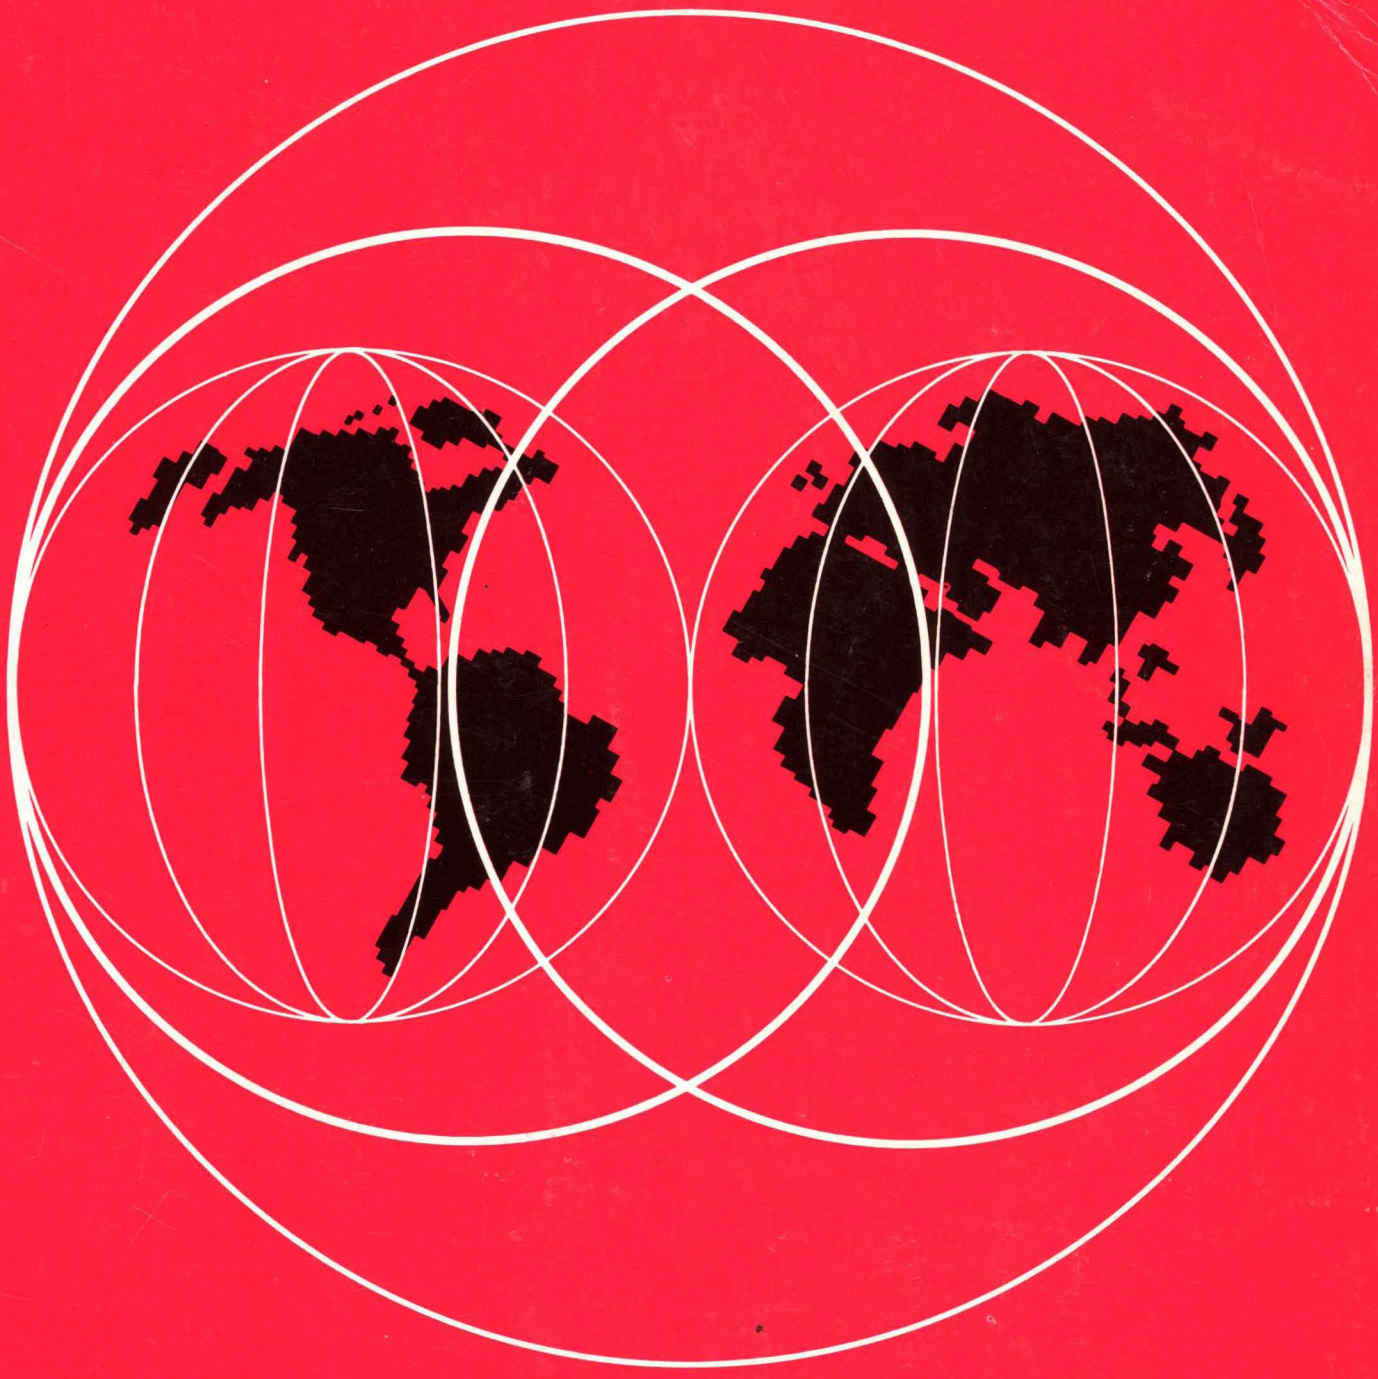

IBM

International Technical Support Centers

IBM MULTISEGMENT LAN DESIGN GUIDELINES

GG24-3398-00

Local Area Network Library
International Technical Support Center
IBM Multisegment LAN Design Guidelines

Document Number GG24-3398

June 1989

International Technical Support Center
Raleigh, North Carolina

First Edition (June, 1989)

This document applies to Versions 2.0 and 2.1 of the Token-Ring Bridge Program (Product Numbers 16F0492 and 16F0493) and Version 1.0 of the PC Network Bridge Program (Product Number 96X5860) for use with the PC/DOS control program.

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 program product in this document is not intended to state or imply that only IBM's program product may be used. Any functionally equivalent program may be used instead.

The information contained in this document has not been submitted to any formal IBM test and is distributed on an 'As Is' basis without any warranty either express or implied. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any performance data contained in this document was obtained in a controlled environment based on the use of specific data and is presented only to illustrate techniques and procedures to assist IBM personnel to better understand IBM products. The results that may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data in their specific environment.

No performance data may be abstracted or reproduced and given to non-IBM personnel without prior written approval by Business Practices.

Publications are not stocked at the address given below. Requests for IBM publications should be made to the IBM branch office serving your locality.

A form for reader's comments is provided at the back of this publication. If the form has been removed, comments may be addressed to:

IBM Corporation, International Technical Support Center
Dept. 985, Building 657
P. O. Box 12195
Research Triangle Park, NC 27709 USA

IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

Abstract

This document is intended to provide assistance for customers and Systems Engineers in the task of planning for a multisegment LAN. It provides design guidelines and management considerations for large LANs with local and/or remote bridges. It primarily contains information about the:

IBM Token-Ring Network Bridge Program, Version 2.0
IBM Token-Ring Network Bridge Program, Version 2.1 and the
IBM PC Network Bridge Program, Version 1.0

General guidance for the logical and physical design as well as network management and traffic control considerations is addressed in a number of LAN configurations.

CSYS PSYS (149 pages)

Acknowledgements

The author of this document is:

Gerard Rousset
IBM France

Project Advisor:

Joan Cavin
International Technical Support Center, Raleigh

This publication is the result of a residency conducted at the International Technical Support Center, Raleigh.

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

Guillermo Diaz	International Technical Support Center, Raleigh
Gerrit Nel	IBM South Africa
Walter Borghi	IBM Australia

Preface

Document Purpose and Scope

The purpose of this document is to provide guidelines in the task of planning for a multisegment LAN, using local and (or) remote bridges. It primarily describes the facilities provided by:

IBM Token-Ring Network Bridge Program, Version 2.0
IBM Token-Ring Network Bridge Program, Version 2.1 and the
IBM PC Network Bridge Program, Version 1.0

It also provides assistance in the logical and physical design of a large LAN. Important bridge parameters are discussed in a number of environments to improve performance and accommodate high availability requirements. General guidance for improved network management is also addressed.

Audience

This document is intended for persons requiring a better understanding of the IBM bridges facilities for planning and installation purposes:

- Customers
  - Local Area Network planning staff
  - Local Area Network installation staff
- IBM
  - Account System Engineers
  - Local Area Network Specialist System Engineers

It is assumed that the reader has prior knowledge of and experience with IBM Local Area Network products and implementations. Knowledge of the contents of *LAN Concepts and Products* and *IBM Token-Ring Network Architecture Reference* is recommended.

Document Organization

The document is organized as follows:

- *Chapter 1 Introduction*

  Chapter 1 gives a brief introduction to the LAN design objectives for large multisegment LANs and describes the LAN management philosophy.

- *Chapter 2 Bridge Concept and Benefits*

  Chapter 2 provides an overview of the bridge concept and architecture. It also provides positioning information versus other LAN interconnection techniques such as gateways and routers. The facilities and benefits provided by the bridges are described in a simple scenario.

- *Chapter 3 IBM LAN bridges*

  Chapter 3 provides an overview of the recently announced bridge products and describes their common functions and characteristics

- *Chapter 4 Routing Support*

  Chapter 4 describes the IBM source routing approach and the use of the corresponding bridge parameters in different scenarios for improved performance. It also provides positioning information between source routing and transparent bridging techniques.

- *Chapter 5 IBM Token-Ring Network Bridge Program V2.0 (Local Bridge)*

  Chapter 5 describes the &brdge20. and the new bridge functions and parameters. It also provides installation guidelines for the bridge products and LAN manager bridges definition.

- *Chapter 6 IBM Token-Ring Network Bridge Program V2.1 (Local or Remote Bridge)*

  Chapter 6 describes new facilities provided by the remote bridge function at the enterprise level. It also provides information on the remote bridge components with emphasis on installation, performance and recovery considerations. The filtering facility which is particularly important in a remote bridge configuration is also described in this chapter.

- *Chapter 7 IBM PC Network Bridge Program*

  Chapter 7 describes the IBM PC Network Bridge Program and provides information on connectivity and network expansion considerations for IBM PC Network (Broadband) segments.

- *Chapter 8 LAN Design Methodology*

  Chapter 8 describes LAN design criteria and the main tasks to perform during the LAN design process.

- *Chapter 9 Logical Design Considerations*

  Chapter 9 addresses the LAN logical design process for user rings and backbone rings. High availability, connectivity and performance considerations are discussed in different scenarios, with local and remote bridges.

- *Chapter 10 Backbone and Bridge Physical Design Guidelines*

Chapter 10 provides guidelines for the LAN physical design process. It addresses issues such as bridge station placement and backbone physical topology in typical environments.

- *Chapter 11 LAN Management Considerations*

  Chapter 11 provides guidelines for improved LAN management with emphasis on centralized network management. by taking advantage of the facilities provided by IBM LAN Manager V2.0 and NetView Release 3.

Related Publications

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

- *IBM Token-Ring Network Architecture Reference* (SC30-3374)
- *SNA Technical Overview* (SC30-3073)
- *IBM LAN Administrator's Guide* (GA27-3748)
- *IBM Cabling System Planning and Installation Guide* (GA27-3361)
- *IBM Cabling System Technical Interface Specification* (GA27-3773)
- *IBM Token-Ring Network Installation Guide* (GA27-3678)
- *IBM Token-Ring Network Introduction and Planning Guide* (GA27-3677)
- *IBM Token-Ring Network Telephone Twisted-Pair Media Guide* (GA27-3714)
- *IBM Token-Ring Network Telephone Twisted-Pair Media Guide (EMEA version)* (GA27-3733)
- *IBM Token-Ring Network Optical Fiber Cable Options* (GA27-3747)
- *IBM Token-Ring Network Problem Determination Guide* (SX27-3710)
- *IBM PC Network Baseband Planning Guide* (S68X-2269)
- *IBM PC Network Baseband Adapter Installation Instructions* Packaged with Adapter
- *IBM PC Network Baseband Adapter/A Installation Instructions* Packaged with Adapter
- *IBM PC Network Baseband Extender Technical Reference* (S68X-2266)
- *IBM PC Network Broadband Planning Guide* (S68X-2268)
- *IBM PC Network Hardware Maintenance and Service* (S68X-2240)
- *IBM PC Network Adapter II - Frequency 1 Installation Instructions* Packaged with Adapter
- *IBM PC Network Adapter II - Frequency 2 Installation Instructions* Packaged with Adapter
- *IBM PC Network Adapter II - Frequency 3 Installation Instructions* Packaged with Adapter
- *IBM PC Network Adapter II/A - Frequency 1 Installation Instructions* Packaged with Adapter
- *IBM PC Network Adapter II/A - Frequency 2 Installation Instructions* Packaged with Adapter
- *IBM PC Network Adapter II/A - Frequency 3 Installation Instructions* Packaged with Adapter
- *IBM PC Network Adapters Technical Reference* (GA27-xxxx)
- *IBM PC Network Adapter II - Frequency 2 Technical Reference* (SC30-3490)
- *IBM PC Network Adapter II - Frequency 3 Technical Reference* (SC30-3491)

- *IBM PC Network Adapter II/A - Frequency 2 Technical Reference* (SC30-3492)
- *IBM PC Network Adapter II/A - Frequency 3 Technical Reference* (SC30-3493)
- *IBM Token-Ring Network Starter Kit Guide, Version 1.1* (SK2T-0303)
- *IBM RT PC Technical Reference Token-Ring Adapter* (SK2T-0291)
- *IBM Token-Ring Network Adapter/A Installation and Testing Instructions* (GA27-3784)
  Packaged with Adapter
- *IBM Token-Ring Network PC Adapter Guide to Operations*
  Packaged with Adapter
- *IBM Token-Ring Network 16/4 Adapter Guide to Operations*
  Packaged with Adapter
- *IBM Token-Ring Network 16/4 Adapter/A Installation and Testing Instructions*
  Packaged with Adapter
- *IBM 8220 Customer Setup Instructions* (GA27-3817)
  Packaged with 8220 Optical Fiber Converter
- *IBM Token-Ring Network Trace and Performance Program User's Guide*
  Packaged with Program Product
- *IBM Token-Ring Network Remote Program Load User's Guide* (GA27-3763)
  Packaged with the Remote Program Load feature, not orderable
- *IBM Local Area Network Host Information* (GC30-3479)
- *IBM Local Area Network Technical Reference* (SC30-3383)
- *IBM Local Area Network Support Program User's Guide, Version 1.1*
  Packaged with Program Product
- *IBM Token-Ring Network NETBIOS Program User's Guide*
  Packaged with Program Product
- *IBM Token-Ring Network PC Adapter Support Program for the 3270-PC*
  Packaged with Program Product
- *IBM PC Network Bridge Program User's Guide, Version 1.1* (SC30-3402-1)
  Packaged with Program Product, not orderable
- *IBM PC Network Bridge Program User's Guide, Version 2.0* (SC30-3402-2)
  Packaged with Program Product, not orderable
- *IBM Token-Ring Network Manager User's Guide, Version 1.1* (LY30-5595-1)
  Packaged with Program Product, not orderable
- *IBM LAN Manager User's Guide, Version 1.0* (LY30-5595-2)
  Packaged with Program Product, not orderable
- *IBM LAN Manager User's Guide, Version 2.0* (LY30-5595-3)
  Packaged with Program Product. not orderable
- *IBM LAN Manager Entry User's Guide, Version 1.0*
  Packaged with Program Product
- *IBM PC 3270 Emulation LAN Management Program User's Guide, Version 1.0* (SC30-3456)
  Packaged with Program Product
- *IBM PC 3270 Emulation Memory Management Enhancement*

Packaged with Product (P/N 8575345)

- *IBM Remote NETBIOS Access Facility Program User's Guide, Version 1.0* (SK2T-0314)
- *IBM Remote NETBIOS Access Facility Program Installation and Configuration Guide, Version 2.1* (SK2T-0323)
- *IBM Asynchronous Communications Server Program User's Guide, Version 1.0* (SC30-3464)
- *IBM LAN Asynchronous Connection Server Program Installation and Configuration Guide, Version 1.0* (SC30-3509)
  Packaged with Program Product, not orderable
- *IBM 8232 LAN Channel Station Installation and Testing* (GA27-3796)
- *IBM 8232 LAN Channel Station Operator's Guide* (GA27-3785)
- *IBM 8232 LAN Channel Station Safety Notices* (GA27-3833)
- *IBM LAN Channel Support Program User's Guide* (SC30-3458)
- *IBM Token-Ring Network Manager and NetView/PC Planning and Installation* (GG24-3128)
- *IBM Token-Ring Network Bridges and Management* (GG24-3062)
- *NetView/PC Primer* (GG24-3115)
- *IBM 3725 Network Control Program Token-Ring Interface Planning and Implementation* (GG24-3110)
- *An Introduction To Programming For APPC/PC* (GG24-3034)
- *IBM Enhanced Connectivity Facilities SRPI Guide* (GG24-3086)
- *Connectivity In 1986* (GG24-3079)
- *IBM Local Area Networks In An Office Systems Environment* (GG24-3071)
- *IEEE 802.3 Local Area Network Considerations* (GG22-9422)

ITSC Publication Structure - LANs

The rapid evolution of Local Area Network products has resulted in the availability of a wide variety of documents, including the reference materials available with each product and additional technical planning and support material available from various development and support groups.

To assist users to locate appropriate, up-to-date information the International Technical Support Center is structuring its local area network documentation into a library of publications.

Each publication is produced to address some technical requirements of a specific audience as described in the abstract and preface of the document. Because the ITSC publications are intended to complement, but not replace reference material available with the products themselves, each document also provides a bibliography of related publications.

The International Technical Support Center publications related to local area networks have been planned with the following structure in mind to simplify the problem of locating up-to-date information.

1. Overview manuals which provide tutorial information and cross product conceptual and planning information.
2. Installation manuals which complement product reference material by describing the experiences of the ITSC in installing particular products within a total system. These documents do not address all installation parameters or options as do the product reference materials, but are intended to highlight those aspects of installation which have the greatest impact on successful use of the product, including the relationship between a specific product and other network or system products.
3. Network design and management manuals which describe trade offs and considerations for managing or planning local area networks.

The current library contains the following publications :

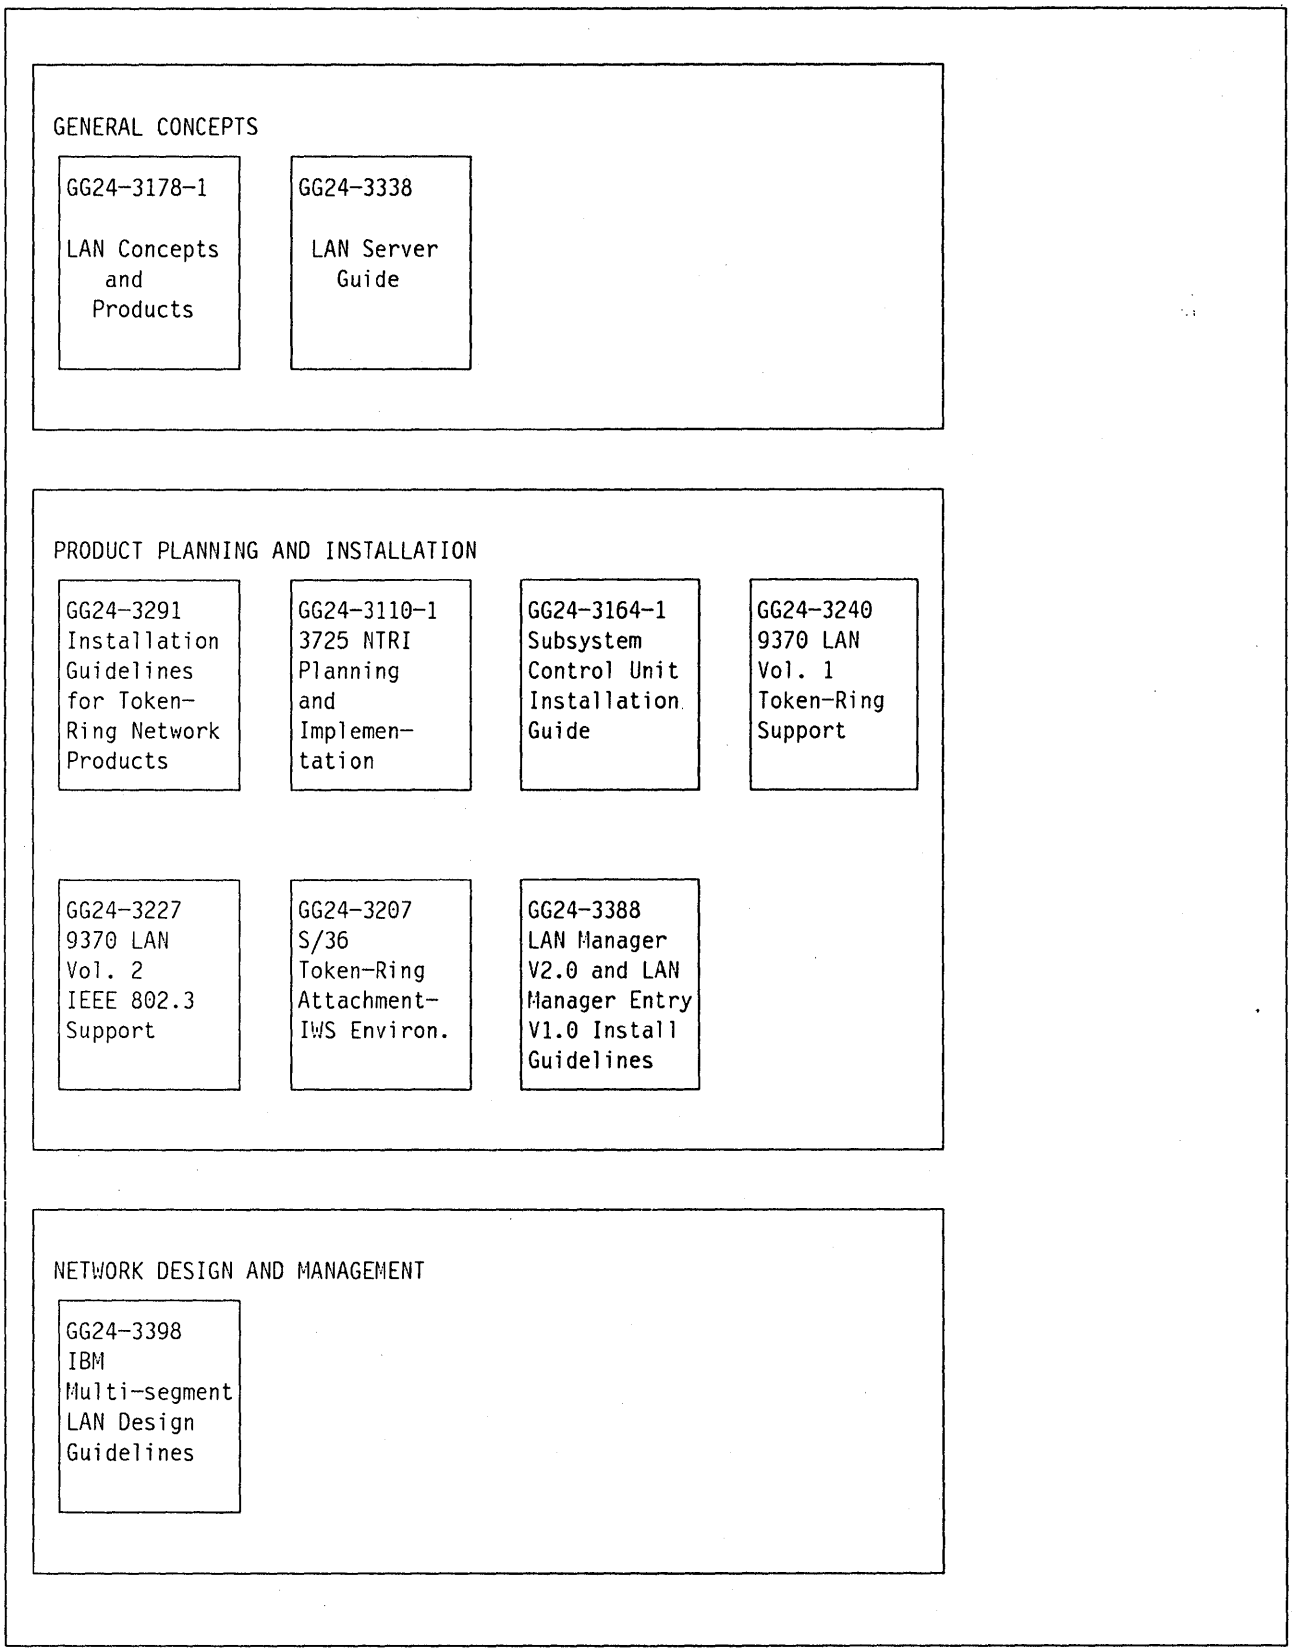

Figure 1. Existing Bulletins in ITSC LAN Bulletin Library.

Contents

Figures

Tables

1. Introduction

In many cases, local area network requirements exceed the capabilities of a single ring or bus and thus need to connect to other segments by means of bridges or routers. This document describes design considerations for supporting interconnection of IBM Token-Ring Network rings and/or IBM PC Network (Broadband) segments using recently announced bridge products.

The new IBM Token-Ring Network Bridge Program V2.0, IBM Token-Ring Network Bridge Program V2.1 and IBM PC Network Bridge Program capabilities provide increased flexibility for users to implement local area network solutions that address rapidly evolving requirements for both host and peer connectivity. Together with the IBM LAN Manager V2.0 these products enable design of a multisegment LAN to address such current or anticipated requirements as:

- Growth in numbers of users or application traffic to provide acceptable availability and application response time
- Need to extend local area network capability beyond the cabling guidelines for a single segment
- Need to provide redundancy for high-availability or capacity
- Need to support token-ring network segments at both 16 and 4 Mbps
- Need to intermix token-ring network segments with PC/Network (Broadband) segments
- Need to interconnect PC/Network (Broadband) segments using different channel frequency ranges
- Need to provide connectivity for autonomous user LAN segments.
- Need to extend LANs into the wide area network as an alternative for some of today's connectivity.

The resulting LAN may involve only two or three segments within a single building, or may involve a complex structure of LANs over one or more buildings in a campus. With the new IBM Token-Ring Network Bridge Program V2.1 it is even possible to design LANs interconnected through dedicated leased communications links.

A wide variety of design alternatives are possible to address each of the above requirements and consequently the design process may appear overly complex. This complexity may be reduced by an understanding of topology and product capabilities, and by an awareness of design trade-offs and recommendations for representative scenarios.

While it is important to simplify the design and implementation process for multisegment LANs it is equally important to provide appropriate management capabilities to ensure that the complexity of the topology or interconnection neither reduces management control of larger LANs nor increases end-user requirements to perform the control function.

When a LAN is part of a larger system complex it is highly desirable to be able to support all aspects of the system with a consistent set of tools that can be used by a smaller group of trained personnel.

IBM communication and system management standards address both these goals through defined alert and command capabilities. For the LAN environment, these capabilities are implemented in the IBM LAN Manager V2.0. IBM LAN Manager V2.0 provides local support for a multisegment LAN from a single workstation (which can also be used for other concurrent tasks). IBM LAN Manager V2.0 also interacts with NetView R3 to enable a host NetView [1] operator to manage the LAN environment as he would other system resources, from a single NetView display.

This connection between the LAN Manager and NetView Release 3 also enables an installation to use enhanced NetView capabilities to automate some LAN operator procedures in the same way as other system or network management processes. For large systems or LANs this may represent a substantial improvement in operations productivity and improved availability of LAN application connectivity. Considerations for using IBM LAN Manager V2.0 are described in "LAN Management Considerations" on page 119, and in LAN Manager V2.0 and LAN Manager Entry V1.0 Installation Guidelines.

[1] NetView is a trademark of the International Business Machines Corporation

2. Bridge Concept and Benefits

2.1 Multisegment LAN and Bridge Concept

Large LANs are usually composed of several rings or buses, called LAN segments. A LAN segment consists of a single common medium and all LAN devices connected to this medium [2]. Due to such physical constraints as signal attenuation, propagation delays or noise susceptibility, LAN segments are limited in the number of stations that can attach to it. For example, a token-ring segment is limited to 260 stations. If you have a large number of stations in your LAN, similar or different type LAN segments can sometimes be combined to form a single large logical LAN.

An example of a LAN composed of four segments is shown in Figure 2. In this configuration, there are three token-ring segments and one IBM PC Network (Broadband) segment interconnected via three bridges.

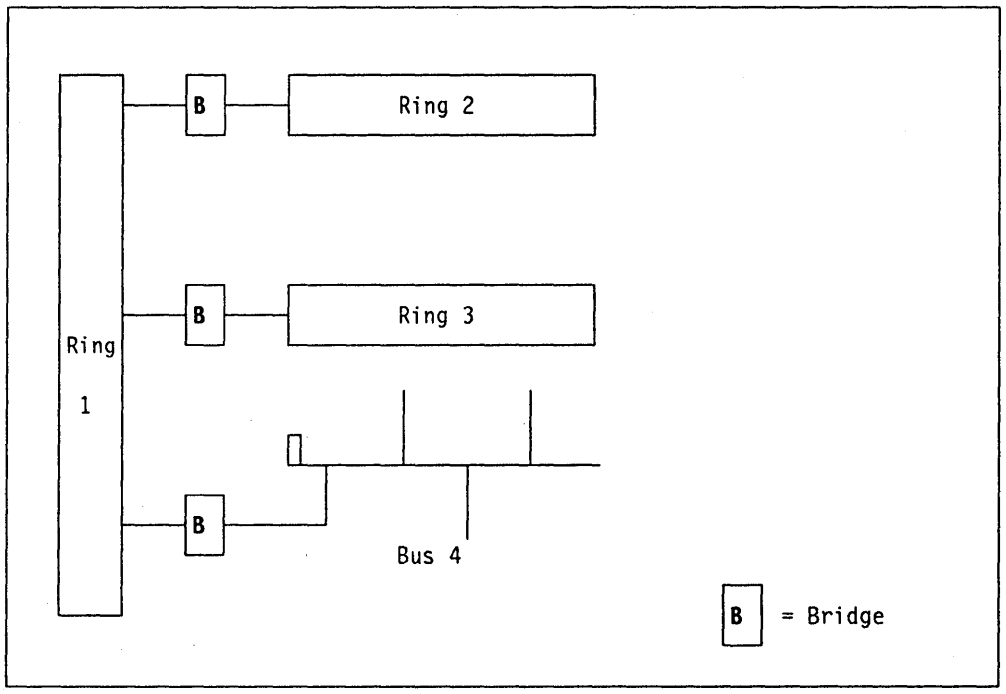

Figure 2. Multisegment LAN

LAN segment interconnection can be achieved through bridges, routers or gateways, as discussed in "LAN Interconnection Techniques" on page 7. Bridges are frequently preferred as they provide protocol independence and potential for optimal performance by interconnecting different LAN segments at the lowest possible level, that is the Medium Access Control (MAC) sublayer, as shown in Figure 7 on page 10.

[2] A common medium is one which has a consistent set of physical attachment specifications supporting operation as a single LAN. In many cases, cable types may not be intermixed. Refer to *IBM Cabling System Planning and Installation Guide*.

2.2 Why do we Need Bridges in a LAN ?

LAN segment interconnection may be considered for various reasons. A prime reason of course is to expand the networking capability by:

- Increasing the number of attached devices or wiring closets above that supported on a single LAN segment.
- Increasing the geographic area covered by the total LAN.
- Providing connectivity between stations attached to different LAN segments so that segment differences (due to use of different MAC protocols, speeds or frequencies) are transparent to higher layer protocols.
- Increasing the bandwidth available to stations on a single segment by splitting a LAN segment into one or more bridged LAN segments. In this way fewer stations have to share the same common medium and access protocol mechanisms.
- Protecting users from the impact of wiring changes or media errors since bridges isolate errors or disruptions on one segment from the other segments they are connected to.
- Permitting smaller ring entities while maintaining full LAN management functions, as the IBM bridges will send notifications and reports to IBM LAN Manager V2.0.

Even when LAN stations use the same MAC protocol, there may be other physical constraints such as wiring or speed preventing a station from joining a particular LAN segment.

- For example, in the case of the IBM Token-Ring Network, bridge interconnection may be required when some stations are attached via unshielded telephone twisted pair wire while all others are attached to data grade media twisted pair wire. Shielded and unshielded wiring can be mixed in a single physical ring only through use of Data Grade Media-to-Type 3 Filters which subject the entire ring to the limitations of telephone twisted pair wiring.

  Similarly, when one station is operating at 4Mbps and another at 16Mbps, they must insert into separate rings.
- Two IBM PC Network (Broadband) segments using different channel frequency pairs to transmit and receive information may use a bridge to provide connectivity between devices attached to the individual segments.

2.3 Single Segment LAN or Multisegment LAN ?

Except in some of the cases mentioned above where you must have bridges in your configuration, a LAN could be designed with very large segments in order to limit the number of bridges and hence the cost of the configuration.

As a matter of fact, there are many other criteria to consider, as explained in "Logical Design Considerations" on page 93.

In the following simple example, we will try to illustrate some of the bridge benefits. Suppose you have to connect 200 stations using token-ring protocols. You can decide to have only one ring with 200 stations, or you can decide to

have several smaller rings, for example two rings (with 100 stations on each) interconnected via a bridge.

- **Single segment approach**

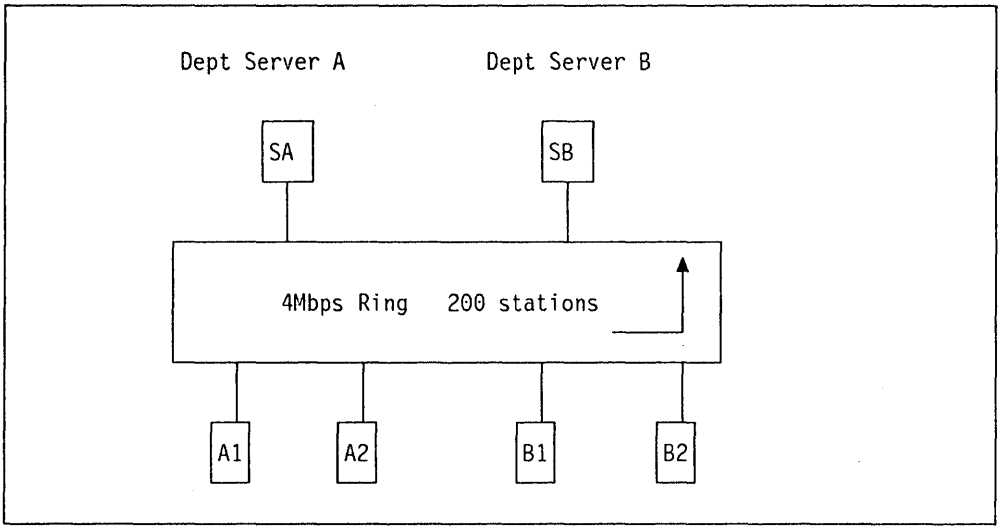

Figure 3. Single Segment LAN Approach

In the first approach, where there is a single large segment, the "apparent" benefits of this configuration are:

- Simplicity of the topology
- Better performance (no bridge delay)
- Lower cost (no bridge).

These apparent benefits could be misleading as the reverse propositions can be made as well. For example, if the stations are spread throughout the building, you may have to install repeaters in order to accommodate physical design guidelines. This will add to the cabling complexity and to the LAN cost.

In addition, despite the usually high bandwidth of a LAN, if the traffic between the stations is very heavy, the load of the single segment (ring or bus) could be very high and could result in performance degradation. This will be particularly true if the protocol used on the segment is a CSMA/CD type protocol.

- **Multisegment Approach**

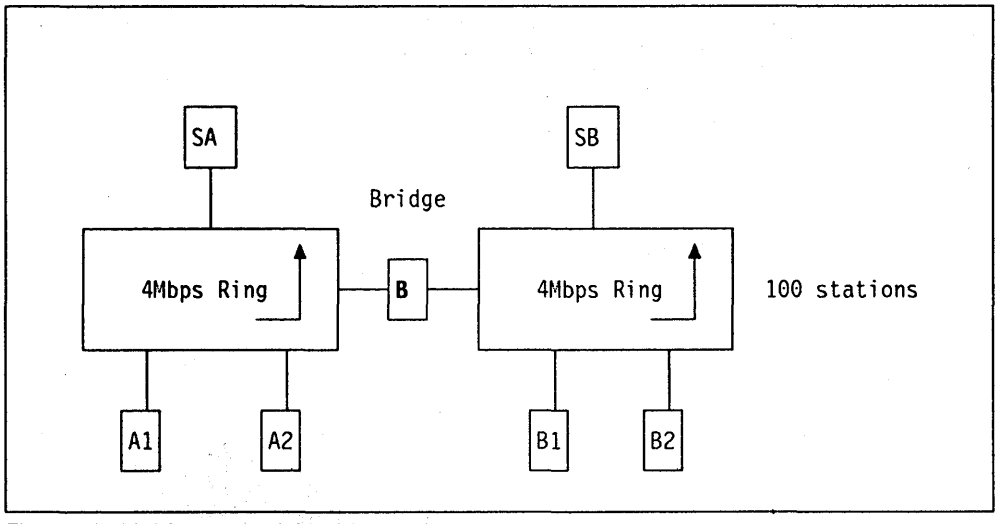

Figure 4. Multisegment LAN Approach

In this second approach, where there are two rings interconnected via a bridge, the following benefits are provided by the bridge:

- Better availability

  One of the main benefits of the bridge is that it will isolate segments A and B from an error point of view.

  For example, if segment A experiences a hard error and enters a beaconing state, segment B users can still communicate with each other. The resulting availability can be translated in terms of obvious cost savings.

- Better performance

  Although the bridge introduces a short delay for inter-ring communication, this configuration can result in better performance than the previous one. As a matter of fact, the total bandwidth of such a configuration is about twice 4Mbps (minus the bridge traffic), as there will be two tokens circulating independently on each ring. If the bridge traffic is low, which may be the case if the two segments correspond to two different departments, the load of each ring will be half of the previous configuration and the response time may be better, especially for future applications requiring high throughput. Note, however, that bandwidth is not the most significant factor in predicting the performance of a local area network.

- Ease of network management

  The use of smaller rings is usually better for network management and applications management if they are based on affinity considerations (an affinity group is a group of users who perform related tasks on the network and have little information interchange with other end-users).

- Future growth

  Using smaller rings increases the cable length available for inter-wiring closet runs or lobe lengths because of the reduced attenuation of the signal.

2.4 LAN Interconnection Techniques

Techniques for interconnection may be classified into three categories as shown in Figure 5 on page 7. The numbered boxes refer to the seven layers of the OSI Reference Model.

Potential optimal performance may be achieved by interconnecting at the lowest possible level, starting from levels 1 and 2, and supporting the same protocol for the network segments to be interconnected.

At a particular interconnection either two or more segments may be interconnected.

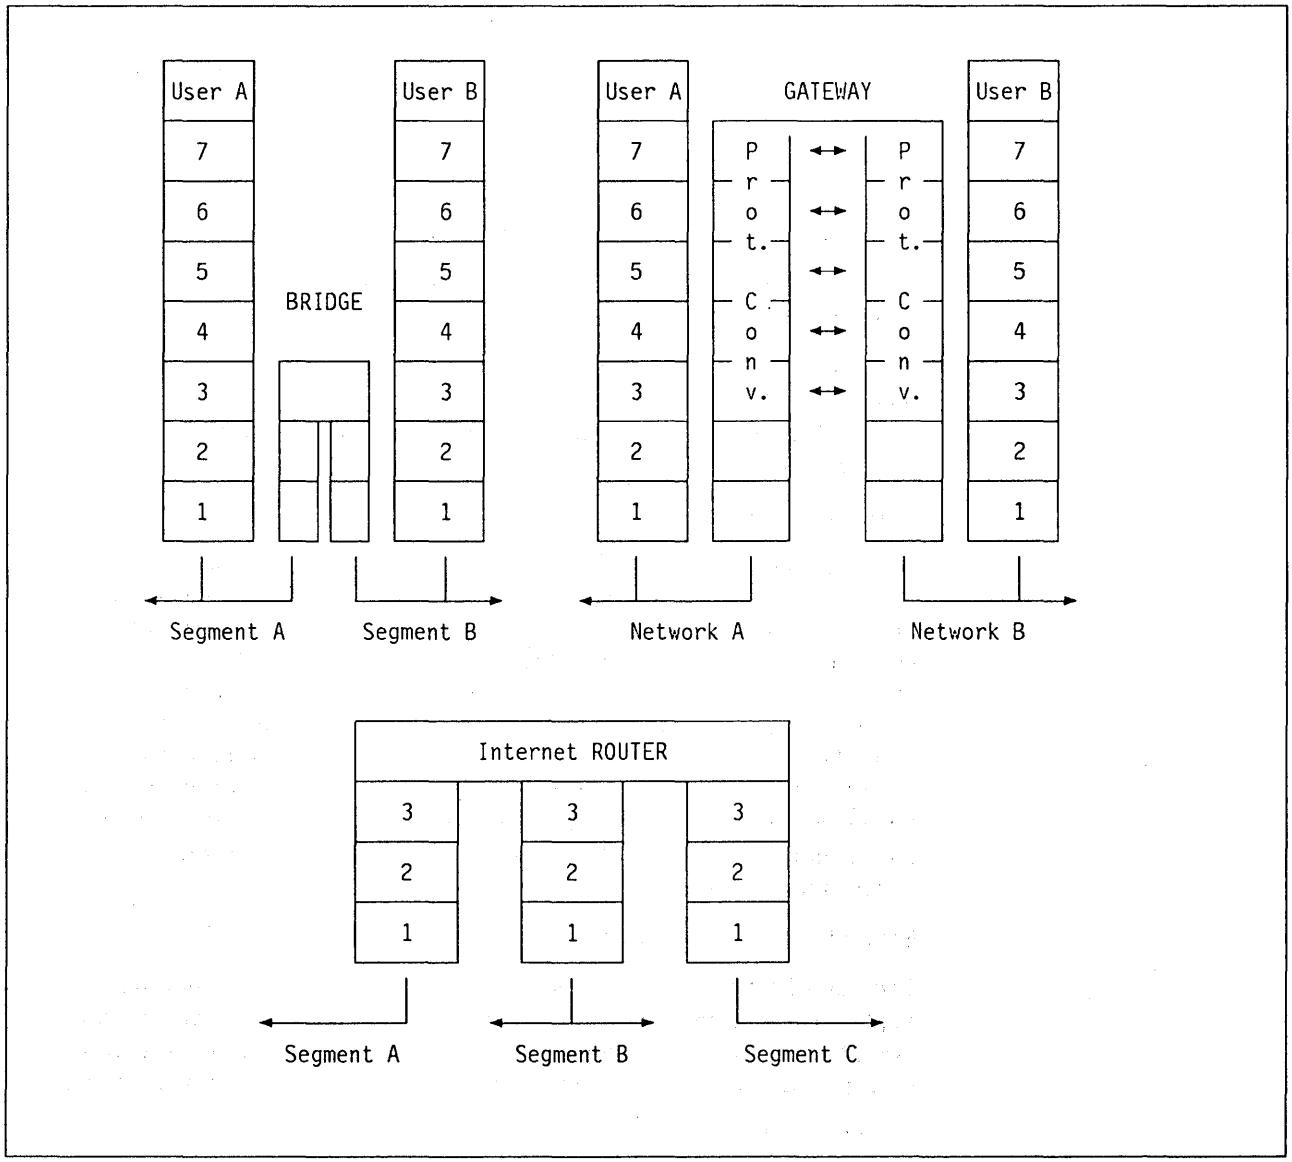

Figure 5. Network Interconnection Techniques (Based upon OSI Reference Model Layers 1 - 7)

2.4.1 MAC Bridge versus Router or Gateway approach

- IBM bridges link segments together at the medium access control (MAC) level.

  The LAN segments bridged together must share the same logical link control sublayer. In general, all stations in a bridged local area network must use the same length MAC address (either all stations use 16-bit addresses or all stations use 48-bit addresses [3]). It is desirable for each specific MAC address to be unique throughout the bridged LAN to permit unambiguous station addressing.

  Bridging may provide dynamic routing, transparent to the higher level protocols.

  Bridged local area networks support communication between stations attached to separate LAN segments as if they were attached to a single LAN. The LAN segments may have dissimilar medium access protocols. Above the DLC layer, bridged LAN segments appear as one single logical LAN.

  Because of the low level at which interconnection is established, bridging is likely to be an optimal solution for connecting LAN segments.

- A router offers segment interconnection at the network layer. The routing service provided by this approach however is not transparent and must be explicitly called if access to remote network segments across the router is required.

  Various router products interconnect more than two network segments. These segments must all share a common network layer protocol, such as Internet [4].

- Gateway interconnection uses quite a different approach. A gateway is a system which supports more than one architecture or network addressing scheme to permit connectivity and interoperability between the devices in the attached environments.

  Gateway products support mapping of addresses from one network to another, and may also provide transformation of data between the environments to support end-to-end application connectivity. Gateways usually link subnetworks at higher layers than bridges or routers. Gateways may operate at layers ranging from layer 3 to layer 7. Depending upon the amount of processing involved, and the number of environments supported by a gateway, a gateway will usually provide less throughput than an equivalently priced bridge or router.

---

[3] All IBM products use 48-bit addresses.

[4] Internet (IP) is a non-proprietary protocol which operates at the level of the OSI Network Layer and is part of TCP/IP. TCP/IP is a set of protocols and specifications of the US Department of Defense that originated with the ARPANET and Defense Data Networks (DDN). It has since become a widely used set of protocols for multi-vendor connectivity, and has been defined by the Department of Defense as the basis from which they plan to migrate to OSI network protocols.

2.4.2 IEEE 802 and MAC Layers

In February 1980, The Institute of Electrical and Electronic Engineers' (IEEE) Computer Society established "Project 802" to draft standards for local area networks. In keeping with the OSI approach, IEEE Project 802 created a reference model with two layers (which correspond to the data link and physical layers of the OSI model). In the IEEE model, however, the data link layer is further divided into two sublayers: the *logical link control* (LLC) sublayer, and the *medium access control* (MAC) sublayer.

The IEEE 802 standards have also been proposed (and in most cases accepted) as OSI standards. The commonly used names for IEEE standards are derived from the project's initial designation "802". Hence, we have:

- IEEE 802.1 - Higher level interface standard, ISO DIS 8802-1
- IEEE 802.2 - logical link control standard, IS 8802-2
- IEEE 802.3 - CSMA/CD standard, IS 8802-3
- IEEE 802.4 - token passing bus standard, IS 8802-4
- IEEE 802.5 - token passing ring standard, IS 8802-5.

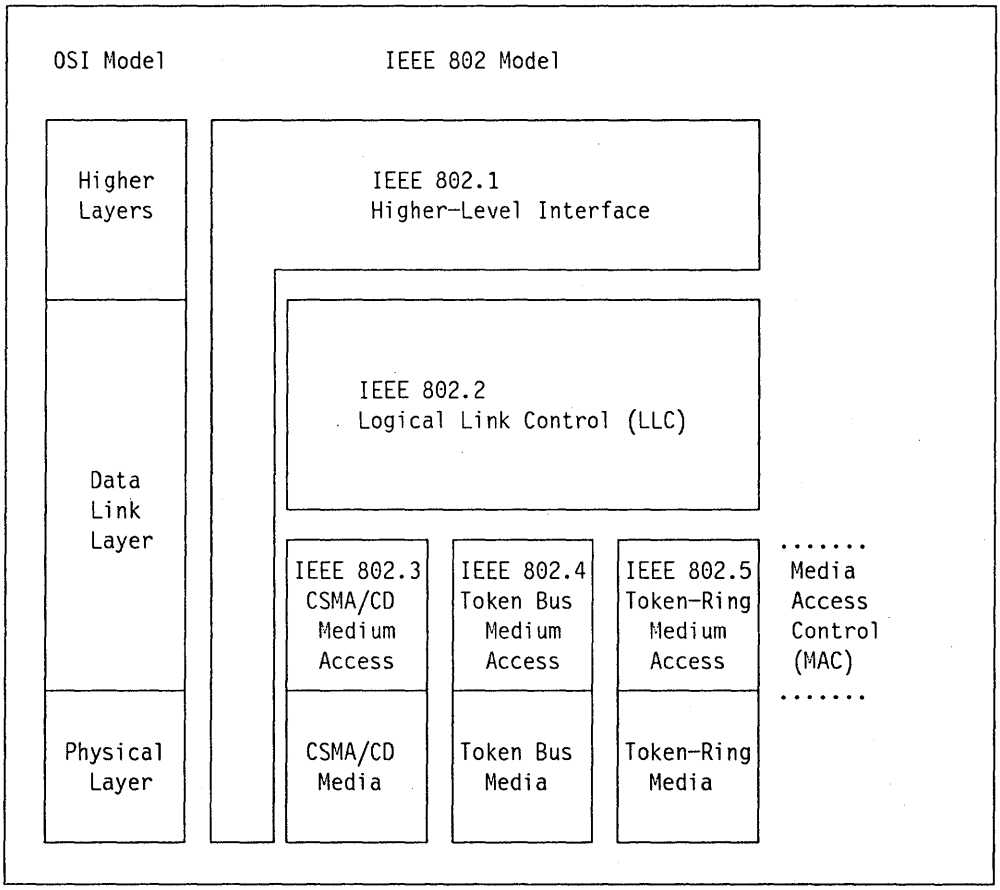

Figure 6. IEEE 802 Model Compared to the OSI Model

2.4.3 MAC Layer Bridges

MAC layer bridges, as the name implies, relay messages from one LAN segment to another at the MAC sublayer level (see Figure 7). MAC layer bridges consist of two (or more) physical and MAC layers (one for each LAN segment they interconnect). The MAC layer functions contained in the bridge are interconnected by a relay function (see Figure 7) which passes frames received from one MAC to the other MAC if certain forwarding conditions are satisfied. The relay function not only passes messages between the MACs but also does the protocol conversion necessary between the different protocols.

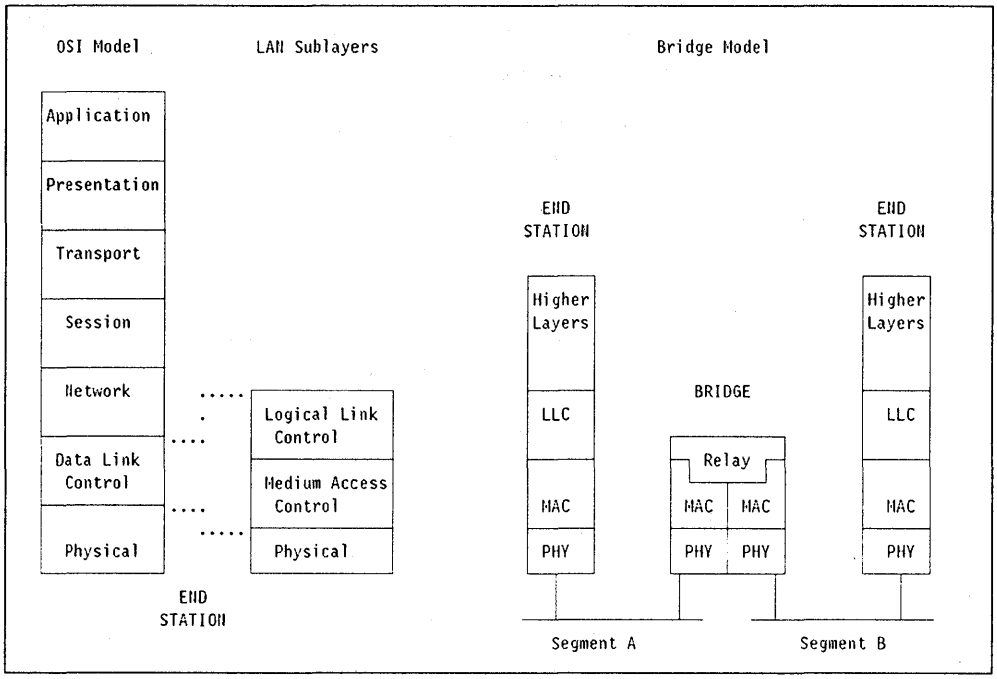

Figure 7. Model for a Bridge. Relationship between the OSI seven layer model, LAN sublayers and a model for a bridge

MAC layer bridges are transparent to the users of the DLC layer (see Figure 7). Note that LLC is an end-to-end communication protocol which is common to all MACs and provides frames numbering, error detection and basic recovery as well as flow control services. Due to this transparency to the higher layers, MAC layer bridges may serve multiple higher layer protocols (for example SNA, NETBIOS, TCP/IP)

2.5 Bridge Topologies

This section introduces some basic different topologies for interconnected LAN segments.

Selecting a logical and physical topology suitable for large multisegment LANs is discussed in detail in "Logical Design Considerations" on page 93 and "Backbone and Bridge Physical Design Guidelines" on page 111.

The main topologies involving bridges are the serial, loop, parallel and backbone topologies.

- **Serial Topology**

  This topology may be used in a smaller multisegment network. It is simple, but is usually limited to three segments because a characteristic of this configuration is that a bridge failure or segment failure will affect the overall LAN connectivity. An example of such a topology is shown in Figure 8.

  In this example, three work groups or departments (A B and C) share similar files or printers, but are attached to separate rings for other reasons. Departments A and C have little inter-ring communication. Ring B can be used as a server ring to reduce the costs of providing such resources as letter quality printers or large disk storage.

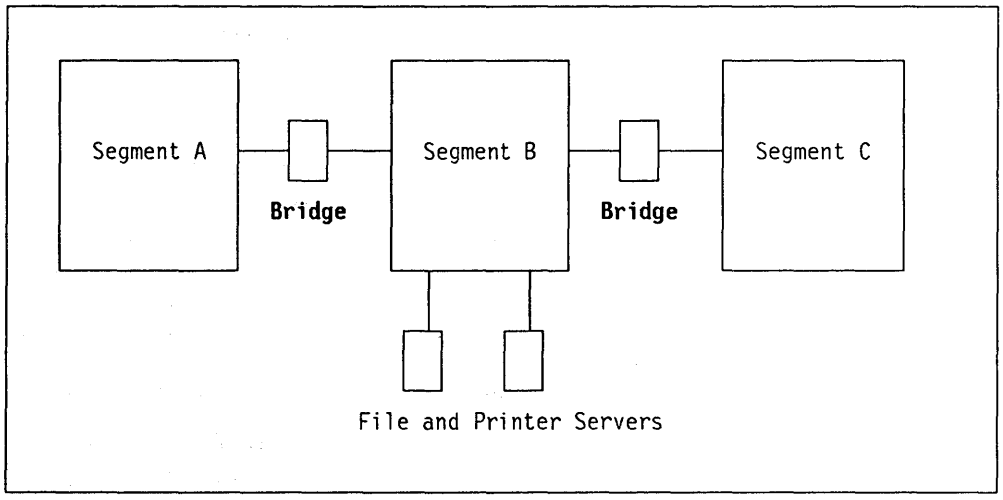

Figure 8. Serial LAN Configuration. Segment A and segment C are two separate departments with little inter-communication. They share servers attached to segment B.

- **Fully-Interconnected and Loop Configurations**

  A fully-interconnected configuration (mesh) provides alternate paths from each LAN segment to another. Should a bridge or path fail, traffic can be routed through an alternate path, thereby increasing availability of the server segment.

  This solution tends to become impractical as the number of LAN segments grows. For *n* LAN segments, one would require *n(n-1)/2* bridges. Therefore,

a loop configuration as shown in Figure 9 on page 12 can be an acceptable compromise between availability versus cost and complexity [5].

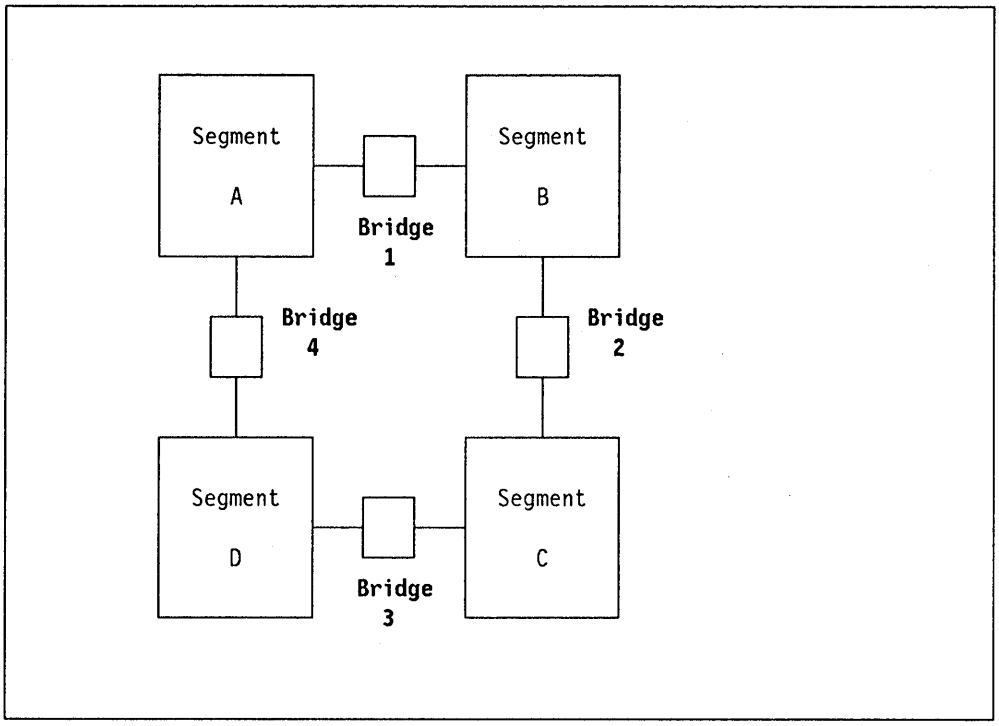

Figure 9. Loop Configuration with Four LAN Segments

- **Parallel Bridge Configuration**

  Parallel bridges can address the problems of heavy traffic flows through particular bridges and/or high availability requirements. While the failure of one bridge would impact connectivity between LAN segments over that bridge, sessions could be recovered via the parallel bridge.

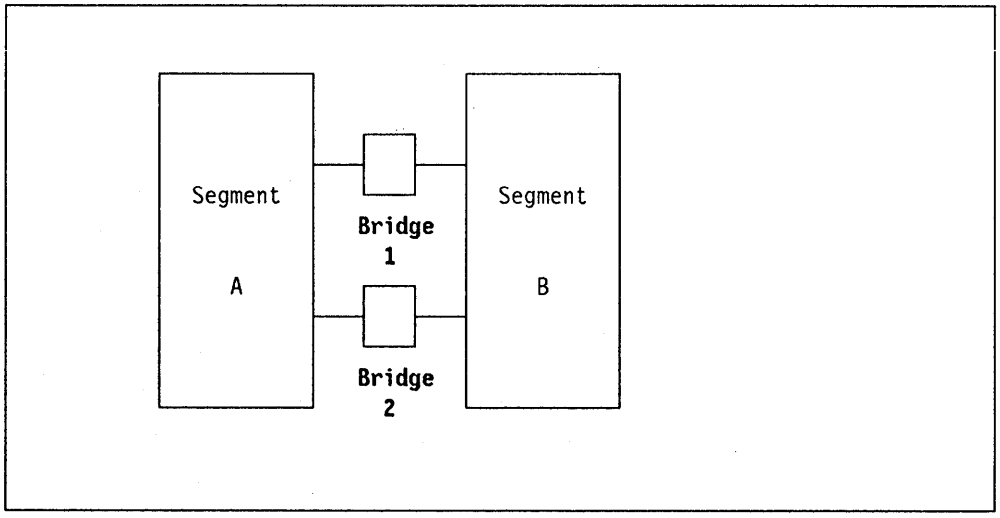

Figure 10. Parallel Configuration

[5] A loop configuration with three LAN segments is also fully-interconnected.

- **Backbone LAN Configuration**

  Backbone configurations are usually recommended for large LANs, as they offer common server and gateway access to a potentially very large number of LAN stations.

  If growth is an important factor, a backbone LAN configuration will provide the necessary flexibility.

  In a backbone configuration, a number of LAN segments, sometimes referred to as departmental LANs (or user rings in the case of a token-ring segments) are all connected to the same backbone LAN segment as shown in Figure 11. This implies that between any two LAN stations attached to departmental LANs, there is always a communications path that includes relatively few bridges, whatever the size of the bridged LAN.

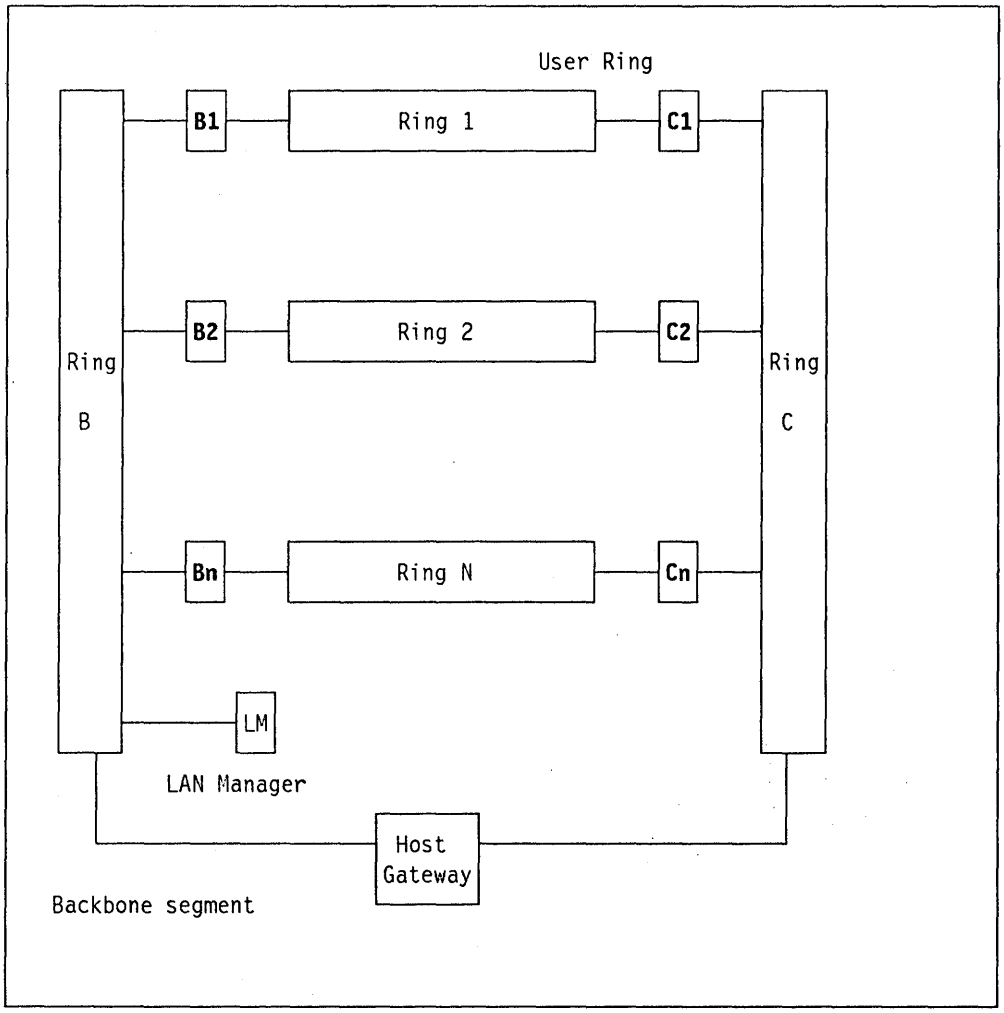

Figure 11. Backbone LAN Configuration

Because of the potentially high concentration of traffic on the common backbone segment, a LAN segment with stable performance characteristics and possible higher bandwidth than the individual segments may be required.

As an example, the backbone segment B in Figure 11 could be a 16 Mbps IBM Token-Ring Network, interconnecting a mixture of IBM PC Network (Broadband) segments and 4 Mbps IBM Token-Ring Networks.

Because of high availability, backup and/or capacity considerations, one might consider implementing a duplex backbone, as a mirror image of the first one.

2.6 Design Considerations

Many factors should be considered when configuring a LAN consisting of a number of interconnected LAN segments, each with its own (perhaps different) MAC protocol.

Bridged LAN topology can be influenced by any of the following:

- The number of workstations physically supported by a particular type of segment or cabling.
- Requirements to balance or alleviate heavy traffic associated with particular applications over one or more interconnected LAN segments.
- Requirements for a geographical approach to interconnecting LAN segments by associating segments with specific areas within a building or campus layout.
- Desire to concentrate communications by connecting users with related information needs within LAN segments, known as affinity groups.
- Requirements for performance, reliability and/or availability which may be addressed through use of parallel bridges [6] or parallel routes [7], providing both increased capacity and backup paths.
- Requirements to separate some LAN stations from others for security purposes by using bridges to provide controllable connectivity paths between secure segments and other stations.
- Requirements to access special function devices such as host gateways or LAN servers which may be best satisfied through use of a backbone topology.

It is evident that there is no "best solution" for every network. However, general guidelines can be given based on the strengths and weaknesses of general topologies, and possible performance expectations.

When the number of workstations is small (typically a network of fewer than 50 workstations) the distribution of workstations within the network will be determined mainly by physical considerations. The existence of departmental groups or affinity groups will also be an important factor in selecting a particular physical layout, as will the number and location of servers.

However, when the number of workstations grows, other factors (such as increased traffic load versus capacity, availability, performance, management and network expandability), increase in importance.

---

[6] Parallel bridges are bridges interconnecting the same two LAN segments.

[7] Parallel routes define end-to-end paths for the same source and destination LAN segments.

Logical and physical design considerations are described in more detail in "Logical Design Considerations" on page 93 and "Backbone and Bridge Physical Design Guidelines" on page 111.

3. IBM LAN Bridges

The most recent IBM Token-Ring Network Bridge Programs are the IBM Token-Ring Network Bridge Program V2.0, IBM Token-Ring Network Bridge Program V2.1 and IBM PC Network Bridge Program. All of them will be available in 1989.

Note that the "local" bridge function provided by IBM Token-Ring Network Bridge Program V2.0 has been replaced by IBM Token-Ring Network Bridge Program V2.1, which provides either "local" or "remote" bridge functions. In addition, the capability of filtering frames to be forwarded by the bridge, which was originally announced and which is particularly important for a remote bridge configuration, has been extended to "local" bridges as well (IBM Token-Ring Network Bridge Program V2.1 only)

The structure and the main functions of these three products are very similar, although they support different environments. Common functions like the bridge adapter interface, LAN Manager support and bridge characteristics are discussed in this chapter.

The differences related to the different environments will be explained in the corresponding sections about IBM Token-Ring Network Bridge Program V2.1 and IBM PC Network Bridge Program.

The bridge installation steps are described in "Installation/Utilization Guidelines" on page 51. In the rest of this chapter, "any bridge program" means one of the three recent products just mentioned above.

3.1 IBM Bridge Programs Overview

IBM local area network bridges are implemented using dedicated workstations attached to two different LAN segments. A bridge workstation contains two LAN adapters (each appropriate to the type of attached LAN segment) and a bridge program running in the workstation's memory.

Communication across any IBM bridge is transparent to applications written to the IEEE 802.2 logical link control interface or higher using source routing.

Since a bridge workstation is entirely dedicated to the bridging function, the different IBM bridge programs operate in a single-task IBM PC/DOS 3.3 or 4.0 environment.

Any IBM bridge program supports the same routing information field as described in "Source Routing Approach" on page 27. Since the routing information field can contain only eight route designators, end stations in a multisegment environment can be separated by a maximum of seven bridges in order to communicate, independent of the bridge program used.

When all bridges in the network are running either IBM Token-Ring Network Bridge Program V2.0, IBM Token-Ring Network Bridge Program V2.1 or IBM PC Network Bridge Program and configured appropriately, they will communicate with each other to automatically configure the network single-route broadcast

path. Refer to the single-route broadcast route determination algorithm in "Single-Route Broadcast (Manual or Automatic)" on page 36 for an introductory discussion of this route resolving technique.

- **IBM Token-Ring Network Bridge Program V2.0**

  The IBM Token-Ring Network Bridge Program V2.0 interconnects token-ring segments operating at either 4 Mbps or 16 Mbps. IBM Token-Ring Network Bridge Program V2.0 communicates with up to four LAN Managers implemented by IBM LAN Manager V2.0 (one of which may control the bridge operations). See "IBM Token-Ring Network Bridge Program V2.0 (Local Bridge)" on page 47 for more details on IBM Token-Ring Network Bridge Program V2.0.

- **IBM Token-Ring Network Bridge Program V2.1**

  The IBM Token-Ring Network Bridge Program V2.1, also referred to as *Split Bridge*, or *Remote Bridge*, includes all the functions and capabilities of IBM Token-Ring Network Bridge Program V2.0. In addition, IBM Token-Ring Network Bridge Program V2.1 provides bridging functions between remote token-ring segments (4 Mbps or 16 Mbps) over a leased teleprocessing (TP) line operating at speeds from 9.6 Kbps to 1.344 Mbps. In this case, a LAN workstation at each end of the TP link will implement one half of the Split Bridge.

  When operating as a remote bridge, the filtering facility provided by IBM Token-Ring Network Bridge Program V2.1 may be particularly interesting. Via a programming interface, a user program can determine which frames are allowed to pass through the split bridge. This feature is especially important to limit the bridge traffic when a relatively slow TP line is being used between the bridge halves (for example, 9.6 Kbps or 19.2 Kbps). See "IBM Token-Ring Network Bridge Program V2.1 "Local" or "Remote" Bridge Function" on page 53 for more details specific to IBM Token-Ring Network Bridge Program V2.1.

- **IBM PC Network Bridge Program**

  A more general LAN segment interconnection solution is offered by the IBM PC Network Bridge Program. A bridge running the IBM PC Network Bridge Program supports interconnection between any two of the following LAN segments:

  - 4 Mbps IBM Token-Ring Network segment
  - 16 Mbps IBM Token-Ring Network segment
  - PC Network (Broadband) operating on channel pair T13 - J
  - PC Network (Broadband) operating on channel pair 2′ - O (Frequency 2)
  - PC Network (Broadband) operating on channel pair 3′ - P (Frequency 3)

  The PC Network segments, operating on different channel pairs, may or may not share the same broadband medium.

  In this way, PC Network (Broadband) attached devices may have access to host gateway devices attached to token-ring segments of the network. See "PC Network Bridge" on page 81 for more details specific to IBM PC Network Bridge Program.

3.2 Common Bridge Characteristics

Before looking at the specific IBM bridge product implementations in the following chapters, Figure 12 shows the major components and their interfaces for a "normal" bridge between two LAN segments.

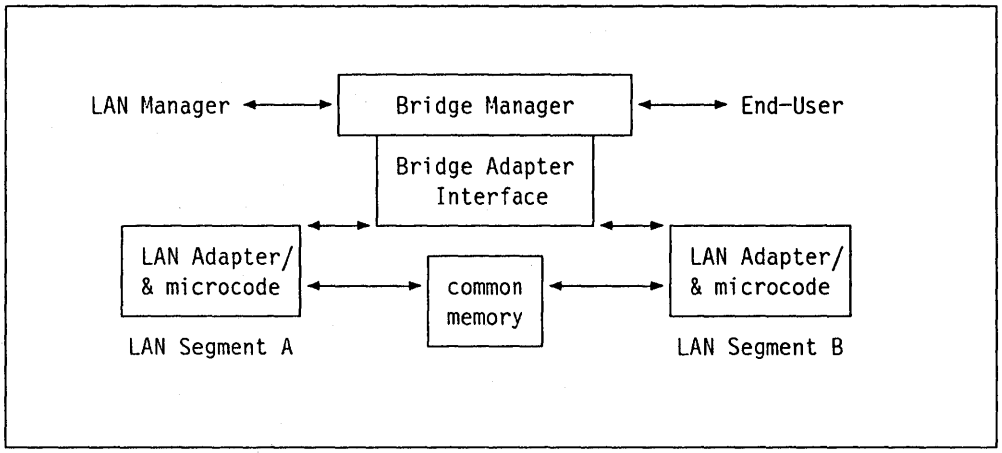

Figure 12. LAN Bridge Structure

Similarly, Figure 13 illustrates the major components and their interrelationship for a split bridge (or remote bridge).

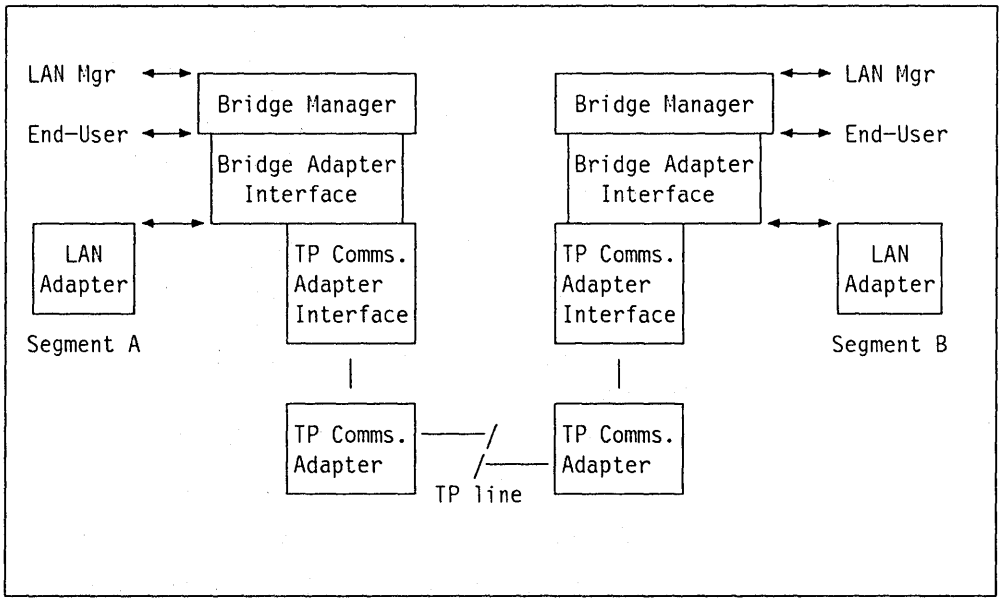

Figure 13. Remote Bridge Structure

As you can see in the figure, the two main components of a bridge are:

1. The bridge adapter interface
2. The bridge manager (or LAN reporting mechanism)

Both components' major functions are described hereafter. Bridge parameters and coexistence considerations are also discussed in the next sections.

3.3 Bridge Adapter Interface

The main function of the adapter interface is to transfer frames between the segments to which the bridge is connected. This process can be divided in two stages:

1. Copy Decision

   Before copying a frame from one segment, the bridge adapter will examine the destination address and the routing information (if any) of that frame. The bridge will copy the frame in the following cases:
   - There is a specific, functional or group address match
   - A non-broadcast frame if the bridge is part of the route identified in the RI (Routing Information) field
   - A broadcast frame as long as it has not already been on the target ring and it meets the adapter's single-route broadcast criteria.
2. frame forwarding

   Once a frame has been copied into the bridge, the adapter microcode will determine if the frame is to be routed through the bridge. Before moving the frame to the output adapter, several checks are performed by the bridge adapter interface. Depending on the nature of the frame (broadcast or non-broadcast), some of the following major checks will be performed:
   - Frame length checking (see "Largest Frame Size" on page 50)
   - Loop checking (see "Loop Check" on page 39)
   - Target ring status checking
   - Hop count checking (see "Hop Count Limit" on page 39)
   - Frame filtering by user appendage (see "The Filtering Facility" on page 71) in "IBM Token-Ring Network Bridge Program V2.1 "Local" or "Remote" Bridge Function" on page 53.

   Once these checks are made, the data will be moved from the receive buffers on the input segment's adapter to the transmit buffer on the output segment's adapter. In split-bridge mode, the data must be sent across the communication link to reach the output segment's adapter on the other half of the bridge.

3.4 LAN Manager Support

The IBM Token-Ring Network Bridge Program V2.0, IBM Token-Ring Network Bridge Program V2.1 and IBM PC Network Bridge Program are all capable of communicating with IBM LAN Manager V2.0 to provide consistent network management for the entire multisegment LAN.

IBM LAN Manager V2.0 will treat a remote bridge (composed of two bridge halves) like a "normal" bridge and will also generate alerts for errors detected in the communications link between the bridge halves.

Any IBM bridge program communicates with IBM LAN Manager Version 2.0 to provide a network management capability for the multi-ring environment. The network management information sent to the LAN Manager consists of:

- Soft error and beaconing notification
- Bridge status and performance data
- Configuration reports.

In addition, any IBM bridge program via a component called the LAN reporting mechanism, is capable of establish a reporting link with up to four IBM LAN Manager V2.0's. Each reporting link is an IEEE 802.2 logical link control Type 2 (connection-oriented) session, dedicated to transport network management information in either direction.

Some LAN management commands change the way in which a LAN segment operates. When a bridge has a reporting link with several IBM LAN Managers to control the same LAN segment, those commands must be reserved to one LAN Manager only in order to avoid conflicts. This LAN Manager is called the *controlling LAN Manager* for a given bridge. All other (up to three) LAN Managers to which this bridge reports are called *observing LAN Managers*.

The link between a bridge and its controlling LAN Manager is reporting link 0, also referred to as authorization level 0. Once a reporting link 0 is established, any IBM bridge program will reject subsequent attempts by other IBM LAN Managers to establish a reporting link 0, thereby avoiding duplicate controlling LAN Managers.

Observing LAN Managers have reporting links 1, 2 or 3 with a given bridge. Reporting links are always initiated by the LAN Manager station, either at LAN Manager initialization or during LAN Manager operation by the network operator through a link-bridge command. The operator may also decide to un-link a reporting link. Setting up a reporting link requires the operator to provide a link password.

Any IBM bridge program may receive requests from IBM LAN Manager V2.0, and send back appropriate responses over a reporting link after executing a given request on an attached ring segment.

In addition, any IBM bridge program may receive commands, reserved to the controlling IBM LAN Manager V2.0, over reporting link 0. Some of these commands include:

- Set/reset single-route broadcast selection mode
- Set/reset single-route broadcast
- Set/reset hop count
- Set/reset ring number
- Set/reset bridge number
- Remove station on a ring segment
- Set soft error logging options for a ring segment
- Perform a path test.

The communication between IBM LAN Manager V2.0 and IBM Token-Ring Network Bridge Program V2.0 is shown in Figure 14 on page 22.

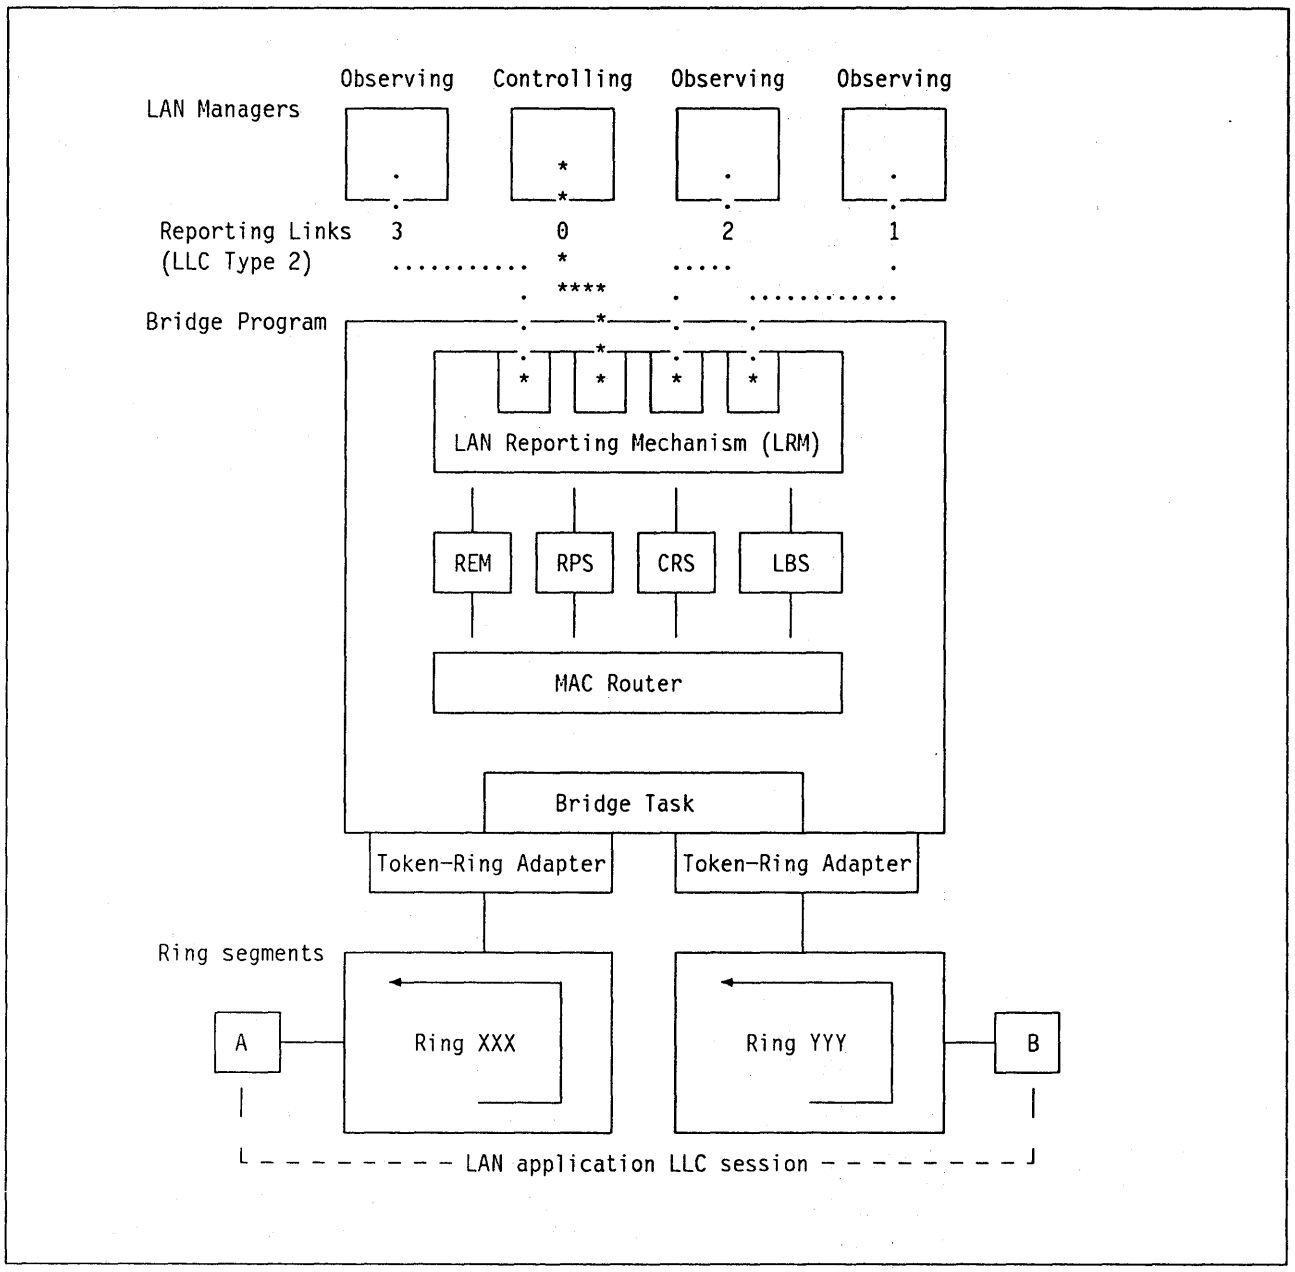

Figure 14. IBM Token-Ring Network Bridge Program V2.0 - IBM LAN Manager V2.0 Communication

This figure summarizes both the LAN application flow of LLC sessions between end stations (carrying data and higher-level protocol header information) and the network management data flow. The latter consists of LLC Type 2 sessions between LAN Manager and the LAN reporting mechanism which will direct a LAN Manager request/command to the appropriate server function in the bridge.

On the ring segments side, LAN Manager requests/commands are translated into the appropriate MAC frames for execution on one of the two attached ring segments. Responses provided by the appropriate server function are directed to the LAN reporting mechanism which packages them into an LLC frame addressed to the originating LAN Manager.

The different bridge server functions which may be enabled optionally on each of the bridge token-ring adapters are also shown in Figure 14.

- **Ring Error Monitor** (REM)

  The REM server function processes both hard and soft errors. It compiles error statistics reported by stations on the ring, analyzes these statistics, and selectively sends reports to the LAN reporting mechanism to notify LAN Manager of critical problems.

- **Ring Parameter Server** (RPS)

  In addition to supplying the ring number to stations as they send Request_Init_Parameters MAC frames to the RPS functional address during the ring insertion process, the RPS function also notifies LAN Manager when a given station has entered the LAN segment. A specific bridge protocol assures that there is only one RPS on each ring segment.

- **Configuration Report Server** (CRS)

  The CRS function forwards configuration notifications to LAN Manager. When receiving a MAC level configuration notification, it will transmit the receive information via the LAN reporting mechanism to LAN Manager. From LAN Manager, CRS accepts such commands as Query Adapter, Remove Adapter and Set Station Parameters for execution on a bridge-attached LAN segment.

- **LAN Bridge Server** (LBS)

  The bridge processing by the LBS function is always present and consists of:

  - Reading and validating bridge parameters from a configuration file at bridge initialization time and whenever a controlling LAN Manager modifies bridge parameters.
  - Performing the bridge self-test during bridge initialization or upon request from an operator through the bridge user interface. This test includes detection of duplicate parallel paths and invalid network configurations (that is, inconsistent ring numbering).
  - Maintaining a set of performance counters for each adapter, including counters for the number of frames discarded, not received or not forwarded for any other reason, and for the number of frames and bytes forwarded (both for broadcast and non-broadcast frames). On request, the accumulated values may be reported to LAN Manager.
  - Accumulating path trace information for frames carrying a system path trace bit set on in the routing information control field. Path trace report frames may be sent to the controlling LAN Manager.

3.5 Bridge Parameters

All IBM LAN bridges are characterized by a set of variables, which is almost the same for the three bridge products mentioned earlier.

As an example, Figure 15 on page 24 lists the configuration parameters for the IBM Token-Ring Network Bridge Program V2.0. All default values may be

modified using the configuration utility ECCCNFG, (except the largest frame size which can be modified only for the remote bridge function) and most of them can be also modified by the LAN Manager or the NetView operator. (The bridge label and path cost can not be modified by the LAN Manager or the NetView operator).

Parameter Description	Default	Range
Bridge number	1	0-9,A-F
Largest frame size	see note	
Frame forwarding active	Y	Y,N
Bridge performance threshold	10	0-9999
Drive for error log	0	0,A,B,C,D
Restart on error	Y	Y,N
Drive for memory dump on error	0	0,A,B,C,D
For each ring segment: (Adapter 0 = Primary, Adapter 1 = Alternate)		
Ring segment number (Adapter 0/1)	001/002	001-FFF
locally administered address	000000000000	4000nnnnnnnn
Shared RAM address (Adapter 0/1)	D800/D400	Adapter Ref.
Hop count limit	7	1-7
Single-route broadcast selection mode	M	M,A
Single-route broadcast	Y	Y,N
Automatic single-route broadcast		
Bridge label	8000	0000-FFFF
Path cost	0000	0000-FFFF
Ring Parameter Server	Y	Y,N
Ring Error Monitor	Y	Y,N
Configuration Report Server	Y	Y,N
Link Password 0	00000000	6-8 chars
Link Password 1	00000000	6-8 chars
Link Password 2	00000000	6-8 chars
Link Password 3	00000000	6-8 chars

Figure 15. IBM Token-Ring Network Bridge Program V2.0 Configuration Parameters

Important considerations related to the bridge parameters are mentioned below:

- The bridge number must be unique for parallel bridges (but only for parallel bridges).
- The largest frame size indicates the frame size this bridge is able to forward. This parameter is not configurable for the local bridge and is determined by the bridge program based on the installed adapters. See "Largest Frame Size" on page 50 for more details.
- Ring numbers must be unique within the entire multisegment LAN.
- Multiple bridges should refer to the same ring with the same segment number, otherwise bridge initialization will fail.
- Bridge performance threshold indicates the tolerance for lost frames (due to target ring inoperable conditions, congestion, etc.) before a notification is sent to LAN Manager (default value is 10 lost frames per 10,000).

For a remote bridge configuration using IBM Token-Ring Network Bridge Program V2.1, there is also a telecommunication link error threshold reflecting the maximum number of frames that can be lost on the telecommunication link before the LAN Manager is notified of a threshold exceeded.

- Restart on error = Y will force DOS to reload and execute AUTOEXEC.BAT automatically if an adapter check, critical bridge resource failure or internal programming error occurs.
- Dump on error = Y will cause a dump of the IBM bridge program's memory and buffers to be written to a file ECCDUMP.DAT if an internal programming error occurs.
- Drive for error log: if an error condition occurs, the error termination message is written to a file ECCLOG.DAT.
- Hop count limit may be a very important performance parameter. A broadcast frame [8] broadcast frame (for example, during route resolution) will not be forwarded by the receiving adapter if its hop count limit is less than or equal to the number of bridges the frame has already crossed (and as indicated by the routing information). See "Hop Count Limit" on page 39 for more information and examples on hop count limit.
- Single-route broadcast selection mode:
  - Manual mode

    If you choose M (manual), you must set the single-route broadcast parameter **manually for each bridge in your network**:

    If you set this parameter value to Y(yes), all frames with single-route broadcast active (BBB='11X') will be forwarded (if hop count limit not reached). If set to N all frames with single-route broadcast active will be discarded.
  - Automatic Mode

    If A (automatic), the bridge program will communicate with other bridge programs to determine how to set the single-route broadcast parameter value to Y or N to compensate automatically for changes in the network configuration. If Automatic mode is selected, **all bridges in the network should also be set to automatic**. You can also specify the bridge label and the path cost increment. These parameters are used by the automatic single-route broadcast algorithm and are discussed in "Automatic Single-Route Broadcast" on page 31.
- REM, RPS and CRS functional addresses are optional. Duplication should be avoided; for example, duplicate REM function provided by two bridges on the same segment may cause duplicate error message to flow to the same LAN Manager.
- Link passwords are used to authorize a LAN Manager to establish a reporting link with a given bridge.

---

[8] Hop count does not apply to non-broadcast frames or single-route broadcast frames.

3.6 IBM Bridge Products' Coexistence and Migration

The IBM Token-Ring Network Bridge Program Version 1.0 and the IBM Token-Ring Network Bridge Program Version 1.1 can coexist in a multisegment network with the IBM Token-Ring Network Bridge Program V2.0, the IBM Token-Ring Network Bridge Program V2.1 and the IBM PC Network Bridge Program.

However, in order to take advantage of the dynamic maintenance capability of the single-route broadcast path, all bridges in the multisegment network must be running either the IBM Token-Ring Network Bridge Program V2.0, IBM Token-Ring Network Bridge Program V2.1 or the IBM PC Network Bridge Program, since only those bridge products can be configured for automatic maintenance. In addition, use of the automatic configuration of single-route broadcast function will require a PTF for early versions of IBM Token-Ring Network Bridge Program V2.0 and IBM PC Network Bridge Program when used in the same network with IBM Token-Ring Network Bridge Program V2.1.

The IBM Token-Ring Network Bridge Program Version 1.0, withdrawn from marketing, has no functional capability to communicate with a LAN Manager.

The IBM Token-Ring Network Bridge Program Version 1.1 may communicate with up to four IBM LAN Managers V1.0. One IBM LAN Manager V1.0 can maintain up to 32 concurrent reporting links with several IBM Token-Ring Network Bridge Program Version 1.1.

The IBM Token-Ring Network Bridge Program V2.0, IBM Token-Ring Network Bridge Program V2.1 and IBM PC Network Bridge Program all communicate exclusively with IBM LAN Manager V2.0, providing extended network management capabilities. One IBM LAN Manager V2.0 can maintain up to 64 concurrent reporting links with IBM Token-Ring Network Bridge Program V2.0, IBM Token-Ring Network Bridge Program V2.1 and IBM PC Network Bridge Program bridges.

As a conclusion of all the above, in order to have the full LAN segment interconnection capabilities currently offered by IBM bridge products while maintaining the network management capabilities available with the previous IBM Token-Ring Network Bridge Program Version 1.1 - LAN Manager Version 1.0 communications, **all bridges and LAN Managers must be upgraded concurrently**.

The IBM Token-Ring Network Bridge Program V1.0 must be replaced by the IBM Token-Ring Network Bridge Program V2.0. The IBM Token Ring Network Bridge Program Version 1.1 may be upgraded at a charge to either the IBM Token-Ring Network Bridge Program V2.0 or IBM Token-Ring Network Bridge Program V2.1. IBM PC Network Bridge Program may be added to provide additional mixed LAN segment interconnection.

The IBM LAN Manager V1.0 may also be upgraded at a charge to IBM LAN Manager V2.0 to support the network management capabilities of the upgraded bridges.

4. Routing Support

4.1 Source Routing Approach

Source routing is the method used in IBM Token-Ring Networks to determine the route a frame must travel through a multi-ring LAN in order to reach its destination. With the new PC Network Bridge Program, source routing will also be used for PC Network segments, as well as Token-Ring segments. In this section of the document, the terms ring, bus or segment are equivalent from a source routing point of view.

The source routing algorithm provides for the acquisition and use of routing information by the ring stations. Source routing eliminates the need for routing tables at intermediate points (bridges) in order to direct a frame to its proper destination. When communicating stations are on the same ring, no routing information is required. In source routing, each frame carries with it the information about the route it is to follow. This *routing information* is acquired at connect time through a search process that originates at the source station and fans out throughout the network to the target station.

The discovery algorithm is contained within the application support software and microcode of the adapter, so that the end user is insulated from the process.

The search process to acquire the routing information is based upon a broadcast search message sent to all rings. The merit of source routing is that it is reliable and places the burden for acquisition and storage of routing information on the communicating applications, rather than on a central control utility.

If several routes are available between ring stations, then the search process will provide several alternative routes. The preferred route will be chosen by selecting the first discovery response to return. If routing information becomes outdated (for example, due to configuration changes on the network), a new route can be sought by the source station after attempts to use the old route have failed.

Routing information is acquired by the originating ring station by broadcasting a search frame throughout the interconnected LAN, which accumulates routing information as it traverses bridges. Each bridge that is traversed by this frame adds information to it indicating the rings that it connects.

A route through a multisegment LAN can be described as a sequence of *segment numbers* and *bridge numbers*. A segment number - bridge number combination is called a *route designator*.

Therefore, before looking at the route-resolving mechanisms used in source routing, Figure 16 on page 28 shows how *routing information* may be contained within a token-passing ring frame, and how route designators fit within the RI field. Figure 17 on page 28 lists the control bit settings involved in source routing.

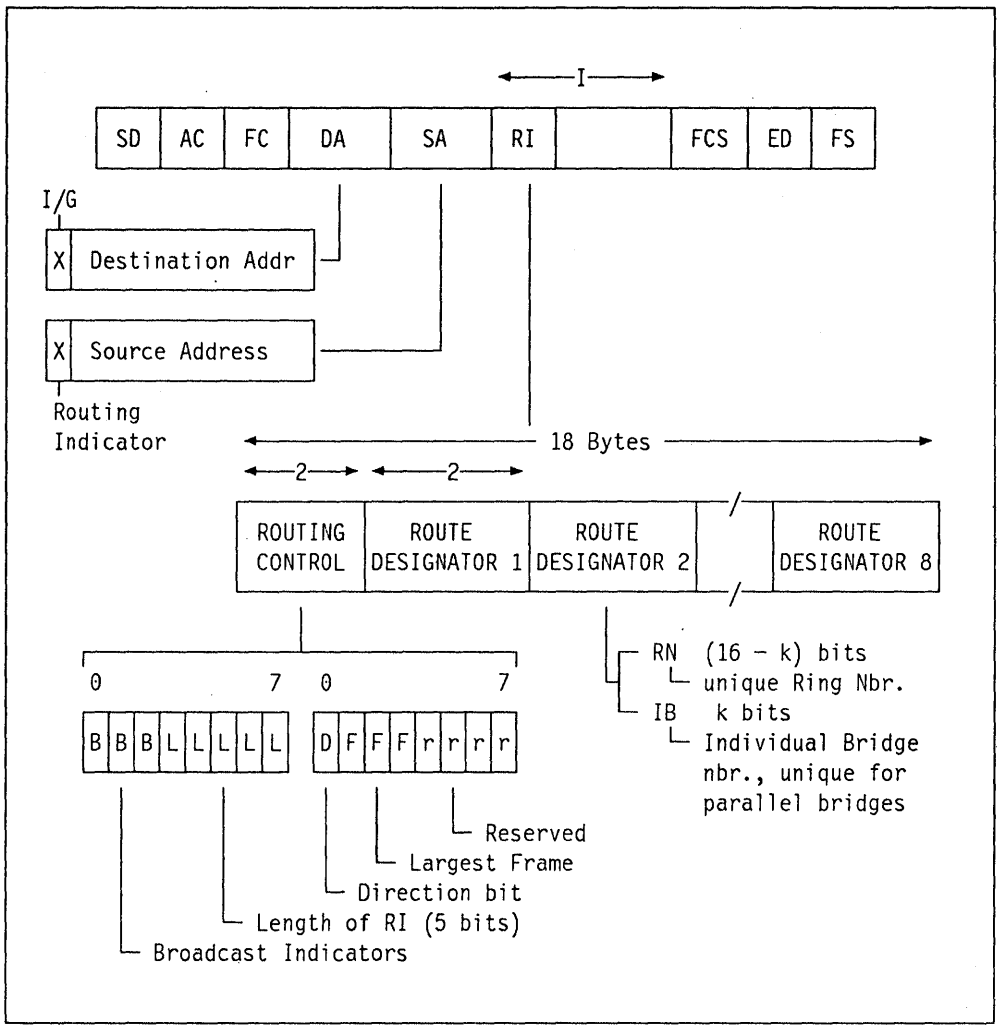

Figure 16. Token Passing Ring Frame - Routing Information Field

Broadcast	Designator	Comments
B'0XX'	Non-broadcast	Used in all-routes broadcast frames
B'10X'	All-routes Broadcast	Frame transmitted along every route in the network to the destination station
B'11X'	Single-route Broadcast	Only designated bridges will relay the frame from one segment to the other

Route Designator	Length (bits)	Comments
Ring-Number	16-k	k is the same for all bridges in one multi-ring LAN, typically k = 4 (Bridge V1)
Individual-Bridge-Nr.	k	

Figure 17. Routing Information Field - Code Bits

Source routing is started at connect time (LLC) by DLC_LAN_MGR (the data link control LAN management component) taking actions to *resolve the route.*

Source routing is normally a two-stage process in which the destination station is first searched for on the *local LAN segment.* If the destination cannot be located on the local segment, the destination will be searched on remote segments connected via bridges using broadcast frames. Not all products search the local ring first as described in the following flows, nor do they always use TEST and/or XID commands. Specific IBM product implementations are described in Figure 18 on page 30.

- On-segment route determination: the source station sends a TEST or an IEEE 802.2 XID command LLC protocol data unit (LPDU) in order to find the destination on the local ring. The destination, if present, responds with a TEST or XID response LPDU. If no response LPDU is returned within a specific amount of time, the destination is not on the local LAN segment, and the second stage of the route discovery is initiated.
- Off-segment route determination: the source station immediately retransmits the TEST or XID LPDU, indicating however the presence of routing information by setting byte 0, bit 0 of the source address field to 1 and appending an initial 2-byte RI field after the SA (see Figure 16 on page 28). At this point the architecture provides two main dynamic route discovery processes:
  - **All-routes broadcast route determination**
  - **Single-route broadcast route determination.**

  The route discovery process used at connect time depends on the type of software (or microcode) used in the station. These techniques are described in the following paragraphs, and have different implications in terms of broadcast traffic amount and control, as discussed in the "Broadcast Traffic Control" on page 35.

4.1.1 All-Routes Broadcast Route Determination

All-routes broadcast route determination is sometimes called *general broadcast* in other documents.

In the TEST or XID command LPDU, broadcasted to all rings by the source station, the first two bits of the RI field are set to B'10'. This triggers all bridges to copy the LPDU frame and while forwarding it, to complete the RI field with additional route designator information.

The destination receives as many command LPDUs as there are available routes, while any received frame contains in its routing information field exactly the route that has been followed. Any received command LPDU will be returned as a response LPDU, setting the first RI-bit to B'0' (= non-broadcast) and another routing control bit, the direction bit, to B'1'. This forces any response frame to flow back to the source station through the bridged LAN following exactly the same route as built in the command LPDU they respond to.

The source station selects the *preferred route* from all returned response LPDUs, by choosing the **first** discovery response to return. All subsequent transmissions to the particular destination follow the preferred route. The destination learns the preferred route from first non-broadcast frame received

PRODUCT		RI	COMMAND TYPE
NETBIOS:	CMD/REQ	SRB	NAME_QUERY
	REPLY	GB	NAME_RECOGNIZED
3270 EMUL. V3.0:	CMD/REQ	none	XID (SAP 0) Local Segment
(PC/DOS)	CMD/REQ	SRB	XID (SAP 0) Multi-Segment
	REPLY	Routed	XID resp from SAP 0
WORKSTATION PGM:	CMD/REQ	none	TEST (SAP 0) Local Segment
Version 1.1	CMD/REQ	GB	TEST (SAP 0) Multi-Segment
(PC/DOS)	REPLY	Routed	XID resp from SAP 0
PERSONAL COMMUNIC/	CMD/REQ	none	TEST (SAP 0) Local Segment
3270 V 1.0	CMD/REQ	GB	TEST (SAP 0) Multi-Segment
(PC/DOS)	REPLY	Routed	XID resp from SAP 0
OS/2 EE COMM MGR:	CMD/REQ	none	TEST (SAP 0) Local Segment
Version 1.1	CMD/REQ	GB	TEST (SAP 0) Multi-Segment
(OS/2 EE)	REPLY	Routed	XID resp from SAP 0

Figure 18. Route Discovery Techniques Used by Various IBM Products.

from the particular source station. The preferred route information between two stations remains valid for the duration of the DLC session between both.

4.1.2 Single-Route Broadcast Route Determination

Single-route broadcast route determination is sometimes called *limited broadcast*.

A value of B'11' in the first two bits of the RI field, marks the broadcasted TEST, XID or UI(NAME_QUERY) command LPDU as a single-route broadcast frame. Depending on the setting of the limited broadcast bridge parameter, the bridge will decide whether or not to copy and forward the frame (leaving the RI-field unchanged).

It is possible to select (manually or automatically) a subset of the bridges to build a **single-route broadcast path** so that there is only one path in the network to forward single-route broadcast frames between two different segments.

In this way, each LAN segment may receive exactly one copy of the TEST or XID command LPDU. **Therefore, the main benefit of the single-route broadcast technique (when used properly) is that each segment and the destination station will receive one and only one command LPDU.**

The TEST or XID command LPDU may be returned as a response LPDU with first two RI-bits indicating all-routes broadcast. Again multiple response LPDUs may be received by a source station, but in this case routing information has been collected in the response flow. However, most IBM Token-Ring Network

gateway products (37XX, 3174) reply with a routed TEST or XID response. Thus with these products, if the source station uses single-route broadcast for the TEST or XID request, then **all** traffic using this method will follow the single-route broadcast path through the network, and parallel bridges which have SRB off will not be used.

As in the first procedure, the source station selects the *preferred route* from all returned response LPDUs. Current implementation chooses the **first** discovery response to return.

Limited broadcast at the bridge level may be set on or off at bridge initialization time by a (controlling) LAN Manager. This requires accurate determination of this parameter for every individual bridge in the multisegment LAN to meet the single-path requirement of this route-resolving mechanism. In addition, whenever a bridge in the path becomes unavailable for any reason, the LAN Manager must take appropriate action to (manually) define an alternate path.

Alternatively, one may configure all bridges to provide **automatic single-route broadcast path determination and maintenance**. This facility, described in detail in "Automatic Single-Route Broadcast," presents several major advantages over the manual setting.

4.1.3 Automatic Single-Route Broadcast

During IBM Token-Ring Network Bridge Program V2.0 configuration, (as well as IBM Token-Ring Network Bridge Program V2.1 and IBM PC Network Bridge Program) the single-route broadcast parameter for each bridge accepts a new value *Auto* for "automatic".

The objective of the automatic single-route broadcast facility is to enjoy the single-route broadcast benefits without having to select the single-route broadcast options manually for each bridge adapter in the network. In order to use the automatic single-route broadcast facility, all bridges in the network must be configured with the "automatic" option.

Since the earlier IBM Token-Ring Bridge Programs do not support dynamic maintenance, they must be upgraded to the IBM Token-Ring Network Bridge Program V2.0 to utilize this feature.

4.1.3.1 Automatic Single-Route Broadcast Route Determination

The bridge can be configured with the configuration utility or by the IBM LAN Manager V2.0 to communicate with other active bridges to dynamically maintain the network single-route broadcast path.

A bridge, configured to participate in the dynamic maintenance of the network's single-route broadcast path, can be in any of three modes:

Blocking The bridge does not forward single-route broadcast frames and does not participate in the bridge protocols.

Listening (or non-designated) The bridge does not forward single-route broadcast frames but participates in the bridge protocols.

Forwarding (or designated) The bridge forwards single-route broadcast frames and participates in the bridge protocols.

A bridge is in blocking mode during initialization. Once the bridge has opened the adapters and has set the appropriate functional address, it is in listening mode. After participating in the protocols long enough to determine if it should forward single-route broadcast frames, the bridge will stay either in listening mode or move to forwarding mode. Bridges participating in the protocols monitor the single-route broadcast path with inter-bridge communication. All inter-bridge communication is sent as logical link control type 1 (connectionless) data to the bridge functional address. The inter-bridge communication is periodically initiated by a "Heartbeat" or "Hello" frame sent from the **root bridge**.

This root bridge is automatically selected through processes specified by the Spanning Tree algorithm. In IBM bridges, the selection of the root bridge is done on the **bridge ID** basis. The **bridge label** (four digits bridge parameter) combined with the **adapter address** of the adapter connecting the bridge to the lowest LAN segment number is the bridge ID. The active bridge which has the *lowest bridge ID* will be the *root bridge*.

The root bridge will periodically send out a "Hello" frame on both LAN segments to which it is connected. Other bridges which provide the best or only path to other LAN segments in the network are selected as **designated bridges**.

This selection is based on the route path cost between a bridge and the root bridge. The route path cost is the sum of the bridge path costs increments involved in the route between the bridge and the root bridge.

The path cost increment is a four-digits bridge parameter which indicates the relative length (cost) of the bridge forwarding process. For example, traversing a remote bridge is longer and more expansive due to the TP link cost than traversing a local bridge.

For a given segment, in case of equivalent path costs, the selection algorithm between the bridges will consider the bridge ID; the bridge with the lowest bridge ID will become the designated bridge for that segment.

When a designated bridge receives a "Hello" frame on its root side port, it will send out a "Hello" frame on its non-root side port (or designated side port). Each LAN segment will have a "Hello" frame circulated on it by either the root or a designated bridge. The root and all designated bridges will have single-route broadcast active for both ports. All other bridges will have single-route broadcast deactivated for both ports. Any bridge that is not selected as the root or a designated bridge provides a parallel path through the network. If the root or a designated bridge is shut down, one of the blocking bridges will automatically be selected to replace it and will turn on single-route broadcast for both ports.

The bridges will also reconfigure the single-route broadcast path to accommodate network changes caused by a new bridge being activated.

It is important to notice that single-route broadcast path maintenance does not affect the way in which a bridge processes all-routes broadcast or non-broadcast frames.

4.1.4 Source Routing Example (Single-Route Broadcast)

In this example, we'll suppose the route discovery mechanism used by a station to connect to another station is the single-route broadcast process.

The figure below is a typical network configuration, with local segments (rings) 1,2,...,N connected through bridges B1 and C1 to duplexed backbone segments B and C.

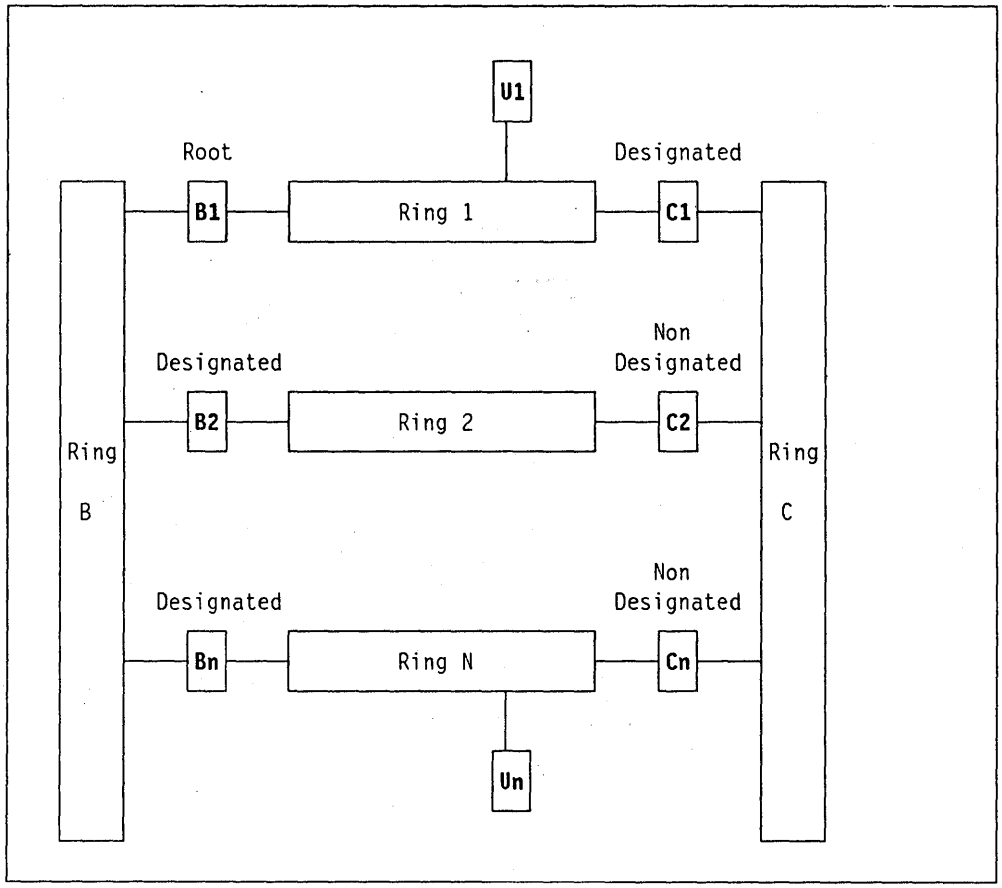

Figure 19. Single-Route Broadcast in a Dual Backbone Configuration

The preferred way for the end station to determine an optimal route is as follows. First, a subset of the bridges is designated "single-route broadcast", so that only **one** route exists between any two segments across the single-route broadcast bridges. As explained earlier, the bridges selection can be done automatically or manually at the LAN Manager or at the bridges. The rest of this section assumes we use the automatic option.

In our example, bridges B1, B2, ..., BN, and C1 could be designated single-route broadcast, depending on the setting of the bridges parameters like bridge label and Path cost. If the bridge label is the name shown on the figure, then B1 will be the root (X'B1' is the lowest bridge label) [9]. For a given user segment like Ring 2, (assuming equivalent adapter and speed settings and default path cost increments) it is clear that the route path cost of bridge B2 is less than the

[9] The lowest bridge ID always becomes the root.

route path cost of C2. That is why bridges B2, ..., BN will be designated "single-route broadcast", while bridges C2, ..., CN won't forward single-route broadcast frames. In addition, bridge C1 will be selected as a designated bridge in order to access to ring C, because bridge C1 has the lowest path cost from the root bridge (compared to C2,..., CN).

Suppose an end station on one segment wants to connect to a target station on another segment. The end station sends a "discovery" frame marked "single-route broadcast" so that all bridges so designated forward the frame. The result is that exactly **one** copy of the "discovery" frame will appear on each segment in the network, including the target segment. The target station is programmed to return the "discovery" frame to the originating station, but marked "all-routes broadcast". An all-routes broadcast frame will traverse **all** routes back to the originating station, including routes using bridges C2-CN.

As the discovery response passes through each bridge, the bridge is programmed to insert the local bridge-segment identification in the routing field in the frame, thereby recording the travelled route inside the discovery response frame. Before forwarding an all-routes broadcast frame, a bridge checks the routing field to see if the max-hop value (optionally) assigned to that bridge would be exceeded. If "yes", the frame will not be forwarded. A max-hop value can be assigned for each direction through the bridge, but the check is only applied to all-routes broadcast frames. The hop count limit bridge parameter is covered in details in "Hop Count Limit" on page 39. The bridge also will not forward a broadcast frame to a segment that it has already traversed; that is, frames will not "loop". This check can further reduce broadcast traffic in a "mesh" network.

For example, if an end station (U1) on segment 1 wanted to connect to a target station (UN) on segment N, then the discovery frame would flow through bridge B1 to segment B and then into segment N through bridge BN. Assume that the max-hop values were all set to 1 for all-routes broadcast frames flowing **out** of the segments 1, 2, ...,N (so that only all-routes broadcast frames **originating** on those segments would be forwarded). The discovery response would multiply and pass into the other segments. All subsequent routes would be eliminated by the max-hop check. For example, routes like "ring N, BN, ring B, B2, Ring2, C2 ..." will be eliminated.

Among the possible routes, the response would pass through BN-B-B1 and CN-C-C1. Depending on the bridge/segment activity at the time, one of the following two frames would arrive first at the originating station on segment 1 and will provide the station with the routing information. The routing information corresponding to the two most probable routes is shown below.

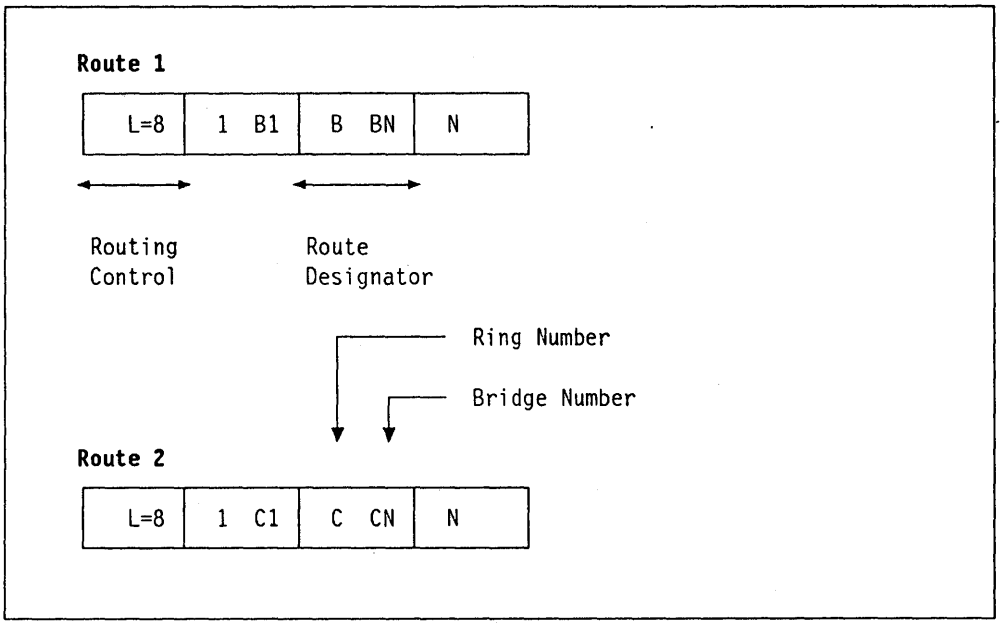

Figure 20. Example of Possible Routes

The preferred route, stored in the routing field, will be used for the subsequent connection.

As you can see, although bridge CN was not selected as a designated bridge (bridge CN does not forward single-route broadcast frames), bridge CN is still an active bridge and can be on the route selected for the subsequent session. This is a major difference with the transparent bridging technique, as explained later.

Over time, multiple connections between stations on any two segments would be statistically distributed across both backbones. If either backbone failed (or one of the bridges), then connections could be re-established over the other backbone. For example, if bridge B2 fails, bridge C2 will automatically be selected by the automatic single-route broadcast process as the designated bridge for segment 2.

4.2 Broadcast Traffic Control

During the route discovery process, broadcast frames are propagated to all segments in the network. Depending on the topology of the network and the type of applications run, this broadcast traffic may generate a lot of frames on each segment.

Simple topology networks with a few bridges and no alternate routes should not be affected by the broadcast traffic. Such examples are described in the first scenarios discussed in "LAN Design Methodology" on page 89.

However, typical dual backbone configurations or meshed topologies with many rings and bridges could generate redundant broadcast traffic if no precaution is taken. In addition, server stations and bridges may experience unnecessary congestion situations if this broadcast traffic is not controlled.

Fortunately, there are several simple ways to control and minimize this broadcast traffic. The most important techniques rely on **bridge parameters or facilities** like:

- Single-route broadcast facility (automatic or manual)
- Hop count limit
- Bridge filtering.

Other ways to minimize the broadcast traffic in **NETBIOS applications** environments include the following:

- Use the RND parameter in LAN Support Program NETBIOS driver
- Avoid using the gateway configuration for PC 3270 Emulation Program Version 3 inside an establishment, whenever possible.

The use of these techniques can reduce the broadcast traffic significantly by specifying a minimum number of parameters at installation time. In addition, all these techniques are complementary and should be combined in order to avoid unnecessary traffic in the LAN and provide the users with excellent response times, at least from a LAN point of view, even during peak logon periods. In the following typical examples as well as in the "High-Availability Design Considerations" on page 96, you will see that the network administrator can easily reduce the broadcast traffic by a factor of 10 or more, depending on the LAN configuration and connectivity requirements.

4.2.1 Single-Route Broadcast (Manual or Automatic)

As described earlier in "Source Routing Example (Single-Route Broadcast)" on page 33, the single-route broadcast technique applies only to the single-route broadcast traffic, which is by the way used by several IBM program products like all NETBIOS applications. The single-route broadcast technique does **not** apply to all-routes broadcast frames. All-routes broadcast frames traffic can be controlled by using the hop count limit described in "Hop Count Limit" on page 39.

4.2.1.1 Manual Setting

In order to use the single-route broadcast facility, one can do it manually by designating a subset of the network bridges as single-route broadcast bridges. Using the manual setting of the corresponding Bridge parameters requires a good understanding of the single-route broadcast. Usually, the LAN administrator will select the appropriate bridges subset, keeping in mind that he should not affect the "any-to-any" connectivity user requirements.

A manual method of designating the subset of the bridges which should have single-route broadcast active is described in the *LAN Administrator's Guide*.

The setting of the single-route broadcast parameter has to be done at installation time for each bridge adapter. Different values can be chosen for each side of the bridge, allowing a full control of the single-route broadcast traffic. In case of a bridge failure, some changes of the single-route broadcast bridge parameters will be required from the LAN administrator or LAN operator, in order to maintain a full connectivity.

Those changes can be done either from the LAN Manager or NetView console, which needs human intervention, or could be done from an appropriate NetView Clist.

4.2.1.2 Automatic Single-Route Broadcast Benefits

On the other hand, the new bridge programs all offer an automatic single-route broadcast technique which does not require from the LAN administrator any selection of the single-route broadcast bridges subset. This selection process is completely handled automatically by the bridges, as explained in "Automatic Single-Route Broadcast" on page 31. In addition, bridge parameters will be **automatically** adjusted by the remaining bridges in case of a bridge failure. This is a major consideration in terms of operations and *we strongly recommend using the automatic single-route broadcast facility provided by the bridge programs.*

The only thing to do to activate the automatic single-route broadcast is to choose the automatic single-route broadcast option at bridge installation. Note however that the LAN administrator can orient the single-route broadcast bridges subset selection by giving different values to bridge parameters such as path cost and bridge label, as explained in "Automatic Single-Route Broadcast" on page 31.

The only advantage of the manual setting is that the LAN administrator can define different single-route broadcast options for each bridge adapter. With the automatic option, both sides will automatically forward the single-route broadcast frames or not, depending on the current network configuration. In addition, in very simple LANs without alternative routes between segments, the automatic single-route broadcast facility is useless and should not be selected.

4.2.1.3 Path Cost and Bridge Label Recommendations

The automatic single-route broadcast function uses *path cost* to decide which parallel path between two LAN segments to use as the single-route broadcast path. You should use the default value of 0000 for the path cost increment for the bridge. The default depends on the type(s) and data rate(s) of the adapters used in the bridge station.

For local bridges, default values for the path cost increment range from 16 to 64. The root bridge path cost increment is always 0.

The remote bridge's path cost increment default value depends on the TP link speed and is significantly higher than path cost default values for local bridges. Default values for a remote bridge path cost increment range from 101 for a 1.344 Mbps link speed to 240 for a 9.6 Kbps link speed. This will prevent using the remote bridge for single-route broadcast frames if there is an alternate path using local bridges.

However, depending on your specific installation criteria, you can specify a path cost increment value based on other factors such as bridge's load or bridge filters appendages. See the appropriate bridge program user's guide or the IBM Local Area Network Administrator for additional information on path cost increment values.

The automatic single-route broadcast function uses the *bridge ID* during the single-route broadcast path selection process. The bridge with the lowest bridge ID will be automatically selected as the root bridge. The root bridge should be in a central location in your network to provide the shortest average path to the connecting LAN segments. In addition, when several parallel bridges connected to the same segment have the same route path cost, the

selection of the designated bridge for that segment is based on the lowest bridge ID.

As the bridge ID first two characters are the *bridge label*, you can influence the selection process depending on your particular configuration. The bridge label default value is 8000.

As a general rule, it is recommended to assign a high bridge label to remote bridges (for example 9000). This could prevent using the remote bridge for single-route broadcast frames if there is an alternate path with an equivalent route path cost using local bridges.

4.2.1.4 Example of Broadcast Traffic Reduction Using the Single-Route Broadcast

Let's assume we have a typical configuration as shown in Figure 21.

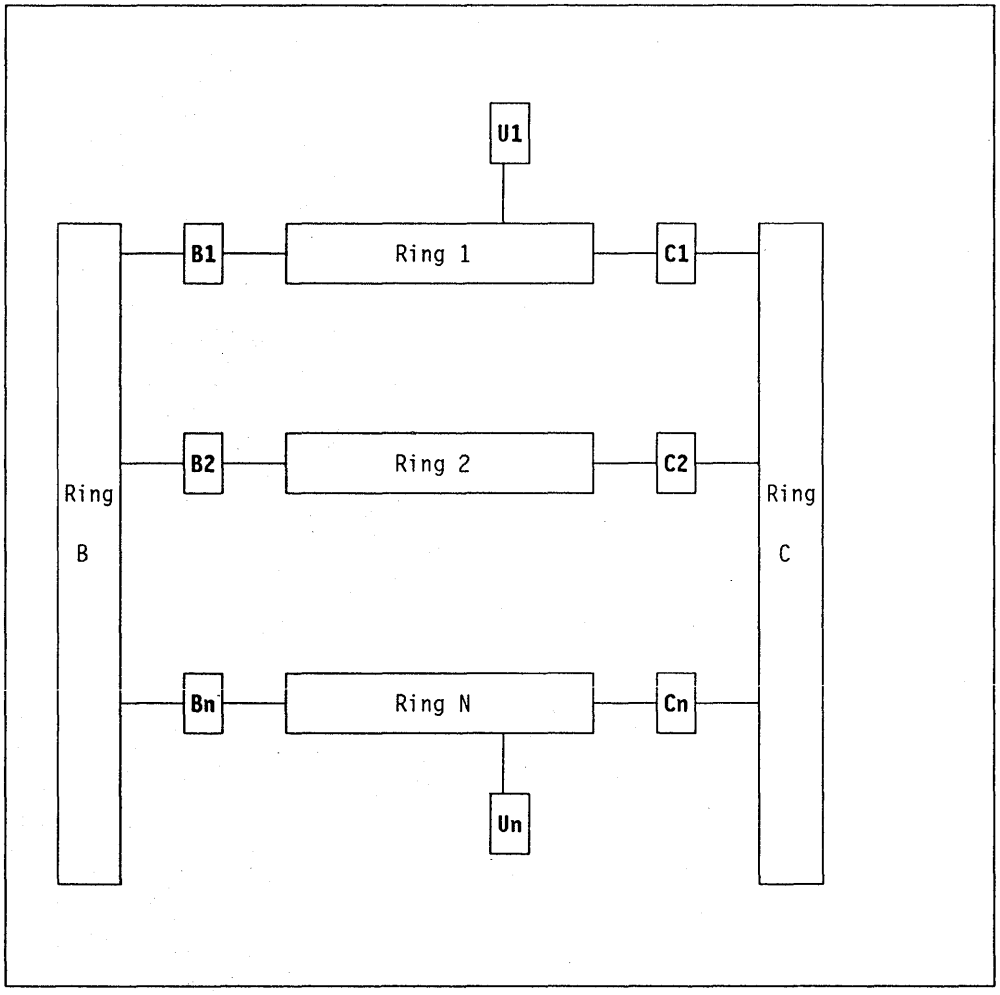

Figure 21. Broadcast Traffic Reduction with Single-Route Broadcast

Suppose an application running on station U1 on segment 1 wants to connect to a target station UN on segment N, using a single-route broadcast frame.

If you don't use the automatic single-route broadcast facility and keep the default single-route broadcast (Yes, Yes) and hop count limit default options (7,7) for all bridges, the U1 to Un connection will generate many frames on the

different segments. In fact, you do not have in this case a single-route broadcast path.

As a result, the target destination will receive:

4 frames if N = 3

18 frames if N = 10.

As the destination will answer to each of these frames using all-routes broadcast, a simple calculation shows that in the case where N=3, the total number of frames generated if hop count limit=(7,7) on the different paths back to the source station will be equal to 60.

On the other hand, **if you choose the automatic single-route broadcast** on the bridges the target destination will receive only:

1 frame if N = 3

1 frame if N = 10.

This is due to the fact that each segment will receive one and only one frame for this particular segment during the single-route broadcast process.

As you can see, using the automatic single-route broadcast facility has "just" reduced the number of frames received by station Un by a factor of 18 where N = 10 and the hop count limit = (7,7). This calculation is valid for every connection involving two different user segments in this particular but typical configuration.

This is the reason why we strongly recommend using the automatic single-route broadcast facility in all complex LAN configurations.

4.2.2 Hop Count Limit and Loop Check

In addition to the single-route broadcast facility, a bridge provides other ways to limit the broadcast traffic in the network.

4.2.2.1 Loop Check

An important bridge feature is called "loop check". Loop check means that a bridge will never forward a broadcast frame to a segment that it has already traversed. Before forwarding an automatic single-route broadcast frame, the bridge will check the existing routing field to see if the segment number on the other side of the bridge has already been traversed by the frame. The net result is that a broadcast frame will not loop in the network.

This function is automatic and does not require any customization.

4.2.2.2 Hop Count Limit

An other interesting feature of the bridges is the hop count limit bridge parameter. This parameter can further reduce the broadcast traffic in a "mesh" network and should always be considered in multiple routes LAN topologies, in addition to the automatic single-route broadcast facility.

The hop count limit indicates to a bridge the maximum number of bridges a broadcast frame can traverse before it is discarded by the bridge. The two adapters of the bridge can have different values specified for the hop count

limit. The default value is (7,7), which means that broadcast frames arriving on both sides of the bridge could have already traversed up to six bridges before traversing it. The number 7 represents the maximum number of bridges that can be traversed. The routing information is 18 bytes long with the first two bytes for routing control; the following 16 bytes are the route designator bytes with the first two route designators used to designate the first bridge crossed and the following route designators added as each additional bridge is traversed. Eight routing designators are used for seven bridges (maximum).

The use of the hop count limit bridge parameter is entirely dependent on the configuration. Before using a value different from the default value, the LAN administrator will have to consider the impact of this parameter on normal flow, but also in backup situations, if a bridge failure occurs. In particular, connectivity requirements should not be limited by the use of the hop count limit even in backup scenarios.

The hop count limit value for the bridges should obviously be set to the lowest possible value depending on the topology, as illustrated in the following examples.

4.2.2.3 Examples of Broadcast Traffic Reduction Using Hop Count Limit

If the configuration of the LAN is hierarchical, hop count limit may be assigned as follows:

- When a broadcast frame has entered a local ring from a backbone ring, that frame should not be allowed to leave the local ring.
- When a broadcast frame has entered a lower-level backbone ring from a higher-level backbone ring, that frame should not be allowed to go to a higher-level backbone ring.

An example of a hop count assignment for a hierarchical LAN is shown on Figure 22 on page 41. The four bridges connected to the second level backbone (ring A) have a hop count limit value of (3,2), 3 being the value for the ring A side. All other bridges have a hop count limit value of (4,1), 4 being the value for the first level backbone side.

Consider a single-route broadcast broadcast frame leaving local ring 1 for a connection to a station located in building 2. Since HC=1 for the local ring side of the bridge, the frame will be forwarded through the first level backbone ring B (or C). Now the frame has already passed one bridge, but the bridge interconnecting first and second level backbone rings has an HC=2 and can forward the frame towards the second level backbone ring. The frame has been forwarded by two bridges. Again, the bridge connecting the second level backbone ring A and the first level backbone ring D forwards the frame (it has HC=3) to the first level backbone ring D. Finally, bridges connecting the first level backbone rings D or E and local rings 4 5 and 6 are able to forward the frame to their local rings, since they have an HC=4. Notice that these hop count assignments do not allow for broadcast frames "falling down" from one ring to another at a lower-level, to go to a higher-level ring.

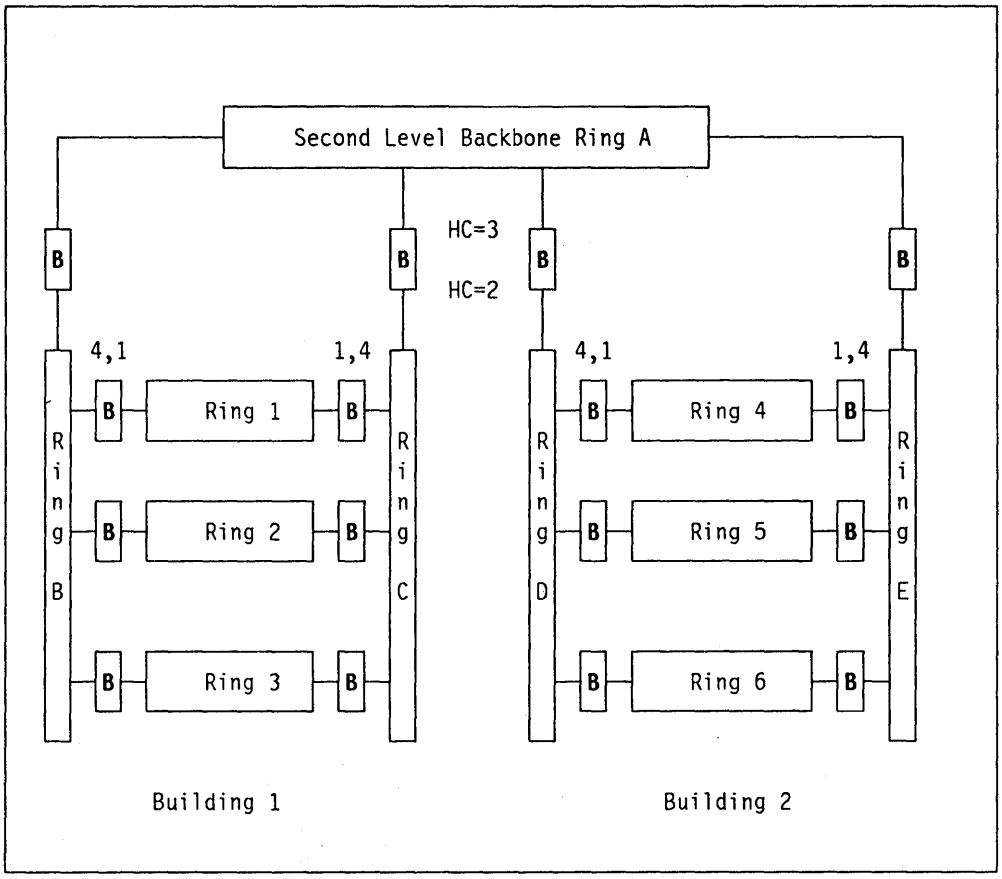

Figure 22. Hop Count Assignment for a Hierarchical Configuration. B are bridges and HC stands for hop count.

If the network topology is a typical dual backbone configuration as shown in the following example, the use of the hop count limit parameter can drastically reduce the broadcast traffic.

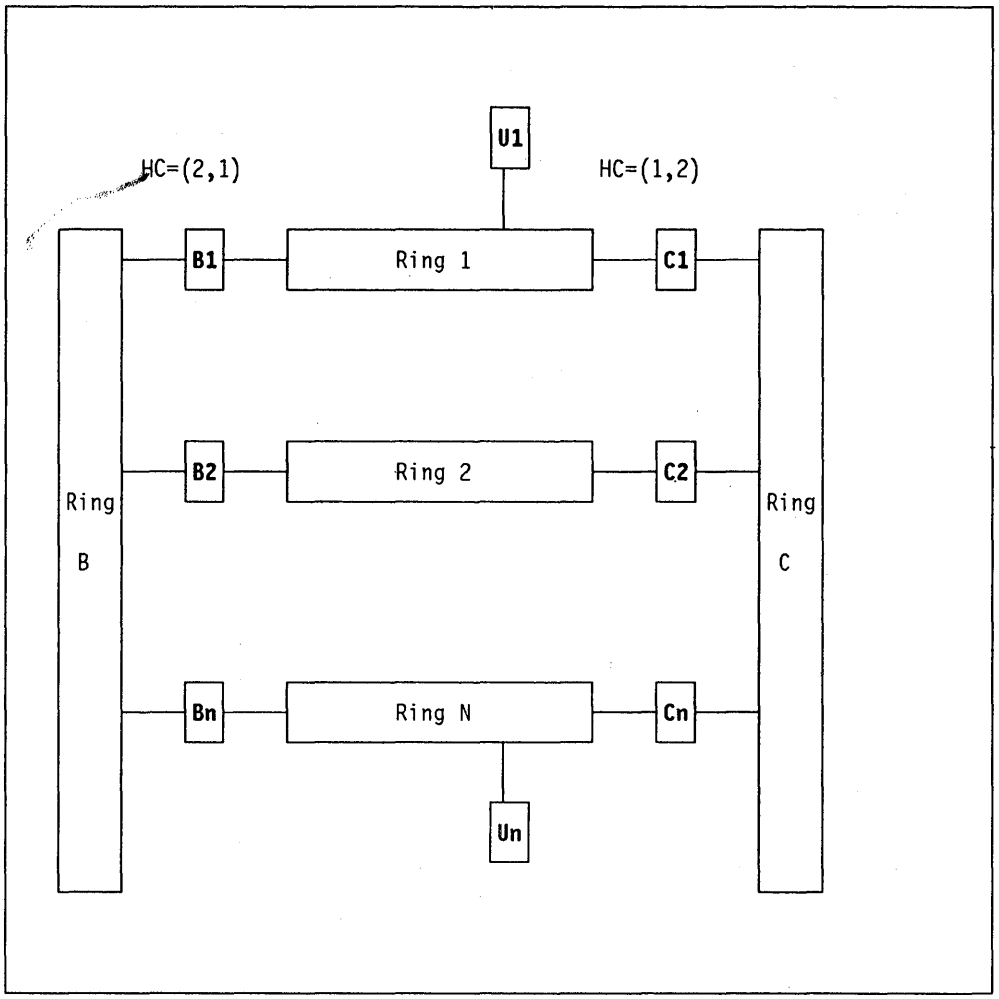

Figure 23. Dual Backbone Configuration and Hop Count Limits

In this example, suppose station U1 wants to communicate with station Un on ring N, using an **all-routes broadcast** frame. We'll assume a (7,7) default value for the hop count limit.

In this case, the number of frames received by the target destination will be equal to:

4 frames if N = 3

18 frames if N = 10.

The target destination will answer to each of these frames using the reverse path route built in the frame by the bridges which have been traversed. As a result, an equivalent number of frames will be sent back to the source station.

On the other hand we could set the hop count limit value for all bridges to (2,1). 2 would be the hop count limit value for the backbone side bridge adapter, 1 being the hop count limit value for the user ring side bridge adapter.

In this case, the target destination will receive only:

2 frames if N = 3

2 frames if N = 10.

As you can see, using the hop count limit facility has divided the number of frames received by station Un by a factor of 9 when N = 10.

Other scenarios illustrate the use of the hop count limit in "High-Availability Design Considerations" on page 96.

4.2.3 NETBIOS Applications (RND)

A characteristic of NETBIOS applications is that the stations will establish NETBIOS sessions across the LAN by using NETBIOS names. In order to find the target station, NETBIOS applications generally issue CALL commands which will generate NAME_QUERY frames sent to the NETBIOS functional address. The NAME_QUERY frames are sent as single-route broadcast by NETBIOS. Although very convenient, this approach can generate a lot of broadcast traffic in complex LAN configurations.

To reduce the unnecessary interrupts to NETBIOS nodes, as well as the number of frames circulating on the different segments, the Remote Name Directory (RND) function is used to send the frame to a specific address, whenever possible.

When RND is used, after the local station has located a remote name, **the remote address is saved** and subsequent messages to that name will be sent to a specific node rather than a broadcast to all NETBIOS nodes.

It is recommended to specify the RND option on NETBIOS nodes which establish sessions or issue a lot of NETBIOS CALL commands.

For more information on NETBIOS protocols and RND option, refer to the *Local Area Network Technical Reference*.

4.2.3.1 Gateway Configuration of the IBM PC 3270 Emulation Program V3

When you define a station as an IBM PC 3270 Emulation Program V3 gateway, you specify in that station a list of NETBIOS names that will be using the gateway to access the host. Those names correspond to network names defined in the network stations using the gateway. The gateway station will continuously "poll" each network station name until the station is active. As explained in the previous paragraph, this NETBIOS "polling" will generate a lot of broadcast traffic and unnecessary interrupts to NETBIOS nodes in complex multisegment LANs. If the stations can be connected to a host gateway via the LAN, the LAN administrator should consider defining all stations running the IBM PC 3270 Emulation Program V3 as "stand-alone" instead of "network stations". There are two major advantages in doing this:

1. Performance will be better, as the traffic between the station and the host gateway won't be affected by the PC defined as a gateway station.
2. The availability will be improved as stations will not have to rely on the PC gateway availability in order to access the host.

- IBM Personal Communications/3270 provides the user with more functions such as several 3270 Emulation sessions.
- The gateway configuration has been implemented using the 802.2 LLC protocols which gives better performance than NETBIOS. In addition the IBM Personal Communications/3270 gateway station does not "poll" the network stations and therefore does not generate a cyclic broadcast traffic for non-active stations as IBM PC 3270 Emulation Program V3 does.

- In addition, IBM Personal Communications/3270 provides an LU pooling function

4.2.4 Bridge Filtering

The filtering facility is provided by the IBM Token-Ring Network Bridge Program V2.1. It requires some programming and applies **both** to remote and local bridges. This feature will reduce the probability of a LAN station or remote bridge being congested by unnecessary broadcast traffic especially in a low-speed TP link environment.

The bridge filtering facility is discussed in "The Filtering Facility" on page 71.

4.3 Source Routing Approach versus Transparent Bridging

4.3.1 Source Routing Benefits

As illustrated in the source routing example and in the design scenarios (see "High-Availability Design Considerations" on page 96), the strengths of source routing include:

1. Multiple, concurrently **active** routes between source/target segments are supported.
2. Bridges need only know the local topology (such as the bridge number and the two segments connected).
3. The discovery process provides for statistical load balancing and route backup.
4. Other route parameters (for example, largest frame size and link timer values) can be negotiated or tuned during the discovery process.
5. End station awareness of the route being used for communication can be used for trace and diagnostic purposes.
6. Hop count limits in bridges can be used to control the "distance" that all-routes broadcast frames travel in the network. In addition, the automatic single-route broadcast facility will reduce the number of frames received by the destination to one, even in very large configurations.
7. Bridges only need to examine frames containing a routing field, as indicated by the first source address bit (=1). Non-broadcast frames containing a routing field will not be copied unless the segment-bridge-segment triple associated with that bridge is in the routing field.
8. Bridges can be implemented on inexpensive, general-purpose hardware (for example, PCs).

4.3.2 Transparent Bridging

As discussed in *Local Area Networks: Concepts and Products*, another architecture, called *transparent bridging*, allows interconnection of two or more segments via bridges. The philosophy of transparent bridging is not to modify the end stations to support bridging. All modifications to the network for bridging must occur in the bridge.

Transparent bridges do not build or inspect routes in passing frames. Instead, transparent bridges copy **all** passing frames and look at the source (SA) and destination (DA) addresses. The SA are used to build two-sided tables (SA is on this/that side somewhere) and the DA are matched against the two-sided tables for routing.

- If the destination address appears in the routing table on the **same** side as the frame-copy, the frame is not forwarded.
- If the destination address appears in the routing table on the **opposite** side as the frame-copy, then the frame is forwarded.
- If the destination address does not appear in the routing tables, then the frame is forwarded anyway.

Obviously, such a routing scheme does not tolerate networks with multiple paths between stations (stations would be on both sides of some bridges, creating network "loops"). Frames forwarded under these conditions would circulate indefinitely, with multiple copies arriving at the target station. Note that the max-hop check and the bridge check to prevent re-entering a segment are **not** available with transparent bridging because these options depend on knowledge of the route travelled. Therefore, a single-route broadcast subnetwork is used for **all** data movement.

The **Spanning-Tree algorithm** is used to maintain this single-route broadcast subnetwork. Details on this algorithm can be found in the transparent bridging section in *Local Area Networks: Concepts and Products*.

As stated earlier, one advantage of transparent bridging is that it can be used without explicit support in the end station, thus allowing non-source routing LANs to be bridged. In addition, the LAN administrator does not have to specify segment or bridge numbers. The bridges will build a tree structure based on default values and end stations will be able to communicate.

Weaknesses of transparent bridging include:

1. Multiple active routes between segments are not supported. Parallel bridges between adjacent segments are possible, but only by using static address filters, which increase processing overhead and must be updated for station location changes. Loss of one of these parallel bridges necessitates modifying or turning off the address filter in the other bridge(s), potentially losing frames. *Alternate, concurrently* **active** *paths for load balancing and high network availability do not exist. Redundancy, if any, is provided by passive, back-up bridges in blocking state, that are not available for data transfer, constituting a non-active, "silent" subnetwork.*

   Therefore, the standard design for large networks of a hierarchical scheme with multi-level, duplexed backbones and multiple, concurrently active routes cannot be supported.

2. Transparent bridges require special-purpose hardware capable of copying **all** frames and examining **large** address tables, at network data rates. Because the table entries in the bridge are periodically refreshed, the bridge must continually re-acquire address information. Note that until a station address appears as a source address, all traffic to that station is broadcast throughout the entire network.

3. End stations are not aware of the "extent" of the station-to-station communication (such as route length) and therefore are not able to make adjustments in protocol timer values, isolate "route" problems, etc.
4. Re-configuration (or failure) in the Spanning-Tree algorithm can lead to lost or circulating frames.

As a summary, each approach has advantages and disadvantages depending upon the size and traffic characteristics of the local area network. Both routing techniques should become standards in a very near future in order to address the requirements for large interconnected LANs and the installed base of stations which don't support source routing.

5. IBM Token-Ring Network Bridge Program V2.0 (Local Bridge)

5.1 IBM Token-Ring Network Bridge Program V2.0 Overview

IBM Token-Ring Network Bridge Program Version 2.0 is intended for establishments requiring connectivity between 4 Mbps and/or 16 Mbps IBM Token-Ring Network LANs. Using the bridge program, combinations of 4 Mbps or 16 Mbps IBM Token-Ring Networks can be connected into a single logical network. IBM Token-Ring Network Bridge Program Version 2.0 provides bridging capability between two rings of an IBM Token-Ring Network with the following data rate combinations:

- 16 Mbps to 16 Mbps
- 16 Mbps to 4 Mbps
- 4 Mbps to 4 Mbps.

Note that this "local" bridge function has been withdrawn from marketing, and that IBM Token-Ring Network Bridge Program V2.0 has been replaced by IBM Token-Ring Network Bridge Program V2.1 which provides "local" or "remote" bridge functions.

The primary functions performed by the local bridge are:

- Connecting multiple rings into a single logical network by transferring frames between the two rings to which the bridge is connected.

  A major extension to the source routing mechanism is provided by the automatic configuration and maintenance capability of the single-route broadcast path, as explained in "Automatic Single-Route Broadcast" on page 31.
- Displaying ring status and fault domain details for hard and soft error conditions
- Maintaining and displaying performance statistics
- Providing network management capability by sending notifications and reports to IBM LAN Manager Version 2.0.

Communication across the bridge is transparent to applications written to the IEEE 802.2 standard logical link control interface or higher using source routing.

5.2 IBM Token-Ring Network Bridge Program V2.0 Architecture

Figure 24 on page 48 shows the general bridge structure of IBM Token-Ring Network Bridge Program V2.0 and its interfaces with the IBM Token-Ring Network Adapters. As with any IBM bridge program, IBM Token-Ring Network Bridge Program V2.0 includes the adapter handler code and Direct (MAC) and LLC application programming interfaces.

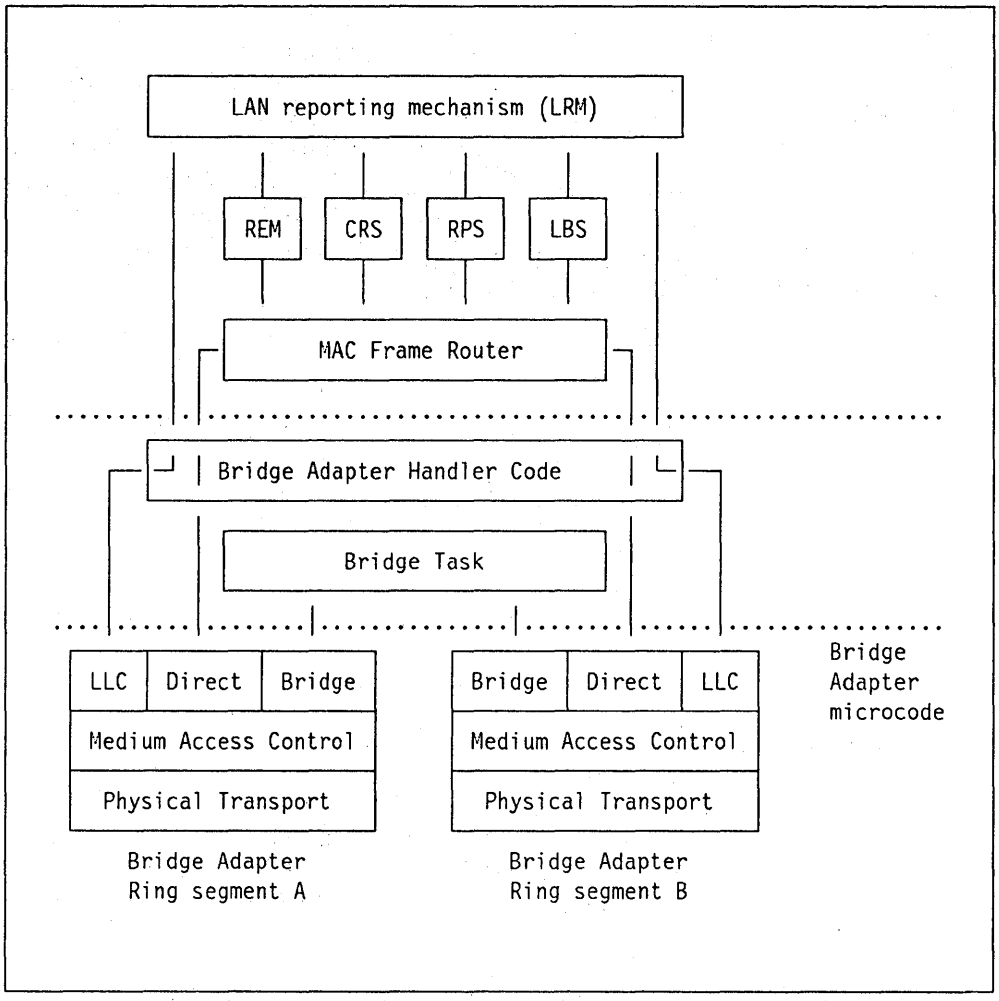

Figure 24. IBM Token-Ring Network Bridge Program V2.0 - General Bridge Structure

Bridge processing by the LAN Bridge Server function now includes *single-route broadcast path maintenance*.

5.3 New Bridge Functions and Parameters

Several new bridge features have been introduced with the three recent bridge programs. The most important are:

- 16 Mbps support
- Automatic single-route broadcast
- LAN Manager support improvements
- Largest frame size increase.

5.3.1 16 Mbps and 4Mbps Support

IBM Token-Ring Network Bridge Program Version 2.0 provides bridging capability between two rings of an IBM Token-Ring Network operating at 4 or 16 Mbps. For example, one bridge's adapter can operate at 16Mbps while the other operates at 4Mbps, allowing different ring speeds in the network to match the different performance requirements.

Backbone rings are ideal candidates for the advantages of 16Mbps speed. User rings with heavy traffic such as image processing should also migrate to the 16Mbps speed.

Like all other stations on a segment, the bridge will run the insertion tests on both segments it is attached to before it can complete opening its adapters. All adapters on a given segment must run at the same speed. If you insert a bridge adapter on a segment with a speed different from the existing adapters, the bridge initialization will fail with an error message, followed by the following:

"Bridge will retry adapter open command until it is successful"

These retries will create temporary beaconing situations on the corresponding ring(s), with alerts sent to the LAN Manager.

The reason the bridge tries its open command indefinitely is that if the ring beaconing is due to another adapter on the ring, the bridge will be able to come back to full operation without user intervention. As a wrong bridge adapter speed should be easily detected at bridge installation setup, it was decided to retry the bridge open command to solve the more probable case of another adapter creating a beaconing situation.

5.3.2 Automatic Single-Route Broadcast

IBM Token-Ring Network Bridge Program Version 2.0 can be configured locally or by an IBM LAN Manager Version 2.0 to communicate with other active bridges in the network to automatically configure the network single-route broadcast path. If this capability is desired, all bridges in the network must be configured to participate in the protocols. Because earlier versions of the bridge program do not support automatic configuration, they must be upgraded to Version 2.0 to utilize this new function.

The bridges will reconfigure the single-route broadcast path to accommodate network changes caused by a new bridge being activated or an active bridge being shutdown.

It is important to notice that single-route broadcast path maintenance does not affect the way in which a bridge processes all-routes broadcast or non-broadcast frames.

We strongly recommend use of the automatic single-route broadcast technique in all complex network configurations (meshed, dual backbone,..,) as explained in "Automatic Single-Route Broadcast" on page 31.

The automatic single-route broadcast algorithm and the new bridge parameters (bridge label and path cost) are explained in "Automatic Single-Route Broadcast" on page 31.

5.3.3 LAN Manager Support

IBM Token-Ring Network Bridge Program V2.0 interfaces with up to four IBM LAN Manager V2.0 stations to provide all network management functions supported by the IBM Token-Ring Bridge Program V.1 plus additional capabilities (see "LAN Manager Support" on page 20 for more details on the

communications between IBM Token-Ring Network Bridge Program V2.0 and IBM LAN Manager V2.0).

With IBM Token-Ring Network Bridge Program V2.0 (as well as IBM Token-Ring Network Bridge Program V2.1 and IBM PC Network Bridge Program), the controlling LAN Manager (V2.0) or the NetView operator now has such facilities as:

- Set/reset hop count
- Set/reset ring number
- Set/reset bridge number
- Set/reset single-route broadcast selection mode

In addition, a bridge executing IBM Token-Ring Network Bridge Program V2.0 can be defined as a critical resource in the IBM LAN Manager V2.0 and can be monitored by the IBM LAN Manager V2.0 This facility should only be used for bridges which are not linked to the LAN Manager, in order to get an alert in case of a bridge failure. Bridges which are linked to the LAN Manager will automatically trigger an alert in case of a loss of the session with the LAN Manager V2.0.

At installation time or during problem determination procedures, it can be useful to define bridges as critical resources, especially if the links with the LAN Manager can't be established.

5.3.4 Largest Frame Size

When you display the IBM Token-Ring Network Bridge Program V2.0 configuration, you will see a parameter called largest frame size (LFS). The frame size does not include the MAC header. This parameter cannot be modified by the bridge operator using the Token-Ring Bridge Program Version 2.0. It is set by the bridge program based on the installed adapters. The frame forwarding size of the most limiting adapter will be used for local bridges:

```
IBM Token-Ring Network Adapter II or
IBM Token-Ring Network Adapter /A                          2052 bytes
IBM PC Network Adapter II/A                                2052 bytes
IBM Token-Ring Network 16/4 Adapter or
IBM Token-Ring Network 16/4 Adapter /A (16K RAM paging)    2052 bytes
IBM Token-Ring Network 16/4 Adapter or
IBM Token-Ring Network 16/4 Adapter /A (32K RAM paging)    4472 bytes
IBM Token-Ring Network 16/4 Adapter or
IBM Token-Ring Network 16/4 Adapter /A (64K RAM paging)    8144 bytes
```

Figure 25. Largest Frame Size Supported by Bridge Adapters

The bridge will inspect the largest frame size (LFS) information contained in the routing information field of the broadcast frames it receives. If the requested LFS is greater than the LFS that can be provided by this bridge, the bridge adapter will replace the requested LFS with its LFS before passing the frame to the attached product.

5.4 Installation/Utilization Guidelines

The bridge installation process for all the new bridge products described in this document is essentially the same. It is fully described in the associated bridge user's guide manuals. Basically, there are three major steps:

1. Bridge planning
2. Bridge physical installation
3. LAN Manager bridge definition.

Some additional specific steps are required for the remote bridge installation. See "Installation/Utilization Guidelines" on page 79 for these specific steps and considerations.

5.4.1 Bridge Planning

Before you physically install the hardware and bridge software, you should perform the following tasks:

1. Assign ring and bridge numbers.

   Draw your whole network on a chart and assign segment numbers and bridge numbers. Remember segment (ring or bus) numbers must be unique. Bridge numbers need not be unique, but are desirable for ease of management, provided the network uses less than sixteen bridges.

2. Assign bridge locally administered addresses (optional).

   If you have decided to use locally administered addresses for stations in your network, you should assign logical addresses to each bridge adapter.

3. Determine bridge parameters like hop count limit and automatic single-route broadcast. The hop count limit parameter is discussed in details in "Hop Count Limit" on page 39. If you decide to use the automatic single-route broadcast facility (which is highly recommended), remember that all bridges must participate in the automatic single-route broadcast process.

4. Fill out the bridge planning chart

   You will find the bridge planning charts in the associated bridge program user. There is one planning chart per bridge station. You should write down in the bridge planning charts the bridge characteristics as well as the bridge parameters you are going to use for each bridge. There are several sections in these charts:

   a. Physical connections

      You specify here physical location and cabling information.

   b. Bridge installation parameters

      These parameters will specify the bridge adapters' characteristics such as adapters' name, speed and address.

   c. Bridge configuration parameters

      These parameters specify all bridge configuration parameters such as ring and bridge numbers, hop count limit, and automatic single-route broadcast. See "Bridge Parameters" on page 23 for the list of these parameters.

d. Communication adapter configuration parameters

   This section is for the remote bridge function only and specifies the communication adapter characteristics (line speed, electrical interface and communication adapter size).

5.4.2 Bridge Physical Installation

After the planning steps are finished, you should perform the following tasks on each bridge:

1. Install the adapters and configure them, according to your bridge planning chart and the bridge user's guide. Write down the physical addresses of the adapters if you don't use locally administered addresses. These addresses will be used during LAN Manager bridges definition.
2. Install DOS and the bridge software

   You need at least DOS 3.3 or DOS 4.0. Installing the bridge software is a simple task, as described in the corresponding product user's guide.
3. Configure the bridge parameters according to the bridge planning chart

   It is very important at this level that the same segment is always defined with the same number in all bridges connected to it. Otherwise, you will not be able to start the bridges (you will get the message: "frame forwarding not active").

   In addition, path cost and label must be defined correctly at the bridge, as the LAN Manager cannot modify them.

   Most other parameters can be changed from the LAN Manager at a later time, which is very useful when the bridges are spread throughout large buildings.
4. Start the bridge program.

5.4.3 LAN Manager Bridges Definition

When all bridges are operational, you should go to the LAN Manager station with your bridge planning charts to define and control them from the LAN manager station (or from NetView). You should then perform the following tasks:

1. Define the bridges' adapters' symbolic names and addresses for each bridge.

   You must specify the locally administered addresses or the physical addresses specified at bridge installation.
2. Define the bridges, using the previously defined adapters' symbolic names.
3. Link the bridges from the LAN Manager.
4. Change the bridge parameters if necessary using the "bridge configure" option.

6. IBM Token-Ring Network Bridge Program V2.1 "Local" or "Remote" Bridge Function

6.1 Remote Bridge Overview

IBM Token-Ring Network Bridge Program Version 2.1 can be configured to provide local or remote bridge function. The local bridge configuration provides all of the functions and capabilities available in IBM Token-Ring Network Bridge Program Version 2.0 (see "IBM Token-Ring Network Bridge Program V2.0 (Local Bridge)" on page 47). Note that IBM Token-Ring Network Bridge Program V2.0 has been withdrawn from marketing and has been replaced by IBM Token-Ring Network Bridge Program V2.1. To avoid repetition of previous descriptions, refer to "IBM Bridge Programs Overview" on page 17 and "IBM Token-Ring Network Bridge Program V2.0 (Local Bridge)" on page 47 for the "local" bridge function of IBM Token-Ring Network Bridge Program V2.1.

In this document a "remote bridge" designates a bridge using the remote bridge function.

The remote bridge configuration extends the access of LAN resources between geographically separated rings by supporting a dedicated communications link between the bridge components on each ring. This remote bridge function can enhance any-to-any connectivity between geographically dispersed workstations, file servers, hosts, and applications.

The remote bridge configuration uses one PC or PS/2 [10] with bridge software at each end of a point-to-point, leased line. When the program is configured as a remote bridge, frames are transferred in full duplex between two rings over the teleprocessing (TP) line at speeds from 9.6 Kbps to 1.344 Mbps.

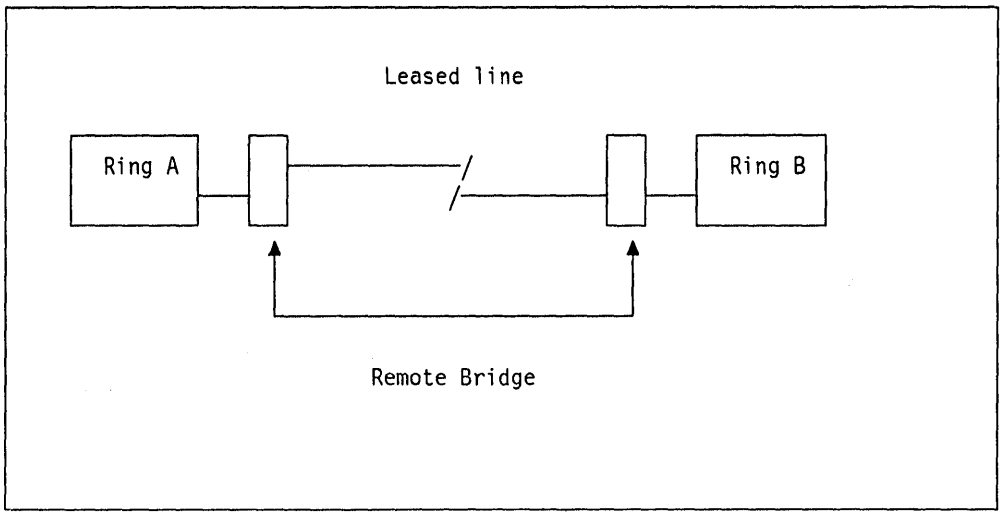

Figure 26. Remote Bridge Scheme

[10] PS/2 is a registered trademark of the International Business Machines Corporation

6.1.1 Line Speeds and Interfaces Supported

A remote bridge (or split bridge) consists of two bridge halves connected by a TP line. Each bridge half requires a token-ring network adapter attached to a LAN segment and a communications adapter attached to the telecommunication network.

As telecommunication network environments vary from country to country, you should refer to the specific announcement letter in your country to verify availability of the following ways 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 1.344 Mbps.
  - X.21 bis/CCITT V.24 at 9.6 Kbps to 19.2 Kbps.
  - X.21 bis/CCITT V.35 at 9.6 Kbps to 1.344 Mbps.
  - X.21 (leased only) at 9.6 Kbps to 64 Kbps.
- Via a multiplexor, such as the Integrated Digital Network Exchange (IDNX) Models 20, 40 and 70 through:
  - the USD or HSD 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 TP link. Valid options are 1 (RS-232), 2 (V.35, default value) 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 Token-Ring Network Bridge Program Version 2.1) can be used.[11]

[11] Note that X.25 circuits are not supported by the remote bridge function

Interface	IBM Realtime Interface Co-Processor	IBM X.25 Interface Co-Processor/2
RS-232C/V.24	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*
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 64 Kbps

Figure 27. Possible interfaces and speeds. * 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.

It should be noted that the IBM X.25 Interface Co-Processor/2 adapter is required for use in the IBM Personal System/2 Models 50, 60, 70, and 80, and the IBM Realtime Interface Co-Processor adapter is required for use in the IBM Personal Computer AT, IBM 7531 and 7532 Industrial Computers, and IBM Personal System/2 Model 30.

6.1.1.1 ISDN Support

Although the remote bridge requires a leased line to operate, it is possible to use the remote bridge functions across a switched network such as ISDN (Integrated Services Digital Network) as long as "the connection looks like a leased line".

In the case of ISDN, it should be possible (depending on your country telecommunications facilities) to use the remote bridge function at 64Kbps once the communication is established between two 7820 connected to the half-bridges using the V.35 interface. The 7820 must be configured in direct call mode to provide the leased line appearance required by the bridge. (Refer to the appropriate 7820 documentation for information on how to establish the communication between the two 7820s)

The IBM X.25 Interface Co-Processor/2 communication adapter with the V.35 cable option must be used in the remote bridge to support this attachment.

One advantage of using the ISDN facility is that you could connect to a backup site if the normal remote site experiences major problems.

An example of a remote bridge configuration operating across an ISDN network is shown in Figure 28 on page 56.

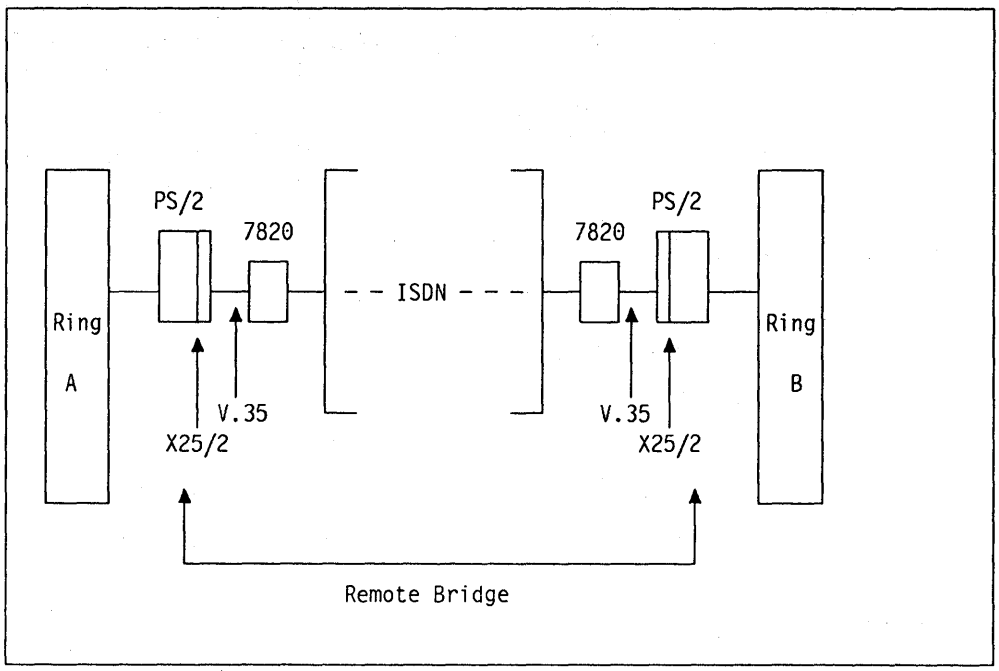

Figure 28. Remote Bridge Function across an ISDN network

6.1.2 Half-Bridge Components

6.1.2.1 Hardware Requirements

IBM Token-Ring Network Bridge Program V2.1 requires two dedicated IBM Personal System/2 Model 30, 50, 60, 70 or 80, or a PC AT, or an Industrial Computer 7531 or 7532, each with 512 KBytes of memory and a 720 KBytes or 1.2 Mbytes 3.5 inch diskette drive.

Each device attaches to a ring segment by means of an IBM Token-Ring Network adapter:

- IBM Token-Ring Network 16/4 Adapter/A
- IBM Token-Ring Network 16/4 Adapter
- IBM Token-Ring Network Adapter/A
- IBM Token-Ring Network PC Adapter II.

In addition, both bridge halves must be equipped with a matching communications adapter to connect to the TP line. A single port of the following communications adapters can be used on each split bridge device:

- IBM X.25 Interface Co-processor/2 (for Micro-Channel PS/2's) with one of the following cable options:
  - Cable option V.24
  - Cable option V.35
  - Cable option X.21
- IBM Realtime Interface Co-processor with 512 KBytes (for PC's and PS/2's with a PC Bus) with the supporting features for one of the following interfaces:

- EIA RS-232C/CCITT V.24
  EIA RS-232C/CCITT V.24 Interface Board with RS-232C Modem Attach Interface Cable or RS-232C Direct Attach Interface Cable.
- CCITT V.35
  CCITT V.35 Interface Board and CCITT V.35 Interface Cable.

6.1.2.2 Software Requirements

IBM Token-Ring Network Bridge Program V2.1 requires for each of the bridge halves:

- IBM PC/DOS 3.3 or 4.0
- IBM Token-Ring Network Bridge Program V2.1
- IBM Realtime Interface Co-processor DOS Support Version 1.0 or higher.

6.1.3 IBM Token-Ring Network Bridge Program V2.1 Architecture

Figure 29 on page 58 shows the general bridge structure of IBM Token-Ring Network Bridge Program V2.1 and its interfaces with the IBM Token-Ring Network Adapter and TP communications adapter. The other bridge half (not represented) has exactly the same structure.

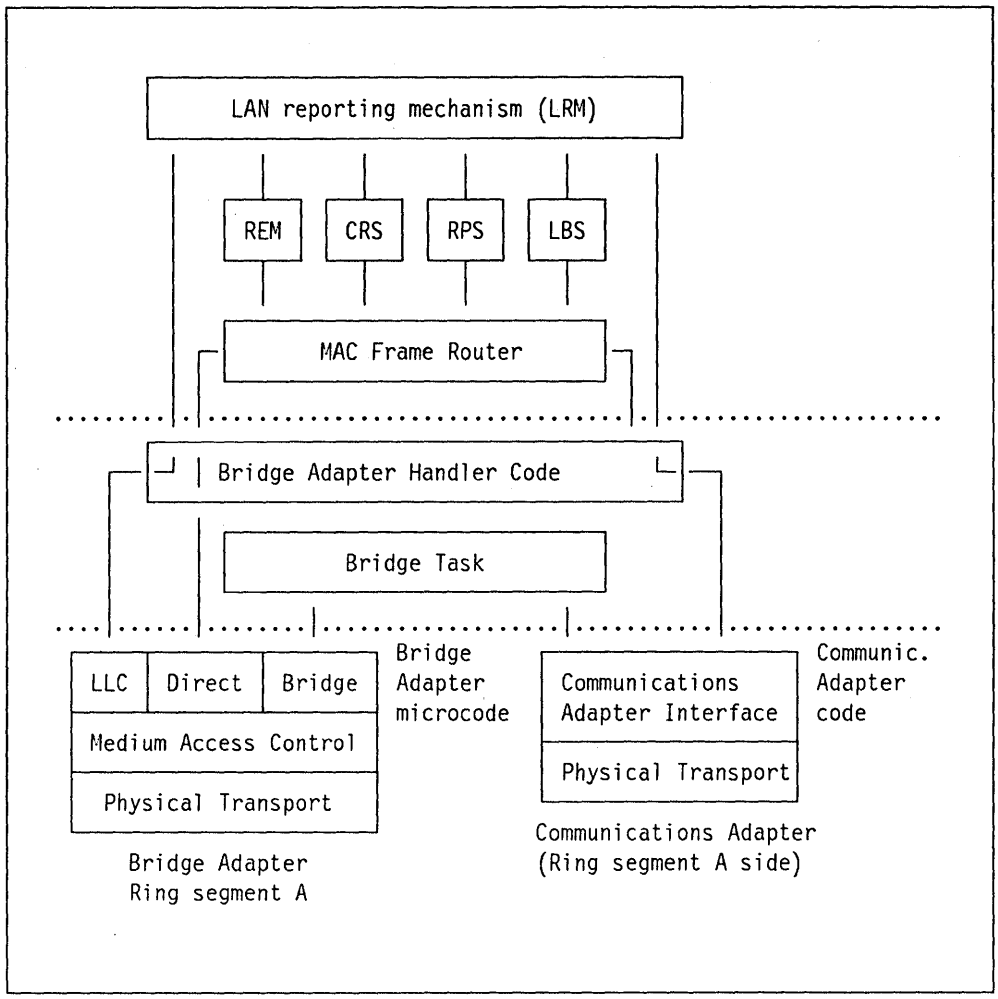

Figure 29. IBM Token-Ring Network Bridge Program V2.1 - General Bridge Structure

6.1.4 LAN Manager Reporting

From a LAN management perspective, the remote bridge is defined and operates like a local bridge. Please refer to "LAN Manager Support" on page 20 for more detail on the LAN management reporting mechanism of a bridge.

As explained in "Installation/Utilization Guidelines" on page 79 configuration information is passed across the communications link directly during bridge initialization. Both bridge halves verify that the common information is consistent (for instance bridge number and ring segment number). As the configuration file must exist at only one side of the remote bridge, it is recommended to have the "primary" side on the central site side whenever possible to make changes easier.

In addition, as explained in "Recovery and Problem Determination" on page 80, both sides of a remote bridge will locally display error messages and terminate their links with the LAN manager in case of a TP line failure.

Finally, when any of the operational bridge parameters is changed by the IBM LAN Manager V2.0, both bridge halves will be updated. The new bridge

configuration parameter values will be permanently recorded by writing them to the ECCPARMS.BIN file on the primary bridge station (half).

6.2 The Enterprise LAN Benefits

The remote bridge introduces many new alternatives for connectivity and operations. LAN functions and benefits which were limited to a building or a campus can now be extended to the enterprise level. Some major benefits of this new approach include:

- Protocol independence
- Integration into an APPN network
- Direct access to remote servers
- PC software distribution
- Remote help desk

6.2.1 Protocol Independence

The remote bridge is a MAC layer bridge which provides **transparent communication across the bridge** to applications written to the IEEE 802.2 logical link control interface using source routing. Therefore, it is possible to establish **direct peer to peer communications across the remote bridge** using a variety of protocols such as:

- IEEE 802.2 LLC (using source routing)
- NETBIOS
- SNA protocols
- TCP/IP

There is no difference from a connectivity standpoint between a remote bridge and a local bridge, except that time-outs and frame sizes might require adjustment to accommodate the link speed.

For the remote bridge function, the default largest frame sizes are based on the link speed, but are settable by the operator.

An example of the advantages provided by the remote bridge configuration is shown in the following scenario. In this example, a company was using a traditional SNA gateway (for example a 3174) connected to a leased line and an IBM LAN Asynchronous Communication Server (ACS) connected to two switched lines to access resources located in an SNA Host and two ASCII Hosts, as shown in Figure 30 on page 60.

The IBM LAN ACS servers provide an efficient way to support link access to ASCII hosts and reduce ASCII devices wiring requirements, as they use the LAN as a transport mechanism inside the building.

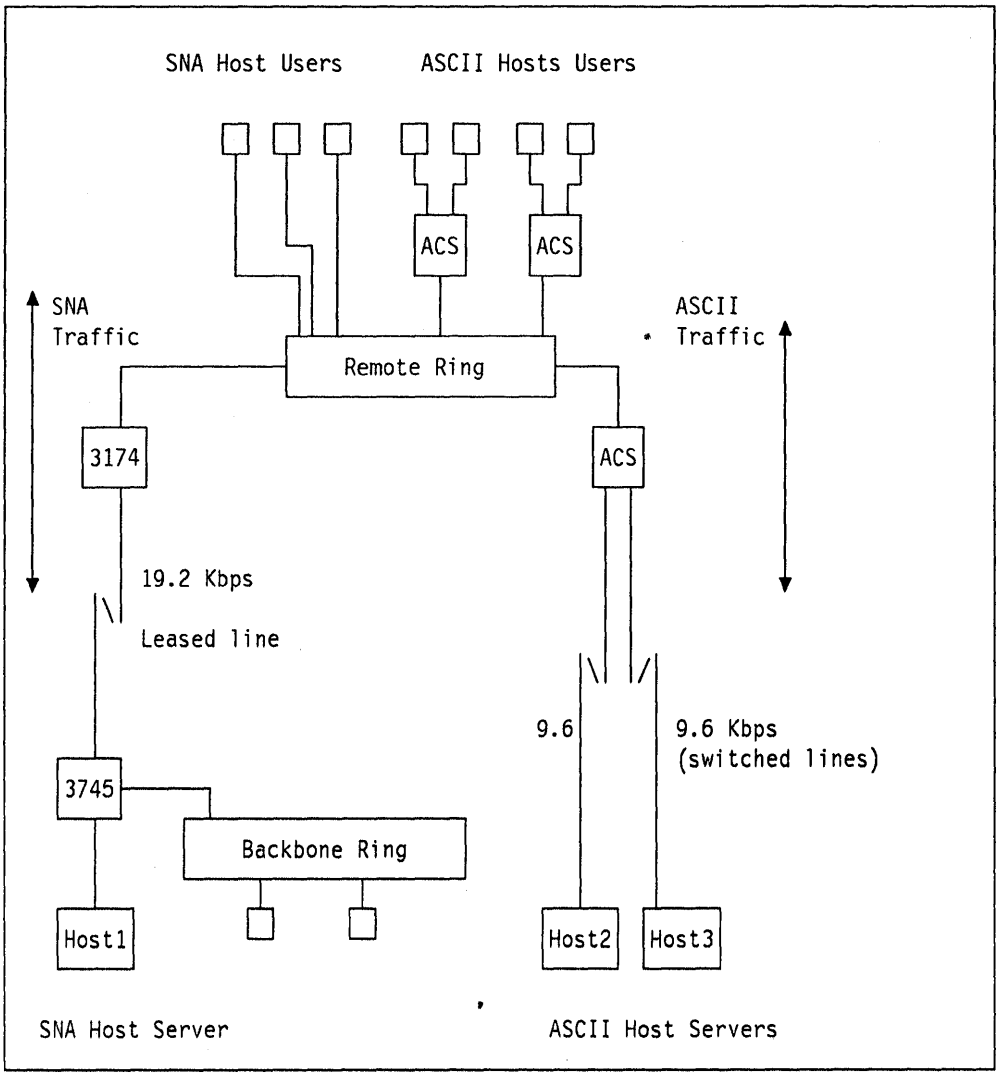

Figure 30. SNA and ASCII Traffic on Separate Links

With the new remote bridge function, it is possible to multiplex the SNA and ASCII traffic on a single leased line. (Note that the ASCII traffic is "enveloped" by the ACS servers in NETBIOS frames).

An example of a new configuration using a remote bridge with a single 56Kbps leased line is shown in Figure 31.

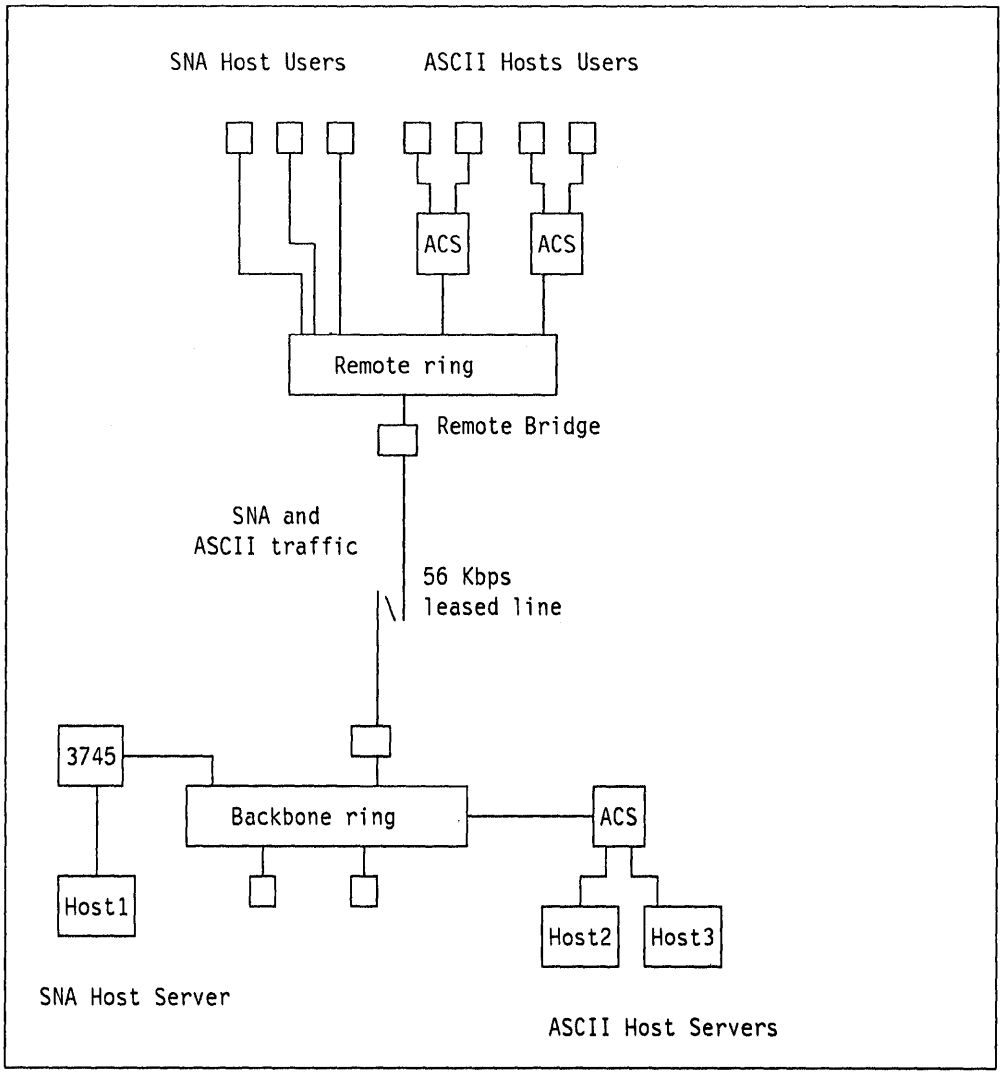

Figure 31. SNA and ASCII Traffic via the Remote Bridge

Some advantages of that new configuration include:

- Higher availability of the TP line for ASCII applications

  As the ASCII flow is multiplexed over the leased line, ASCII terminals won't have to wait for an available switched port on the ACS server. The maximum number of concurrent ASCII "sessions" will not be limited anymore by the number of switched lines and modems on the ACS server, but only by the number of ASCII ports on the remote ASCII hosts. In addition, the quality of the leased line will usually be better than switched lines.

- Better performance of the TP line for ASCII and SNA applications

  The maximum TP line speed supported on the public switched network (and by ACS) is 19.2 Kbps but many switched lines operate today at speeds below or equal to 9.6 Kbps. For a remote 3174/1R gateway, the maximum line speed supported is 64Kbps.

The remote bridge removes this constraint by supporting TP links at speeds up to 1.344 Mbps. As a result, ASCII applications could now use the highest speed supported by the stations instead of being limited in our example by the modem speed (9.6 Kbps). Likewise, SNA applications can take advantage of the faster leased line, especially when the ASCII traffic is low. The 56 Kbps line will provide a better response time or support additional terminals. (Note however that dependent upon traffic, a lower speed might be sufficient for the SNA and ASCII traffic.)

- Lower overall network cost

  It is very likely in this example that the cost of a remote bridge (and a single fast leased line) is less than the total cost of an SNA gateway (such as a remote 3174 or 3720), one leased line and two switched lines with the appropriate number of modems.

There is now another approach to providing connectivity for the ASCII devices and hosts in this scenario. With the new 3174 configuration support S release 5, the ASCII emulation adapter (AEA) can coexist with the gateway feature of the 3174. So an alternative is to attach all ASCII devices and hosts to the 3174 gateway using the AEA feature as illustrated in Figure 32 on page 63.

In this case, no ACS servers are required. This requires however a leased or switched link between each ASCII device/host and the 3174 gateway, as shown in Figure 32 on page 63, which may be difficult or costly to provide in large buildings where distances between the ASCII devices and the 3174 gateway are considerable.

On the other hand, the 3174 approach has the advantage of supporting ASCII passthru between the devices and hosts or of supporting 3270 protocol conversion. (3270 coax attached to the 3174 can also access ASCII ports with reverse protocol conversion).

Finally, it is also possible to have a mixed scenario and to use the IBM LAN ACS in conjunction with the 3174 gateway equipped with the asynchronous emulation adapter.

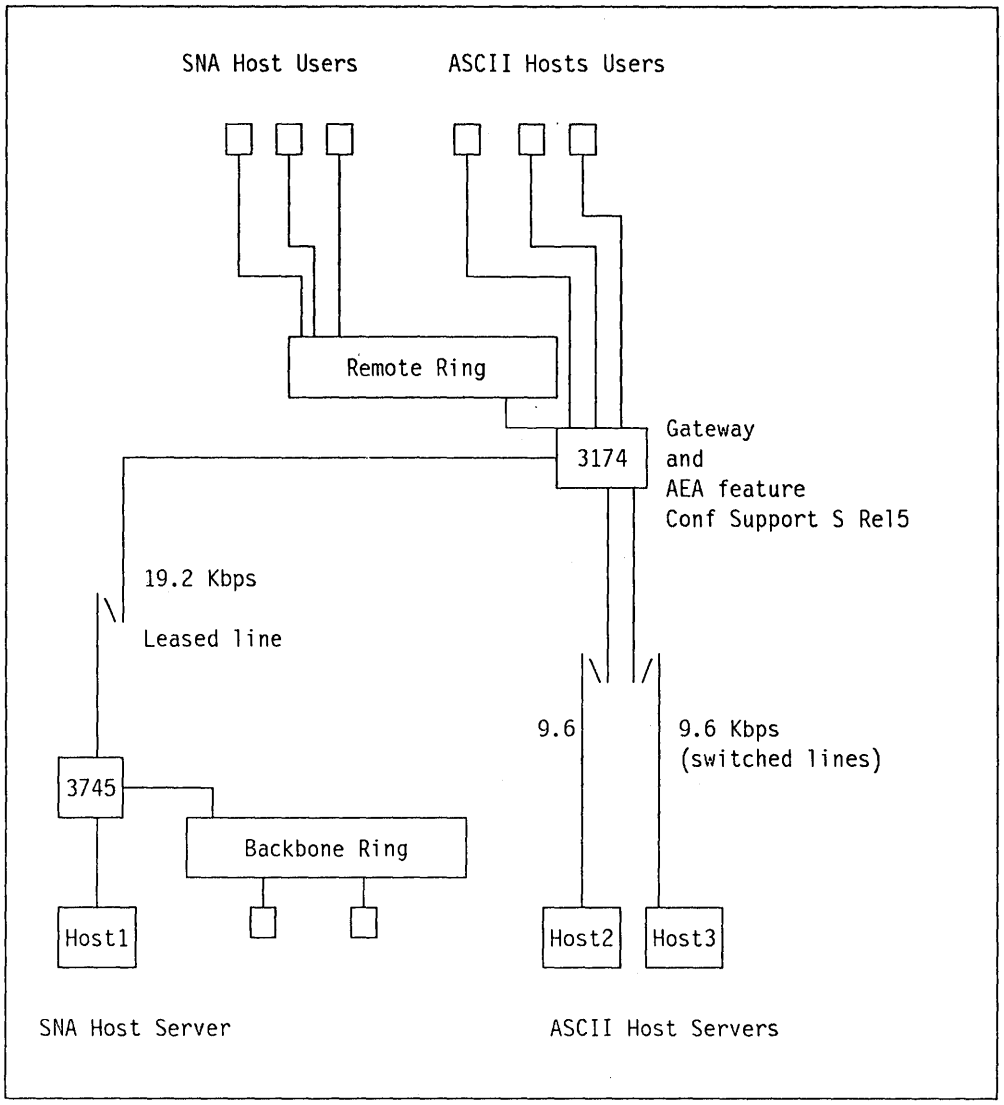

Figure 32. SNA and ASCII Traffic via 3174 LAN Gateway

6.2.2 Integration into an APPN network

A consequence of transparent communication across the remote bridge is that it is easy to integrate remote rings into an APPN network without restriction.

The token-ring is fully supported by APPN nodes and therefore it is possible to use all LU6.2 functions and APPN facilities between APPN nodes communicating across a remote bridge.

There are no restrictions such as the ones you might have with SNA T2.0 node gateways or even composite nodes such as communication controllers.

6.2.3 Direct Access to Remote Servers

One of the major benefits of the remote bridge is that PC users can access remote NETBIOS servers such as IBM PC LAN Program V1.3 or IBM OS/2 LAN Server. With the traditional SNA gateway approach, users on remote rings could access 3270 type host applications but could not access central NETBIOS servers' resources. In the following example, U1 on remote ring 1 can access resources shared by the central LAN server (station LS), as if U1 was connected to the backbone ring.

Likewise, it is possible for stations connected to the backbone ring (or other local segments) to access the remote ring servers' resources. For example, station HD (help desk) may directly access resources shared by servers S1 and S2.

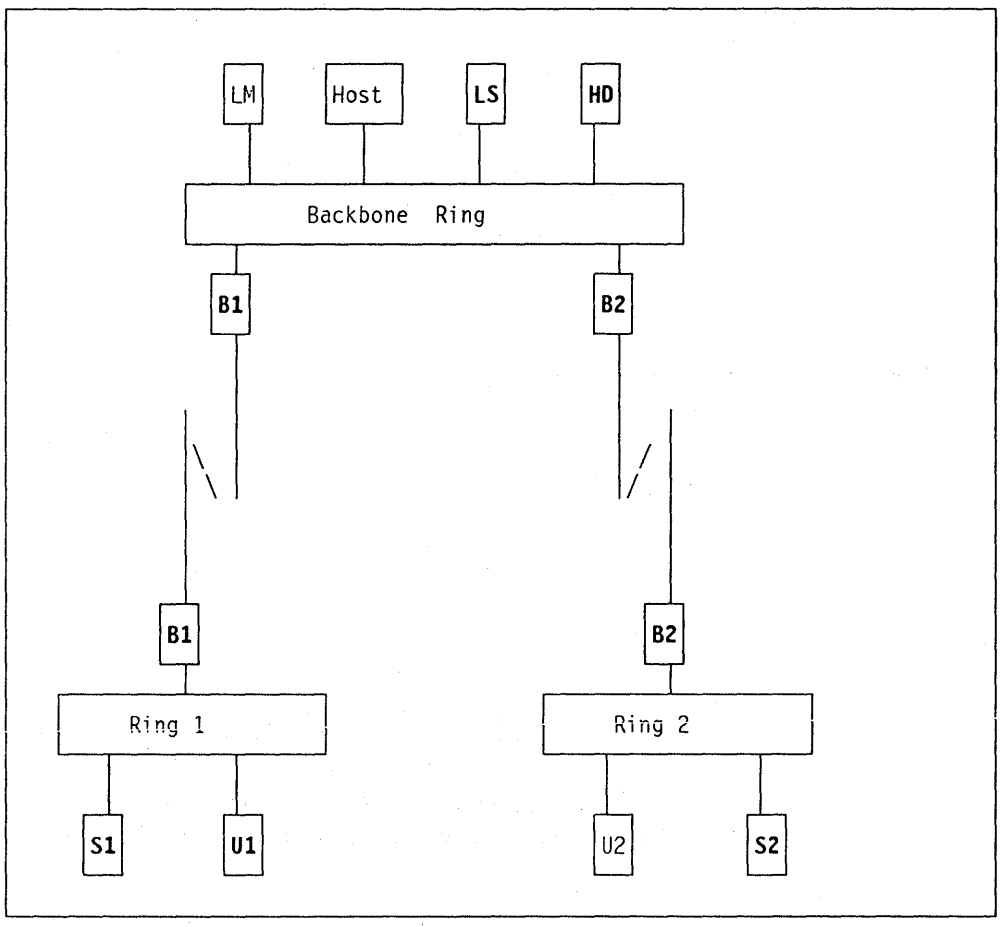

Figure 33. Remote Bridge Configuration

6.2.4 Software Distribution

When the number of the LAN programmable workstations increases, PC software distribution and maintenance can be very time consuming and become a major concern if not properly planned. There are several techniques and IBM products such as NetView/DM or DSX to automate and fully manage software distribution to remote servers and remote stations from a central host.

If these techniques are not appropriate to your environment (for example if you have no host), you can consider another alternative based on PC to server or

server to PC file transfer. Appropriate DOS or OS/2 procedures can be developed to automate the process using the DOS (or Operating System/2 Extended Edition V1.1) COPY function.

With the remote bridge function, procedures which were initially developed for local rings might also be used for the remote rings, subject to time-out or performance constraints resulting from the link speed.

6.2.5 Help Desk New Facilities

Help desk functions are essential at the local and enterprise level of the network to provide assistance to the end users. A difficulty inherent to remote buildings is that help desk personnel do not always have the capability to check the precise status and contents of remote resources such as PC servers. If an end user has difficulty accessing a server on a remote LAN, problem determination procedures can be greatly simplified if help desk personnel can access that particular server to check its status and directory contents.

This is now possible with the remote bridge facility. With the correct password and access level, help desk personnel will be able to access the information they need and take the appropriate actions to correct the situation. As shown in Figure 33 on page 64, station HD (help desk) may access all resources and servers in the network, including remote servers S1 and S2.

6.3 Performance Considerations

Because the performance of the remote bridge is constrained by the TP line, network traffic flow should be considered when selecting the line speed. For example, the broadcast traffic generated by a large network could cause severe performance problems at the 9.6 Kbps TP line speed. Because the effective throughput at the slower line speeds is low, the network traffic flow is a determining factor in the number of concurrent sessions that can be supported by the remote bridge.

Additional information on restrictions on line speed, frame size, number of active concurrent sessions through the remote bridge, and network configuration will be provided when the program becomes available.

The common rational behind these different restrictions is to avoid end-to-end session time-out due to excessive delays experienced while crossing the remote bridge.

When using connection-oriented service across a remote bridge, the following factors must be considered to determine whether a specific application will work:

- TP line speed and quality
- Link protocol parameters (timers, frame size, window size)
- Application parameters (timers, frame size)
- Remote bridge parameters (communication adapter transmit buffer size, frame size)
- Network traffic.

As far as link protocol parameters are concerned, time-outs could be experienced at different levels such as LLC, NETBIOS or higher application layers.

To compensate for timing delays created by the TP line, parameters such as frame size and protocol timer values may need to be adjusted in the application. Applications not allowing such adjustments may not be able to communicate across the remote bridge at the slower line speeds.

For lower layers like LLC and NETBIOS, it is possible to adjust NETBIOS and LLC acknowledgment timers (such as the T1 timer) and window values by modifying the corresponding parameters in IBM LAN Support Program V1.10 or Operating System/2 Extended Edition V1.1 configuration panels. Likewise you can adjust the LLC window size and T1 parameter values in VTAM and NCP definitions. For example, you could set the window size to one or raise the LLC T1 value in order to accommodate the additional delay due to the line speed.

However, you should realize that a low window size will decrease performance for file transfer types of application and that a high T1 value can also result in performance degradation due to link error recovery. You should also check that the changes you want to make are possible and consistent on *both ends* for proper operation of connection-oriented protocols.

For example, LAN station adapters maintain a response timer (T1) to detect failure of a link station to receive a required response or acknowledgement from the partner link station. The recommended time value for this timer, based upon typical LAN bandwidth and quality is one second. However, the link stations are not aware that the link between the halves of a bridge may be slower or less reliable. Thus, unless the link provides adequate reliability and capacity, time-outs between link stations are likely to occur.

The response timer value (T1) can be customized in the LAN Support Program Version 1.1 or in OS/2 Extended Edition Communications Manager. For link speeds from 9.6 Kbps to 64 Kbps it should be set to a value greater than 1.6 seconds for each remote bridge in the path between partner stations. Thus, if two stations will be connected by a single remote bridge pair (one link), then the T1 value in each station should be set to greater than 1.6 seconds. If they are to be connected via cascaded bridges with two bridges in the path, then they should be customized with values greater than 3.2 seconds.

Another way to reduce the probability of time-out is to reduce the maximum frame size that the partner link stations will transmit, thereby reducing the probability of queuing and congestion at the bridge. This may increase the processing requirements in the end stations because the same amount of traffic will require transmission of more frames. Note, however, that if the bit error rate of the communications link is even slightly raised, a reduction of frame size will have a significant impact on session availability and link throughput. See Figure 34 on page 67.

Error Rate	Frame Size	Ratio of Lost Frames	Theoretical Throughput
1 in 10**4	100 Bytes	0.07688	0.86773
	516 Bytes	0.33821	0.65410
	1470 Bytes	0.69149	0.30725
	2052 Bytes	0.8063	0.1937
1 in 10**5	100 Bytes	0.00797	0.93251
	516 Bytes	0.04044	0.94840
	1470 Bytes	0.11095	0.88542
	2052 Bytes	0.15139	0.84613
1 in 10**6	100 Bytes	0.000800	0.93925
	516 Bytes	0.004120	0.98430
	1470 Bytes	0.011691	0.98428
	2052 Bytes	0.016282	0.98084
1 in 10**7	100 Bytes	0.000080	0.93992
	516 Bytes	0.000413	0.98796
	1470 Bytes	0.001175	0.99418
	2052 Bytes	0.001640	0.99544

Figure 34. Effect of Frame Size With Different Line Error Rates.

The source routing support provided by the Token-Ring Bridge Program allows paths between stations of up to seven MAC level bridges. In larger LANs or with heavier traffic, there is a potential for congestion frame loss at every bridge. This is increased to the extent that one or more of the bridges in question is a remote bridge with a slower speed link between the two halves. To ensure that such loss does not impact the reliability and manageability of a LAN, each bridge maintains counters and reports when defined traffic threshold failure values are exceeded.

Each bridge keeps a count of the number of frames forwarded and the number of frames lost due to congestion or other causes. **Bridge performance reporting** has been enhanced to reflect remote bridge specific information. As mentioned earlier, a special *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 TP communications link.

If the ratio of frames lost to frames forwarded exceeds a frame loss threshold (default = 0.0010), the bridge sends a warning message to the LAN Manager. Thus, in the worst case of seven bridges with fourteen hops, (seven over and seven back) there could be a ratio of 0.014 frames lost (0.0010*14) before the LAN Manager would be informed. This warning does not reflect loss of

sessions, which may be preserved by retries which are attempted by Logical Link Control. Up to eight retries will be attempted, thus reducing the probability of failing to deliver a message (and consequent session loss) to 1.5E-15 (0.014*8).

This retry support is especially valuable when using a relatively low speed transmission link, since token-ring adapter buffering capacity may be insufficient to store incoming frames from the ring segment before they can be transmitted to the other bridge half. In this case, the *Frame Not Received (adapter congested)* counter may reach its threshold much faster than it would for the same segments and same traffic in a normal bridge configuration.

To determine appropriate values to define for remote bridge configurations, an estimate can be made of the probability of path loss based upon the size of the LAN for different threshold values and retry counts. The calculation is based upon the following assumptions:

1. Frame loss = frame size * line error rate (under 1 %) or

   Frame loss = 1 - exp (- frame size * line error rate) over 1%
2. Bridge loss = congestion loss + TP Link loss
3. Path loss = sum of the bridge losses
4. Frame failure rate = path loss ** retry count
5. Frame success rate = 1 - (path loss ** retry count)
6. Link success rate = (1 - (path loss ** retry count)) ** frame count

From results such as those illustrated in Figure 35 on page 69, tradeoffs in values can be identified and parameters set with greater confidence.

For example, assume a single backbone network using the default retry count of eight, and acceptability of a 0.999999 link reliability rate during transfer of 10**6 frames. The 316 lost frames per 10000 transferred can be divided evenly across four bridge hops, and the traffic thresholds set for 79 with confidence. On the other hand, if all traffic is directed to a particular station such as a host or server on the backbone, then the thresholds could be set for 158.

Assume for another example, that a path has two remote bridges, each with a link bit error rate of 1 bit in 10**5. The defaults for retry produce a frame loss rate of 2048*8/10**5*2 or 0.32768 over the route. With a retry of eight, this represents a link failure in about 10000 frames or perhaps an hour and a half. If the frame size had been set to 512 bytes, the path loss would be about 0.08192, or a link failure in about 10 years of one shift per work day. Another way to reduce the risk is to increase the retry count to 14 or more. This would increase the time between expected failures to about six months.

Notice that the effect of reducing the frame size is that the real throughput is increased. Increasing the retry count overcomes probable link loss by spending more time, and consequently reduces real throughput.

Retry Count	Link Success	Path Loss For 10,000 Frames	Path Loss For 10**6 Frames	Path Loss For 10**8 Frames
8	0.99999999	0.03162	0.01013	0.01013
8	0.99999990	0.04217	0.02371	0.01333
8	0.99999900	0.05623	0.03162	0.01778
8	0.99999000	0.07499	0.04217	0.02371
8	0.99990000	0.10000	0.05623	0.03162
8	0.99900000	0.13360	0.07499	0.04217
10	0.99999999	0.39812	0.03981	0.02354
10	0.99999990	0.50116	0.05012	0.03162
10	0.99999900	0.63064	0.06310	0.03981
10	0.99999000	0.79040	0.07943	0.05120
10	0.99990000	0.95517	0.10000	0.06310
10	0.99900000	0.999995	0.12590	0.07944
12	0.99999999	0.46416	0.06813	0.04682
12	0.99999990	0.56232	0.08254	0.05623
12	0.99999900	0.68101	0.10000	0.06813
12	0.99999000	0.82200	0.12115	0.08254
12	0.99990000	0.96250	0.14678	0.10000
12	0.99900000	0.999996	0.17784	0.12115
14	0.99999999	0.51795	0.10000	0.07215
14	0.99999990	0.61052	0.11788	0.08483
14	0.99999900	0.71943	0.13895	0.10000
14	0.99999000	0.84534	0.16379	0.11788
14	0.99990000	0.96777	0.19307	0.13895
14	0.99900000	0.999997	0.22759	0.16379

Figure 35. Relationship of Retry Count and Link Success Rate

6.3.1 Cascading Considerations

Apart from the time-out restrictions and performance considerations mentioned earlier, remote bridges may be cascaded as local bridges. An example of two level cascaded bridges is illustrated in Figure 36 on page 70.

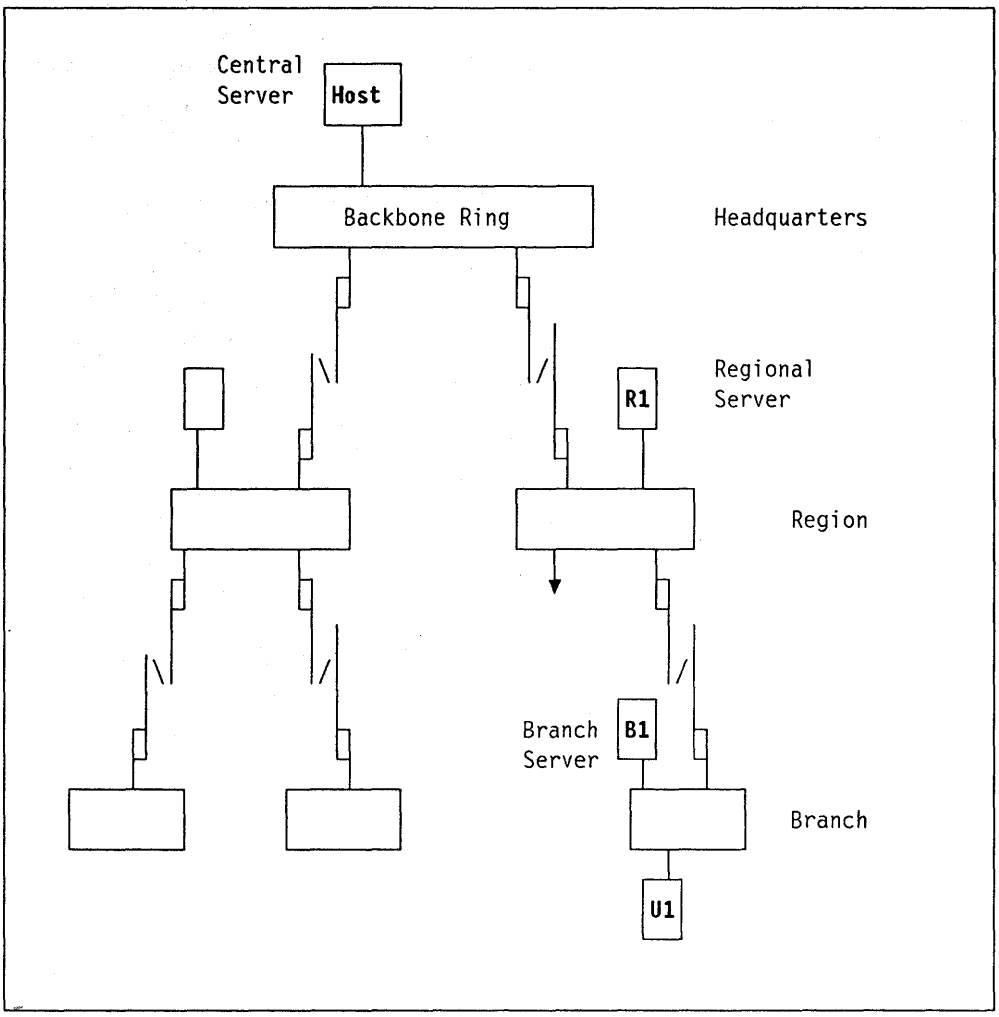

Figure 36. Example of a Cascaded Bridges Configuration

Due to the number of factors which can impact network delays in a cascaded configuration, it is the customer's responsibility to determine whether specific applications will work in this configuration.

From a connectivity and operations standpoint, a cascaded bridges configuration introduces new alternatives. Many large companies have a three level organization at a country or international level. If the primary communication vehicle is the token-ring at each level, cascaded remote bridges might allow full interconnection of the different entities, subject to the performance constraints discussed earlier in this section.

As shown in Figure 36 a station in a branch (like U1) may be able to access concurrently resources located in a local server (Branch server B1), a regional server (R1) and a central server (Host).

Note that in such a configuration, it is possible to have a three level hierarchy without any single session having to cross two cascaded bridges. For instance, station U1 could be restricted to access resources located only in servers B1 and R1. Server R1 could have downstream sessions with B1 and upstream sessions with the Host, which involves only one hop in both cases from a remote bridge standpoint. This could be very useful if the performance factors

such as low line speeds or application timers allow you to traverse one remote bridge but not more. Besides having appropriate definitions in the user and server stations, an efficient way to control the traffic and to avoid people experience time-outs due to the cascaded bridges is to use the filtering facility provided by the remote bridge.

6.3.2 Largest Frame Size (Remote Bridge Function)

For a remote bridge, the largest frame size that can be forwarded by the bridge is user settable. For local bridge configurations it is usually recommended to use default values as the size of the frames has a definite impact on performance. However, as indicated in Figure 34 on page 67, the frame size may be a significant factor in addressing time-out due to telecommunication link errors. The default values for the largest frame size forwarded by a remote bridge are the following:

TP Line Speed (Kbps)	Maximum Frame Size (bytes)
9.6 ≤ TP line ≤ 19.2	516
19.2 < TP line < 56	1028
56 ≤ TP line =≤ 1344	2052

Figure 37. Largest Frame Size for a Remote bridge

As explained in "Largest Frame Size" on page 50 for local bridges, the remote bridge will inspect the largest frame size (LFS) information contained in the routing information field of the broadcast frames it receives. If the requested LFS is greater than the LFS that can be forwarded by this bridge, the remote bridge adapter will replace the requested LFS with its LFS before passing the frame to the attached product.

Note that the AS/400 [12] supports Token-Ring communications through the remote bridge when the bridge is configured for the maximum frame size of 2052 bytes.

6.4 The Filtering Facility

An important additional feature of IBM Token-Ring Network Bridge Program V2.1 called *Frame-Forward Filtering*, provides a mechanism which is especially interesting in a remote bridge configuration and which accomplishes the following:

- Limit the volume of traffic across a bridge
- Filter frames for security or naming convention reasons across a bridge.

[12] AS/400 is a trademark of the IBM Business Machines Corporation

At bridge start-up, it is possible to specify one or several filter program names. The bridge will invoke these *filter user appendages* during the frame forwarding decision process. If the filter appendage returns indicating not to forward the frame, the frame will be discarded and a specific counter (frame not forwarded, filtered) is incremented.

6.4.1 Bridge Filters

A bridge filter is a TSR (Terminate and Stay Resident) program which registers an appendage with a bridge adapter handler. The appendage is invoked by the frame forwarding process of the bridge. The criteria that the appendage uses to discard/forward frames is controlled by the user when the filter is started.

The filters have the following structure in order to minimize their use of storage:

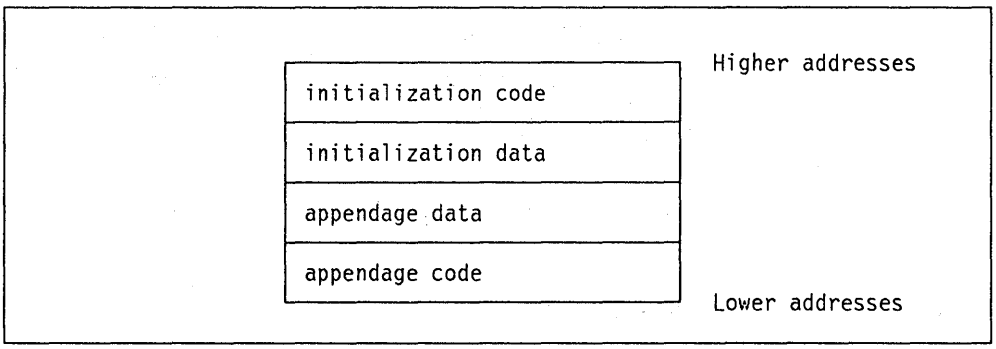

Figure 38. Filter Layout

When the initialization code of the filter is complete it terminates and returns all the storage beginning at the initialization data to DOS. The appendage code and data remain resident.

The interfaces of the filter such as registering with the adapter handler and exchanging information with the frame forwarding process is described in the *IBM Token-Ring Network Bridge Program V2.1 User's Guide*. The Token Ring Network Bridge Program V2.1 contains the following filter related files:

- FILTER.ASM - A sample filter program.
- FILTER1.COM - The link limiting filter.
- FILTER2.COM - The NETBIOS filter.
- FILTER3.COM - The address filter.

6.4.2 Sample Filter

The FILTER.ASM file contains the (MASM) code for a sample filter. This program is meant to be used only as a guideline for a customer to develop a filter program. It demonstrates the correct usage of the DIR.SET.FILTER.APPENDAGE command and the interface between the appendage and the frame forwarding process. This filter leads to forwarding of frames if either the source or destination node address lies within a specified range. The range is supplied by parameters passed to the filter program. Addresses used as parameters can be physical addresses or locally administered addresses.

6.4.3 Link Limiting Filter

The FILTER1.COM file is the link limiting filter. The purpose of this filter is to limit the number of links established when the speed of the TP line is 9600 bps. By entering "FILTER1 ?" the following information is displayed.

```
The format of this filter command is: FILTER1 LINKS=1-255 (TIME=1-3600)
(CONT)
LINKS: The number of unique source and destination address pairs that will
       be allowed to establish links through the bridge.  Default is 2.
 TIME: Time interval in which the link must contain activity to guarantee
       the link remains intact.  Default is 60 (seconds).

 CONT: Do not prompt the user if an error occurs

 NOTE: Frames with a source address of the bridge and frames with a
       group or null destination address are forwarded.
```

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 (frame is filtered). The discarded address pair will go through this same procedure when resuming traffic on its link. If it can not re-enter the table, the link will be disconnected (station not found). Note, even though a link may not have traffic in the specified time period, it will not be removed from the table until an attempt to establish a new link.

There are "special" frames that are always forwarded by the appendage. These consist of frames that originate at the bridge (e.g. bridge self test), frames to a group address (e.g. NETBIOS functional address) and frames to the null address. Upon recognition of one of these frames, the appendage returns control to the frame forwarding process and indicates that the frame be forwarded (not filtered).

The filter should reside on only one side of the bridge. Consider a remote bridge connecting rings A and B with the filter installed in the bridge local to ring A. Requests that originate on ring A are passed through the filter. Requests that originate on ring B use frame forwarding code in the bridge local to ring B, these frames do not pass through the filter. However the responses to ring B requests do enter the filter in the bridge local to ring A. If the response is filtered (not forwarded), the request from ring B will time out.

6.4.4 NETBIOS Filter

The FILTER2.COM file is the NETBIOS filter. The purpose of this filter is to restrict the proliferation of NETBIOS frames throughout the network. By entering "FILTER2 ?" the following information is displayed.

```
The format of this filter command is:

 FILTER2 ADP=PRI|ALT ACTION=DISCARD|FORWARD|DISCARDDB|DISCARDALL (NAME=name|FILE=fil

    ADP: The adapter that will invoke the appendage

 ACTION: The destination NETBIOS name of the frame will be compared to the
       specified name(s).  If a match is found the frame will be discarded
       (DISCARD) or forwarded (FORWARD). If a match is NOT found, the opposite
       action will be taken.  Also, all DATAGRAM_BROADCAST frames can be
       discarded (DISCARDDB) or ALL NETBIOS frames can be discarded
       (DISCARDALL).

 NAME: A NETBIOS name (up to 16 characters).  An "*" or "?" may be used as a
       wildcard.  This name will be compared to the name in the NETBIOS
       UI-frame header.

 FILE: A file name or path to a file which contains NETBIOS names.  The file
       must contain 1 NETBIOS name per line (up to 50).

 CONT: Do not prompt the user if an error occurs


EXAMPLE: FILTER2 ADP=PRI ACTION=DISCARD FILE=DISCARD.LST
```

Unlike the link limiting filter, FILTER2 contains no default parameters. The use of this filter ALLOWS duplicate NETBIOS names to coexist on the network.

If ACTION=DISCARDALL or DISCARDDB, all other parameters are ignored. If ACTION=DISCARDDB is used, all NETBIOS DATAGRAM_BROADCAST frames are discarded. If ACTION=DISCARDALL is used, all frames containing a destination address of the NETBIOS group address are discarded.

The filter maintains a table of NETBIOS names. If the NAME parameter is used, the table contains only 1 name. If the FILE parameter is used, the table may contain more than 1 name.

If ACTION=DISCARD or FORWARD is used, the NETBIOS names in the table and the frame are compared. The appendage looks in the NETBIOS frame header to determine the name to be compared to the name(s) in the table. If the NETBIOS frame is an ADD_GROUP_NAME_QUERY or a ADD_NAME_QUERY, the source name field is used. If the NETBIOS frame is a NAME_QUERY or DATAGRAM (non-broadcast), the destination name field is used. If the NETBIOS frame is any other type, the frame is forwarded.

The following scenarios show some uses of the NETBIOS filter.

In this scenario LAN1 and LAN2 use a centralized server on the backbone ring.

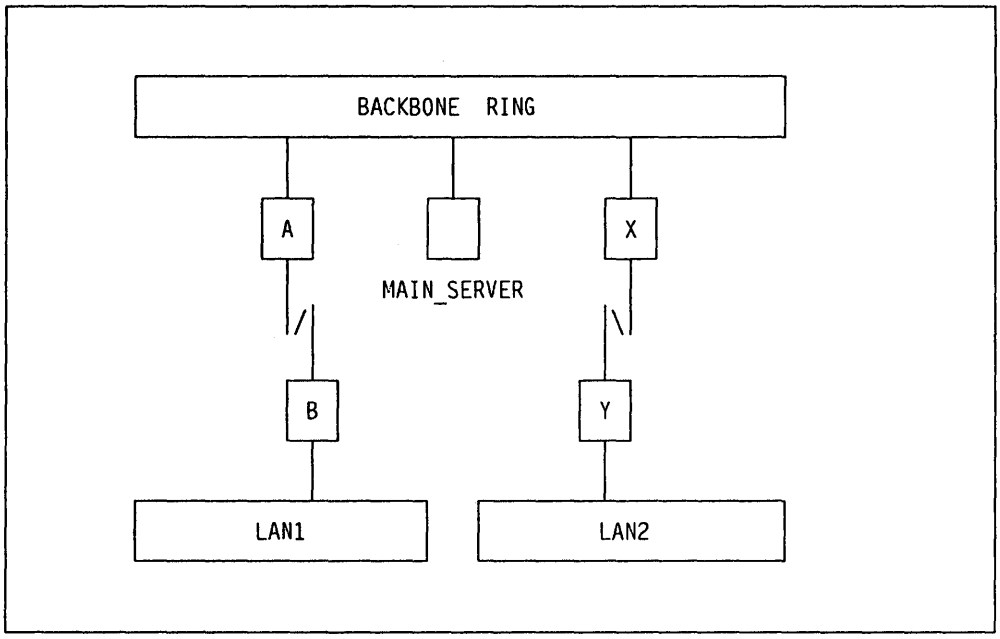

Figure 39. Central Server

Assume that the primary adapter in local bridges A and X forwards frames from the backbone.

To stop any frame with a destination address of the NETBIOS functional address from being forwarded to LAN1 or LAN2, the filter command entered at bridges A and X should be:

```
FILTER2 ADP=PRI ACTION=DISCARDALL
```

To allow frames destined for MAIN_SERVER to be forwarded from LAN1 or LAN2, the filter command entered at bridges A and X should be:

```
FILTER2 ADP=ALT ACTION=FORWARD NAME=MAIN_SERVER
```

In this case, all frames with a destination other than MAIN_SERVER will be discarded.

Another scenario may have LAN1 and LAN2 each using a local server. Also, support personnel on the backbone ring may need access to the remote servers for problem determination or maintenance. In this case, a remote bridge is used to communicate to the backbone.

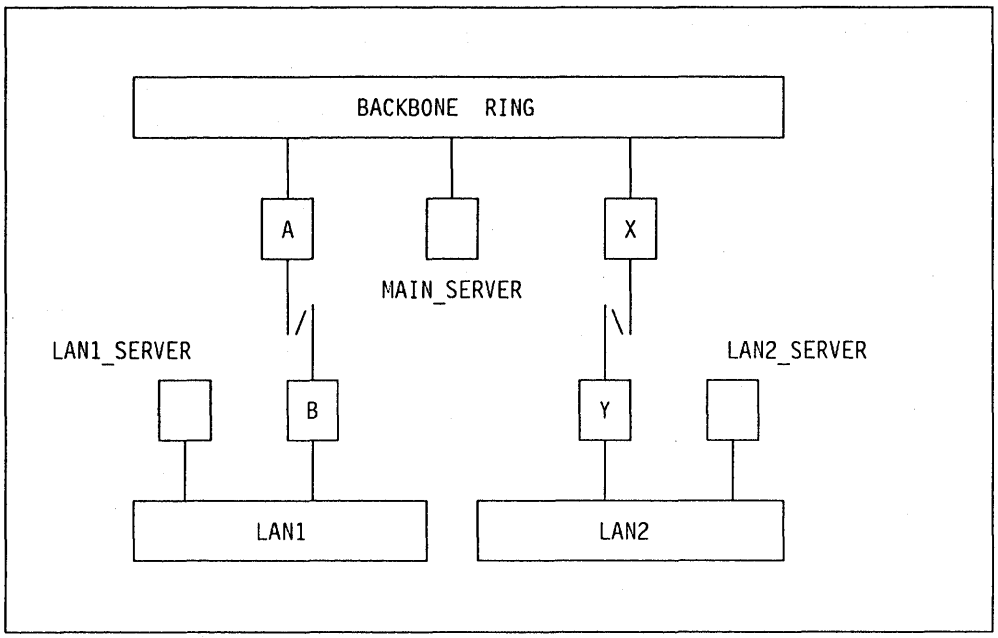

Figure 40. Local Server

Bridges A and X control the frames that flow from the backbone to LAN1 and LAN2. Bridges B and Y control the frames that flow from LAN1 and LAN2 to the backbone.

To stop any frame with a destination name beginning with "LAN1" from being forwarded from LAN1, the filter command entered at bridge B should be:

```
FILTER2 ADP=PRI ACTION=DISCARD NAME=LAN1*
```

This assumes that there are names on LAN1 with a prefix of "LAN1" other than LAN1_SERVER. The same holds true for LAN2. The filter command entered at bridge Y should be:

```
FILTER2 ADP=PRI ACTION=DISCARD NAME=LAN2*
```

To filter frames at bridges A and X we have several options. The best choice depends on the naming conventions of the backbone. If all the NETBIOS names on the backbone begin with "MAIN" and are followed by a certain number of characters (e.g. 2), we could use the following filter command:

```
FILTER2 ADP=PRI ACTION=DISCARD NAME=MAIN??
```

If the NETBIOS names on the backbone are not consistent, it may be easier to explicitly specify the destination. This can be done by the following filter command:

```
FILTER2 ADP=PRI ACTION=FORWARD FILE=SERVERS
```

Where SERVERS is a file containing the names of LAN1_SERVER and LAN2_SERVER. If ACTION=FORWARD is specified and the name is not found in the file, the frame is discarded.

The names file must contain 1 name per line. The file can contain up to 50 names, any extras will be ignored. The names in the file can also use the wildcard feature (e.g. LAN1* and LAN2*). The "*" and "?" wildcards are used to represent characters identical to their definition in DOS.

6.4.5 Address Filter

The FILTER3.COM file is the address filter. The purpose of this filter is to limit access across a bridge according to the source and destination addresses of a frame. By entering "FILTER3 ?" the following information is displayed.

The format of this filter command is:

FILTER3 ADP=PRI|ALT SA=addr1(-addr2) DA=addr3(-addr4) (CONT)

ADP: The adapter that will invoke the appendage

SA: A MAC address or range to be compared to the SA in the frame.

DA: A MAC address or range to be compared to the DA in the frame.

CONT: Do not prompt the user if an error occurs

NOTE: A MAC address is exactly 12 HEX characters (0-9 or A-F) in length. A range is 2 MAC addresses separated by a hyphen. The value in the frame is compared to the specified value or range. If the value in the frame equals the specified value or is in the range, the frame is discarded. If both SA and DA are specified, BOTH conditions must be satisfied for the frame to be discarded.

EXAMPLE: FILTER3 ADP=ALT DA=4000A0001000-4000A0002000

The appendage compares the SA and DA of a frame being forwarded to the values specified on the command line. If the SA and DA meet the specified conditions, the frame is discarded. There are no exceptions made for group addresses or the address of the bridge. There is no validity checking performed on the user specified range(s).

This filter is meant to be used when a network contains well maintained locally administered addresses. An example of the address filter command is:

```
FILTER3 ADP=PRI SA=400000001000 DA=4000AAAA0000-4000BBBB0000
```

The SA and DA of a frame must satisfy both conditions for the frame to be discarded, otherwise it is forwarded. This implies an "and" condition. To create an "or" condition the filter can be activated twice with different parameters.

```
FILTER3 ADP=ALT SA=400000001000

FILTER3 ADP=ALT DA=4000AAAA0000-4000BBBB0000
```

If the DA of a passing frame falls within the range, the discard-flag is set and all the other filters return to the caller, immediately. Remember, filters are executed in the reverse order they are activated. If the DA did not fall within the range, the SA of the frame is compared to the specified SA. If the SA values are equal, the discard-flag is set.

If there is more than 1 range which needs frames discarded, the filter can be activated multiple times.

```
FILTER3 ADP=PRI DA=400000000000-400000004FFF

FILTER3 ADP=PRI DA=40000000A000-40000000FFFF
```

The filter can discard an errant all-stations broadcast in order to relieve network traffic.

```
FILTER3 ADP=ALT DA=FFFFFFFFFFFF

FILTER3 ADP=PRI DA=FFFFFFFFFFFF
```

6.4.6 Filter Combinations

It may be necessary to have multiple filters (maximum of 10) active concurrently. Filters are executed in the REVERSE order in which they were activated. *Thus care should be taken in planning and using filters* . It is advantageous to execute filters that discard more frames before filters which 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 execute all the filter code. The filters shipped with the bridge product DO NOT avoid executing other filters (i.e. they do not set AH = 1 on exit). The filters return immediately if a previous filter has decided to discard a frame (i.e. they check that AL = 0 on entry).

The link limiting filter and the NETBIOS filter can be used together. For example, the user can enter the following commands.

```
FILTER1 LINKS=6 TIME=120

FILTER2 ADP=PRI ACTION=DISCARDDB
```

The link limiting filter should not be executed more than 1 time. Otherwise, the tables of address pairs would be redundant.

If the user wishes to explicitly identify the filtering attributes for a group of NETBIOS names, the filter can be invoked a number of times. Assume FILTER.LST contains NETBIOS names to be discarded and FORWARD.LST contains NETBIOS names to be forwarded. The user can enter the following commands:

```
FILTER2 ADP=ALT ACTION=FORWARD FILE=FORWARD.LST

FILTER2 ADP=ALT ACTION=DISCARD FILE=DISCARD.LST
```

6.5 Installation/Utilization Guidelines

6.5.1 Installation Hints

The bridge installation methodology described in "Installation/Utilization Guidelines" on page 51 applies to the remote bridge and is not repeated here. However, there are some important additional considerations specific to the remote bridge which are discussed hereafter (please refer to the *IBM Token-Ring Network Bridge Program V2.1 User's Guide* for more information on the physical installation procedure).

Before you start the software installation process (using SETUP), you must ensure you have the following three diskettes:

- Bridge program backup copy
- Realtime Interface Co-Processor DOS Support Program
- RTIC diagnostics, Realtime Control Program or RTIC/2 X.25 and Multiport/2 Option diskette.

Before you start the next configuration step (using INSTALL), you should have decided which half of the bridge is the primary or the secondary side, as the installation procedure is slightly different for each side.

After the installation, both sides will have the same files, except the ECCPARMS.BIN file, which contains all bridge configuration parameters values. **The ECCPARMS.BIN file must exist only on one side of the remote bridge.** The same configuration file is used for both halves of the bridge (primary and secondary). The primary side of the bridge is the only half which contains the ECCPARMS.BIN configuration file. Therefore, when installing the Bridge Program, you must configure only the primary half of the bridge, not the secondary.

If you change later on your configuration, for example if you swap the primary and secondary roles of the two halves, be sure that you have the ECCPARMS.BIN file only on one side (by renaming or erasing the other one). Otherwise, the bridge initialization will fail with the following messages:
"Only one half of the bridge can have an ECCPARMS.BIN configuration file"
"Bridge initialization has failed"
"Shutdown is complete"

The reason for this is that the ECCPARMS.BIN configuration information is passed across the communications link during bridge initialization. If there are two such files, one on each side, the remote bridge logic has no reason to select one rather than the other. Once this unique file has been transferred to the other half, both bridge halves verify that the common information matches (for instance bridge number and ring segment number).

Likewise, the communication adapter configuration files (ECCSBPRM.BIN) must contain the same parameter values on each half of the bridge, in order to be able to communicate.

In addition, when any of the operational bridge parameters is dynamically changed by the IBM LAN Manager V2.0, both bridge halves will be

automatically updated and the changes will be permanently recorded in the ECCPARMS.BIN file.

6.5.2 Recovery and Problem Determination

Recovery can be considered in the two following cases:

- Bridge failure

  In the case of a bridge failure, the automatic restart option will usually be successful, for example after a short power failure. A dump of the bridge's memory may also be automatically written at the bridge. If a bridge which is already linked to a LAN manager is powered off, the following alert will appear on the LAN manager console:
  "Management Server Reporting Link Error"

  When the bridge is restarted, the link between the LAN Manager and the bridge will have to be reestablished either manually (from the LAN manager or NetView console) or automatically via an appropriate NetView Clist.

- Telecommunication link failure

  In the case of a link failure, both sides will detect that the link has failed as they periodically exchange messages and counters, even if there is no actual activity on the link. As a result, both halves will terminate their link with the LAN Manager(s) in order to generate an alert at the LAN Manager's or NetView console.
  "Management Server Reporting Link Error"

  In addition, the following messages will be displayed on both halves of the remote bridge:
  "Telecommunication line is down"
  "Network manager links are terminated"

  The alert displayed on the LAN manager console is the same as in the case of a bridge failure and the alert does not tell you that it is a telecommunication link problem.

  To determine if the actual problem is due to the bridge or to the telecommunication link, you can check if the bridge is still alive by querying the status of the remote bridge token-ring adapters.

  It is important to notice that the link will be recovered from the bridge point of view as soon as the link is operational again, as both halves continue to try periodically to contact each other. No intervention is required on either bridge half, and you will automatically see a "normal" status for the telecommunication link on the bridge's screen, once the line is operational.

7. PC Network Bridge

7.1 PC Network Bridge Overview

The IBM PC Network Bridge Program provides MAC level interconnection between the following LAN segments:

- Two IBM Token-Ring Network segments operating at 4 or 16 Mbps.
- Two IBM PC Network (Broadband) segments.
- An IBM Token-Ring Network segment operating at 4 or 16 Mbps and an IBM PC Network (Broadband) segment.

Although the IBM PC Network Bridge Program may interconnect two Token-Ring Network segments, its primary purpose is to integrate PC Network (Broadband) segments into a larger (token-ring based) LAN environment. It alleviates the limitation on the number of devices attached to the network and provides increased network design flexibility and integration of LAN management.

Devices attached to a broadband PC Network segment being bridged by the IBM PC Network Bridge Program require IBM Local Area Network Support Program Version 1.0 with PTF UR22583 or Version 1.1.

For any type of LAN interconnection, the IBM PC Network Bridge Program uses the concepts of source routing.

The bridge function is transparent to applications written to the IEEE 802.2 logical link control interface using source routing, whether the MAC protocol at either side of the bridge is identical or dissimilar.

IBM PC Network Bridge Program also provides the automatic single-route broadcast facility.
If the automatic single-route broadcast facility is desired, all bridges in the multisegment LAN must be configured to participate in the automatic single-route broadcast bridges protocols.

7.2 PC Network Bridge General Structure

Figure 41 on page 82 shows the general structure of IBM PC Network Bridge Program in the case of an interconnection between a PC Network (Broadband) segment and an IBM Token-Ring Network segment. The IBM PC Network Bridge Program includes the adapter handler for both PC Network and the IBM Token-Ring Network, as well as logical link control and bridge protocol software to support PC Network (Broadband) adapters. IBM Token-Ring Network Adapters contain those protocols as microcode in Read Only Memory (ROM).

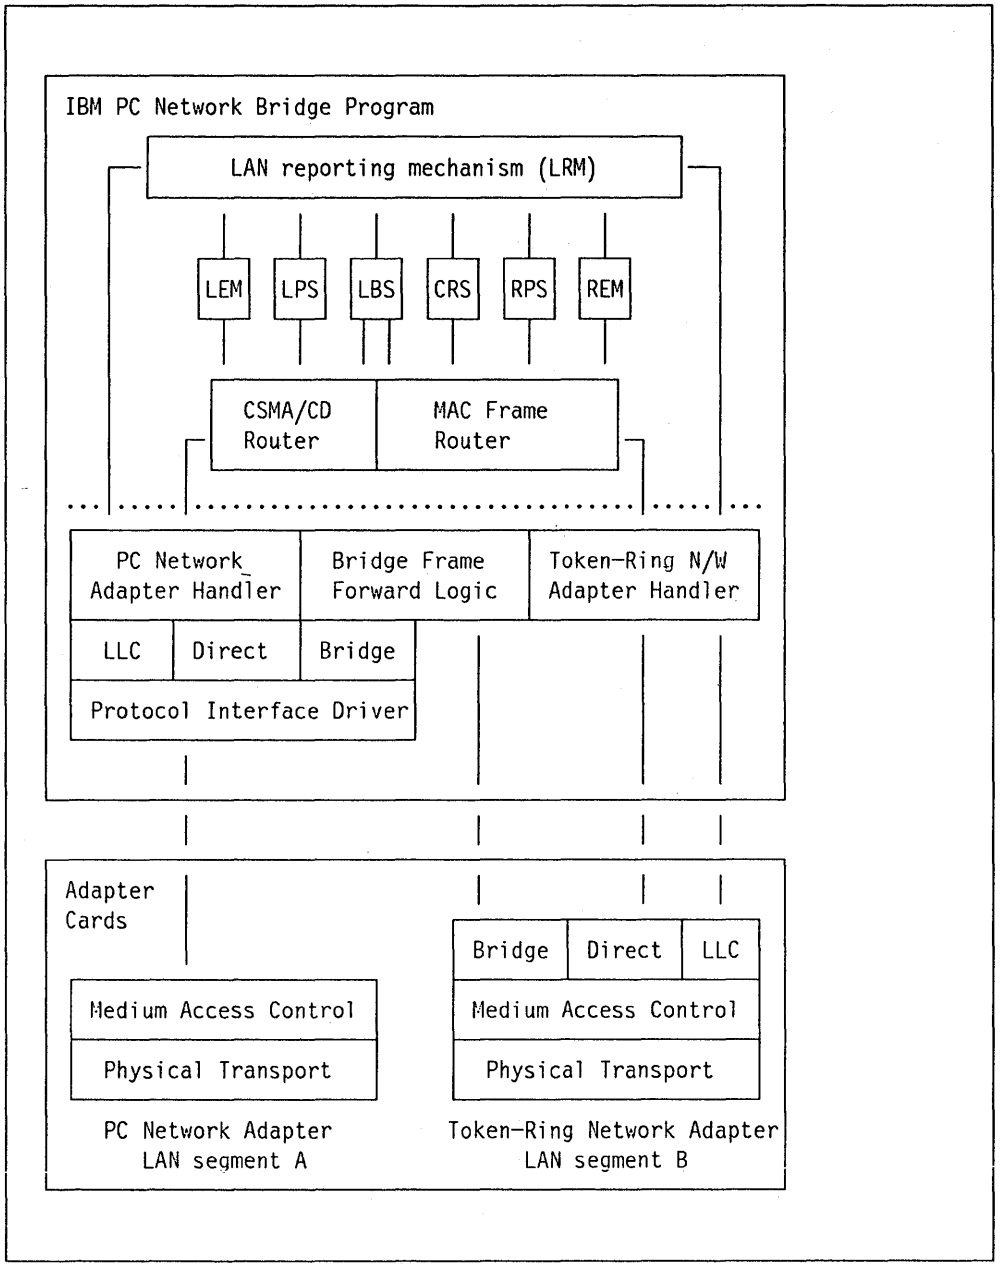

Figure 41. IBM PC Network Bridge Program - General Bridge Structure

The LAN management server functions have been defined in "LAN Manager Support" on page 20 and are referenced explicitly in Figure 42 on page 83.

7.3 LAN Management Facilities for PC Network Segments

The IBM PC Network Bridge Program includes all the network management support provided by IBM Token-Ring Network Bridge Program V2.0 and supports also network management data from PC Network (Broadband) segments.

Like IBM Token-Ring Network Bridge Program V2.0 and IBM Token-Ring Network Bridge Program V2.1, IBM PC Network Bridge Program interfaces with up to four LAN Managers running IBM LAN Manager V2.0.

The following IBM Token-Ring Network management information is sent to IBM LAN Manager V2.0:

- Soft error reports and beaconing notification.
- Bridge status and performance data.
- Ring configuration reports.
- Path trace reports.

The following PC Network (Broadband) management information is sent to IBM LAN Manager V2.0:

- Continuous carrier and no carrier notifications.
- Bridge status and performance data.
- Ring configuration reports.
- Path trace reports.
- Topology reports (although the sequence between stations on a CSMA/CD LAN segment is irrelevant).

Some IBM PC Network Bridge Program server functions support only a token-ring network segment, some support only a PC Network (Broadband) segment, and others support both types of network segments.

IBM PC Network Bridge Function	Token-Ring N/W Segment	PC Network Segment
LAN reporting mechanism (LRM)	Yes	Yes
Ring Error Monitor (REM)	Yes	No
LAN Error Monitor (LEM)	No	Yes
Ring Parameter Server (RPS)	Yes	No
LAN Parameter Server (LPS)	No	Yes
Configuration Report Server (CRS)	Yes	n/a
LAN Bridge Server (LBS)	Yes	Yes

Figure 42. IBM PC Network Bridge Program LAN Management Functions

As there is no equivalent to token-ring MAC frames in a IBM PC Network (Broadband), some network control functions have been defined using LLC protocols.

Therefore the LAN Manager can communicate directly with IBM PC Network (Broadband) stations for network control reasons. For example, the LAN manager will issue a "Remove adapter" command directly to a IBM PC Network (Broadband) station via an LLC session, while this is done with a MAC frame sent by the bridge on request of the LAN manager in a token-ring segment. In

addition, PC Network (Broadband) segments do not need a CRS server function provided by the bridge.

7.4 Network Expansion for PC Network (Broadband) Segments

Each individual IBM PC Network (Broadband) segment can operate at either of three supported channel pairs, referred to as frequency 1, frequency 2 and frequency 3.

PC Network (Broadband) segments may share the same broadband medium (at different channel pairs) or have totally separate media.

With the IBM PC Network (Broadband) translator unit and cables, the number of stations on a IBM PC Network (Broadband) is limited to 72 stations.

Bridges running IBM PC Network Bridge Program can interconnect multiple IBM PC Network (Broadband) segments into a single logical LAN.

An example of an expanded IBM PC Network (Broadband) LAN is shown in Figure 43.

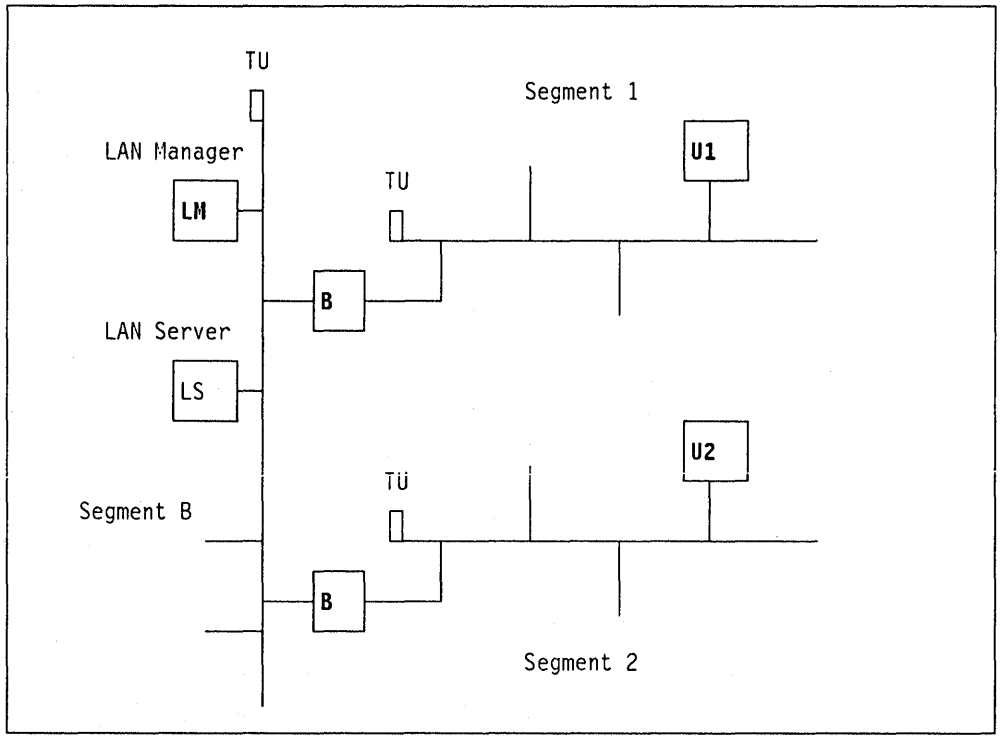

Figure 43. LAN with Multiple PC Network (Broadband) Segments. TU stands for Translator Unit. B states for PC Network Bridge

In this example, all stations can communicate with each other and can access shared resources such as the LAN server connected to segment B.

Although the backbone shown in this example is an IBM PC Network (Broadband) segment, a token-ring backbone is recommended because of its higher speed and for host connectivity, as shown in the next section.

7.5 Host Connectivity via a Token Ring Backbone

Another benefit of the IBM PC Network Bridge Program is that it improves host connectivity for IBM PC Network (Broadband) stations.

Because host systems cannot connect directly to an IBM PC Network (Broadband) segment, IBM PC Network (Broadband) stations should be able to access the host directly at LAN speeds if the IBM PC Network (Broadband) segments are connected to a token-ring segment via the IBM PC Network Bridge Program. With the current implementation of IBM PC 3270 Emulation Program V3, IBM PC Network (Broadband) stations need an IBM PC 3270 Emulation Program V3 gateway on the token-ring to access the host.

Another alternative is to have an IBM Personal Communications/3270 gateway connected to the token-ring backbone as well as the IBM PC Network (Broadband) segment.

An example of a mixed LAN is shown in Figure 44.

In this example, stations U2 and U3 on segments 2 and 3 are defined as network stations and can access the host via station GW, defined as a gateway station. Token-ring network stations can access the host directly as stand-alone stations.

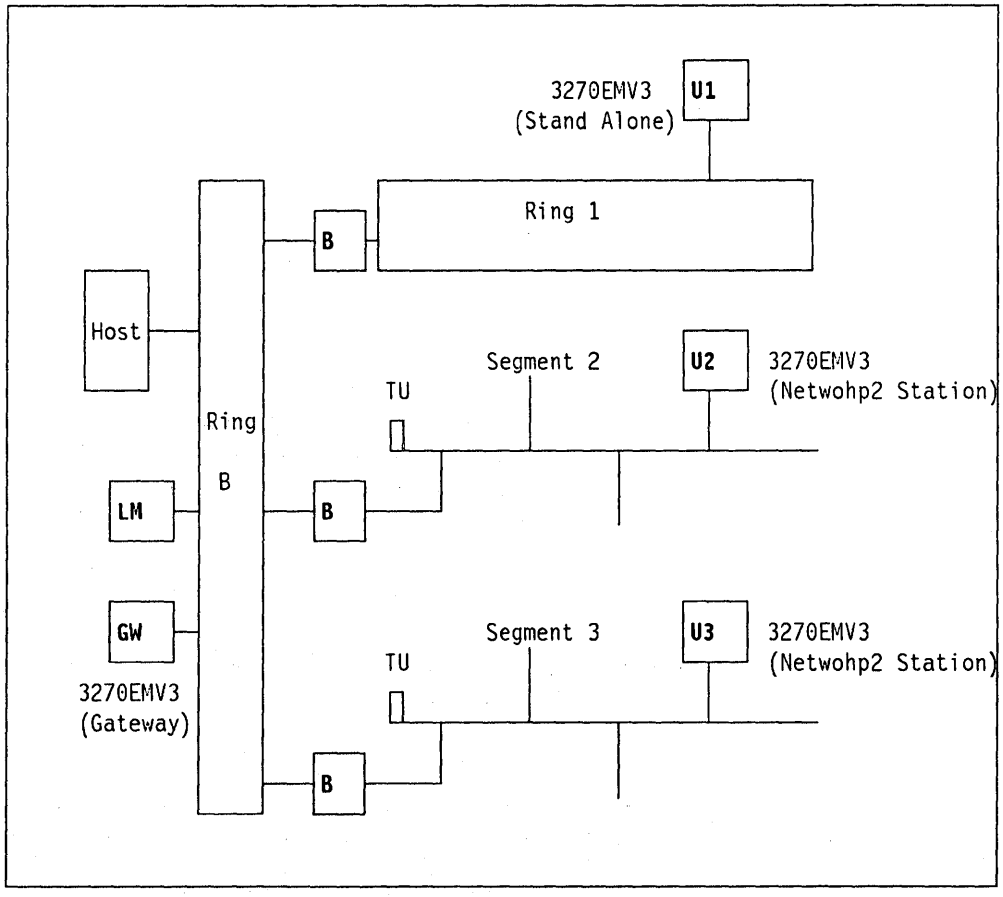

Figure 44. Mixed LAN with a Token-Ring Backbone and PC Network segments. TU stands for translator unit

7.5.1 Hardware and Software Requirements

IBM PC Network Bridge Program requires a dedicated IBM Personal System/2 Model 50, 60, 70 or 80. It requires 512 Kbytes of memory and a 720 Kbytes or 1.2 Mbytes 3.5 inch diskette drive.

Supported local area network adapters are shown in Figure 45.

AFE	P/N / FC
IBM Token-Ring Network Adapter /A	69X8138 / 4790
IBM Token-Ring Network 16/4 Adapter/A	16F1133 / 1133
IBM PC Network Adapter II/A	1501222 /
IBM PC Network Adapter II/A - Frequency 2	96X5647 /
IBM PC Network Adapter II/A - Frequency 3	96X5648 /
(EMEA)	P/N / FC
IBM Token-Ring Network Adapter /A	/
IBM Token-Ring Network 16/4 Adapter/A	/
IBM PC Network Adapter II/A	90X6969 /
IBM PC Network Adapter II/A - Frequency 2	96X6061 /
IBM PC Network Adapter II/A - Frequency 3	96X6066 /

Figure 45. IBM PC Network Bridge Program - Bridge Supported Adapters

Note that the limitation requiring the IBM PC Network (Broadband) adapters identified in the original Programming Announcement 288-488, dated September 20,1988 has been removed.

IBM PC Network Bridge Program's only software prerequisite is IBM PC/DOS 3.3 or 4.0.

Devices attached to a broadband PC Network segment being bridged by the IBM PC Network Bridge Program require IBM Local Area Network Support Program Version 1.0 with PTF UR22583 or Version 1.1.

7.6 Installation/Utilization Guidelines

The installation of the IBM PC Network Bridge Program is very similar to the installation of the IBM Token-Ring Network Bridge Program V2.0.

The methodology and the main installation steps of the IBM Token-Ring Network Bridge Program V2.0 are described in "Installation/Utilization Guidelines" on page 51. That methodology applies to the IBM PC Network Bridge Program as well. The only difference is that you may have to install IBM PC Network (Broadband) adapters instead of token-ring adapters. While there is only one speed for a IBM PC Network (Broadband) segment (2Mbps), there is a variety of adapters with different frequencies as mentioned earlier. Adapter parameters such as shared ram address and early token release are not relevant for the IBM PC Network Bridge Program.

Refer to the *IBM PC Network Bridge Program User* for details on the IBM PC Network Bridge Program installation and customization procedure.

As far as performance is concerned, the expected bridge throughput depends on the type of interconnected MAC segments:

- About 1 Mbps for PC Network (Broadband) to PC Network (Broadband) or PC Network (Broadband) to Token-Ring Network.
- About 2 Mbps or more for Token-Ring Network to Token-Ring Network, depending on the bridge processor and the bridges' adapter speeds.

It should be noted that on the IBM PC Network (Broadband) side, a bridge adapter will copy all frames from the network. Then the bridge process task will determine if the frame is to be routed or not. This is not the case for the token-ring side, where the adapter microcode only copies frames that must be handled by the bridge station.

8. LAN Design Methodology

This chapter introduces the design criteria and main tasks a LAN planner should consider during the LAN design process.

8.1 LAN Design Criteria

Many factors need to be considered when designing a multisegment LAN. A partial list includes:

- The number of stations
- The connectivity requirements (hosts, departmental servers)
- The physical layout of the establishment
- The existence of affinity groups
- Performance requirements
- Reliability and availability (alternate paths, backup gateways)
- Cost
- Network management
- Expected network growth in a long term view

All these factors influence the decision of what topology to select for a particular installation. Most of them are discussed in "Logical Design Considerations" on page 93.

Considering the number of factors, it is obvious there is no "best solution" for every network. However general guidelines can be given on the design methodology and on the servers' location.

8.2 LAN Servers Considerations

Before you start the LAN design process, you should decide what type of servers you will need and how the users are going to access them.

There are mainly two categories of servers:

- Central servers

  These servers are usually hosts and might be accessed by all LAN users. They are usually connected to the backbone and completely managed by the data processing department.

- Local servers

  These servers are normally connected to a single segment and accessed by a smaller group of users, or affinity group. Examples of such servers are departmental servers which provide disk sharing or printing facilities.

  It is important to decide where the local servers will be located and who and how they are going to be maintained, depending on your organizational environment.

For example, print servers are usually placed in an area close to its users to avoid time loss or requirements for a print delivery service.

Disk servers should be placed in a secure area to avoid accidental damage or intentional misuse of the server.

In addition, from a performance point of view, the load and number of local servers such as PCs should be evaluated to avoid bottlenecks and provide good response time to the users.

Last but not least, the servers (and users) software level maintenance strategy must be determined. In particular, if you have many stations, disk or print servers in your installation, automated software maintenance procedures should be considered to alleviate the LAN personnel (or end-user) maintenance workload.

8.3 Design Methodology

The LAN designer should go through the following steps in an iterative way to select the right LAN topology for a particular installation. As cabling considerations are beyond the scope of this document, we will assume that the physical cabling has already been done and documented, or will be done in accordance with the design.

1. **Collect the required information**

   Before designing a LAN, it is important to know the user requirements as well as the physical constraints of that installation. The following information must be carefully collected to achieve a good design.

   - A detailed **physical layout "blueprint"** and the cabling related information is absolutely essential to accommodate the physical design rules. In particular, the following data must be documented:
     - The precise locations (and size) of the wiring-closets
     - The **exact cable lengths between the wiring-closets**
     - The **maximum lobe lengths** for each wiring-closet
     - The **number of stations** to be connected to each wiring-closet.

     This type of information is required to calculate the adjusted ring length and accommodate the physical design rules.

   - Connectivity requirements

     It is important to know the type of applications and the number and type of the hosts or servers that must be accessed by the different user groups. If there are several buildings in the establishment, you should also know if inter-building communication is required at the end-user level.

   - Performance objectives and traffic statistics

     To select the appropriate number and type of servers or gateways that you need, you must have a good idea of the host traffic and of the local server (for example a print server) traffic.

     For example, you should have a good estimation of:

     - The frequency and size of file transfers per host or server

- The number and type of interactive transactions per host or local server

Peak periods volumes have to be taken into consideration to guarantee good response times.

- Affinity groups

  The designer should be aware of the existence, size and locations of the eventual affinity groups if he wants to structure the LAN segments according to these particular entities.

- Availability and security requirements

  The objectives of the particular installation in terms of high availability requirements must be clearly defined, in order to select the appropriate topology with for example alternate paths, dual backbones and backup gateways.

2. **User ring design**

   The "horizontal" or "user ring" design is usually the task you will perform first when you have collected the above necessary information. Most of the design criteria's listed earlier have to be considered in designing each LAN segment. How to design a user ring is discussed in "User Ring Design" on page 93.

3. **Backbone ring design**

   After the "horizontal" design, the next task to perform is the "vertical" or "backbone" design. At this stage, you must choose the best topology to interconnect the previously defined user rings. How to design a logical backbone is described in "Backbone Ring Design" on page 94. Additional considerations on backbone physical design can be found in "Backbone and Bridge Physical Design Guidelines" on page 111.

4. **Gateway selection**

   Once the backbone design has been made, you should select the appropriate type and number of gateways depending upon the traffic characteristics, performance objectives, availability, cost and network management considerations.

5. **Network management and performance**

   At this stage, you should determine requirements or values for the following options or parameters:

   - Number and location of LAN Managers
   - Automated LAN operations with NetView
   - REM function activation in specific stations
   - Bridge parameters such as automatic single-route broadcast and hop count limit.

6. **Backup scenarios**

   Depending on the availability requirements in your particular installation, you should check the effect of a component failure and adapt your configuration if necessary. In particular, you should consider the impact of a failure of a component such as a bridge, a backbone or user ring, repeaters or cables, gateways, and local or central servers.

As stated earlier, the design process is iterative and you will sometimes have to change previous choices due to the number of criteria to be considered in a large LAN design.

8.4 Broadcast Traffic Control Techniques

As mentioned in the network performance design step, several bridge parameters or options are available to minimize broadcast traffic and alleviate the server's workload in complex topologies. The most important options, which have been explained in detail in "Source Routing Approach" on page 27, are:

1. The automatic single-route broadcast facility
2. The hop count limit parameter
3. The RND NETBIOS option.

Use of these options is highly recommended in most cases. Several examples are discussed in "Logical Design Considerations" on page 93.

9. Logical Design Considerations

9.1 User Ring Design

As explained in "LAN Design Methodology" on page 89, the designer of the user rings will have to take several factors into consideration:

- Physical topology constraints

  The design will usually be greatly influenced by the physical topology of the building. For example, you could consider a LAN design with one user ring per floor, or even one user ring for several floors, as long as the number of stations is less than 260.

  However, distances between the wiring closets, the number of 8228s, ring speed and the maximum lobe lengths might sometimes oblige you to use two rings (or more) per floor, to accommodate the ring physical design rules. This will add some additional bridges to your configuration.

  An alternative is to install repeaters, which support very large rings. This is particularly true with the new 8220 optical repeaters which additionally provide a high level of availability (automatic wrapping on both ends to the backup path due to a problem in the optical segment (including the repeaters) and backup path monitoring).

  Depending on other criteria's discussed below, you will make a choice between large user rings with repeaters, or smaller rings with more bridges.

  In addition, even if the adjusted ring length of the ring is compatible with the maximum lobe length, you should also consider the growth potential of the user ring in terms of number of stations and the number of 8228s to support them. For example, if you have many 3270 type terminals connected to 3174/3R, many of those terminals could migrate to PC type terminals directly attached to the ring and therefore would need more 8228s in the user ring. So make sure your design is made in the long-term range.

- Number of stations

  The maximum number of stations on a single ring is 260. You can design a user ring with 260 stations with no concern for distance or performance.

  However, although a LAN segment and especially a token-ring segment is very reliable, a question to consider is: "How many users do you impact if a ring is down?". There is no general answer to that question as it is highly dependent on your environment. Many operational rings have between 50 and 150 stations connected to them.

  A general guideline is to avoid breaking the network into too many small rings (less than 50 stations) without good reason, as you will multiply the number of bridges, under utilize the ring capacity and raise the LAN overall cost.

- Affinity groups

  An affinity group is a group of users who perform related tasks on the network and have little information interchange with other end users.

Each of those affinity groups can be attached to different rings within a network. This could simplify the design and maintenance of applications running in the servers of the respective affinity groups.

- Organization factors

  In some cases, some departments want "their own ring". This can be explained for management, control or security reasons. Those departments could be completely responsible for selecting and buying the equipment (workstations and servers) and installing and maintaining their local applications.

- Moving (relocation) factor

  If user groups are often subject to relocation inside the building, use of affinity groups as a basis for ring structure should be avoided as you could have to change your topology very frequently. So, if the relocation ratio is high, the design of the user rings should be based essentially on geographical considerations.

- Performance and ring speed

  With traditional 3270-type interactive applications, load on the user ring is unlikely to be a bottleneck, even with 260 stations on a 4Mbps ring. However, as your design considers long-term views, it is recommended to use 16 Mbps ring design guidelines, even if you start with adapters running at 4Mbps. 16Mbps will be used initially for backbone rings, and for rings with users who need very fast high volume data transfer like image-oriented applications.

- Management and software maintenance

  Although some departments want "their own ring" and sometimes their own ring administrator, some coordination is required in terms of naming conventions (to avoid duplicate names, addresses, and ring numbers within the network), LAN management, problem determination, performance control and common applications maintenance. Therefore, the LAN manager should be able to access all bridges and the LAN administrator should be able to give at least some naming convention standards or software guidelines to all user groups.

9.2 Backbone Ring Design

One of the very first steps in a LAN design is to select the most appropriate topology for a particular installation.

Serial topologies are usually inadequate for multisegment LANs, as they don't provide alternate paths. Any bridge or ring problem will affect all other rings.

Loop and mesh configurations become very complex as soon as you exceed four rings and you want to maintain any-to-any connectivity with good performance. In addition, the servers location is not obvious at all.

Backbone ring topology is generally recommended as the best approach to multi-ring networks, and can be the best solution for configurations ranging from four to hundreds of rings.

One-level hierarchy or multi-level hierarchy can be achieved with the backbone approach. In all cases, the backbone ring should use either IBM Cabling System type 1 (or equivalent) cables and preferably optical fiber to support high speed adapters and high traffic requirements. In this section, we'll consider first single-level hierarchy topologies. An example of a single-level backbone topology is shown in Figure 46.

9.2.1 Backbone Configuration Benefits

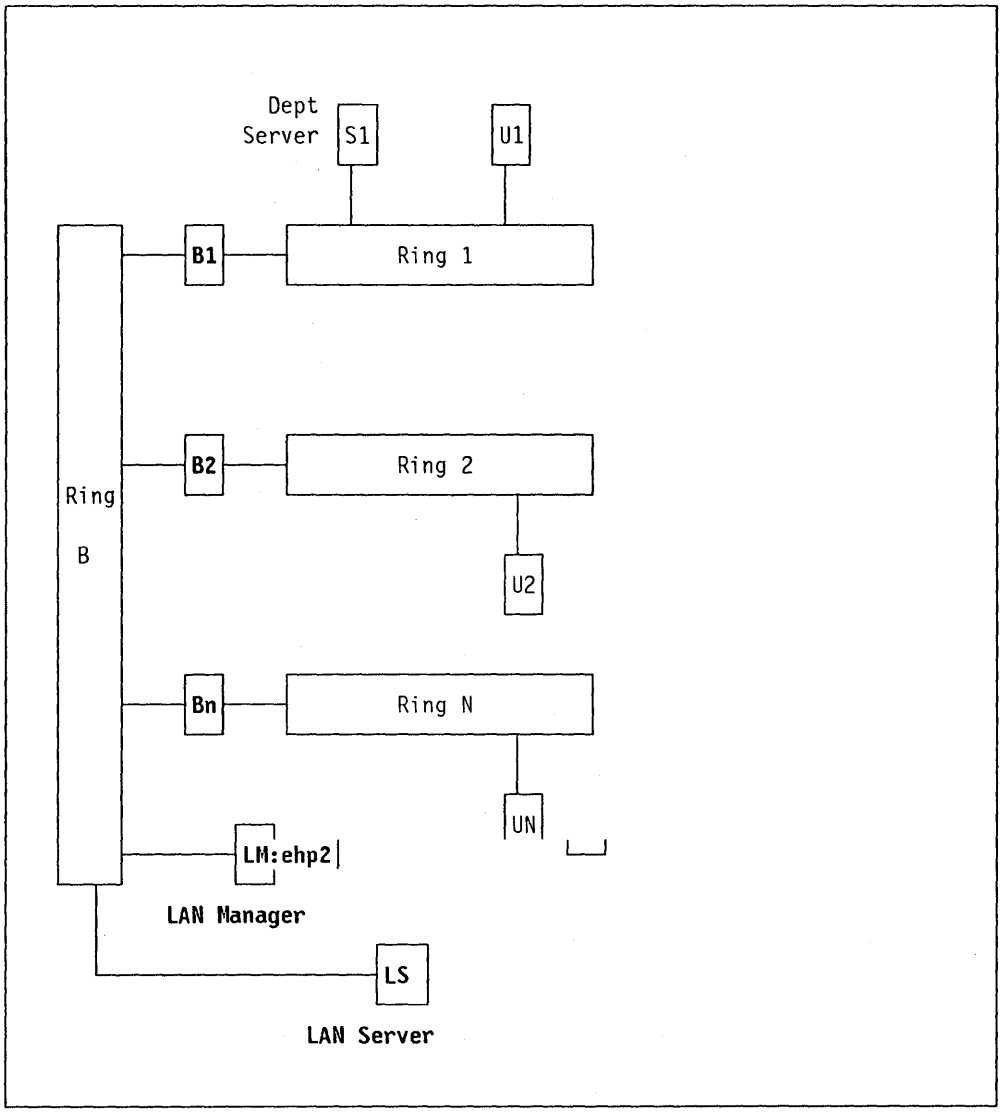

Figure 46. Simple Backbone Configuration

Compared to other topologies like serial, loop, or mesh topologies, a backbone design presents several characteristics and advantages:

- Each bridge in the LAN attaches one user ring to the backbone.
- All inter-ring communications flow through the backbone.
- Communication between two devices will traverse a maximum of two bridges, regardless of the number of rings in the LAN.

- Installation of an additional ring is simple and can be achieved by adding a single bridge per ring, without interference with the other user rings.
- Resources shared or accessed throughout the LAN such as LAN servers or hosts will usually be connected to the backbone.
- The physical attachment or relocation of rings or workstations is simplified.
- Evolution to higher speed backbones is simple and does not impact the user rings.
- Network management and traffic analysis is easier.

The more rings you have in your configuration, the more desirable is the backbone approach.

However, **the single** backbone topology has also some weaknesses:

- Each user ring may rely on a single bridge for its connectivity with servers or other workstations attached to different rings.
- The backbone ring is a critical component and could be a bottleneck from a performance point of view.

That is why the backbone configuration is very often enhanced by adding a second backbone ring, as explained in the next section.

From a performance point of view, a single backbone topology does not need any customization of the automatic single-route broadcast and hop count limit bridge parameters:

- As there is only one natural route to go from a workstation to another, there is a "natural" single-route broadcast path in the network. Therefore, in this particular configuration, the automatic single-route broadcast facility is absolutely useless and is not recommended.
- Likewise, there is no need to set the hop count limit parameter to any particular value, assuming any-to-any connectivity is required (that is any user should be able to access the LAN server connected to the backbone as well as other workstations connected to other rings). As a matter of fact, any hop count limit value less than (2,1) would affect the connectivity and should be used only in very particular cases (2 is the hop count limit value for the backbone side adapter).

9.3 High-Availability Design Considerations

To satisfy high-availability requirements, the LAN designer should discard all configurations with a single point of failure.

In particular, any bridge, repeater or cable failure should not inhibit the any-to-any connectivity user requirements which are the fundamental reasons for the popularity of LAN solutions.

Any component failure in the LAN should be automatically bypassed or at least have a short and minimal impact on the user's activity.

9.3.1 Dual Backbone Approach

A dual backbone approach addresses high-availability requirements and is therefore the recommended topology for most large LANs.

A typical dual backbone configuration is shown in Figure 47.

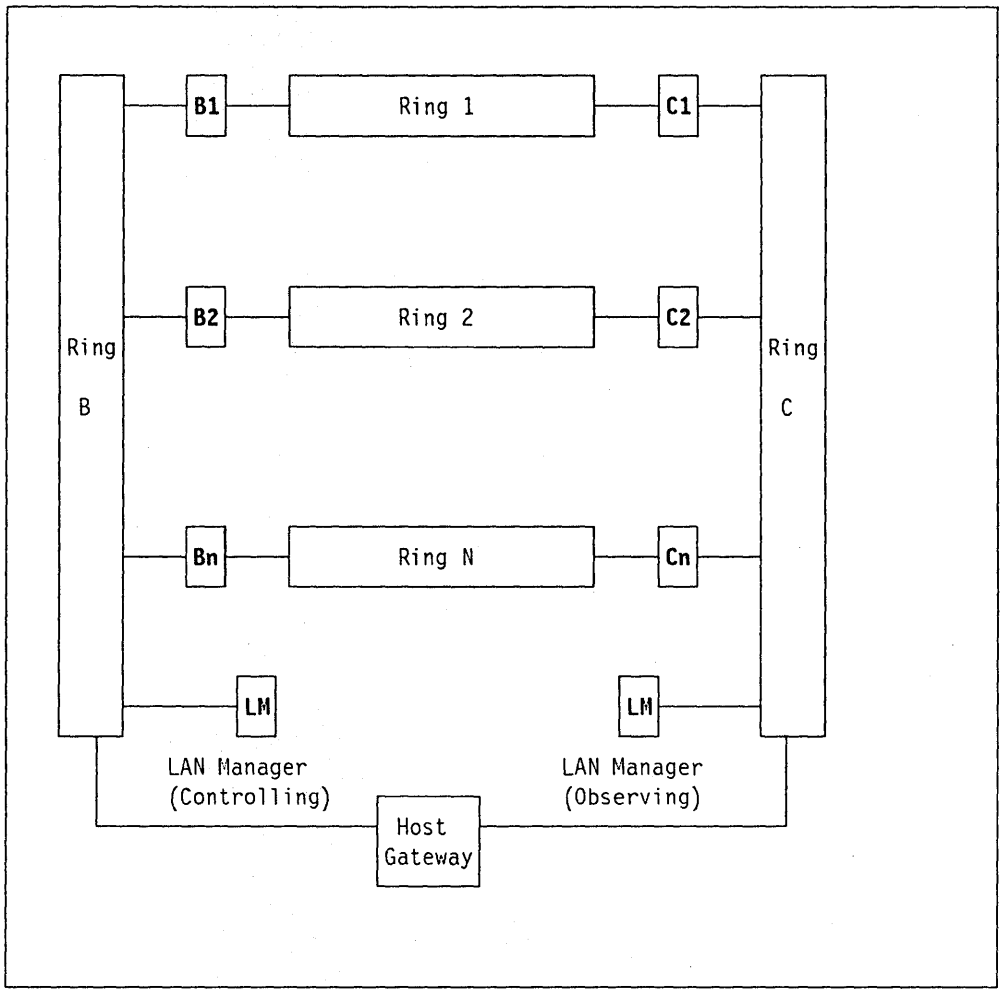

Figure 47. Availability Using a Dual Backbone Configuration.

As you can see in this figure, there are two backbones in this configuration.

Each user ring is connected via two bridges to each backbone.

Although this topology involves more bridges and an additional backbone, there is no single point of failure in such a configuration. The major benefits of this topology are:

- If a bridge fails, the user will be able to reestablish his sessions with his partners via an alternate route without operator intervention.
- Likewise, a backbone failure will not prevent a user from using an alternate route to communicate with other stations.
- Finally, from a performance point of view, the load of the backbones will usually be automatically and naturally balanced during session establishment thanks to the source routing algorithm.

Regarding servers or host connections to the LAN, you have several solutions:

- Connect the LAN servers or hosts to the backbone(s).

  This is the most usual configuration.

- Connect the host directly to all user rings.

  This can be done when the traffic to the host is very important but the number of rings is not too high.

The following scenarios illustrate these solutions.

9.3.2 Host-Connected Backbones

The most usual configuration for a multisegment LAN with a host connection is a dual backbone topology, with the host (or servers) connected to the backbone. Connecting the host to the backbones is usually the best solution to share host resources in an efficient way. However, as we will see in the section "Host Connected User Rings" on page 104, there is an alternative when this traffic is very heavy.

A typical dual backbone configuration is shown in Figure 48 on page 100.

Although the selection and positioning of different host gateways is beyond the scope of this document, some typical configurations are illustrated in the following scenarios.

The topology shown in Figure 48 on page 100 provides you with the following benefits:

- The traffic flow to the host is spread over N bridges or 2xN bridges, depending on the number and type of gateways to the host.

  If the host gateway is a 37xx communication controller with several TICs (as shown in Figure 48 on page 100) then the traffic to the host will flow via all bridges and the two backbones.

  If the host gateways are IBM 3174/1L type gateways, you should connect at least one of them to each backbone, in order to spread the traffic over the two backbones.

- The availability of access to the host is very high. If a backbone or one of the TICs fails, the users will still access the host via the other backbone and TIC. (This can be done automatically depending upon the products in use.)

  In case of 3174/1L gateways, the automatic backup can be implemented by having two active 3174/1L with duplicate address, each of them being connected to a single backbone. An example is shown in Figure 50 on page 102. Note that this requires duplication of host definition for complete backup and that some network management command is required.

- The reliability of the host-connected backbone rings is high, because there are no user workstations on those rings (with resulting frequency in changes or insertion processing).

9.3.3 Host Attachment via a 37xx Gateway

With its numerous hosts and Token-Ring attachments, the 37xx communication controller (for example the 3745 or 3725) is certainly the best gateway for large networks with many hosts from a connectivity point of view. In addition, 37xx gateways offer excellent backup and recovery facilities. A typical host attachment via a 37xx gateway is shown in Figure 48 on page 100.

As you can see in this figure, each 37xx is attached to both backbones via two TICs (Token-Ring Interface Couplers). NCP [13] 1 TIC 1 is connected to backbone ring B while NCP1 TIC 2 is connected to backbone ring C. On each 37xx, both TICs can be operational, provided there are no duplicate TIC address on each backbone. (Both TIC 1 have the same ring address and both TIC 2 have the same ring address, but different from TIC 1.)

This configuration currently offers the highest availability and completely automatic backup facilities with appropriate NetView CLISTs. In case of a failure of a backbone ring, a bridge, a TIC, or even an NCP, a station will always be able to reestablish its sessions with the host via an alternate route.

---

[13] NCP is a trademark of the International Business Machines Corporation

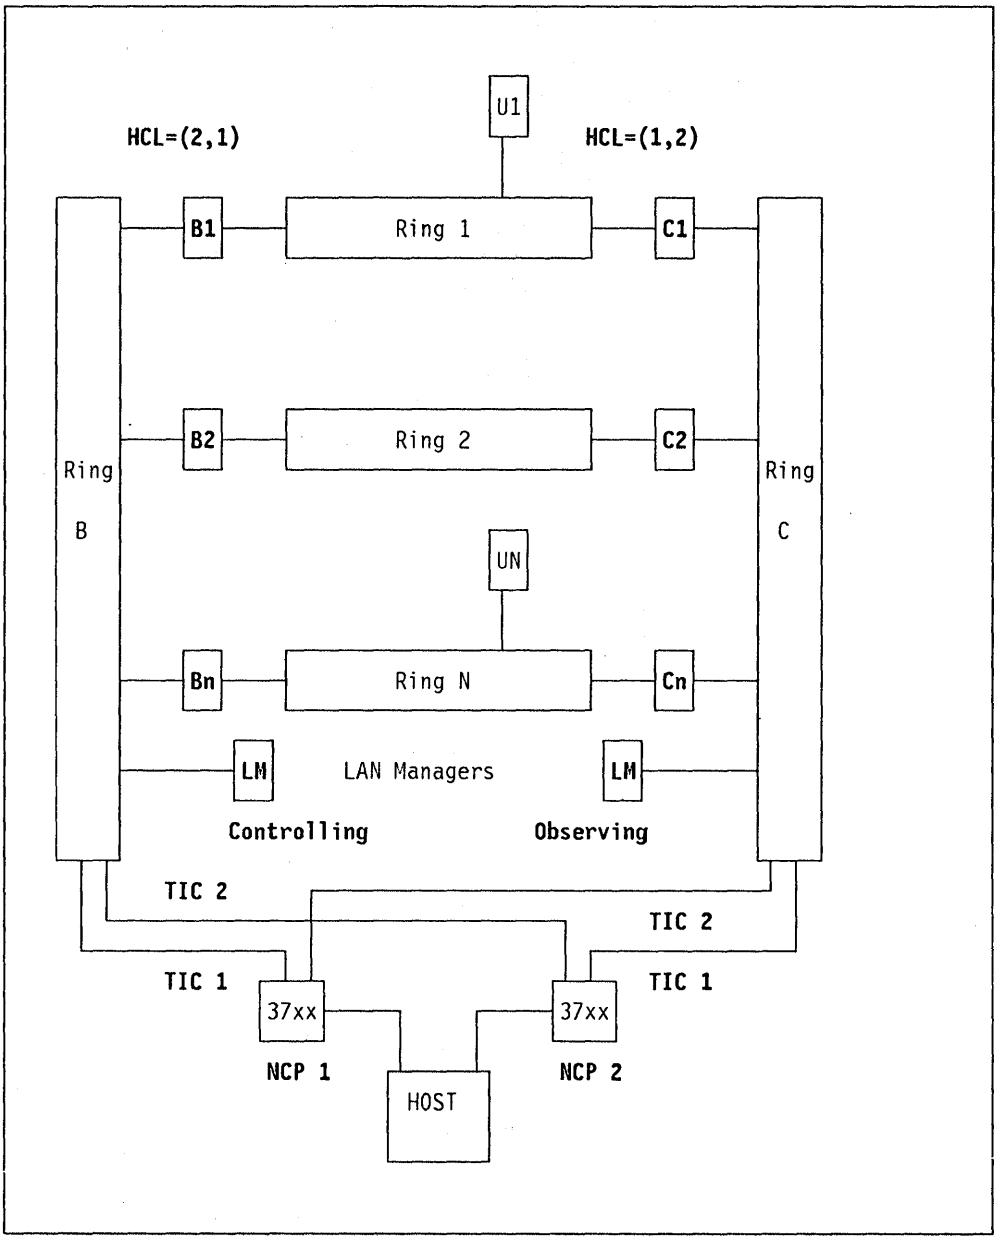

Figure 48. Host Connection via Two 37xx and Four TICs.

A controlling LAN Manager is usually connected to one of the backbones and another observing LAN Manager can be connected to the other backbone. That way, if a backbone or a LAN Manager fails, the other LAN Manager will always be available for LAN management purpose.

In this configuration, you could set the bridge parameters as follows:

- Select the automatic single-route broadcast option on all bridges
- Set the hop count limit value to (2,1) on all bridges.

  The value of 2 is for the backbone side bridge adapter, the value of 1 is for the non backbone side adapter.

With these options and values, the broadcast traffic is reduced to a minimum.

However, if you have a slightly different configuration like the following, with only one TIC 1 and one TIC 2 active, you still should use the automatic single-route broadcast option, but the hop count limit value should then be set to (3,3) for each bridge.

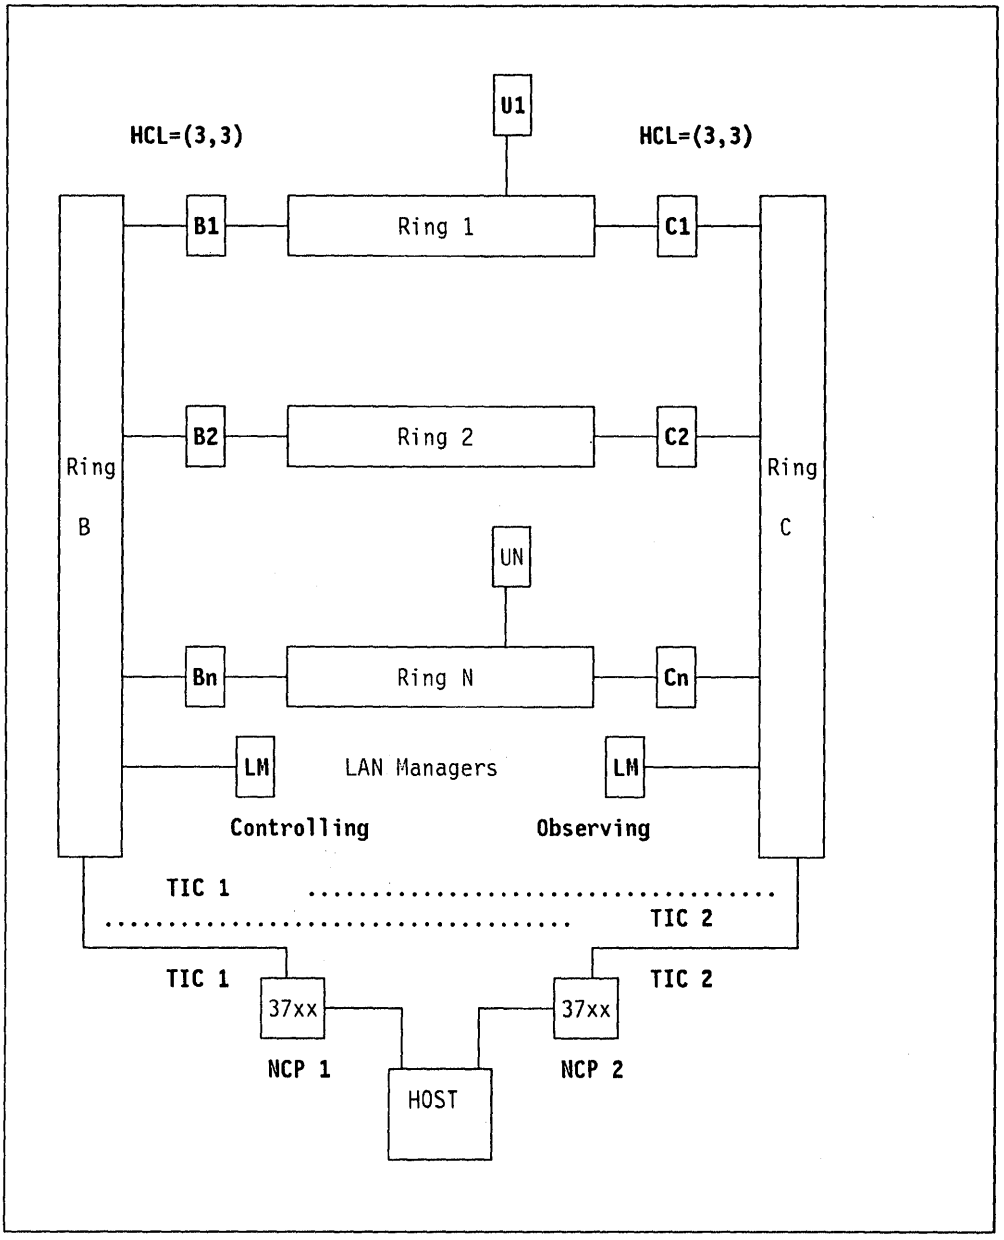

Figure 49. Host Connection via Two 37xx and Two TICs.

For example, suppose station U1 uses the TIC 1 ring address to connect to the host. A probable route between U1 and TIC 1 would be (Ring 1, Bridge B1, Ring B). If bridge B1 fails, station U1 will have to use a route like (ring 1, bridge C1, ring C, bridge CN, ring N, bridge BN) to access TIC 1 of NCP 1. As you have three bridges in such a route, the hop count limit value must be at least (3,3), and not (2,1) as in the previous example.

In case of an active TIC failure (for example TIC 1 NCP 1), it is necessary to activate the backup TIC (TIC 1 NCP 2) of the other communication controller (this activation can be automated with appropriate host definition)

Note that in this configuration, the active TIC addresses are unique in the network, which is not the case in the previous scenario.

9.3.4 Host Attachment via Two 3174/01L Gateways

The 3174/01L gateway offers very good price performance solution for medium size LANs. Although it does not offer all the flexibility of the 37xx (especially in terms of device definition requirements and host connectivity), the following configuration also offers fully automatic recovery capability. In this example, both 3174 gateways are activated with the same ring address. If a backbone or gateway fails, no intervention is required and sessions can be reestablished with the host via the other backbone and the other 3174/01L (only if duplicate PU definitions and all PUs/LUs activated).

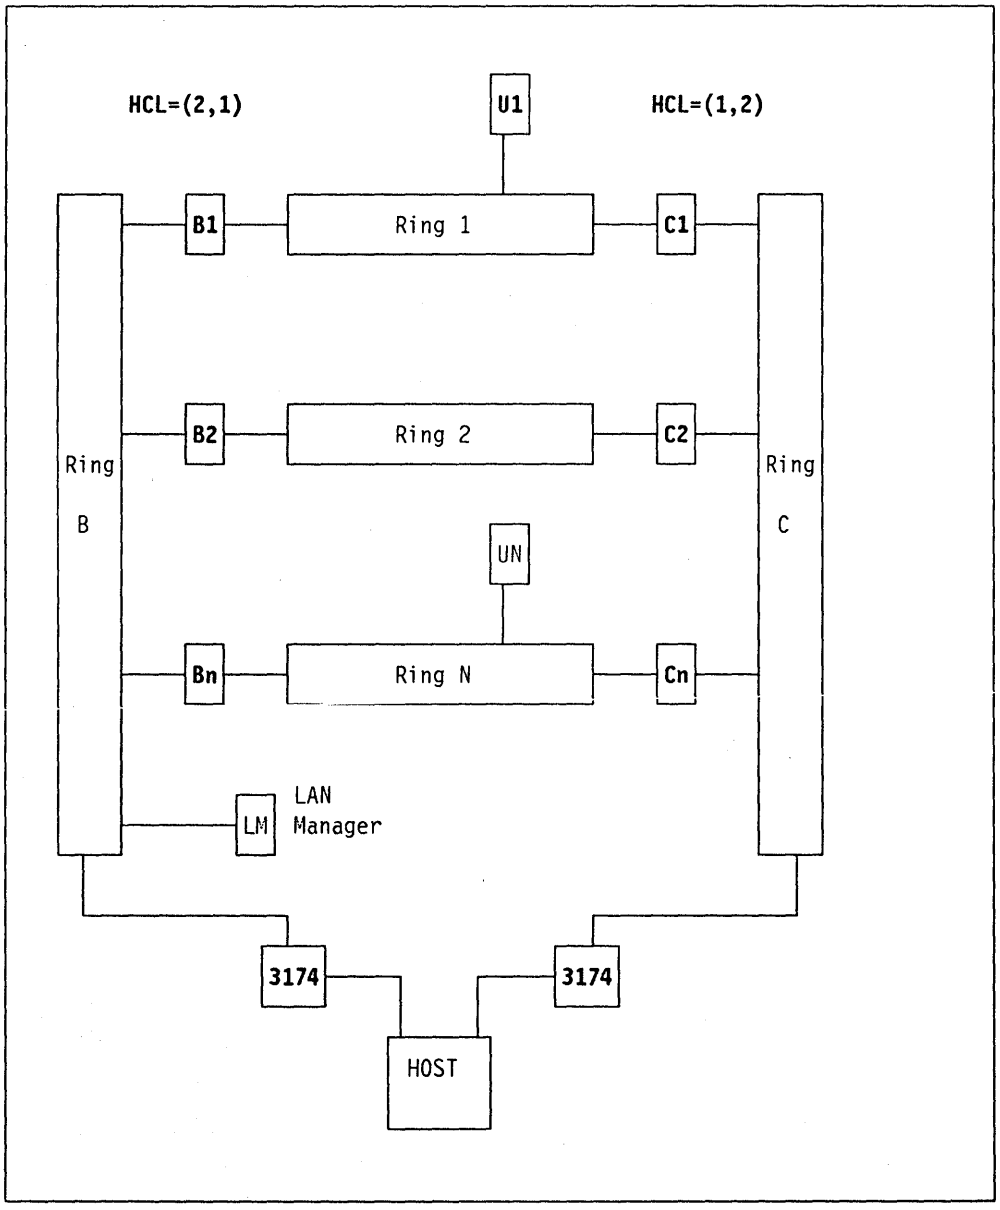

Figure 50. Host Connection via Two 3174/01L

9.3.5 LAN Server or Host Attachment via a Single Gateway

The following configuration will apply if your LAN server is a PC type server, or if your host gateway (or host in the case of a 9370 or AS/400) is connected to a single backbone as shown in Figure 51. In this case, you lose many benefits mentioned in the previous scenarios, because you break the symmetry of the configuration.

If the LAN server (or gateway) or the backbone attached to it fails, there is no automatic backup capability.

In addition, the total traffic to the server flows through a single backbone.

Finally, the hop count limit value can not be set in this case to (2,1), as in the other scenarios. If a bridge like B1 fails, then station U1 must traverse bridges C1 C2 and B2 in order to access the server. Thus the hop count limit value must be (3,3) in this configuration.

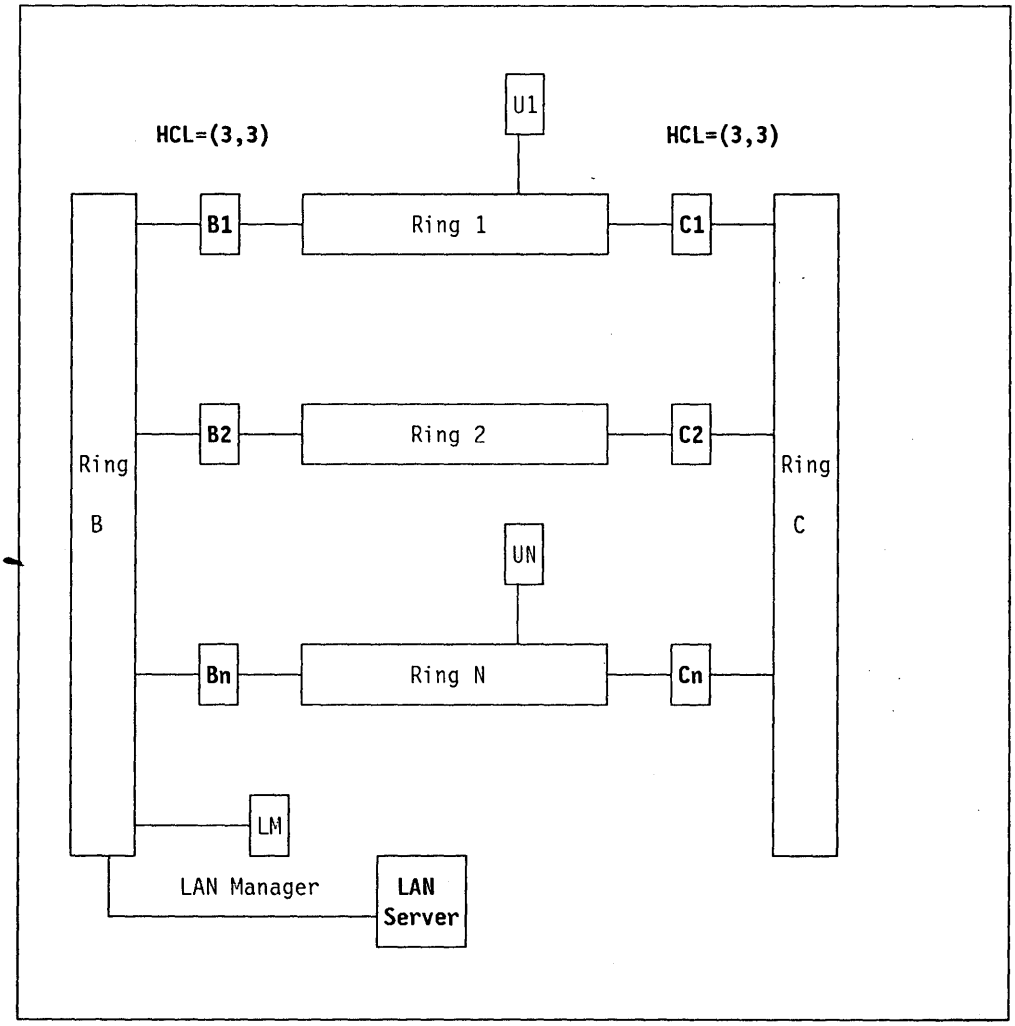

Figure 51. LAN Server Attached to a Single Backbone

9.3.6 Host Connected User Rings

When significantly heavy and important LAN traffic is host-oriented, for example with a lot of 3174/3R control units or stations running IBM 3270 Workstation Program V1.1 with many file transfers with the hosts, you can also consider a solution such as shown in Figure 52. In this example, the host gateways are 3745 communication controllers with four TICs. So each gateway can be directly connected to four user rings.

On each gateway, two TICs are active and each gateway handles the host traffic coming from two user rings. For example 3745 A's TICs 1 and 2 manage the host traffic coming from rings 1 and 2, while 3745 B's TICs 3 and 4 manage the host traffic coming from rings 3 and 4.

The two others TICs are backup TICs for the two other user rings. These TICs are normally not active and can be automatically activated by NetView Clists in case of the other 3745's TICs or NCP failure. For example, 3745 A TICs 3 and 4 will be activated if 3745 B fails.

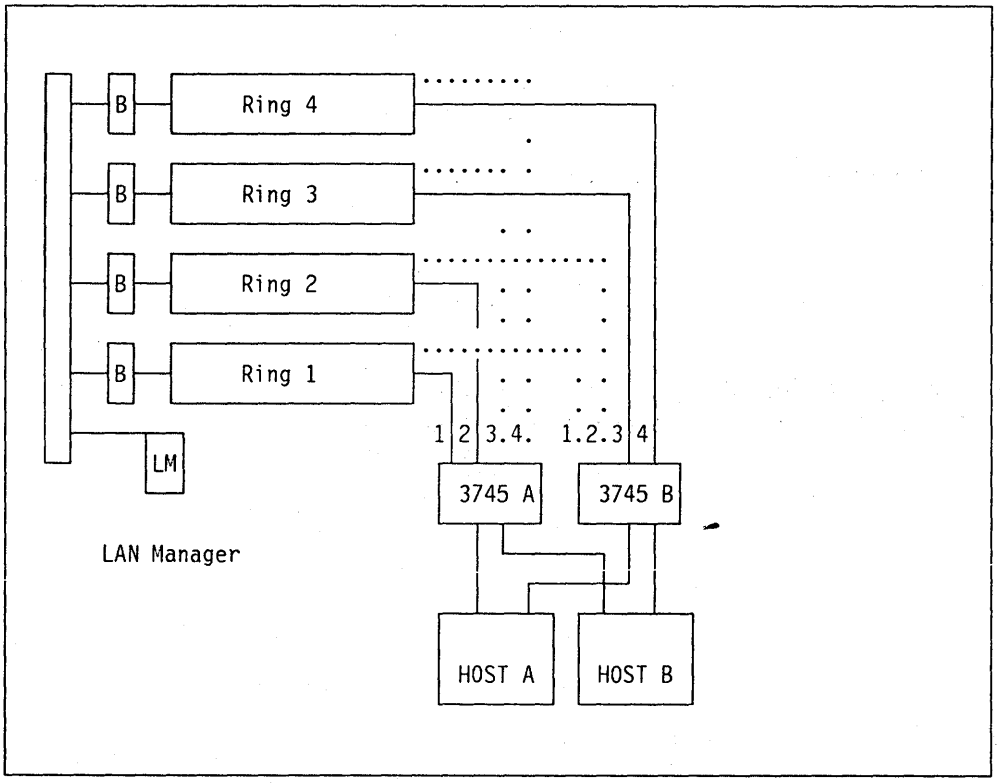

Figure 52. Host Connected User Rings. Solid lines indicate active TICs. Dotted lines are non-active backup TICs.

Some characteristics of this configuration are the following:

- The host traffic does not flow via the backbone.

  If the backbone or a bridge fails, the host traffic is not impacted.
- If a 3745 fails, it is possible to have an automatic backup recovery.
- The backbone is used only for inter-ring communication.

As we assume that the main traffic is host oriented, the backbone is not likely to be a bottleneck and is not duplicated in our example. The LAN Manager is connected to the backbone to control all stations in the network.

- The bridge parameters could be set in this case to default values:
  - The automatic single-route broadcast facility is not useful as there is already a "natural" single-route broadcast path in this particular configuration.
  - The hop count limit value could also be set to the default, or set to (2,1), 2 being the hop count limit value on the backbone side.
- The network response time for host oriented applications is minimized, as there are no intermediate bridges and rings to traverse to access the host.
- This approach could imply many TICs (or host gateways) if the number of rings is high.

9.3.7 Three Level Hierarchy Sample Scenario

For very large networks or campus environments, topologies with only first level backbones are sometimes not sufficient. In this case, your design should include one or more second-level backbone rings, and even third-level backbones if necessary.

An example of a second-level backbone topology is shown in Figure 53 on page 106.

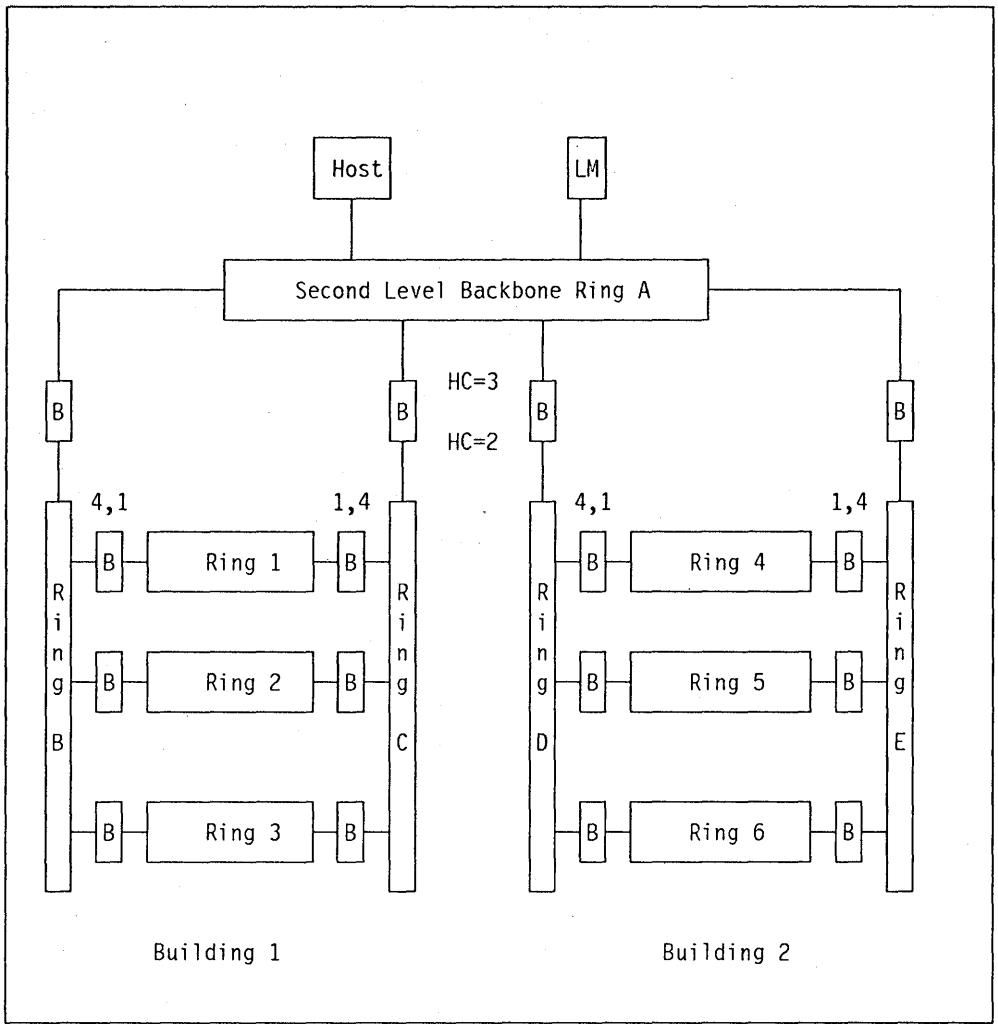

Figure 53. Three-level Hierarchy with Any-to-Any Connectivity

As you can see in this example, there are two different buildings, each of them having a dual first-level backbone like rings B and C for building 1. The figure would be similar if you had more than two buildings in your network.

All first-level backbones are connected to a second-level backbone ring (ring A). Shared resources like 370 hosts are connected to the second-level backbone ring. This second-level backbone ring could be an optical fiber backbone running through a campus between the different buildings.

All users from all buildings can access the host(s) resources.

If you also want "any-to-any" connectivity between all buildings, then you should set the hop count limit values for the bridges as indicated in the figure. HCL=(3,2) is appropriate for the four bridges connected to the second-level backbone ring, and a value of (4,1) should be selected for all other bridges (4 being the hop count limit value for the first-level backbone side bridges' adapters). For more explanations on the hop count limitbridge parameter, see "Hop Count Limit" on page 39 where this example is discussed.

In addition, as in all complex networks, it is highly recommended to use the automatic single-route broadcast facility. This will reduce significantly the

single-route broadcast traffic and hence the number of frames received by the servers.

However, if for some reason you don't want the users to communicate between buildings, you could set the hop count limit parameters to lower values, as indicated in the following figure.

In this example, all users from all buildings can access the host(s) resources. But with these hop count limit values, a user located in building 1 can not communicate with another user or server located in building 2.

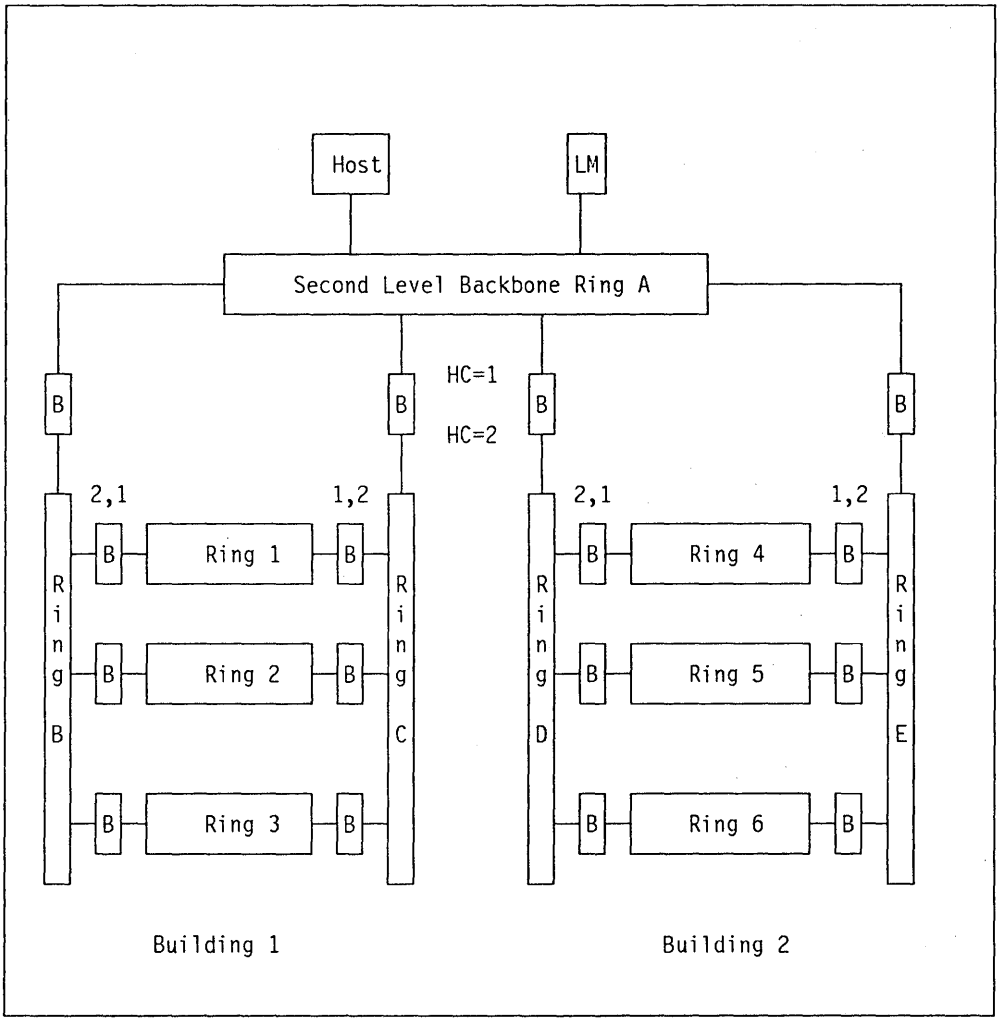

Figure 54. Three level Hierarchy with No Inter Building Connectivity

9.4 Remote Bridge Considerations

If you have a remote bridge in your configuration, your topology and bridge parameter settings must be considered with special attention.

An example of a topology including a remote bridge is shown in Figure 55 on page 108.

From a performance point of view, it is important to avoid unnecessary traffic flow via the TP link.

The bridge parameters could be set in this case to the following values:

- The automatic single-route broadcast facility should be selected to reduce the single-route broadcast traffic. As a result, remote ring 3 should receive only one frame for a single-route broadcast connection.
- The hop count limit value could be set to (4,1) for the remote bridge and (3,3) for all other bridges, in order to maintain a full connectivity. The value of 4 has to be specified for the remote bridge on the backbone side in case of a bridge failure. If bridge B2 fails (for example), a ring 2 attached station has to traverse 4 bridges to communicate with a ring 3 attached station (C2, C1, B1, B3).
- The filters should be used depending on your applications or address naming conventions to filter unnecessary traffic as explained in "The Filtering Facility" on page 71.
- Frame size, response timer (T1), and retries parameters should be modified to reflect appropriate values for link speed and quality as described in "Performance Considerations" on page 65.

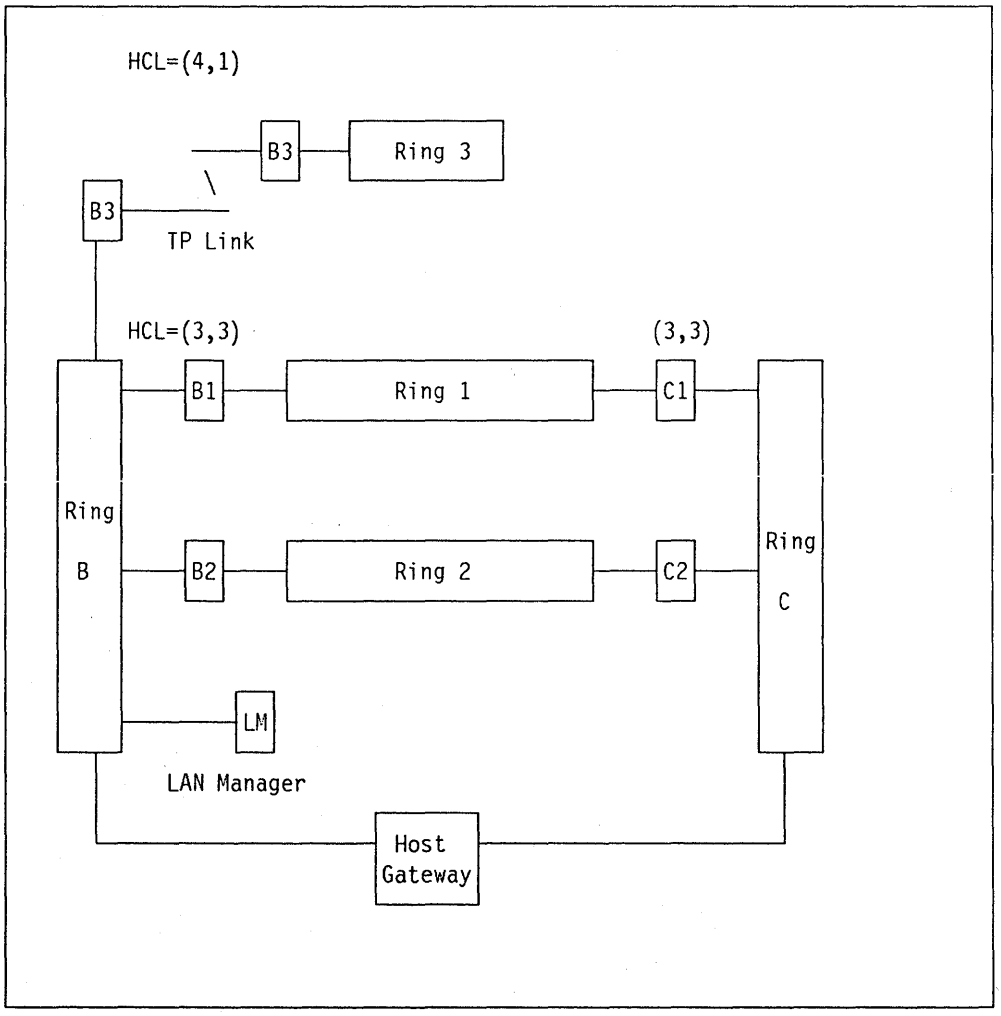

Figure 55. Remote Bridge Configuration

9.4.1 Parallel Remote Bridges

An example of a topology including "parallel" remote bridges is shown in

Remote bridges B3 and C3 are not true parallel bridges as they do not connect ring 3 to the same ring. But they do provide alternate "parallel" paths from local rings 1 or 2 to ring 3. In any event, it is recommended to connect these two remote bridges to the two different backbones for performance and availability reasons, as shown in Figure 56.

- If a backbone or a TP link fails, the traffic will flow via the other link and backbone.
- If the two TP links have different speeds, most of the traffic will be naturally routed via the fastest link because of the source routing algorithm based on the first response received by the source station.

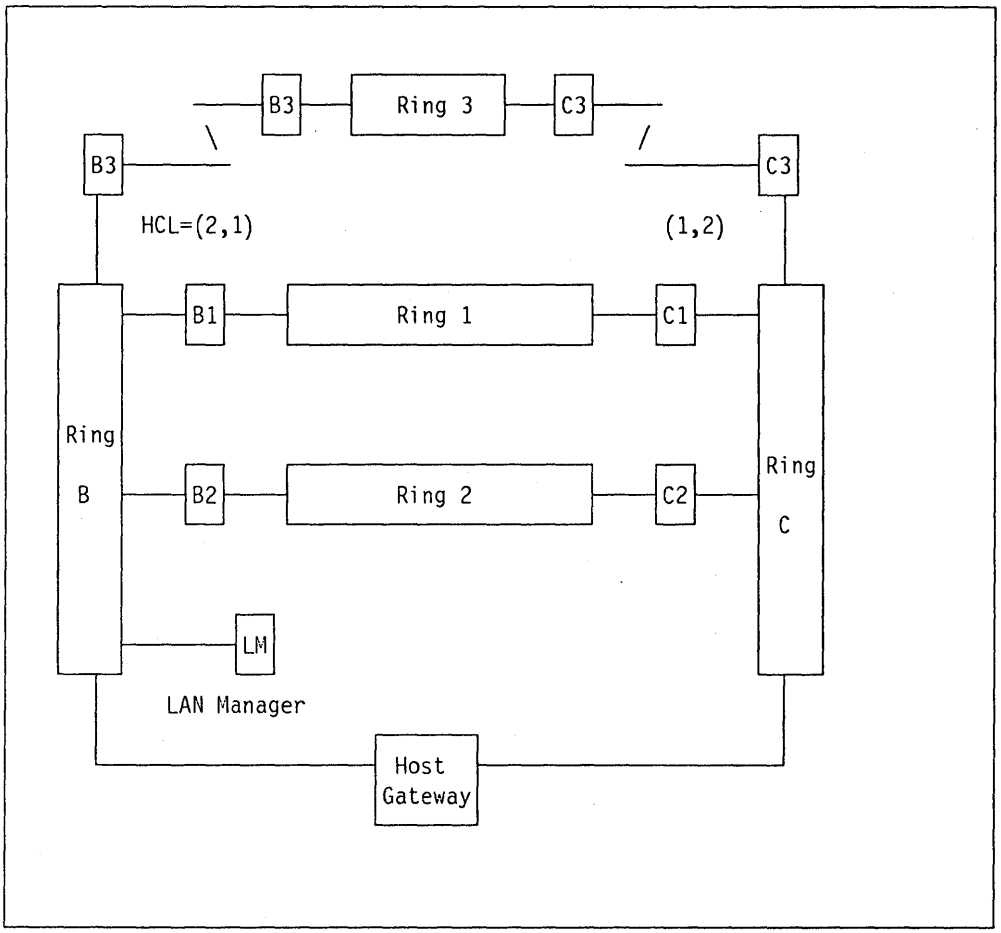

Figure 56. Parallel Remote Bridges Configuration

In this configuration, it is very important to avoid unnecessary traffic across the TP links because of the potential degradation in performance due to link speed. All advices given in "Broadcast Traffic Control" on page 35 apply. As a result, you should specify the following options for the bridge parameters:

- The automatic single-route broadcast facility should be selected in all bridges to reduce the single-route broadcast traffic. As a result, remote ring 3 should receive only one frame for a single-route broadcast connection.

- As the configuration is symmetric, you can set the hop count limit value to (2,1) for all bridges, including the remote bridges.
- In order to be sure that a remote bridge is not selected as the "root bridge", you should give them a bridge label higher than the other local bridges (8000 is the default bridge label value).
- One of the remote bridges will be automatically selected as the designated bridge for ring 3. The one with the fastest TP link speed should be automatically selected, as the default path cost value depends on that speed (the higher the speed, the lower the path cost).
- The filters should be used depending on your applications or address naming conventions to filter unnecessary traffic on the TP links as explained in "Bridge Filters" on page 72.

9.4.2 TCP/IP Considerations

In complex networks where there are multiple routes with different speeds, some TCP/IP implementations choose the slowest route, due to the ARP (address resolution protocol) mechanism. When an IP station receives an ARP command for which it is the intended destination, it updates its IP location cache with the source IP address and route. Therefore, the last message arriving via the slowest path will trigger an update of the cache, which will be used during the duration of the subsequent connection.

For example, suppose you don't use the hop count limit value indicated in the figure and let the default value (7,7) in all bridges. A TCP/IP station on ring 1 which wants to communicate with a VM TCP/IP application in the host may be routed via the slowest path, that is for example via bridges B1, B3, and C3, instead of just traversing bridge B1 (or C1). If you specify a (2,1) value for the hop count limit of all bridges, the slowest path previously mentioned (via the remote bridges) is not possible any more.

Another solution to force the routes used by TCP/IP stations is to use TCP/IP "router stations" which can be attached to different segments in the LAN. This approach implies additional costs and manual definitions in TCP/IP stations to specify the router address to be used to access to the destination station.

10. Backbone and Bridge Physical Design Guidelines

During the design process, physical topology considerations play an essential role. While cabling and wiring considerations are beyond the scope of this document, we will address in this section general guidelines related to:

- Bridge station hardware considerations
- Backbone and bridges physical design.

Several typical physical topologies are given as examples in various environments like high buildings or campus.

10.1 Bridge Station Considerations

Many LAN applications are dependent upon the high bandwidth provided by the LAN environment and will experience time-outs or similar errors if their stringent response-time requirements are not met. Consequently, bridges must provide sufficient throughput to support these applications and to avoid response-time delays for the end users. This throughput is provided by:

- Use of dedicated bridge workstations
- Use of faster workstation processors (80286 or 80386-based PC/AT or PS/2 workstations)
- Efficient bridge logic which reads only those frames which need to be handled by the bridge
- Transmission of frames at a higher priority (priority 4).

Use of the IBM Token-Ring Network 16/4 Adapters is recommended to support transmission of larger frames (up to 8K) to further improve performance potential and to support future LAN expansion.

In a local environment, use of a bridge is typically transparent to the end user. In a remote LAN environment, bridge throughput and transparency to applications and users is highly dependent upon the speed of the communications link interconnecting the bridge "halves" as discussed in "IBM Token-Ring Network Bridge Program V2.1 "Local" or "Remote" Bridge Function" on page 53.

Since IBM bridge workstations described in this document are entirely dedicated to the bridging function, IBM bridge programs have been designed to operate in the single-task IBM PC/DOS 3.3 or 4.0 environment.

Disk and memory requirements are minimal for a bridge station. A bridge will almost never access its disk under normal circumstances.

The remote bridge throughput and additional delay is highly dependent on the TP link speed as discussed in "IBM Token-Ring Network Bridge Program V2.1 "Local" or "Remote" Bridge Function" on page 53. IBM PS/2 equipped with the IBM X.25 Interface CoProcessor/2 are recommended as remote bridge stations as they can sustain much higher speeds than PC stations equipped with the IBM Realtime Interface Co-processor Multiport Adapter

Bridges should be located in a secure area (for example a wiring closet) or operations area where they will not be misused. As a bridge may need to operate 24 hours a day, it is recommended to have adequate ventilation in very small bridge rooms or wiring closets to stay within the allowed temperature range. In hostile environments, one should consider selecting the IBM 7531 or 7532 Industrial Computer for the bridges.

Bridge distances, management and space considerations are discussed in the next section and have a strong impact on the bridges' location and backbone topology.

10.2 Centralized Backbone Approach

The best approach to design a backbone ring is to keep it as simple and as short as possible. An ideal backbone consists of a single 8228, placed in a wiring closet with all bridges connected to it. There are several advantages to do so:

- The backbone is very reliable and has minimal probability of cabling problems or failures.
- As the backbone ring includes one (or very few) 8228(s), you can use longer lobes in order to connect the host gateway(s) or even the bridges if they can not be centralized, as shown in Figure 58 on page 114.
- If the backbone fails, you can replace it by changing only one 8228.

10.2.1 Centralized (Collocated) Bridges Example

The example below shows a centralized backbone with centralized bridges configuration.

The backbone consists of a single 8228. B stands for bridge and S means it is a spare bridge. LM stands for LAN Manager and Gwy stands for gateway.

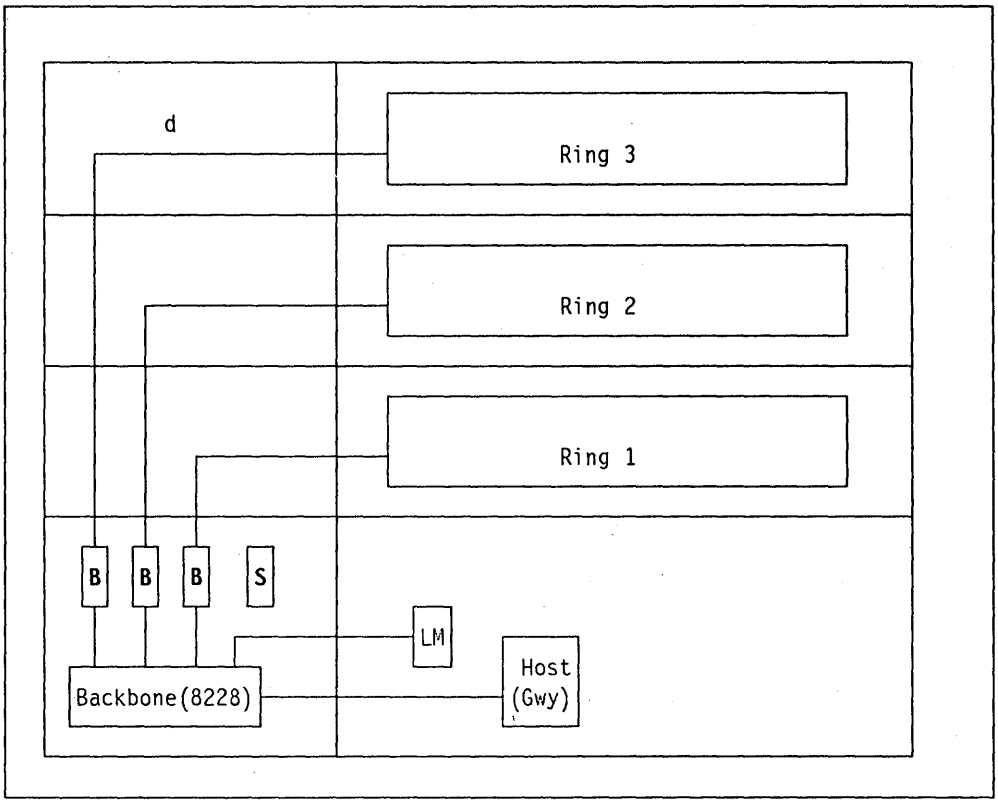

Figure 57. Centralized Bridges and Central Backbone Topology

10.2.2 Centralized Bridges Advantages

Collocation of bridges should be considered for the following advantages:

- The whole backbone ring is situated in a single room (that is, the wiring closet). There is no need to go to another wiring closet in case of a bridge problem.
- The maintenance of the bridges is easy, and the backup of the bridges will usually be obtained by putting a spare bridge in the "backbone" room, as shown in Figure 57.
- Placing all bridges in the backbone room will save some space in the wiring closets located in the other floors, which is sometimes very helpful in old buildings where it is difficult to find a place for the wiring closets. The counterpart is that you will need more space in the backbone room, although you can save some space by mounting the bridges in appropriate racks.

  The maximum distances between the bridges and the centralized backbone must be compatible with the maximum lobe length allowed on each user ring. For example, as shown in the figure, if d is the distance between the bridge and ring 3, you must have:
  $d \leq$ ring 3 maximum lobe length
- The cost of the backbone is very low as there are no cables on the main path and you don't need any repeaters.
- In a centralized bridges configuration, there should be no problem to migrate to a 16 Mbps backbone ring as the bridges' lobe lengths are very

short. However, centralization of bridges as described above precludes use of fiber, and if future use of fiber (for example FDDI) is likely to be required, alternative approaches should also be considered.

10.2.3 Distributed Bridges Example

When bridges cannot be centralized for such reasons as topology constraints or the lack of space in the "backbone room", it may still be possible to have a small centralized backbone whenever possible.

Distributed bridges are usually put in wiring closets located on each floor on a vertical path. We recommend using direct cables between the 8228 and the bridges to avoid signal attenuations and potential cabling problems in the wiring closets located in intermediate floors.

An example of distributed bridges with centralized backbone is shown in the following figure:

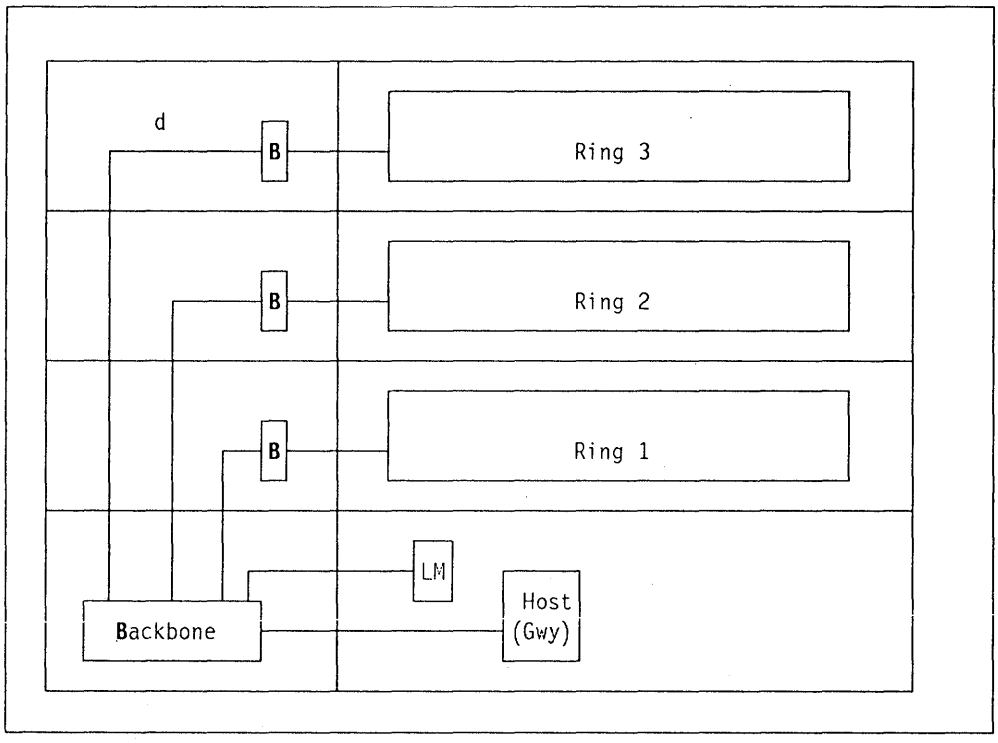

Figure 58. Distributed Bridges and Central Backbone Topology

If the backbone consists of only one 8228, the maximum distance you can use to reach a bridge with IBM cables type 1 or 2 is depending on the ring speed:

- If the backbone ring speed is 4 Mbps, the actual maximum lobe length is about 350 meters, although we recommend not to go over 100 meters for future growth of the backbone ring.
- If the backbone ring speed is 16 Mbps, the actual maximum lobe length is about 160 meters, with the same recommendation as above.

In both cases, if we limit the maximum lobe length to 100 meters, the distances allowed by a centralized backbone will permit you to cover most single building

configurations, unless the number of floors is very high. If d is the longest distance between the backbone and a bridge, you must have:
$d \leq$ Backbone ring maximum lobe length

Although not shown in the figure, two bridges connected to a dual backbone are usually recommended for each user ring in high availability configurations.

10.3 Distributed Backbone

You usually will consider a "distributed" backbone because the topology of the building(s) is incompatible with a centralized backbone.

A distributed backbone is composed of several 8228s spread over several wiring closets interconnected via copper cables or optical fibers(see Figure 59) The backbone ring can either be routed through all wiring closets, as shown in our example, or routed through only some of them with some concentration effect in these selected wiring closets.

10.3.1 High Building Example

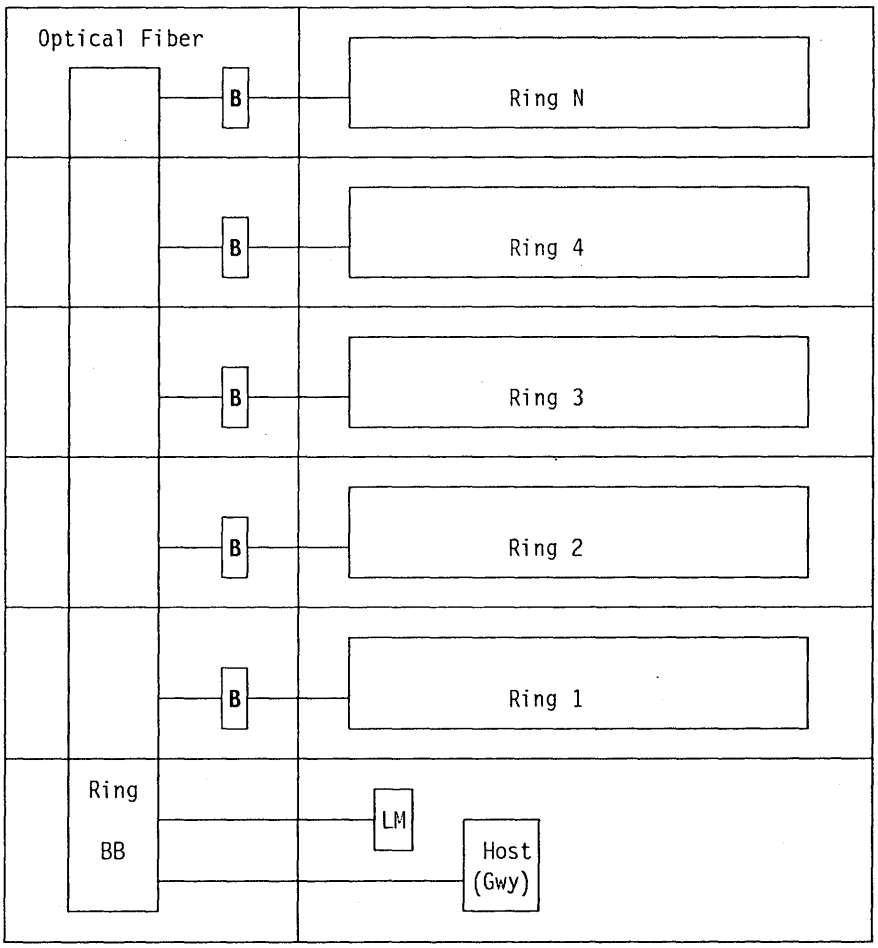

Figure 59. Distributed Backbone Topology in a High Building Environment

As the main path is much longer than a centralized backbone and as you can have many components like 8228s and repeaters, a distributed backbone may be more exposed to cable problems or component failures. Therefore, although only one backbone is shown for clarity, a dual backbone configuration with two bridges per user ring is highly recommended for availability.

The two backbones can be located in the same wing or in opposite wings as shown in Figure 60

One advantage of a distributed backbone is that the number of "vertical" cables is reduced to a minimum. For example, as shown in Figure 60 each backbone may consist of only two vertical cables. With a centralized backbone approach, you would have needed at least N cables (one for each bridge).

Note that it is sometimes possible to implement a multi-floor ring when the number of users per floor is very low. In addition, the number of 8228s belonging to the main backbone path and located in selected floor wiring closets can be reduced as a single 8228 can connect up to eight bridges.

In many buildings, optical fiber and repeaters are strongly recommended in order to support greater distances speed or future migration to an FDDI backbone. In addition, use of an IBM 8220 converter to support such fiber cabling provides additional availability and manageability of the cabling.

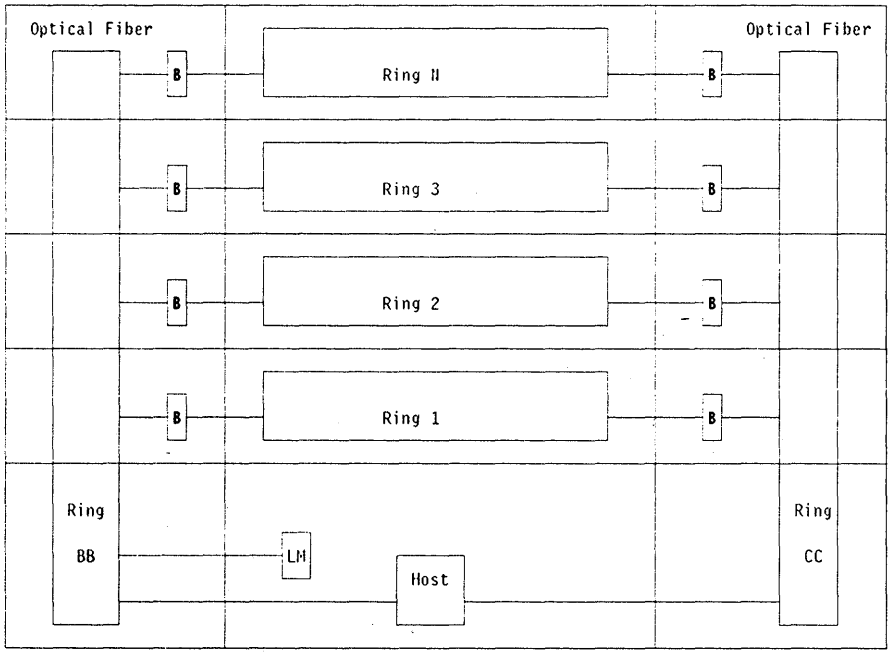

Figure 60. Distributed Dual Backbone Topology (High Building Environment)

10.3.2 Campus Example

Another example of a typical distributed backbone in a campus environment is shown in Figure 61 on page 117.

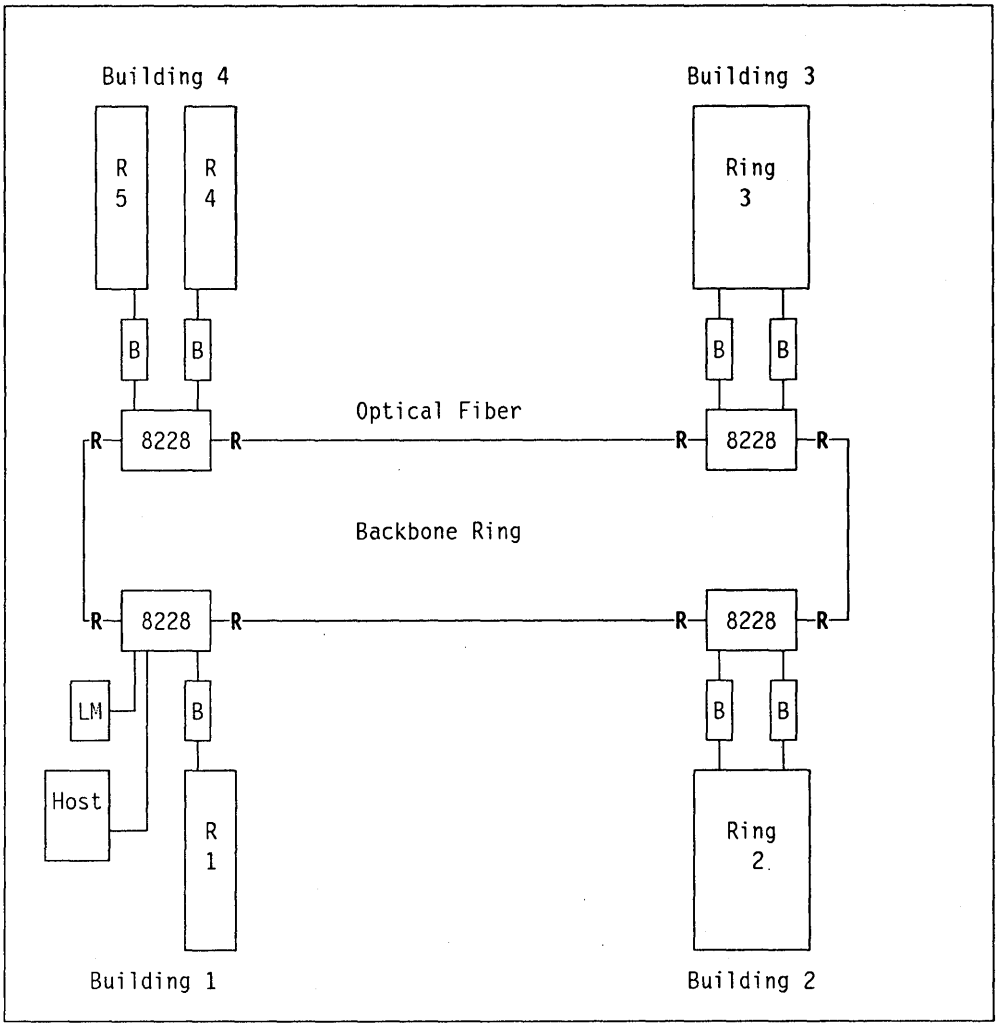

Figure 61. Distributed Backbone in a Campus Environment. (**R** stands for repeater)

In this example, there are four different buildings separated by long distances. The usual topology is to have a distributed optical fiber backbone interconnecting the four buildings. A pair of optical fibers is the ideal medium to interconnect different buildings, and the new 8220 repeaters offer a very high level of reliability. Even if an optical fiber cable is broken, the 8220s will automatically wrap at both ends to the backup path.

Depending on the size of the user rings and the availability user requirements, each ring should be connected to the backbone (or better to a dual backbone) via two bridges, as shown for rings 2 and 3.

A single LAN Manager (or NetView operator) located for example in building 1 can efficiently manage all LAN resources in the campus.

11. LAN Management Considerations

In larger local area networks such as a multisegment LAN, management assumes increased importance from the standpoint of problem awareness and recovery procedures. The bridge products described in this document use the facilities of IBM LAN Manager Version 2.0 to provide enhanced control of attached segments and of the bridges themselves.

The LAN Manager Version 2.0 is described in detail in *LAN Manager V2.0 and LAN Manager Entry V1.0 Installation Guidelines* This chapter will highlight those aspects of multisegment design and operations which are impacted by these management capabilities, and will point to appropriate references for additional information.

11.1 LAN Management and IBM Network Management

Management requirements for local area networks have most of the functions of wide area networks, and consequently it is important that IBM LAN management implementations be consistent with IBM system management architectures, and that they provide support for local or central operator control. The latest version of NetView incorporates support for local and remote control of the LAN via the LAN Manager products.

Network management can thus be implemented at various levels depending on the network design and the requirement for centralized management. For instance, the management of a remote Local Area Network can be done in three ways :

1. At a local level with IBM LAN Manager V2.0
2. At the central site with NetView R3 and IBM LAN Manager V2.0 or IBM LAN Manager Entry V1.0 at the Local Area Network
3. A combination of centralized and local management with NetView R3 and IBM LAN Manager V2.0.

Because IBM LAN Manager V2.0 is required to support the IBM bridges, we will not further consider use of IBM LAN Manager Entry V1.0.

IBM LAN Manager V2.0 supports network management for a multitude of LAN configurations, for 4Mbps rings to 16 Mbps rings, for local and remote bridged segments and for Token-Ring and PC Network segments.

IBM LAN Manager V2.0 is an integral part of IBM's total network management strategy. In a host connected environment the LAN Manager acts as a service point, the LAN stations act as the targets and NetView R3 acts as the focal point. In a standalone environment (no host-connection), the LAN Manager itself performs a number of focal point functions.

Local LAN management server functions (like those in the bridges) can be integrated with a service point (LAN Manager) when the LAN is part of a larger wide area network, and thus relocated to a distant focal point. While the LAN subsystem is then controlled by the overall network management system, the local LAN management capability still remains if needed. LAN operation and

management can be shared, if desired, between a local operator via a local operator interface and the focal point network manager.

11.1.1 LAN Management Relationships

With the latest LAN management product announcements in the Operating System/2 [14] Extended Edition V1.1 environment, IBM increases the range of management function and connectivity options between service points and focal points.

Figure 62 on page 121 presents an overview of the various IBM LAN Management products, their hierarchy within the network management system and the coverage within the LAN subsystem.

LAN Management Products Hierarchy Notes:

a This LAN Manager refers to **IBM LAN Manager V2.0**, which runs in an Operating System/2 Extended Edition V1.1 environment. It is capable of maintaining up to 64 *reporting links* **e**, with bridges. Reporting links are network management LLC sessions which are implemented between IBM LAN Manager V2.0 and IBM Token-Ring Network Bridge Program V2.0, IBM Token-Ring Network Bridge Program V2.1 or IBM PC Network Bridge Program.

LAN Manager, as a Service Point, communicates with NetView as the focal point residing in a System/370 host system via a 3174, 372x or 3745 token-ring gateway or via a 9370 token-ring attachment if NetView resides in this system. Alerts from the LAN Manager can be transmitted up to NetView via this link. This also allows LAN network management to be controlled from NetView by using Service Point Command Services so that the NetView R3 operator can issue commands to the LAN Manager and receive responses.

If only single LAN segment management is required from a focal point and no user interface is needed locally, IBM LAN Manager Entry V1.0 (also running under Operating System/2 Extended Edition V1.1) can be used. IBM LAN Manager Entry V1.0 only functions on a single segment since it has no bridge support and because it has no local user interface it can only be operated by using Service Point Command Services based LAN Manager commands from NetView R3.

b This LAN Manager would typically be **IBM LAN Manager V2.0**. If none of the previous communication capabilities with the focal point are available, the SDLC communications facility included in the Communications Manager of Operating System/2 Extended Edition V1.1 may support the communications link between LAN Manager and NetView in the host. Again, IBM LAN Manager V2.0 will provide multisegment LAN management from a focal point (NetView R3) and/or from the local LAN Manager station. This LAN Manager could also be IBM LAN Manager Entry V1.0 if no local user interface is required and it is a single-segment LAN.

---

[14] Operating System/2 is a trademark of the International Business Machines Corporation

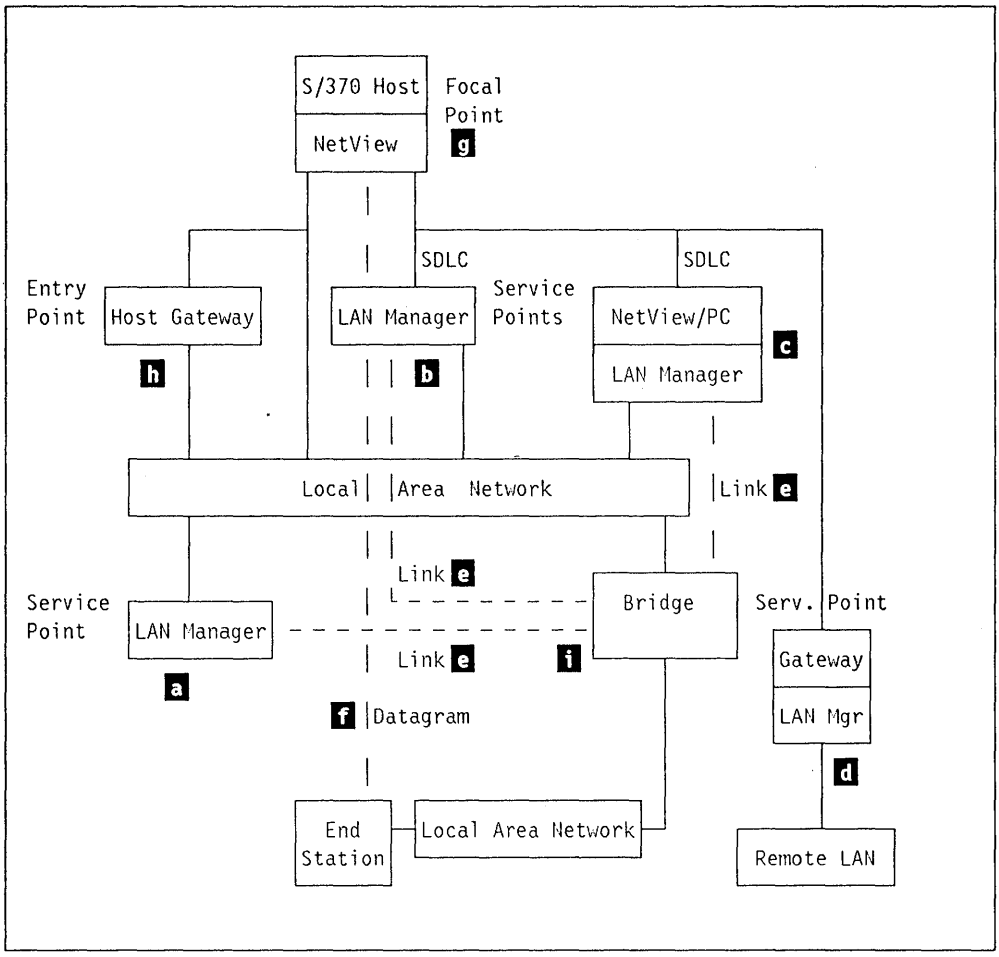

Figure 62. LAN Management Products Hierarchy

c The third LAN Manager shows IBM LAN Manager V2.0 implemented as a NetView/PC application. In this case NetView/PC provides the service point function which is connected to the NetView focal point in the System/370 Host. IBM LAN Manager Entry V1.0 cannot run as an application under NetView/PC.

IBM LAN Manager V2.0 can communicate with a NetView host using the OS/2 Communications Manager. If it is running as a NetView/PC V1.2 application, it can use the NetView/PC V1.2 Communications Manager to communicate with the host.

NetView/PC V1.2 with IBM LAN Manager V2.0 should be considered whenever LAN management from a central site management location is combined with management of other non-SNA parts of the network, taking advantage of the NetView/PC V1.2 function to achieve this.

d This refers to management of small, remote LANs connected to a central host system via a gateway function where central site network management should also cover the remote LAN subnetwork. Usually there will be no local network management expertise at the remote location. Therefore the LAN management solutions proposed in this environment are operated exclusively from the NetView console (focal point) at the host. The service point function resides in the gateway associated with IBM LAN Manager Entry V1.0 (OS/2 Communications Manager).

e This reference indicates the dedicated LLC sessions, called reporting links, between IBM LAN Manager V2.0 and one of the bridge products. A bridge is capable of maintaining up to four concurrent links with four different LAN Managers. However, only one link connects to a *controlling* LAN Manager, while the remaining links connect to *observing* LAN Managers for this bridge.

11.2 IBM LAN Manager V2.0 Overview

IBM LAN Manager V2.0 operates under Operating System/2 Extended Edition V1.1, thereby permitting coexistence with other tasks and removing the 640 Kbytes PC/DOS memory constraint of previous LAN management products.

IBM LAN Manager V2.0 can operate in three different environments as shown in Figure 63 on page 123. Parts of this figure will be referred to by the alphabetic reference numbers throughout this section.

11.2.1 Features

Stand-Alone LAN Manager. This refers to a LAN Manager **a** for which no host communication is available. All the management of the network is done locally at the LAN Manager workstation and the local full-screen operator interface is used. This LAN Manager can control and manage single and multiple segments via communication links maintained with the bridges interconnecting different LAN segments.

LAN Management Agent for NetView As a LAN management agent **b** for NetView running in the host, IBM LAN Manager V2.0 uses the Operating System/2 Extended Edition V1.1 Communications Manager facilities to access the NetView host via:

- Existing LAN gateways **d** attached directly to a token-ring segment. In this case, an exclusive communications link to the host for LAN management is not needed. Such a host gateway device, for example a 3174-01L or a 3745, can only be attached to an IBM Token-Ring Network segment.
- Alternatively, LAN Manager **b** may communicate with the host via an SDLC protocol provided by Operating System/2 Extended Edition V1.1's Communication Manager through an SDLC adapter or a PS/2 Multiprotocol adapter.

NetView R3 is recommended in the host system to provide command functions in addition to alert capabilities. The command capabilities are implemented via Service Point Command Services (SPCS). For more detail on this facility refer to LAN Manager V2.0 and LAN Manager Entry V1.0 Installation Guidelines.

A NetView/PC Application. When NetView/PC is used to provide Service Point functions for other applications or subsystems, you might prefer to run LAN Manager as a subsystem application under NetView/PC. IBM LAN Manager V2.0 can run as an application of NetView/PC V1.2 **c**. In the Operating System/2 Extended Edition V1.1 environment, NetView/PC V1.2 can now support up to 128 concurrent applications using one host reporting interface. If you want to run LAN Manager on the same machine as other applications using this host interface you must execute it as a NetView/PC V1.2 application because the host SSCP-PU session cannot be shared.

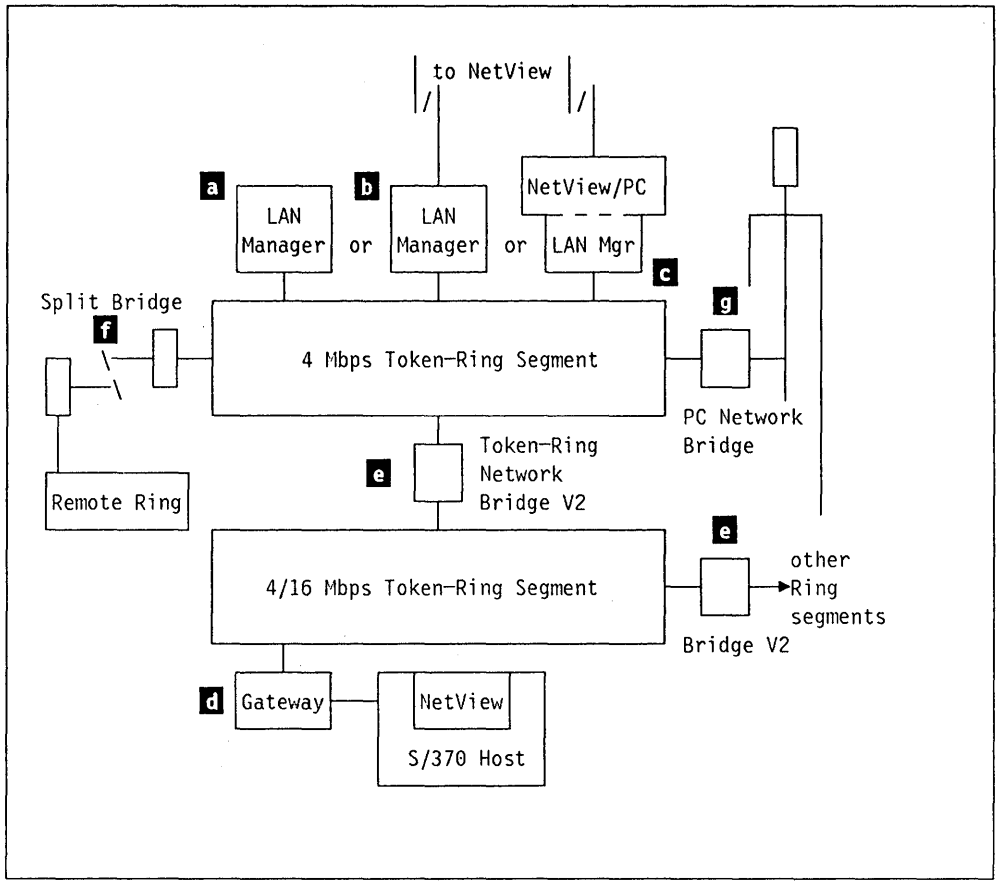

Figure 63. IBM LAN Manager Version 2.0 Overview

A NetView/PC V1.2 remote console facility is not yet available[15] , but the LAN Manager application subtask can be remotely operated via the SPCI LAN Manager command support offered in NetView R3.

The next few sections will provide more information about the extended network management capabilities offered by IBM LAN Manager V2.0.

Mixed LAN Management IBM LAN Manager V2.0 can maintain up to 64 concurrent LAN management LLC links with 64 different bridges to control a maximum of 65 LAN segment (the local segment plus 64 remote segments). These links may be dynamically updated by the network management operator.: The bridges would normally consist of IBM Token-Ring Network Bridge Program V2.0 **e** , IBM Token-Ring Network Bridge Program V2.1 **f** or IBM PC Network Bridge Program **g** . While bridges executing IBM Token-Ring Network Bridge Program V1.1 may coexist in the LAN environment and will operate properly as far as the basic bridge forwarding protocol is concerned, they cannot be configured fully from IBM LAN Manager V2.0. Such coexistence would therefore limit some of function provided by IBM LAN Manager V2.0 and the newer bridges. Migration considerations are discussed in "IBM Bridge Products' Coexistence and Migration" on page 26.

[15] A remote console facility equivalent to that of NetView/PC V1.1 is not available in the initial release of NetView/PC under OS/2 EE. However, the intent to provide remote console facility in the OS/2 EE environment has been announced in a statement of direction.

In this way, IBM LAN Manager V2.0 offers LAN management of almost any LAN environment consisting of 16 Mbps or 4 Mbps IBM Token-Ring Network segments, and PC Network (Broadband) segments, as shown in Figure 63.

In the mixed multisegment LAN environment, each LAN segment has a segment identifier (assigned by the bridges) which consists of a three digit hexadecimal number (range 001 - FFF). This three digit number is used in the operator interface to refer to specific segments, as well as in the routing information field of frames traveling over multiple segments.

Critical Resource Monitoring This IBM LAN Manager V2.0 feature enables the LAN Manager operator to define any workstation, identified by its unique universal or local MAC address, as a critical resource.: IBM LAN Manager V2.0 will monitor this station and send an alert whenever this critical resource is lost on the LAN. This feature is especially useful to monitor essential resources in an operational LAN environment such as print servers, large file servers, and gateway devices etc. It can also complement the operation of an intelligent 8220 Optical Fiber Converter.

The 8220 Optical Fiber converters provide to the LAN Manager a means for an additional level of ring problem determination. The downstream 8220 periodically sends a downstream converter present frame on the backup ring, so that when the LAN begins to operate on the backup ring the LAN Manager will receive this frame and generate an alert notifying that the LAN is now operating on the the backup ring.

Multiple LAN Manager Support Up to four LAN Managers running IBM LAN Manager V2.0 can be in communication with the same LAN bridge. Their relationship with the bridge is set by the network operator, identifying the reporting link to be either *controlling* or *observing*. Only one link to a given bridge can be a controlling link. The request to link to a bridge is verified against a password defined in the bridge for each of the four reporting link and the password provided by the particular LAN Manager requesting the link.

A controlling LAN Manager may issue any IBM LAN Manager V2.0 command on the remote LAN segment for execution by the bridge, even commands that change the way in which LAN components operate or the topology of the remote segment (for example the REMOVE ADAPTER command). Only controlling LAN Managers will log and transport alerts from application programs using the *Alert Transport Facility*.

An observing LAN Manager may not issue a command altering the way in which remote LAN segments function. In general, an observing LAN manager is restricted to *Query* and *Status* type commands and therefore is not capable of removing adapters or reconfiguring bridges.

Communication between LAN Managers is not supported.

NetView Support NetView R3 provides two complementary services for LAN Managers. One is to act as a *focal point for alerts* passed from the LAN Manager as a service point in the IBM Open Network Management Implementation. The other service is a *LAN Manager command service*, to control LAN Managers remotely from a NetView R3 console. This latter service is a new function implemented for IBM LAN Manager V2.0 and IBM LAN

Manager Entry V1.0 in NetView R3 based on the Service Point Command Interface.

Generic Alert Support: Alerts logged locally and transferred to the host use the architected code points and alert subvector frames published by IBM. The LAN Manager implements all the code points of the new LAN alerts documented in *SNA Format and Protocol Reference Manual: Management Services*. The objective is to support as many relevant code points as possible for the environment and therefore those not supported locally at the LAN Manager will be flagged as *unknown*. They will still be transported intact to the NetView R3 host where they may be recognized.

Note: When the IBM LAN Manager V2.0 receives or generates an alert it will pass the alert on to the OS/2 Communications Manager for transmission up to NetView at the host. This is indicated by an asterisk (*) next to the alert in the LAN Manager alert list. However if the host session is not active at the time the alert does **not** get transmitted nor is it queued for later transmission. The user is not notified that the alert cannot be sent to the host.

The following new alerts are supported by IBM LAN Manager V2.0:

- Loss of a critical resource on a Token-Ring Segment
- Receipt of a downstream converter present frame from an 8220 Fiber Converter. This frame is received by the LAN Manager when the Token-Ring backup path becomes part of the main ring path.
- A beaconing condition on the Token-Ring backup path - indicated by an 8220 Fiber Converter.
- Alerts sent from applications using the Alert Transport Service
- Unauthorized presence of the trace function of the IBM Trace and Performance Program on the network.

Application Alert Transport Services This important new feature of IBM LAN Manager V2.0 allows applications to generate and send alerts to a focal point. In this context, the application is referred to as a *reporting entity*.: This capability relies on the *alert transport service* architecture, by which a reporting entity uses the services of the LAN manager to allow it to send alerts to NetView.

Alerts are sent to the LAN Manager station (functional address C00000002000). The LAN Manager then logs the alert, acknowledges receipt of the alert to the sending station, and builds an NMVT frame (which imbeds the alert) and sends it off to the host focal point if required.

A detailed discussion on this service with a programming example can be found in LAN Manager V2.0 and LAN Manager Entry V1.0 Installation Guidelines.

Service Point Command Interface (SPCI).: The NetView *Service Point Command Service (SPCS)* is an integrated function that provides the facility to implement command flow between the *focal point* and a *service point* in the network. To enable applications (or services) to make use of this facility an interface to code these commands is provided: the *Service Point Command Interface (SPCI)*. These two terms (SPCS and SPCI) are sometimes intermixed but refer to the same capability of sending commands and receiving responses from a service point.

The NetView R3 operator is able to query various LAN entities and exercise a high degree of control over a LAN by using the Service Point Command Services and IBM LAN Manager V2.0. Automated operations can be implemented by NetView R3 using Service Point Command Services to perform an action upon receiving particular alerts.

Bridge Configuration Support: The ability of a LAN Manager to configure bridges has been enhanced so that subsequent configuration of bridges can be done once a communications link has been established with the bridge. The following parameters can now be changed from the LAN Manager:

- Hop count limits
- Link passwords
- Bridge number
- LAN segment number per adapter
- Frame forwarding
- Percent frames lost threshold
- Performance notification interval
- Single-route broadcast settings per adapter
- Single-route broadcast mode of operation (automatic/manual).

Previously such parameters had to be set at the bridge during configuration which required a bridge to be shutdown and restarted. If the bridge is running IBM Token-Ring Network Bridge Program V1.1 then only the Single-Route Broadcast settings per adapter and percentage lost threshold can be changed.

11.3 Design Considerations for LAN Management

Some of the factors that influence or determine LAN management requirements are:

- Size of the LAN
- Complexity of the configuration
- Level of control required
- Level of availability required
- Level of skills available
- Security requirements.

The size and complexity of the network are the most common and fundamental factors that determine the LAN management issues. Therefore we will be examining use of a LAN Manager in various network sizes and configurations.

11.3.1 Small LAN with No Host Connection

A small stand-alone LAN can be supported only by IBM LAN Manager V2.0 even when it does not require multisegment support, because IBM LAN Manager Entry V1.0 does not provide a local operator interface. This means that at least one user will require some knowledge of LANs and will require training to use the IBM LAN Manager V2.0.

A main advantage of using IBM LAN Manager V2.0 under Operating System/2 Extended Edition V1.1 is that in this environment it is not necessary to dedicate a PS/2 to the LAN management function because the LAN Manager can run concurrently with other applications.

To improve operator accuracy and productivity in this environment it is strongly recommended that all potential LAN adapter addresses be defined in the symbolic adapter names file with symbolic names. Comments and whether or not they are a critical resource should also be specified. Note that it is not necessary to define bridge adapters in this way, because the LAN Manager sessions with the bridges provide equivalent support. Up to 1,000 adapters can be defined in this file. This has two major benefits:

1. All alerts associated with a named adapter will contain the symbolic name of the adapter involved in the alert and therefore the alerts can more easily be traced to their source.
2. Any adapters defined as critical resources are continuously monitored for problems.

Consideration should also be given to authorization or restriction of trace facilities. Whenever an IBM Trace Adapter inserts itself into the ring it announces itself to a LAN Manager so that the LAN Manager can verify its authorization. IBM LAN Manager V2.0 provides the option of controlling any trace activities on the ring, preventing or permitting tracing in general or allowing only specific adapters to trace.

The alert transport facility can also be useful in this environment because IBM LAN Manager V2.0 will post an alert generated by an application into its alert log and notify the user that an alert has been received. For more information on the alert transport facility refer to LAN Manager V2.0 and LAN Manager Entry V1.0 Installation Guidelines.

11.3.2 Small Host-Connected LANs

A small host-connected LAN is assumed to connect to the host via a host gateway (3174, 37XX) or an SDLC link from the OS/2 Communications Manager.

The basic activities with regards to symbolic names definition and critical resources described in the "Small LAN with No Host Connection" on page 126 obviously become more important in a centralized network management environment. Now that a LAN has become part of a larger network it is more important to quickly identify and react to problems. So the critical resource monitoring and symbolic names play a more significant role in this environment.

The ability to extend centralized NetView network management to remote LANs creates alternative management scenarios. Two main factors should be considered in selecting an appropriate LAN Management capability for small host-connected LANs:

1. The level of local expertise in LAN management
2. The capability to do remote centralized LAN management.

Whenever the LAN environment is connected to a larger SNA network with central site management tools (NetView R3), integration of LAN management into total network management is highly recommended.

Because local systems management expertise is typically limited, use of IBM LAN Manager Entry V1.0 for a single-segment LAN or IBM LAN Manager V2.0 for a small multisegment LAN permits control of the LAN by a NetView R3

operator using the Service Point Command Services based LAN Manager commands. Any alerts generated on the LAN will be forwarded to NetView and can be displayed and attended to by the operator.

Centralized and Local LAN Management: If the LAN includes multiple segments or for some other reason requires local as well as centralized management, then IBM LAN Manager V2.0 is the ideal solution. IBM LAN Manager V2.0 has a user interface enabling a local network administrator to manage the network or allowing a remote NetView operator to provide backup assistance or after-hours support.

With IBM LAN Manager V2.0 it is a simple process to expand a single-segment LAN into a multisegment environment. Thus if LAN expansion is planned, it would be logical to install IBM LAN Manager V2.0 from the beginning.

Remotely Bridged LAN: A remote LAN can also be one connected via the Token-Ring Bridge Program Version 2.1. In this instance the LAN management function can be performed by IBM LAN Manager V2.0 on the central LAN segment. The remote LAN would be monitored by this LAN Manager and all alerts and problems would be reported through it as if the remote LAN were local. However it may be worthwhile to consider having a backup IBM LAN Manager Entry V1.0 with a host SDLC link from the remote site in case the remote bridge fails.

Mixed Token-Ring Network and PC Network (Broadband) LANS: Token-Ring and PC Network segments can be easily inter-connected with the IBM PC Network Bridge Program. Mixing network types as part of a single multisegment LAN creates little difference from a LAN Management perspective. All the bridges are treated in a similar manner by the LAN Manager.

11.3.3 Large MultiSegment LANs

As mentioned previously, if a LAN is large or has a need for local management capability, then IBM LAN Manager V2.0 is necessary.

In the multisegment environment, the bridges themselves become remote management agents that can report problems or soft errors to the LAN Manager. Since the bridges may be critical resources, the LAN Manager should be configured to link to the bridges upon startup. Then if there is any failure of the bridge, the LAN Manager is immediately notified by an alert that a bridge has failed. A bridge adapter need not be defined as a critical resource because the link to the bridge ensures the LAN Manager will monitor it. In designing support for a multisegment LAN, some consideration should be given to the number and location of LAN Manager stations.

Why Multiple LAN Managers?: In a large multisegment LAN, it may be worthwhile to have multiple LAN Managers for the following reasons:

- Distribution of management and control (when decentralized management is required as may occur in a large LAN that spans departments or buildings).
- Distribution of load (when high logging activity or monitoring limitations require multiple LAN Managers).
- Backup LAN Managers controlling tracing activities.
- Backup stations in case of failure or availability of reporting links.

How Do You Implement Multiple LAN Managers?: A LAN Manager can function only in either *controlling* mode or in *observing* mode. The way in which the LAN Manager links to a bridge is what determines its mode of operations. As mentioned in "Multiple LAN Manager Support" on page 124 up to four LAN Managers can link to a bridge but only one can be a controlling LAN Manager, since bridges support one controlling link and three more observing links to any LAN Manager.

It is possible to have multiple controlling LAN Managers which control different bridges. These controlling LAN Managers could monitor each other (as a critical resource). If a failure is detected, the monitoring manager could take over the responsibility of the failed LAN Manager through operator action to link to the bridges previously linked to the failing manager.

An observing LAN Manager could also monitor a controlling LAN Manager in this manner, but to take over the role of the failed controlling Manager it would have to be redefined as a controlling LAN Manager. This process needs a *reset* of the LAN Manager to re-start as a controlling LAN Manager. The reset process is very fast, but valuable information could potentially be lost in a busy LAN. After the reset, those bridges defined to be linked automatically at startup will now be linked in controlling mode. Others not automatically linked will need operator action to manually link them. For this reason, an observing LAN Manager should be linked to all bridges so that a reset to controlling mode would result in automatic linkage to the bridges in controlling mode.

Functions Supported by Manager Mode: The mode the LAN Manager is operating in will determine the functional behavior of that LAN Manager. Some restrictions apply to an observing manager. It will also determine which type of alerts will be transmitted to the host NetView. Only controlling LAN Managers will recognize an alert transport frame, post it to its alert log and send it up to the host. Although an observing LAN Manager will not recognize an alert from an application sent to the alert transport facility, it will still receive and log alerts from bridges it is linked to. However, it will not transmit these alerts to NetView itself, leaving this function to the controlling LAN Manager. For a list of functions and alert transportation capabilities for a LAN Manager in either controlling or observing mode see Table 1 on page 130.

Table 1. Controlling/Observing LAN Manager Functions		
LAN Manager Function	**Controlling**	**Observing**
Bridge Functions		
Link to bridge	X	X
Change bridge parameters	X	
Monitor remote segments	X	X
Adapter Functions		
Remove adapter from LAN	X	
Query adapter	X	X
Alert Management		
Log application generated alerts sent to Alert Transport Facility	X	
Transmit to host application alerts sent to Alert Transport Facility	X	
Log alerts if a critical resource fails	X	X
Send alerts to host if a critical resource fails	X	X
Log LAN alerts	X	X
Send LAN alerts to host	X	
Tracing Control		
Remove trace adapters anywhere on multisegment LAN when tracing is prevented	X	
Log trace adapter insertion anywhere on multisegment LAN	X	
LAN Status Functions		
Do a path test	X	
Do a LAN segment test	X	X
Display adapters with soft error conditions	X	X
Set soft error logging for local token-ring segments	X	X
Set soft error logging for remote token-ring segments	X	
Set configuration change logging	X	

Placement of LAN Managers: Aside from considerations dictated by physical cabling requirements, and requirements for access by operations staff, other logical factors should be considered in deciding where to place a LAN Manager in a particular LAN.

As a general guideline for complex multisegment LANs, it is recommended to place the LAN Manager in a segment such that:

1. It will not be easily isolated in case of bridge failure (that is place the LAN Manager on a backbone segment or where the critical functions are).
2. It is on a segment that has a high probability of receiving alerts with minimum propagation.
3. It is centrally located to minimize path lengths to critical resources, bridges, host connections and generally any node on the LAN.

4. If there is a host connection, it will connect to the same multi-station access unit as the gateway adapter to minimize the likelihood of isolating the NetView operator in the event of segmentation.

In large installations with many segments there may be multiple LAN Managers distributed through buildings and floors (because of physical restrictions). If centralized operations is required then NetView R3 with the Service Point Command Services provides a solution. This allows a high degree of remote operation of those LAN Managers.

The positioning of multiple LAN Managers in relation to each other is an important issue. Careful planning is needed to establish the role of LAN Managers (controlling/observing and backup) and their logical placement on segments of the LAN. The following discussions evaluate some factors which should be considered.

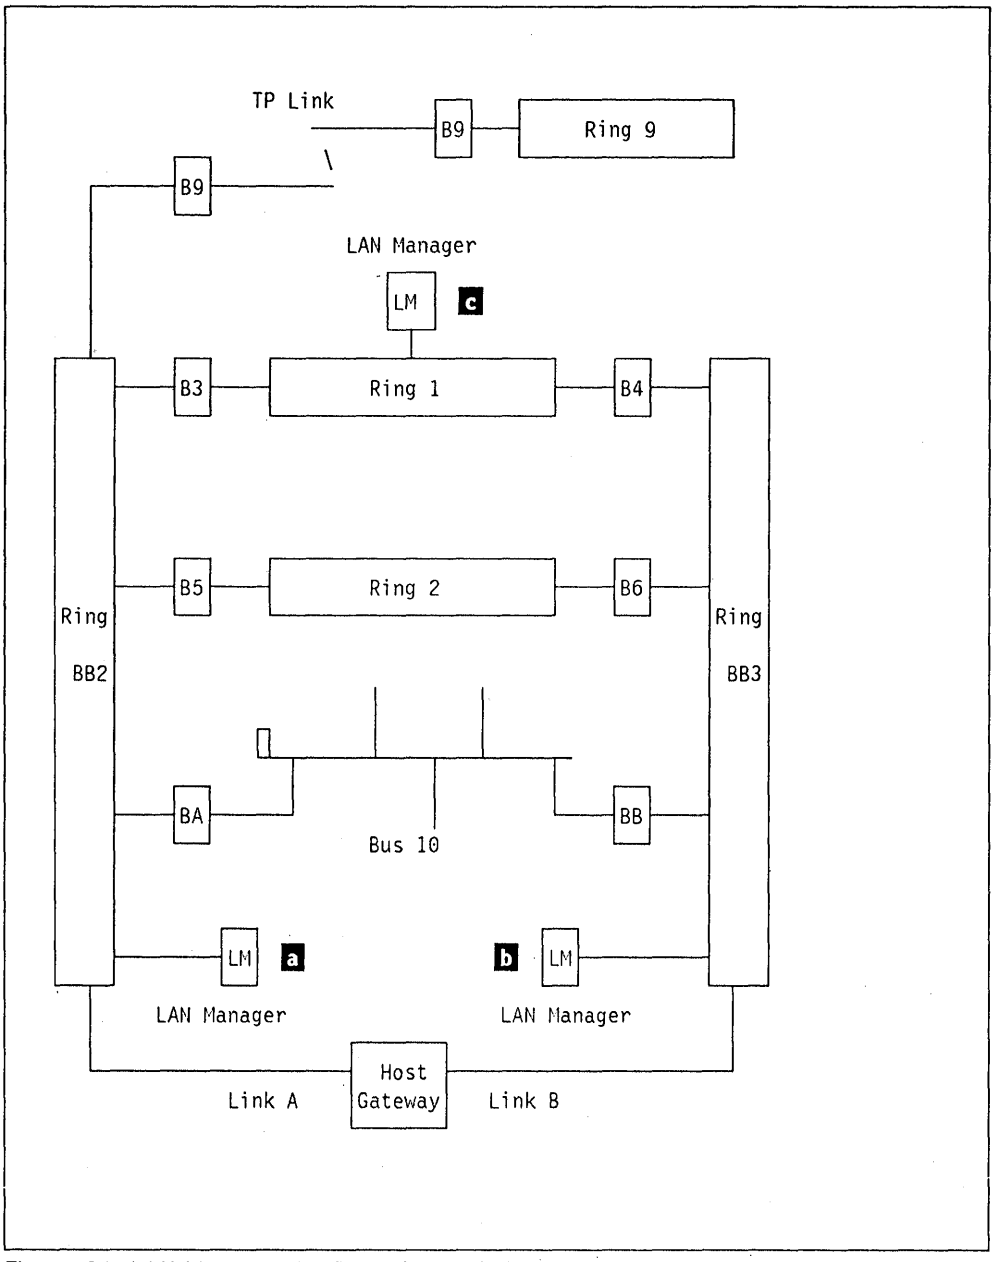

Figure 64. LAN Manager Configuration and Placement

One Controlling LAN Manager: The simplest way to have multiple LAN Managers in a LAN is to have only one controlling LAN Manager, linked to all the bridges, and one or more observing LAN Manager also linked to all the bridges. This way only one set of alerts will be generated and transmitted to the NetView at the host. If the controlling LAN Manager fails then the observing LAN Manager needs to be redefined to controlling mode and then *reset* by an operator.

Note: The changing of a LAN Manager from observing to controlling mode can only be done at the local LAN Manager user interface.

Thus when there is one controlling and multiple observing LAN Managers, it must be taken into consideration that if the controlling LAN Manager fails then backup is straightforward but requires operator intervention.

Multiple Controlling LAN Managers: In a multisegment environment, with multiple controlling LAN Managers determination of the controlling or observing mode for each manager should consider their position and role as backup for each other.

With multiple controlling LAN Managers and redundant bridge paths as described in "How Do You Implement Multiple LAN Managers?" on page 129, it is possible for each LAN Manager to manage the whole network and effectively provide coverage for each other. Some issues would be the possibility for redundant logging and transmission of alerts and/or the requirements for keeping bridge definitions consistent.

Controlling LAN Managers can send alerts to NetView. Thus if there are multiple controlling LAN Managers, the host can receive multiple copies of the same alert. One way to control this is to define one of the controlling LAN Managers without an active host connection. Thus alerts received by this manager will be logged locally and will not be transmitted to NetView. This feature can only be changed at the local user interface and not by a NetView R3 operator.

Since only one controlling LAN Manager can link to any one bridge, any changes to bridges will have to be done carefully between the controlling LAN Managers to maintain consistency. For instance it is advisable to have all bridges control single-route broadcasting in a consistent manner. One method for changing bridge definitions may be for one controlling LAN Manager to link to all bridges in the LAN for the duration of route reconfiguration.

Tracing Control: Only controlling LAN Managers can control tracing on a Token-Ring LAN. Therefore if security control is a high priority in a multisegment LAN, the LAN administrator should consider having multiple controlling LAN Managers to monitor all segments of the LAN and provide backup for each other. Similarly, alternate routes to segments from the controlling LAN Managers is recommended to reduce the chance of a having an isolated segment and thus having no management control over it.

The IBM Token-Ring Network Trace and Performance Program announces itself periodically when tracing. So any controlling LAN Manager that does not allow tracing in general or for this adapter specifically will force the trace adapter off the network and thus terminate the trace. This control is achieved using a general broadcast message. The LAN Manager does not rely on a bridge to control tracing but does it all itself. So the LAN Manager need not be linked to a bridge on the segment with the trace adapter to control its tracing activities. It can still control the trace activities of that adapter provided there is a path to that segment.

Having multiple LAN Managers controlling this function means that all of them must authorize adapters in the same manner. For example:

- If one LAN Manager allows only a certain adapter to trace and other LAN Managers prevent all tracing (or another selection of adapters not including this one), then it would not be authorized.
- When some LAN Managers allow tracing and other LAN Managers prohibit tracing then the prohibition would be overriding and no adapter will be allowed to trace.

Notes:

1. This access control for tracing is only implemented for the IBM Token-Ring Network Trace and Performance Facilities (adapter and program) and for any other manufacturer's adapters that implements the same trace announcement function. This would not prevent adapters that do not announce themselves to the network to trace information from the LAN. If this is an exposure in your installation only encryption of the data by the applications would provide the necessary security level required (this would be the case in any telecommunications environment where open access to the transmission media cannot be controlled to the level that is required).

REM on Other LAN Stations: The Ring Error Monitor (or LAN Error Monitor in the case of PC Networks) has a network management functional address implemented in LAN architecture. It is a function that observes, collects, and analyzes hard-error and soft-error reports send by stations on a *single* segment. It also assists in fault isolation and correction. The REM functional address is the destination address for all soft-error reports (*Report Soft Error* MAC frames) generated by any LAN station. The hard-error reports (*Beacon* MAC frames) are sent to the *all stations* MAC address, and examined by the REM hard-error analysis function. A detailed description of this function can be found in the *IBM Token-Ring Network: Architecture Reference*.

It is important to note that the functional address for REM (X'C00000000008') can be invoked on any adapter, to be the target for any error MAC frames. But the degree of analysis and processing if the error conditions is determined by the particular network management product (program or microcode) residing on that station. There are many products available that offer REM function to different degrees of usefulness. Some capture and log events only in a local repository to be retrieved by operator commands, where others might format the events into generic alert, log locally and/or forward these alerts to a central focal point. Since there could be multiple reporting points of alerts to the focal point the operator should be aware of this and try and customize the environment such that less redundant information arrives at this focal point.

The LAN Manager implements REM function in full. It will log error conditions (as a result of a MAC frame forwarded to its REM functional address) and format certain events as *generic alerts*, log them locally and forward them to NetView if customized to do so. LAN Manager will do this for the segment he is connected to and is capable of doing this for remote segments for which reporting links exist to bridges on those segments. This is done via error reports he receives from those the REM on those bridges which implement event collection capabilities (with limited local statistics). Bridges can only communicate with LAN Managers over their reporting link and have no capability of sending alerts to NetView.

This capability of collecting error MAC frames from all segments in the LAN Manager's domain (via bridge reporting links) allows the LAN Manager to forward alerts to NetView for error conditions on all those segments.

Other products implementing REM function (for example, 3174 9370, or AS/400 LAN Gateway configurations) do not support links with bridges and therefore provide alert support for the local segment only. If these products are customized to process REM error conditions and forward alerts to a focal point

and are on the same segment as a LAN Manager or other REM station, then multiple alerts could be sent to NetView.

The recommendation would be not to invoke the REM function in the other stations (3174, 9370 and AS/400) when they are in a LAN controlled by the LAN Manager. They should only be invoked when needed (as backup or if local information on that subsystem is required).

IBM bridge products also provide REM function, and in large LANs, configuration of this function on specific bridges would be determined by the topology. Each adapter in the bridge could be configured individually to support REM on each segment. If LAN Manager is on a segment with many other bridges (with REM active for this same segment) it would use its own REM function for the local segment and not request reports from the bridges. Similarly when having reporting links to other bridges on a common remote segment, it would predetermine which bridge on that common segment would be used for error reporting, in order to avoid redundant alerts and to use the most efficient reporting path.

11.4 LAN Management Scenarios

Using Figure 64 on page 132 the following discussions describe the implications of the above design factors.

A LAN Manager can only monitor critical resources that reside either on the local segment or on segments that the LAN Manager can access via a link to a bridge on that segment.

Consider the case where LAN Manager **a** links to bridges BA, B5, B3 and B9 and LAN Manager **b** links to bridges B4, B6 and BB. All these link are in controlling mode. In this case they cannot monitor the segments where the other LAN Manager resides; that is, LAN Manager **a** cannot monitor segment BB3. So if LAN Manager **b** fails then LAN Manager **a** cannot learn about the failure because it is not able to monitor LAN Manager **b** as a critical resource.

Therefore for multiple LAN Managers to maintain control over an entire multisegment LAN **each LAN Manager must link to bridges in a LAN such that each manager is linked to at least one bridge on every segment of the LAN**. It is not necessary for a LAN Manager to be linked to intermediate bridges on the path to a linked bridge.

In our example it would be better if LAN Manager **a** were linked to bridges B9, B4, B6 and BA in controlling mode and if LAN Manager **b** were linked to bridges B3, B5 and BB also in controlling mode. Being in controlling mode they both retain full functional capability and in this configuration can both monitor the entire LAN (except ring 9). If bridge B3 goes down, LAN Manager **b** can still monitor Ring BB2. The link to bridge B5 makes ring BB2 visible to LAN Manager **b**. Also if LAN Manager **b** fails then the LAN is still totally covered by LAN Manager **a**.

LAN Manager **c** could have been an observing manager for all bridges. In normal circumstances he is perfectly positioned to serve as a secondary backup for either of the two controlling LAN Managers. LAN Managers **a** and

b control half the bridges each. If either LAN Manager **a** or **b** fails and LAN Manager **c** is reconfigured to controlling and then *reset*, it will try to link to all bridges in controlling mode (if linking was defined to be automatic) but will only be able to link to bridges which have been "orphaned". This will generate error events for the link requests for those bridges owned in controlling mode by the still active LAN Manager. These error messages and events would not do any harm although they may be disturbing.

In the multisegment environment described in this figure, we have looked at the concept of alternate LAN Managers providing general backup for a failing LAN Manager. Another option to increase availability of the LAN Manager is to place two LAN adapters in the workstation. The LAN Manager can use only one at a time. However a second adapter would be useful if the LAN Manager is operational but lost contact because of:

- Failure of the LAN segment (for example beaconing) or
- Lost sessions to bridges (segmentation on invocation of backup path)

A second LAN adapter in this LAN Manager (attached to segment BB3), could be useful to recover from this temporary problem. The operator must reconfigure the LAN Manager to use the alternate adapter and issue a *reset*. This will close the primary adapter and insert the second adapter (if not already in use by another application) into a healthy segment from which bridge links could be re-established. This technique is only feasible when alternative paths to other segments are available as in our diagram.

One final consideration is to ensure that your backup planning does not overlook linking to all bridges accessed by a failing LAN Manager. Otherwise this could leave holes in your reporting process. Also be very careful to accommodate changes in the network (reconfiguration of the bridges, or removal/additions of bridges). Make sure to change the bridge definitions on the LAN Managers accordingly and to maintain your backup procedures.

When the LAN Manager operator displays the *configuration list* of a Token-Ring segment then all current live adapters will be displayed in nearest active upstream neighbor (NAUN) sequence. The adapter's symbolic name from the symbolic name table would also be displayed. A meaningful naming convention which assists in identifying the location of a workstation could be very valuable in identifying when a whole set of devices are missing because of cable problems resulting in ring segmentation. Refer to LAN Manager V2.0 and LAN Manager Entry V1.0 Installation Guidelines for a description of symbolic naming conventions.

11.5 Centralized Network Management - NetView

With an SSCP-PU link between a S/370 host and the LAN Manager, alerts can be transmitted to a NetView console for attention and a NetView R3 operator can then initiate an action in response to the alert using the SPCI. This facilitates centralized network management from a NetView R3 console.

11.5.1.1 Alerts

A *controlling* LAN Manager transmits alerts to NetView via a defined host connection. The NetView R3 operator can then issue commands using the Service Point Command Services to take corrective action. With IBM LAN Manager V2.0 the alerts are also posted locally on the LAN Manager Alert display and the user is notified by an indicator on any other LAN Manager display panel that an alert has been posted. The local user can then browse the alerts log and initiate corrective procedures if required. This also applies to application generated alerts sent by the alert transport facility of IBM LAN Manager V2.0.

Indication that IBM LAN Manager V2.0 has forwarded an alert to NetView via the host link is provided an asterisk (*) next to the alert in the LAN Manager alert list. However if the host session is not active at the time the alert does **not** get transmitted nor is it queued for later transmission. The user is not notified that the alert cannot be sent to the host. Similarly with IBM LAN Manager Entry V1.0 there is no notification if the alert is not transmitted to the host.

Care should be taken to minimize the number of duplicate alerts sent to NetView such as in the case of multiple controlling LAN Managers. One method of doing this is to define the host communication option of LAN Manager to none. This will still allow the NetView R3 operator to execute Service Point Command Services commands but will stop the LAN Manager from transmitting alerts to the host while retaining full controlling manager capabilities. At least one LAN Manager should have the capability of sending alerts up to NetView at the host.

11.5.1.2 Automated Operations

With the introduction of Service Point Command Services support by the LAN Manager products and NetView R3 some automated operations can now be implemented. The alerts received by NetView R3 can be monitored and if particular LAN alerts are received a Service Point Command Services command could automatically be executed. This could speed recovery, and minimize the time required by an operator to manage particular types of events. For instance if a bridge is disabled the LAN Manager may lose links with other bridges that went through the failing bridge. If alternate routes to those bridges are available it is a simple process to re-link to the bridges. This can be done automatically upon receiving an alert saying there is a problem with the bridge link.

11.5.1.3 Host Gateways on the LAN

For a LAN Manager to have an SSCP-PU session with the host, a link must be established either through a host Gateway (3174, 3745 or 3725) or through an SDLC link.

Host gateway placement follows similar guidelines to LAN Manager placement which are based on maximizing availability and accessibility to workstations on the LAN. Therefore the host gateway, like the LAN Manager, should be placed on a backbone segment and if possible on the same MSAU as the LAN Manager to minimize the possibility of the central network manager becoming isolated from the LAN in the case of segmentation.

If access to the host is critical there may be multiple gateways to access the host. Host gateway availability and ease of reconfiguration is dependent on the

gateway type. Therefore backup and alternate gateways should be carefully planned. The host gateway should be monitored as a critical resource so that the LAN Manager is notified in the event that it fails. Of course there are other events which may cause the LAN Manager to lose its host session that cannot be monitored so easily.

The Parameter Server, Error Monitor and Configuration Report Server in the bridges should all be enabled.

The remote bridges such as B9 are treated the same way as local bridges by IBM LAN Manager V2.0 and do not require special operational consideration by a LAN Manager. But the operator must be aware that if the bridge remote link fails the LAN Manager will only generate an alert stating that the reporting link to the remote bridge is lost. There is nothing to indicate the cause of the problem. To determine if it is the remote link which has failed, the operator must go to the bridge and read the messages there.

To ensure total coverage by both LAN Managers, the LAN Manager on segment BB2 links to bridges B4, B6 and B9 while the other LAN Manager on segment BB3 links to bridges B3 and B5. This allows them to monitor every segment in the LAN and have a backup path to the backbone segments. Also every segment except Ring 9 is monitored by two LAN Managers. This provides management backup in the event of a LAN Manager failure or bridge failure.

List of Abbreviations

CCITT	International Telegraph and Telephone Consultative Committee
DCE	Data Communications or Circuit Terminating Equipment
DOS	Disk Operating System
DSU	Data Service Unit
EIA	Electrical Institute of America
IDNX	Integrated Digital Network Exchange
ISDN	Integrated Services Digital Network
KB	Kilobyte
Kbps	Kilobits per second
LAN	Local area network
LLC	Logical Link Control
MAC	Medium Access Control
MB	Megabyte
Mbps	Megabits per second
PC	Personal computer
PS/2	Personal System/2
RAM	Random access memory
ROM	Read-only memory
RS	Recommended Standard

Glossary

This glossary defines abbreviations and terms used in this manual in describing the &brg., the IBM Token-Ring Network, IBM PC Network, and local area networks in general.

It includes information from the *IBM Vocabulary for Data Processing, Telecommunications, and Office Systems*, GC20-1699.

Definitions from draft proposals and working papers under development by the International Standards Organization, Technical Committee 97, Subcommittee 1 are identified by the symbol **(TC97)**.

Definitions from published sections of the *ISO Vocabulary of Data Processing*, developed by the International Standards Organization, Technical Committee 97, Subcommittee 1 and from published sections of the *ISO Vocabulary of Office Machines*, developed by subcommittees of ISO Technical Committee 95, are preceded by the symbol **(ISO)**.

A

active monitor. In the IBM Token-Ring Network, a function in a single adapter that initiates the transmission of tokens and provides token error recovery facilities. Any active adapter on the ring has the ability to provide the active monitor function if the current active monitor fails.

adapter. The circuit card within a communicating device (such as an IBM Personal Computer) and its associated software that enable the device to be attached to a network.

application program. A program written for or by a user that applies to the user's work.

B

beaconing. An error-indicating function of adapters that assists in locating the problem causing a hard error on the Token-Ring Network.

bridge. A functional unit that connects two local area networks (LANs) that use the same logical link control (LLC) procedures but may use different medium access control (MAC) procedures. A bridge consists of the bridge computer, two adapters and their cables, and the Bridge Program.

bridge computer. The dedicated computer in which the Bridge Program is loaded.

bridge ID. The bridge label combined with the adapter address of the adapter connecting the bridge to the LAN segment with the lowest LAN segment number; it is used by the Bridge Program automatic single-route broadcast function.

bridge label. A two-byte hexadecimal number that you can assign to each bridge. See bridge ID.

broadcast frame. A frame that is to be forwarded by all bridges, unless otherwise restricted.

bus. A network configuration where a series of nodes (attaching devices, such as IBM Personal Computers) are connected to a main cable.

bypass. To eliminate a station or an access unit from a ring network by allowing the data to flow in a path around it.

C

configuration. (1) (TC97) The arrangement of a computer system or network as defined by the nature, number, and the chief characteristics of its functional units. The term may refer to a hardware or a software configuration. (2) The devices and programs that make up a system, subsystem, or network.

configuration file. The collective set of item definitions that describe a configuration.

controlling link. The reporting link between a bridge and a network manager program that is authorized to change bridge configuration parameters and to disable and enable certain bridge functions.

D

data communications equipment (DCE). The equipment installed at the user's premises that provides all the functions required to establish, maintain, and terminate a connection, and the signal conversion and coding between the data terminal equipment (DTE) and the telecommunications link.

data service unit (DSU). Device which provides an interface to a telecommunications network by performing conversion of the user 2-level data signals into the standard bipolar line format.

data terminal equipment (DTE). That part of a data station that serves as a data source, data sink, or both, and provides for the data communication control function according to protocols.

designated bridge. A bridge in a network using automatic single-route broadcast that forwards single-route broadcast frames.

diagnostics. Modules or tests used by computer users and service personnel to diagnose hardware problems.

disk image. A representation of a disk or diskette containing files and programs. The image resides in computer storage and is used by the computer as though it were a physical disk or diskette.

Disk Operating System (DOS). A program that controls the operation of an IBM Personal Computer or IBM Personal System/2 computer and the execution of application programs.

dump. (1) Computer printout of storage. (2) To write the contents of all or part of storage to an external medium as a safeguard against errors or in connection with debugging. (3) (ISO) Data that have been dumped.

E

enabled. Active, operational, and can receive frames from the network. (Servers and functional addresses may be enabled by programs running on the Token-Ring Network.)

establishment. A user's premises that does not extend across public rights of way (for example, a single office building, warehouse, or campus).

F

formatted diskette. A diskette on which track and sector control information has been written and which may or may not contain data.

Note: A diskette must be formatted before it can receive data.

frame. The unit of transmission in the Token-Ring Network. It includes delimiters, control characters, information, and checking characters.

H

hard error. An error occurring on the network that makes it inoperative. See beaconing.

"hello" message. A message used by automatic single-route broadcast to detect what bridges enter and leave the network and to reset single-route broadcast parameters accordingly. The root bridge sends a "hello" message on the network every two seconds.

help panel. Information displayed by a program or system in response to a help request from a user. An on-line display that tells you how to use a command or another aspect of a product.

hop count. The number of bridges through which a frame has passed on the way to its destination.

Note: Hop count applies to all broadcast frames that are not single-route broadcast frames.

hop count limit. The maximum number of bridges through which a frame may pass on the way to its destination.

K

kilobyte (KB). 1024 bytes.

L

LAN segment. Any portion of a local area network (for example, a single ring or bus) that can operate independently, but is connected to the establishment network via bridges, controllers, or gateways.

LAN segment status. The condition of the LAN segment (ring or bus).

link. The combination of physical media, protocols, and programming that connects devices on a network.

lobe. In the IBM Token-Ring Network, the section of cable (which may consist of several segments) that attaches a device to an access unit.

local area network (LAN). A data network located on the user's premises in which a serial transmission is used for direct data communication among data stations.

local bridge. A function of the Bridge Program that allows a single bridge computer to connect two LAN segments (without using a telecommunication link).

M

megabyte (MB). (1) A unit of measure for storage capacity. One megabyte equals 1,048,576 bytes. (2) Loosely, one million bytes.

N

NAUN. Nearest active upstream neighbor. For any station on a ring, the station that is sending frames or tokens directly to it.

network. A configuration of data processing devices and software connected for information interchange.

network manager. A program or group of programs that is used to monitor, manage, and diagnose the problems of a network.

non-broadcast frame. A frame containing routing information specifying which bridges are to forward it. A bridge will forward a non-broadcast frame only if that bridge is included in the frame's routing information.

P

page. (1) The portion of a panel that is shown on a display surface at one time. (2) To move back and forth among the pages of a multiple-page panel. See also *scroll*.

panel. (1) A formatted display of information that appears on a terminal screen. See also *help panel*. (2) In computer graphics, a display image that defines the locations and characteristics of display fields on a display surface.

path. (1) The route traversed by the information exchanged between two attaching devices in the network. (2) A command in DOS that searches specified directories for commands or batch files that were not found by a search of the current directory.

path cost. A value, maintained by each bridge, that indicates the relative length of the path between the root bridge and another bridge.

path trace. A function that may be requested of a bridge by a received frame. The request is for a record of the bridges through which the frame has passed.

R

RAM paging. RAM paging is a technique that allows the computer software to access all the RAM on the adapter, without having to map the entire shared RAM into the computer's memory map. The shared RAM on the adapter is paged into the computer's memory map one area at a time.

RAM size. The amount of RAM that is directly mapped into the computer's memory map.

random access memory (RAM). A computer's storage area into which data may be entered and retrieved in a nonsequential manner.

read-only memory (ROM). A computer's storage area whose contents cannot be modified.

remote bridge. A function of the Bridge Program that allows two bridge computers to use a telecommunication link to connect two LAN segments.

ring (network). A network configuration where a series of attaching devices, such as IBM Personal Computers, are connected by unidirectional transmission links to form a closed path. A ring of a Token-Ring Network is referred to as a LAN segment or as a Token-Ring Network segment.

root bridge. The bridge in a network using automatic single-route broadcast that sends the "hello" message on the network every two seconds. Automatic single-route broadcast uses the message to detect when bridges enter and leave the network, and to change single-route broadcast parameters accordingly.

S

scroll. To move all or part of the display image vertically to display data that cannot be observed within a single display image. See also *page*.

server. A device, program, or code module on a network dedicated to a specific function.

shared RAM. Random access storage on the adapter that is shared by the computer in which the adapter is installed.

single-route broadcast. The forwarding of specially designated broadcast frames only by bridges which have single-route broadcast enabled. If the network is configured correctly, a single-route broadcast frame will have exactly one copy delivered to every LAN segment in the network. Also called limited broadcast.

soft error. (1) An intermittent error on a network that requires retransmission. The adapters are able to retransmit the data that had the difficulty and communication continues. (2) An error on a network that affects the network's performance but does not, by itself, affect its overall reliability. If the number of soft errors reaches the ring error limit, reliability is affected.

stand-by bridge. A bridge in a network using automatic single-route broadcast that does not forward single-route broadcast frames. A stand-by bridge is a parallel bridge or is in a parallel path between two LAN segments.

T

threshold. (1) In the Bridge Program, bridge performance threshold refers to a value set for the number of frames per 10,000 that can be lost before an entry is made in the bridge Performance Statistics, and a notification is sent to any network manager program that has requested such reports. (2) The telecommunications link error threshold value expresses the maximum allowable number of frames

per 10,000 not received across a bridge due to errors on the telecommunications link connecting the two stations of a bridge before the Bridge Program counts a "threshold exceeded" occurrence in the Performance Statistics. The Bridge Program also sends a notification to any network manager programs that have requested such reports.

V

virtual drive. A direct access storage device that does not physically exist. It exists logically in computer memory.

W

working disk(ette). A computer fixed disk or diskette to which files are copied from an original diskette for use in daily operation.

Index

T

W

IBM MULTISEGMENT LAN DESIGN GUIDELINES
GG24-3398-00

READER'S COMMENT FORM

You may use this form to communicate your comments about this publication, its organization, or subject matter, with the understanding that IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

Your comments will be sent to the author's department for whatever review and action, if any, is deemed appropriate. Comments may be written in your own language; use of English is not required.

Note: *Copies of IBM publications are not stocked at the location to which this form is addressed. Please direct any requests for copies of publications, or for assistance in using your IBM system, to your IBM representative or to the IBM branch office serving your locality.*

Possible topics for comment are:

Clarity Accuracy Completeness Organization Coding Retrieval Legibility

If you wish a reply, give your name, company, mailing address, and date:

_____

_____

_____

_____

What is your occupation? _____

Cut or Fold Along Line

Reader's Comment Form

Fold and tape | **Please Do Not Staple** | Fold and tape

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST CLASS | PERMIT NO. 40 | ARMONK, N.Y.

POSTAGE WILL BE PAID BY ADDRESSEE:

IBM International Technical Support Center
Department 985A, Building 657
P.O. Box 12195
Research Triangle Park
Raleigh, North Carolina 27709
U.S.A.

Fold and tape | **Please Do Not Staple** | Fold and tape

IBM®

GG24-3398-00

IBM-MULTISEGMENT LAN DESIGN GUIDELINES GG24-3398-00 PRINTED IN THE U.S.A.

IBM®