3745 Communication Controller Models A
3746 Nways Multiprotocol Controller
Models 900 and 950

Planning Guide

(Part 1/3)

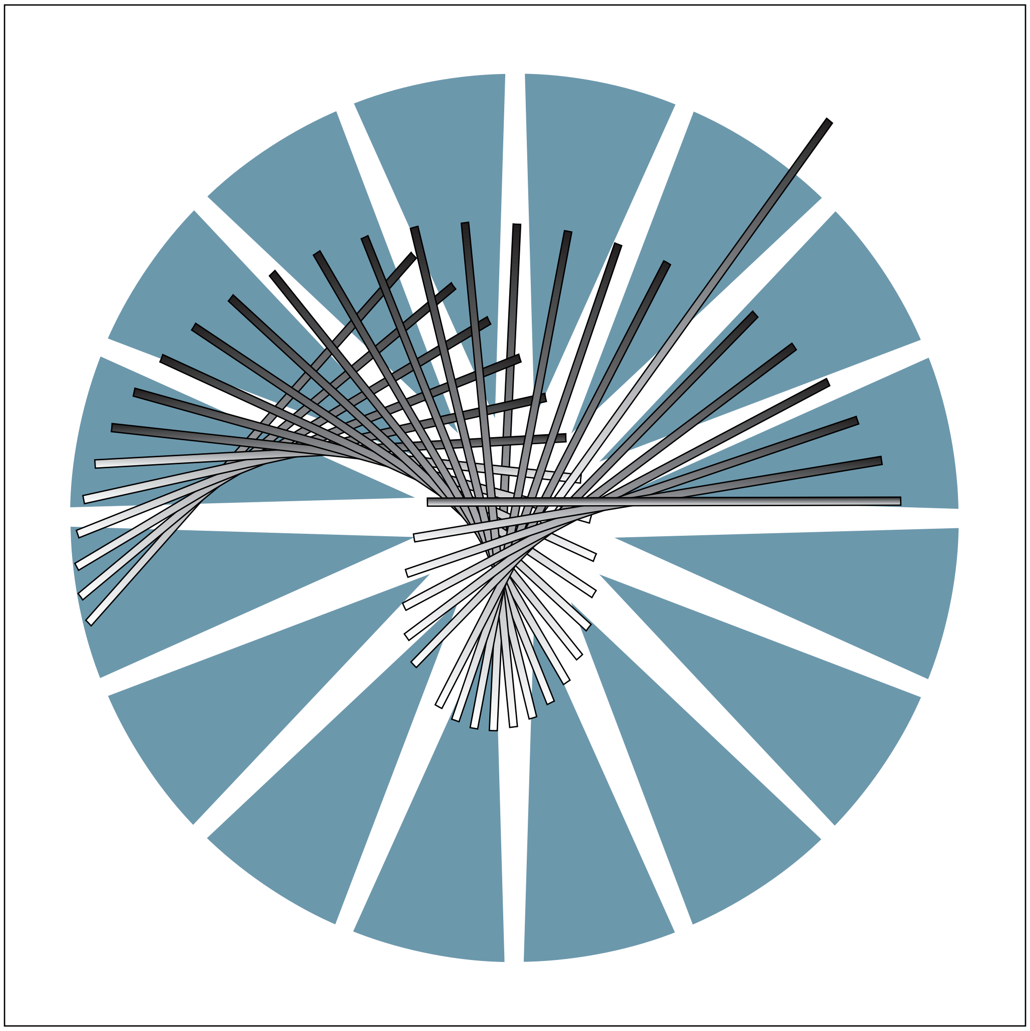

GA33-0457-0A

3745 Communication Controller Models A
3746 Nways Multiprotocol Controller
Models 900 and 950

Planning Guide

(Part 1/3)

GA33-0457-0A

Note!

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

| **Preliminary Second Edition (December 1997)**

| Changes or additions to the text and illustrations are indicated by a vertical line to the left of the change.

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

A form for readers' comments appears at the back of this publication. If the form has been removed, address your comments to:

IBM France
Centre d'Etudes et Recherches
Service 0798 - BP 79
06610 La Gaude
France

- FAX: (33) 4 93 24 77 97
- E-mail: FRIBMQF5 at IBMMAIL
- IBM Internal Use: LGERCF at LGEPROFS
- Internet: rcf_lagaude@vnet.ibm.com

When you send information to IBM, you grant IBM a non-exclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Contents

Part 1. Planning Guide Overview

Part 2. Planning Guide Configuration and Management

Part 3. Physical Planning Guide

Figures

Tables

Notices

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, is the user's responsibility.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, 500 Columbus Avenue, Thornwood, New York 10594, U.S.A.

Product Page/Warranties

The following paragraph does not apply to the United Kingdom or to any country where such provisions are inconsistent with local law.

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore this statement may not apply to you.

Electronic Emission Notices

Federal Communications Commission (FCC) Statement

Note: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and

(2) this device must accept any interference received, including interference that may cause undesired operation.

Industry Canada Compliance Statement

This Class A digital apparatus meets all requirements of the Canadian Interference-Causing Equipment Regulations.

Avis de conformité aux normes d'Industrie Canada

Cet appareil numérique de la classe A respecte toutes les exigences du Règlement sur le matériel brouilleur du Canada.

Japanese Voluntary Control Council For Interference (VCCI) Statement

This equipment is in the 1st Class category (information equipment to be used in commercial and/or industrial areas) and conforms to the standards set by the Voluntary Control Council for Interference by Information Technology Equipment aimed at preventing radio interference in commercial and industrial areas.

Consequently, when used in a residential area or in an adjacent area thereto, radio interference may be caused to radios and TV receivers, and so on.

Read the instructions for correct handling.

Korean Communications Statement

Please note that this device has been approved for business purpose with regard to electromagnetic interference. If you find this is not suitable for your use, you may exchange it for a non-business one.

New Zealand Radiocommunications (Radio) Regulations

Attention: This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

Other Special Notices

Notice to UK Users:

The IBM 3746 Model 900 and IBM 3746 Model 950 are manufactured according to the International Safety Standard IEC950 and, as such, is approved in the UK under the General Approval number NS/G/1234/J/100003.

The Active Remote Couplers (ARCs) and the X.21 Interface, housed within the 3746 Model 900 and 3746 Model 950, are approved separately, each having their own independent approval number. These interface adapters, supplied by IBM, do not contain excessive voltages. An excessive voltage is one which exceeds 42.4 V peak ac or 60 V dc. They interface with the 3746 Model 900 or 3746 Model 950, using Safe Extra Low Voltages only.

In order to maintain the independent approval of the IBM adapters, it is essential that other optional cards, not supplied by IBM, do not use mains voltages or any other excessive voltages. Seek advice from a competent engineer before installing other adapters not supplied by IBM.

Trademarks and Service Marks

The following terms, denoted by ™ used in this publication, are trademarks or service marks of IBM Corporation in the United States or other countries:

ACF
AIX
ACF/VTAM
Advanced Peer-to-Peer Networking
APPN
Enterprise Systems Connection Architecture
ESCON
ES/9000
Hardware Configuration Definition
IBM
IIN

Intel
Pentium

LPDA
Operating System/2
OS/2
NetView
Nways
PS/2
RETAIN
S/390
System/370
VTAM
3090

Safety

This product meets IBM™ Safety standards.

For more information, see the following manual: *IBM 3745 Communication Controller All Models*, *IBM 3746 Expansion Unit Model 900*, *IBM 3746 Nways Multinetwork Controller Model 950 Safety Information*, GA33-0400.

About this Guide

This guide applies to the IBM 3746 Nways™ Multiprotocol Controllers: the 3746 Model 950, 3746 Model 900 and 3745 Communication Controller Models A. Use it for configuration and installation planning, and to gather the information needed during the installation and network integration of 3746 Nways Multiprotocol Controllers operated in the APPN/High Performance Routing (APPN® /HPR) and Internet Protocol (IP) environments. There is also a general description of installation and upgrade scenarios.

Note: This guide contains information about major new features and their effect on existing features, and has been restructured since the last edition. If you are already familiar with the last issue, you are recommended to read this entire guide, not just the new features.

Information is provided about various parameters that have to be available to the IBM service representative and your network specialists for installation or upgrade of your machine. They relate to:

- 3745 Communication Controller Models A
- 3746 Models 900 and 950
- Controller Configuration and Management (CCM)
- Network Node Processor (NNP)
- Multiaccess Enclosure
- Service processor
- Distributed Console Access Facility (DCAF) and TME10 remote consoles
- Network management

There is also an introduction to the Controller Configuration and Management application (CCM), which is required for the definition of 3746 Nways Multiprotocol Controller resources.

Note: Your IBM marketing representative can obtain this 3746 publication, the related 3746 Controller Configuration and Management application (CCM), and the *IBM 3746 Nways Multiprotocol Controller Model 950, IBM 3746 Model 900 Network Node Controller Configuration and Management: User's Guide* from the “900NN950” package, which is available through the IBM worldwide source of marketing materials.

The 3746 *Planning Guide* available in this package may be a preliminary version of the next printed edition and, therefore, may be required for information about the latest 3746 enhancements.

Who Should Use this Guide

This guide is intended for network planners, network specialist, and system programmers responsible for preparing the information that is needed for the installation and network integration of 3745 Communication Controller Models A and 3746 Expansion Unit Model 900 in an SNA environment, as well as the 3746-950 and 3746-900 as APPN/HPR network nodes and IP routers.

Your Task Responsibilities as a Customer

These are not IBM tasks!

The tasks in Table 0-1 are not performed by IBM personnel as part of the machine installation and basic operations. They can be performed by IBM on a fee basis.

Table 0-1 (Page 1 of 3). Customer Tasks

Task	Where to Find Information
Network design	Network design is not covered in this guide. Refer to the following IBM books for SNA, APPN/HPR, and IP network planning guidance: • *Planning for Integrated Networks*, SC31-8062 • The following IBM "redbooks": – *Subarea Network to APPN Network Migration Guide*, SG24-4656 – *IBM 3746 Nways Multiprotocol Controller Model 950 and IBM Model 900: APPN Implementation Guide*, GG24-2536 – *IBM 3746 Nways Multiprotocol Controller Model 950 and IBM Model 900: IP Implementation Guide*, SG24-4845 – *IBM Nways 2216 Multiaccess Connector Description*, SG24-4957 – *IBM 2216 Multiaccess Connector ESCON Solutions*, SG24-2137.
Physical planning: Before the IBM service representative arrive to install your controller, make sure that you have met the necessary requirements for the following: • Electric power • Floor space with service clearances • Space for the cables • The RSF switched line • The Controller Expansion (Feature 5023) • Other components (such as the service processor).	Chapter 44, "Physical Planning Details"

Table 0-1 (Page 2 of 3). Customer Tasks

Task	**Where to Find Information**
Controller hardware configuration definitions: Decide on what type of attachments (lines) and how many of each type you need.	This is input for the IBM ordering system (CF3745). Helpful information is found in “Line Weights and CLP Load” on page 18-9.
Software definitions and tuning:	Refer to:
• ESCON port, host link, and station definitions; ESCON resource, TCP/IP, and VTAM tuning	• Chapter 15, “ESCON Adapters” • Chapter 22, “ESCON Channel Adapter” • Chapter 27, “3746 Base Frame ESCON Configuration Examples”
• Token-ring port and station definitions; PU and LU maximum limits; port sharing with NCP-controlled traffic; duplicate addresses; token-ring APPN, IP, and/or NCP resource tuning and VTAM tuning	• Chapter 16, “Token-Ring Adapters”
• Serial line (SDLC, PPP, frame-relay, and X.25) port and station definitions; location of CLPs, LICs, LCBs, and ARCs; maximum CLA line connectivity; CLP backups	• Chapter 18, “Serial Line Adapters” • Chapter 19, “3746 SDLC Support”
• Multiaccess Enclosure: hardware planning and configuration; software configuration and tuning	• Chapter 20, “Multiaccess Enclosure” • Chapter 21, “Multiaccess Enclosure Adapters Overview” • Chapter 23, “Multiaccess Enclosure ISDN Support” • Chapter 26, “Multiaccess Enclosure Configuration” • Chapter 28, “Configuring the MAE ESCON Channel Adapter”
• Use of the Controller Configuration and Management (CCM) application.	• Chapter 25, “Welcome to the CCM” on page 25-1 • *IBM Controller Configuration and Management User's Guide*, SH11-3081. Also see: • *IBM 3746 Nways Multiprotocol Controller Model 950 and IBM Model 900: APPN Implementation Guide*, GG24-2536 (an IBM “redbook”) • *IBM 3746 Nways Multiprotocol Controller Model 950 and IBM Model 900: IP Implementation Guide*, SG24-4845 (an IBM “redbook”).
Filling out:	Refer to:
• 3746 plugging sheets To keep a record of the processors and couplers (and their addresses) installed in the 3746 frame.	• Chapter 43, “Plugging Sheets for the 3746 Nways Multiprotocol Controller”
• CCM worksheets To plan the non-MAE logical resource definitions. They can then be used when configuring the 3746 via the CCM.	• Chapter 40, “CCM Worksheets for Controller Configuration Definitions.”
• Multiaccess Enclosure worksheets To plan the MAE logical resource definitions. They can then be used when configuring the MAE.	• Chapter 41, “ Multiaccess Enclosure Worksheets.”

Table 0-1 (Page 3 of 3). Customer Tasks

Task	Where to Find Information
NetView definitions in VTAM, the MOSS-E, NPM, CCM, MAE, NetView/360, NetView/AIX for: • APPN traffic • IP traffic • NetView alert path.	Refer to: • Chapter 29, "3746 Management Overview" on page 29-1 • Chapter 30, "3746 APPN/HPR Network Node Management" on page 30-1 • Chapter 31, "3746 IP Router Management" on page 31-1 • Chapter 32, "MAE APPN/HPR Network Node Management" on page 32-1 • Chapter 33, "MAE IP Router Management" on page 33-1.
Controller, service processor, and network node processor definitions. Some examples: • Link IPL port information • Service processor token-ring and IP LAN addresses • Password management • NetView alert reporting path definitions • DCAF LU definitions • Ethernet port definitions for SNMP.	Refer to Chapter 34, "Controller and Service Processor." Fill out the Chapter 38, "MOSS-E Worksheets for Controller Installation (3745)," which are used by the IBM service representative during the installation.
Remote console definitions (using DCAF): • Insure that the necessary hardware and software is available for the type of console attachment chosen • Service processor definitions for DCAF • DCAF installation and configuration on the remote console.	Refer to: • Chapter 35, "Customer Consoles and DCAF." • For the 3746-900, refer to the *3745 Console Setup Guide*, SA33-0158 • For the 3746-950, refer to the *IBM 3746 Nways Multiprotocol Controller Model 950 User's Guide*, SA33-0356.
Connection to the IBM remote support facility (RSF): • Service processor connection (modem) definitions • Customer definitions for RSF records.	Refer to Chapter 36, "Connecting to the IBM Remote Support Facility."
Problem determination through the MOSS-E and NetView	For the 3746-900, refer to: • *Problem Analysis Guide* accessed online from the MOSS-E • *3745 Models A: Alert Reference Guide*, SA33-0175 • *3745 All Models: Advanced Operators Guide*, SA33-0097.

Where to Find More Information

During your migration planning, it may be necessary to use, in addition to this guide, the following documents:

- *IBM 3745 Communication Controller Models A, IBM 3746 Nways Multiprotocol Controller, Models 900 and 950: Overview*, GA33-0180.
- IBM *3746 Nways Multiprotocol Controller Models 900 and 950: Controller Configuration and Management User's Guide*, SH11-3081.
  Preparing controller definitions prior to installation of your 3746 Nways Multiprotocol Controller is recommended. To obtain a stand-alone version of the Controller Configuration and Management that runs on an OS/2 workstation, contact your IBM marketing representative.
- *IBM 3746 Nways Multiprotocol Controller Model 950: User's Guide*, SA33-0356.
  For information about routine operations, installing and testing the communication line adapters, service processor, and remote consoles.
- *IBM 3745 Communication Controller: Console Setup Guide*, GA33-0158
  For information about remote console access to 3745/3746-900(s) via an SNA/subarea path.

Be sure to use the latest editions of these documents. This will ensure that you have the necessary information about the 3746 Nways Multiprotocol Controllers.

Also helpful is:

- *Planning for Integrated Networks*, SC31-8062.

The following *IBM International Technical Support Organization* “redbooks”, are generally helpful for 3746 Nways Multiprotocol Controller implementation:

- *APPN Architecture and Product Implementations Tutorial*, GG24-3669
- *IBM 3746 Nways Multiprotocol Controller Model 950 and IBM Model 900: APPN Implementation Guide*, GG24-2536
- *Subarea Network to APPN Network Migration Guide*, SG24-4656 (an IBM “redbook”)
- *IBM 3746 Nways Multiprotocol Controller Model 950 and IBM Model 900: IP Implementation Guide*, SG24-4845 (an IBM “redbook”)

The following Enterprise Systems Connection Architecture documentation may also be helpful:

- *Enterprise Systems Connection Migration*, GA23-0383
- *Planning for Enterprise Systems Connection Links*, GA23-0367

For the Distributed Console Access Facility (DCAF) for Version 1.3:

- *DCAF: Installation and Configuration Guide*, SH19-4068
- *DCAF: User's Guide*, SH19-4069
- *DCAF: Target User's Guide*, SH19-6839.

See also the “Bibliography” on page X-9.

World Wide Web

You can access the latest news and information about IBM network products, customer service and support, and information about microcode upgrades via the Internet at `http://www.ibm.com/`.

CD-ROM Online Documentation

Starting with engineering change F12380, the Licensed Internal Code (LIC) is shipped on a CD-ROM. Also included on this CD-ROM is:

- 3745/3746 documentation
  (For example, the 3745 Model A and 3746 *Planning Guide*, 3746 NNP and service processor installation and maintenance guides, CCN *User's Guide*, 3746-950 *User's Guide*, and others. Refer to the “Bibliography” on page X-9 for the complete name and form number of these books.)

  The documentation is in the .PDF format. The Acrobat Reader™ for OS/2 is included on the CD-ROM. It lets you easily read the .PDF files and print all or part any book.

- 3746 presentations
  (For example, the latest Announcement, the Multiaccess Enclosure, Processors Type 3, and others.)

  They are available in:

  – .PRE format for Freelance for Windows™

  – .PRS format for Freelance for OS/2™

  – .PDF format for the Acrobat Reader.

- 3746 information pages
  (For example, details about the available presentations.)

  They are available in the .HTM format for use by any WEB browser. (The Netscape Navigator™ for OS/2 browser is available on the service processor.)

- .PRS format for Freelance for OS/2™
- .PDF format for the Acrobat Reader.

To Access the CD-ROM Information

The CD-ROM can be used on a service processor[1] .

To access the CD-ROM from the **service processor**:

Step 1. Install the CD-ROM in the service processor CD-ROM drive.

[1] The following service processors can be used:

- Service Processor 9585 (feature code 5021) equipped with:
  – Feature code 5051, a CD-ROM drive
  – Feature code 5028, 96 MB of memory
  – Feature code 5026, 2 GB hard disk drive.
- Service Processor, type 2 (feature code 5052).

Step 2. In the MOSS-E main window, open the **View** menu and select **Information**.

Step 3. Double click on **CD-ROM documentation**. Netscape Navigator automatically opens and displays the documentation home page.

Step 4. Click on any highlighted text (blue and underlined) to go to the material that interests you:

- Click on **File repository** to access the Freelance and .PDF presentations
- Click on **Documentation** to access 3745/3746 books.

  Then click on the icon marked PDF that corresponds to the item that interests you.

  The Acrobat Reader automatically opens and displays the file in the full screen mode. Use the **Page Up** and **Page Down** keys to move through the document.

  Press the **ESC** key to display the Reader menus that allow you to print all or part of the file.
- When you **Close** the Acrobat Reader, you return to the Netscape Navigator browser.
- When you **Close** the browser, you return to the MOSS-E **Documentation** menu.

Each presentation and book file has one or more of the following identifiers:

- Date
- Form number
- Engineering change level
- Revision code.

Check these identifiers on future releases of the CD-ROM to see if the documents that you use have been updated.

Part 1. Planning Guide Overview

Chapter 1. 3745 and 3746 General Information

Note

The book refers in various chapters to IBM's previewed enhancements to Nways Controllers.

The announcement of Nways Controller enhancements will be based on IBM's business and technical judgment. All information being previewed in this book represents IBM's intent at the time of writing, is subject to change or withdrawal, and represents only goals and objectives. This information is included here to assist in your planning.

Evolution of the 3745 and 3746 Controller Family

In the past 20 years, IBM has designed and developed a large number of communication controllers, particularly the 3745 family, and more recently, the 3746 Nways Multiprotocol Controllers. They are the:

- 3745 Models 130, 140, 150, 160, and 170
- 3745 Models 210, 310, 410, and 610
- 3746 Models Axx, Lxx
- 3745 Models 17A, 21A, 31A, 41A, and 61A (3745 Models A)
- 3746 Model 900 (3746-900)
- 3746 Model 950 (3746-950)

The 3745 family of communication controllers, a symbol of the SNA world, took a big step towards APPN with the 3746-950, which operates as an APPN network node and native IP router independent of any 3745 running the Network Control Program (NCP).

The IBM 3746 Model 900 has been enhanced to support the same routing functions as the 3746-950. The 3746-900 can operate as an APPN network node and IP router, while simultaneously keeping its more traditional role as an NCP-controlled SNA subarea node or APPN composite network node (CNN).

Later the 3746 was enhanced to support APPN High Performance Routing (APPN/HPR) and IP routing. This evolution of the 3746 Models 900 and 950 formed a new generation of controllers, *the 3746 Nways Multiprotocol Controllers*.

The 3746 Nways Multiprotocol Controllers were recently enhanced to provide a greater range of high-speed, high-availability connectivity options through the 3746 Multiaccess Enclosure. Using IBM 2216 hardware technology, the 3746 Multiaccess Enclosure provides interface support for additional ESCON host connectivity, worldwide ISDN primary, for 155 Mbps ATM (multimode and single mode fiber) as a forum-compliant ATM LAN emulation client and ATM classical IP routing. IBM has also previewed Fiber Distributed Data Interface (FDDI), Fast Ethernet (100 Mbps), High Speed Serial Interface (HSSI) for T3/E3 speeds, TN3270e server and native HPR over ATM for later availability.

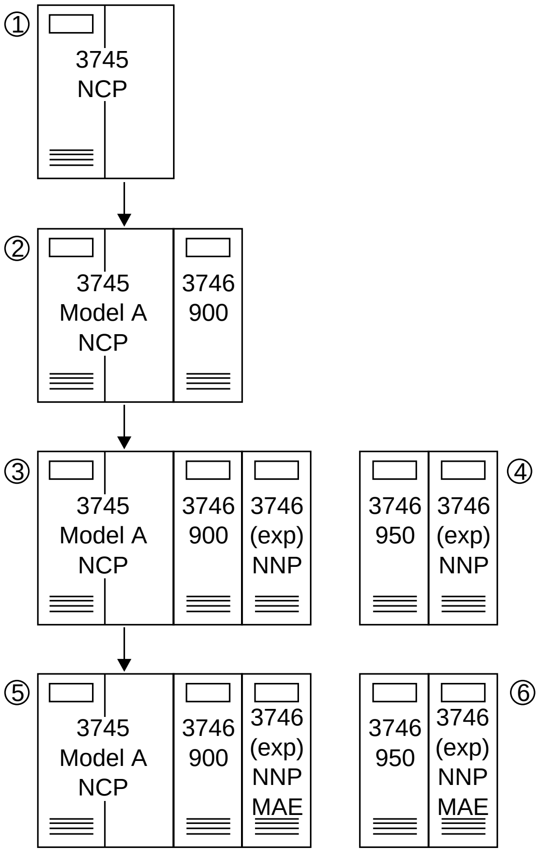

Figure 1-1. 3745 and 3746 Evolution

Figure 1-1 summarizes the 3745 and 3746 Controller evolution. The following list describes the models shown in the diagram:

1. 3745 with NCP

   The main characteristics of all the 3745 models are:

   - All resources are NCP controlled
   - Can support the 3746 Model 900 (Models xxA only)
   - Up to 16 megabytes of main storage per 3745 central control unit (CCU) for the 3745 Models 310, 610, 31A and 61A
   - A service processor, which provides the maintenance and operator subsystem (MOSS) and extended (MOSS-E) functions for the models xxA.

   The 3745 Models A, with or without an attached 3746-900, provide functions that increase the possibilities for designing or expanding 3745-based SNA or APPN/HPR networks. The 3745, with or without a 3746-900, provides APPN/HPR network node support. This requires VTAM to operate as a composite network node with one or multiple 3745/3746-900(s) and NCP(s).

2. 3745/3746-900 with NCP

   As an extension to the 3745 A models, the 3746-900 was introduced to increase the 3745 throughput and connectivity. Initially, high-performance ESCON and token-ring adapters were supported. Later, high-connectivity adapters for communication links, high-performance frame relay capabilities, including frame relay boundary support and primary rate access to Euro-ISDN were added.

3. 3746-900 with Network Node Processor (NNP)

   The addition of the *Network Node Processor* (NNP) allowed the 3745 attached 3746-900 to operate as an APPN/HPR network node, independent from VTAM and NCP, and as an APPN/HPR composite network node (CNN) controlled by NCP and VTAM.

   In these modes, the adapters of the 3746 Model 900 are shared between traffic controlled by the 3746 network node, traffic controlled by NCP, and traffic controlled by the 3746 IP router.

   The 3746-900 allows the 3745 and associated SNA network to evolve and grow in capacity, function, and performance. It also provides upgrades to protect current and future investments in 3745-based networks. As an APPN/HPR network node, the 3746-900 offers a flexible and cost-effective evolution path from SNA networking to APPN/HPR and IP networking.

4. 3746-950 with Network Node Processor (NNP)

   The 3746 with NNP can also operate stand-alone as a 3746-950. This can function as an APPN/HPR network node and IP router concurrently.

   The 3746-950, under control of its network node processor, supports APPN/HPR network node and IP routing functions over communication line adapters, token-ring adapters (type 2), and ESCON adapters (type 2). All the 3746 type 2 adapters can run 3746 network node traffic and 3746 IP traffic traffic simultaneously.

   The 3746-950 supports existing SNA traffic, such as 3270 flows, via its dependent logical unit requester (DLUR) function.

   The 3746-950 also provides high-connectivity and high-throughput IP routing, along with an important set of IP functions, such as support of OSPF, RIP, BGP. The 3746-950 supports IP routing over token-ring and Ethernet LANs, frame relay, X.25, and PPP lines, and ESCON channels.

5. 3746-900 with NNP and MAE (MultiAccess Enclosure)

6. 3746-950 with NNP and MAE (MultiAccess Enclosure).

   The *3746 Multiaccess Enclosure* provides interface support for additional ESCON host connectivity, worldwide ISDN primary, for 155 Mbps ATM (multimode and single mode fiber) as a forum-compliant LANE emulation client, and classical IP routing as well as additional TR, Ethernet and low-speed lines (up to T1/E1).

   The 3746 Multiaccess Enclosure support announced and previewed in 1997 can be seen as a multiple phase enhancement to the 3746.

   Multiaccess Enclosure available 06/30/97

   The 3746 Multiaccess Enclosure is loosely coupled to the 3746. Data can be transferred between the MAE and 3746 via two token-ring connections. Two TIC3 token-ring couplers transfer data over dedicated token-ring connections to the token-ring adapter in slot 1 of the MAE. No other devices may connect to these rings. Configuration of these adapters is done by the MAE configuration program. MAE adapters cannot be activated or controlled by NCP.

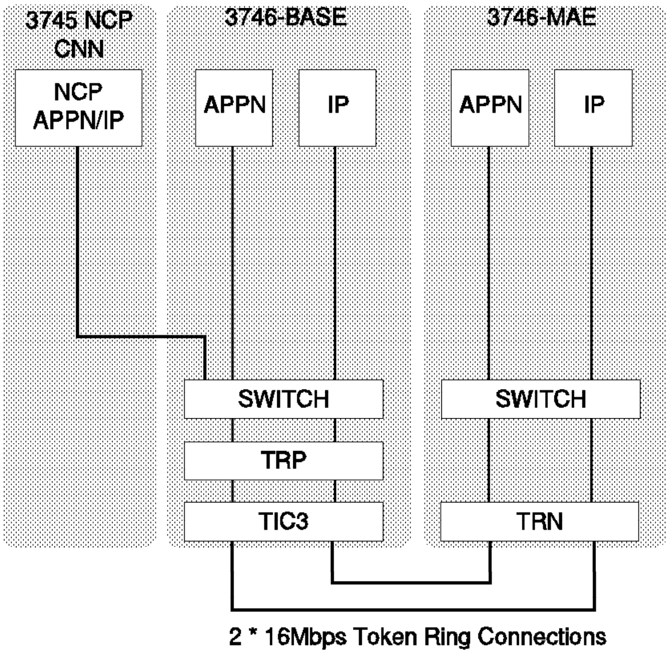

Figure 1-2. 3746 Multiaccess Enclosure Release 0

Multiaccess Enclosure--Preview

The 3746 Multiaccess Enclosure will now be attached directly to the connectivity switch of the 3746. This will require a new card in slot A next to the system card in the Multiaccess Enclosure. This will remove the need for two TIC3 token-ring adapters in the 3746, which are dedicated to data transfer between the 3746 and MAE. Also, the token-ring card in slot 1 of the MAE is now available for general use. The MAE adapters will be under control of the 3746 APPN/HPR and IP control points. They will be configured and controlled by the CCM software in the same way as native adapters.

Multiaccess Enclosure Follow-on Releases

The 3746 Multiaccess Enclosure IP and APPN control points will be merged with the 3746 APPN and IP control points to give a single APPN and a single IP control point image for the 3746.

Further 3746-900 and 950 Enhancements

In addition to the MAE, further enhancements were announced for the 3746 in early 1997. These include:

- Frame Relay Frame Handler (FRFH) functions are now supported by the Network Node Processor of the 3746 Models 900 and 950.
- Frame relay links now support bandwidth management via Committed Information Rate (CIR) and Bandwidth Reservation System (BRS).
- IP traffic flowing on PPP links also can use the Bandwidth Reservation System (BRS) to assign different transmission priorities to different applications.
- APPN adapter connectivity was increased by 100%.
- APPN/DLUR throughput was increased by up to 40%.

- The backup function available on the Euro-ISDN primary adapter of the 3746-900 for SNA (NCP) traffic flowing on frame relay links has been enhanced with automated operations and NetView Performance Monitor (NPM) support.
- Box Connectivity improvements.
  - 5000 PUs supported
  - 15000 ISR sessions supported

Networking Overview

Section 1 of this book introduces a series of topics that provide basic planning and migration guidance for using the 3745/3746-900 and 3746-950 in the following environments:

- Systems Network Architecture/APPN® (SNA/APPN)
- APPN/High Performance Routing (HPR):
  - Rapid Transport Protocol (RTP)
  - Automatic Network Routing (ANR)
  - Adaptive Rate-Based (ARB) congestion control
- Internet Protocol (IP) routing
- Frame-relay Data Link Connection (DLC)
- Point-to-Point Protocol (PPP) Data Link Connection (DLC)
- X.25 Data Link Connection (DLC)
- ISDN Connections
- ATM LANE Emulation and Classical IP

Figure 1-3 on page 1-6 shows the connections available with a 3746 Nways Multiprotocol Controller (3746-900 or 3746-950).

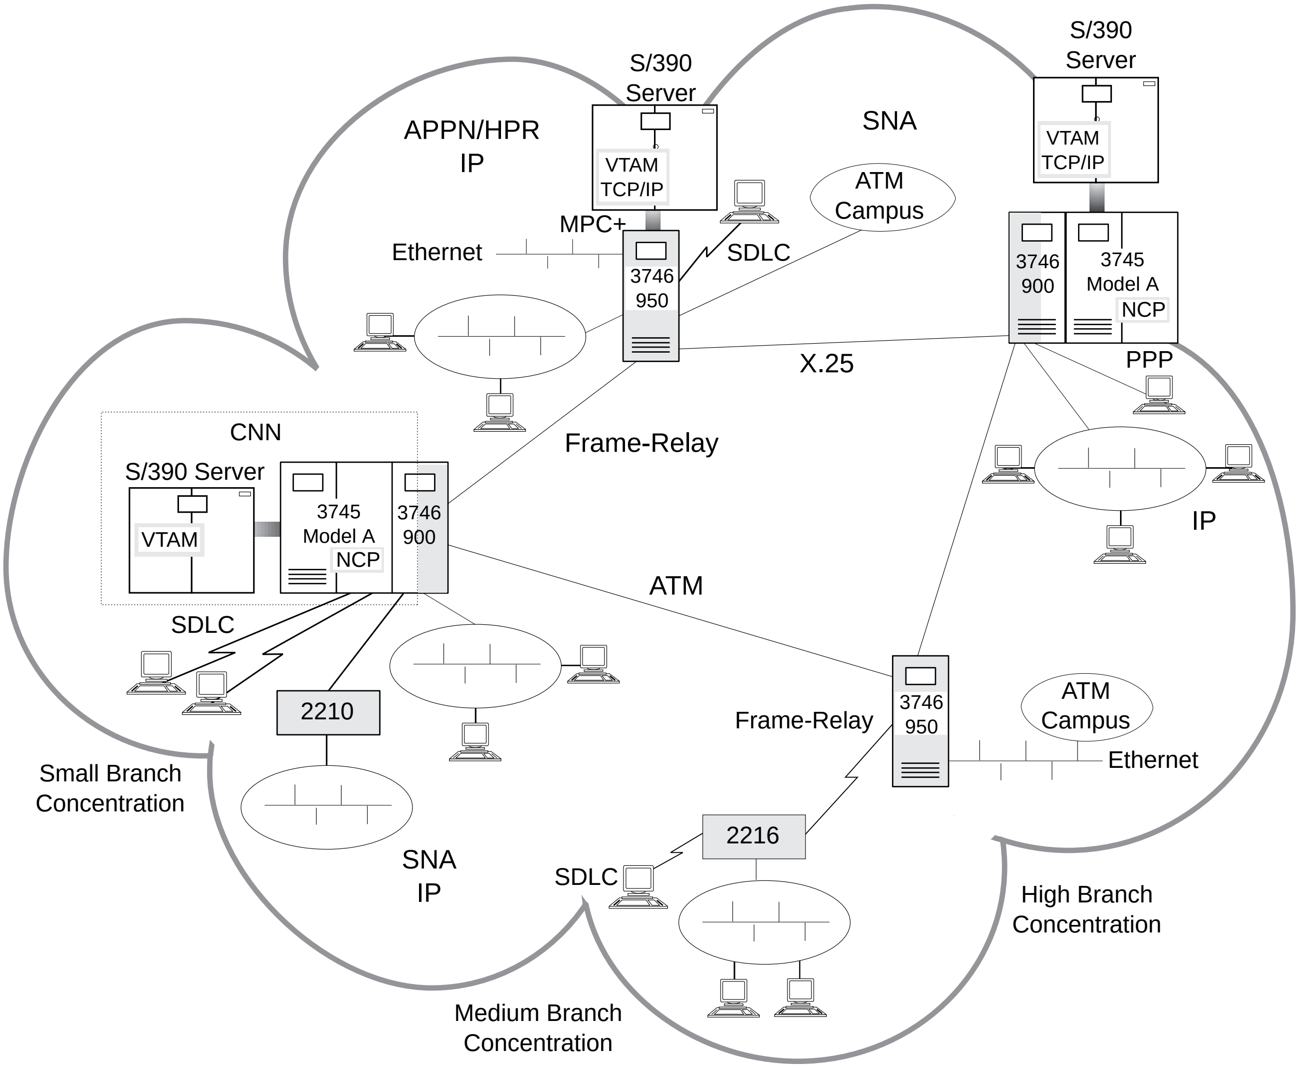

Figure 1-3. 3746 Nways Multiprotocol Controller Connectivity

3745 and 3746 Model Overview

The following is an overview of the 3745 and 3746 models discussed in this book.

The IBM 3746 Models A11, A12, L13, L14, and L15 are expansion units that provide the IBM 3745 Communication Controller with additional parallel channel adapters, token-ring adapters, low-speed scanners, and communication line interfaces. The IBM 3746 Nways Multiprotocol Controller Model 900 and 950 are stand-alone APPN/HPR network nodes and/or Internet Protocol (IP) routers with ESCON channel links, token-rings, communication lines, Ethernet, ISDN-PRI and ATM attachments.

Model 3746-900

The IBM 3746 Expansion Unit Model 900 provides the IBM 3745 Communication Controller Models 17A, 21A, 31A, 41A, and 61A with ESCON® channel adapters, additional SDLC and X.25 communication lines, token-ring, Ethernet, ISDN-PRI (EMO ISDN only) and frame relay attachments. The native support of the ESCON architecture provides flexibility in the design of host front-end installations. Very fast access to S/390 servers can be obtained on WAN and Campus networks when attached to the 3746 Model 900. The 3746 Model 900 can be operated as an APPN/HPR network node and/or an IP router independently from NCP, and a SNA/subarea or Composite Network Node (CNN) controlled by NCP.

Model 3746-950

The IBM 3746 Nways Multiprotocol Controller Model 950 is a host-independent network node, primarily designed for Advanced Peer-to-Peer Networking (APPN)/High Performance Routing (HPR) and IP routing across SDLC and X.25 lines, token-ring, Ethernet, ISDN-PRI and ATM attachments, frame relay links, and ESCON channels, including support for pre-APPN SNA devices (3270 type) through the Dependent Logical Unit Requester (DLUR) function. Routing is distributed in the various (up to 16) adapters, while a Network Node Processor performs the control point services. The 3746 Model 950 operates without being controlled by ACF/NCP running in a 3745 or even having a 3745 attached. (This reduces floor space requirements, maintenance fees and NCP license fees.) A service processor supports the maintenance and operator procedures.

Model 3746-A11

The 3746-A11 is a modular unit that attaches to a 3745. The Input Output Control (IOC) buses attached to the Central Control Unit (CCU) of the 3745 are extended to the 3746 Expansion Units Model A11 in order to attach additional adapters: channel adapters, possibly fitted with a two-processor switch and low-speed scanners.

Model 3746-A12

The 3746-A12 is a modular unit that attaches to a 3745. The Input Output Control (IOC) buses attached to the Central Control Unit (CCU) of the 3745 are extended to the 3746 Expansion Units Model A12 in order to attach additional low-speed scanners.

Model 3746-L13

The 3746-L13 is a modular unit that attaches to a 3745. The 3746 Expansion Unit Model L13 mainly comprises up to four LIC enclosures called LIC units, housing up to 16 LICs.

Model 3746-L14

The 3746-L14 is a modular unit that attaches to a 3745. The 3746 Expansion Unit Model L14 mainly comprises up to four LIC enclosures called LIC units, housing up to 16 LICs.

Model 3746-L15

The 3746-L15 is a modular unit that attaches to a 3745. The 3746 Expansion Unit Model L15 mainly comprises up to four LIC enclosures called LIC units housing up to 16 LICs.

Chapter 2. APPN / HPR Overview

This section gives you a short introduction to Advanced Peer-to-Peer Networking and HPR, serving as a conceptual base to understand the 3746-900 and the Nways Controller in an APPN network. In order to have an in-depth knowledge of the APPN and HPR architecture and to know about the details of specific subjects, we strongly recommend that you read *APPN Architecture and Product Implementations Tutorial*, GG24-3669.

For VTAM and NCP implementations of HPR refer to:

- *ACF/VTAM V4R3 HPR Early User Experiences*, SG24-4507
- *ACF/NCP V7R3: New Functions*, SG24-2592.

Introduction

With the advancements in client/server and peer-to-peer technologies associated with the increasing power of workstations and midrange computers, it became essential to extend the Systems Network Architecture (SNA) in order to address these new networking requirements. That is why Advanced Peer-to-Peer Networking (APPN) was created.

Basing APPN on SNA protocols gives APPN a set of important characteristics, such as:

- Class-of-service: depending on the nature of data (interactive, batch, file transfer, etc.), APPN will select a physical path and assign the corresponding traffic priorities.
- Segmenting: A message can be divided in several smaller units. This allows the application to be independent of the characteristics of the network, such as the maximum blocksize of a line. This also helps provide fairness among the various sessions sharing one physical link.
- Flow control: pacing mechanisms prevents one node from flooding another one with excessive data.

When using SNA, customers were required to configure networks in a hierarchical design. Such topology often lacks the flexibility to address varying network geographies, sizes, and workgroup relationships. Moreover, the application's design has changed; instead of having code running in just one central processor, now there can be many processors running code.

APPN provides the flexibility to meet modern requirements for various dynamic topologies users need in their networks. For example, each networked computer can be directly connected to every other computer (known as a "mesh") or they can all connect through a single routing network node. Alternatively, some customers will choose to continue to use a hierarchical network design. Mesh, network node, hierarchical networks, as well as mixtures of these, are all possible using APPN.

APPC (Advanced Program-to-Program Communication) is usually provided as system software. The APPC software provides two interfaces. The first, a programming interface "at the top," responds to requests from application programs that need to communicate. One example of such an interface is the CPI-C. The second interface, "at the bottom," exchanges data with the communications hardware.

To determine where partner LUs are located in the network, the nodes in an APPN network exchange different types of messages, known as APPN control information. At each node in an APPN network, there is a control point (CP) that is responsible for managing its resources. Each node that has a CP establishes CP-CP with its *adjacent nodes* for both data and control traffic.

In the following sections, we give a description of each node type showing the differences between them.

Node Types

Each node type in the APPN architecture has a specific function set. Following, you have a description of how each node type works. Table 2-3 on page 2-66 and Table 2-4 on page 2-69 summarize all the option sets defined in the APPN architecture, and shows their implementation in the 3746.

APPN Network Node

An APPN network node provides distributed directory and routing services for all LUs that it controls. These LUs may be located on the APPN network node itself or on one of the adjacent LEN or APPN end nodes for which the APPN network node provides network node services. Jointly, with the other active APPN network nodes, an APPN network node is able to locate all destination LUs known in the network.

A facility known as *central resource registration* allows an APPN network node to register its resources at a central directory server. Once a resource is registered, APPN network nodes can locate the resource by querying the central directory server instead of using a broadcast search, thus improving network search performance during session establishment.

After the LU is located, the APPN network node is able to calculate the route between origin and destination LU according to the required class of service. All network nodes exchange information about the topology of the network. When two adjacent network nodes establish a connection, they exchange information about the network topology as they know it. In turn, each network node broadcasts this network topology information to other active and adjacent network nodes with which it has CP-CP sessions.

Alternatively, if the connection between network nodes is deactivated, then each network node broadcasts this change to all other, active and adjacent, network nodes. An APPN network node that is taken out of service will be declared inactive and, after some time, removed from the topology information in all network nodes, together with its routing capabilities to other nodes.

The APPN network node is also capable of routing LU-LU sessions through its node from one adjacent node to another adjacent node. This function is called intermediate session routing.

APPN End Node

An APPN end node provides limited directory and routing services for LUs local to it. The APPN end node can select an adjacent APPN network node and request this network node to be its *network node server*. If accepted by the network node, the APPN end node may register its local resources at the network node server. This allows the network node server to intercept Locate search requests for resources that are located on the APPN end node and pass the request on to the APPN end node for verification.

Without a network node server, an APPN end node can function as a LEN end node and establish LU-LU sessions with a partner LU in an adjacent APPN or LEN node.

The APPN end node sends Locate search requests, for resources unknown to the APPN end node, to its network node server. The APPN network node uses its distributed directory and routing facilities to locate the LU (via directed, central directory, or broadcast searches) and calculates the optimal route starting at the APPN end node toward the destination LU.

The APPN end node may have active connections to multiple adjacent network nodes; however, only one of these network nodes at a time can act as its network node server. The APPN end node selects its network node server by establishing CP-CP sessions with an adjacent APPN network node.

On APPN network nodes, the APPN end nodes are categorized as either *authorized* or *unauthorized*. An authorized APPN end node may send registration requests to register local network accessible resources at a network node server, a facility known as *end node resource registration*, and may, in addition, request that these resources be registered with the central directory server. If during session establishment a network node server does not know where an LU is located, the network node server queries authorized APPN end nodes within its domain that have indicated they are willing to be queried for unknown resources. Network accessible resources on unauthorized nodes require explicit definition at the network node server as part of its system definition or dynamically by the network node server's operator. To avoid explicitly defining resources of authorized nodes at their network node server, the APPN end node should either register its resources or allow the network node server to query the APPN end node for unknown resources.

An APPN end node can attach to any LEN or APPN node regardless of its network ID.

LEN End Node

A LEN end node provides peer-to-peer connectivity to other LEN end nodes, APPN end nodes, or APPN network nodes. A LEN end node requires that all network-accessible resources, either controlled by the LEN end node itself or on other nodes, be defined at the LEN end node. LUs on adjacent nodes need to be defined with the control point name of the adjacent node. LUs on nonadjacent nodes need to be defined with the control point name of an adjacent network node, as LEN end nodes assume that LUs are either local or reside on adjacent nodes.

Unlike APPN end nodes, the LEN end node cannot establish CP-CP sessions with an APPN network node; therefore, a LEN end node cannot register resources at a network node server, nor can it request its network node server to search for a resource, or, to calculate the route between the LEN end node and the node containing a destination resource.

However, indirectly a LEN end node uses the distributed directory and routing services of an adjacent network node by predefining remote LUs, owned by nonadjacent nodes, with the CP name of an adjacent APPN network node. The session activation (BIND) request for that remote LU is sent by the LEN end node to the adjacent network node. The network node, in turn, automatically acts as the LEN end node's network node server, locates the actual destination of the LU, calculates the route to it, and uses this route to send the BIND to its final destination.

A LEN end node can attach to any LEN or APPN node regardless of its network ID.

Other Node Types

Additionally, there are some other node types that will be used in this publication. They are synonyms for nodes as seen from a subarea network, represent a specific junction in the network, or represent an APPN node with additional functions:

- Boundary and peripheral node
- Composite node
- Interchange node
- Virtual routing node
- Border node
- Extended border node

Boundary and Peripheral Node

Within traditional subarea SNA, the resources in a domain of a subarea SNA network are controlled through a hierarchical structure. The nodes that play a role in these networks are categorized as subarea and peripheral nodes. An example of such an SNA network is an S/390 type mainframe running VTAM and a 3745 communication controller running the network control program (NCP). Both VTAM and NCP are referred to as subarea nodes. The VTAM subarea node includes the control point function, hereafter called the System Services Control Point (SSCP). Like the APPN control point, the SSCP controls all the resources that are in its domain.

Attached to these subarea nodes, or *boundary nodes*, are the *peripheral* nodes. The peripheral node is either a PU T2.0 or an APPN or LEN node. The PU T2.0 node is a traditional hierarchical node that requires the support of an SSCP to establish sessions. Traditional subarea SNA allowed LEN connections only; CP-CP sessions could not be established between VTAM and the APPN nodes.

With the introduction of APPN VTAM, a VTAM or a composite network node (subarea network consisting of one VTAM and one or more NCPs) is able to present an APPN image to other APPN nodes. APPN VTAM allows CP-CP sessions with APPN nodes attached to the VTAM or NCP boundary function, to get full APPN connectivity. The term “peripheral” node has lost its value in a network that is truly peer-to-peer.

Composite Node

The term *composite node* is used in some publications to represent a group of nodes that appear as *one* APPN or LEN node to other APPN or LEN nodes in an APPN network. For example, a subarea network that consists of one VTAM host and one or more NCPs consists of multiple nodes, but when connected to an APPN node, appears as *one* logical APPN or LEN node.

A subarea composite node may appear as either a LEN end node or as an APPN network node. In the former case, the term composite LEN node is used; in the latter case the term composite network node (CNN) is used.

Interchange Node

A VTAM host acting as an interchange node (ICN) can be a stand-alone APPN VTAM node or a composite network node. The ICN routes sessions from APPN nodes into and through the subarea network using subarea routing, without exposing the subarea implementation to the APPN part of the network. This is accomplished by making the APPN VTAM node, plus all its owned resources, appear to other nodes as a single APPN network node with multiple connections. At the same time the ICN, and the NCPs it owns, will maintain their subarea appearance to other subarea nodes.

The ICN supports SSCP-SSCP sessions with other VTAM nodes as well as CP-CP sessions with adjacent APPN network nodes and end nodes. This support allows the ICN to use both APPN and subarea data flows to locate LUs and to provide the best route between nodes. APPN session setup protocols, which flow on CP-CP sessions, are converted to the corresponding subarea protocols that flow on SSCP-SSCP sessions, and vice versa.

To an ICN multiple VTAMs and NCPs may connect using subarea protocols (see for example VTAM1/NCP in Figure 2-1 on page 2-6). Session establishment is possible between any LU in the subarea network and any LU in the APPN network. The VTAM host to which APPN nodes attach, or the VTAM host owning the NCPs to which APPN nodes attach, must have implemented APPN VTAM, as it is responsible (as an “interchange node”) for the conversion of subarea to APPN protocols and vice versa; other VTAMs within the subarea network may be back-level VTAMs. From the viewpoint of the APPN nodes, LUs owned by VTAMs (for example, VTAM2 or VTAM3) other than the VTAM providing the interchange function, are considered to reside on APPN end nodes.

Note: Figure 2-1 on page 2-6 shows the basic form of connecting APPN and subarea networks using a composite network node acting as an interchange node.

Virtual Routing Node

APPN allows APPN nodes to reduce the addressing information stored at each node connected to a shared-access transmission facility (SATF), such as a token-ring, by allowing each node to define a virtual routing node (VRN) to represent its connection to the shared facility and all other nodes similarly configured. The SATF and the set of nodes having defined a connection to a common virtual routing node are said to constitute a *connection network*.

A virtual routing node (VRN) is not a node, but it is a way to define an APPN node's attachment to a shared-access transport facility. It reduces end node definition requirements by relying on the network node server to discover the common connection and supply necessary link-level signaling information as part of

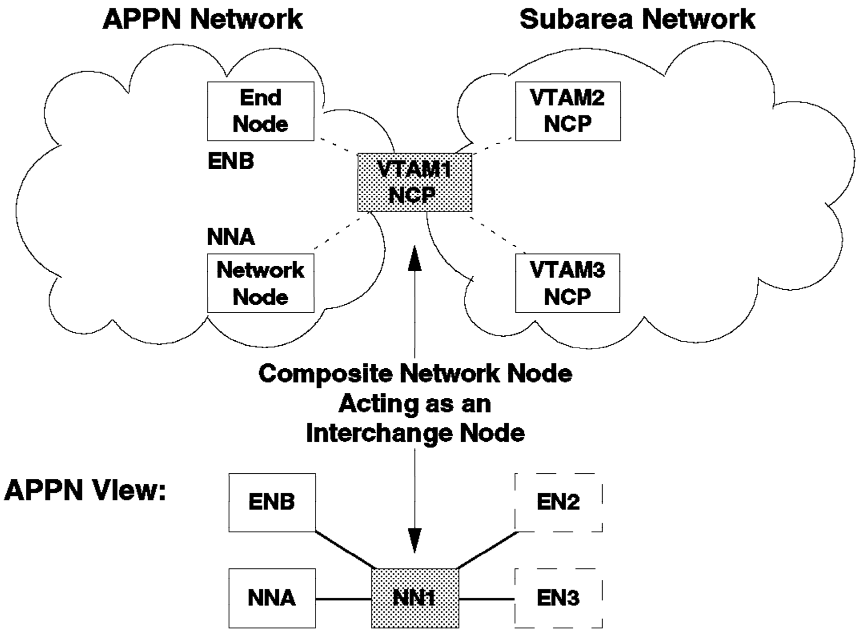

Figure 2-1. Composite Network Node Acting As an Interchange Node

the regular Locate search process; LU-LU session data can then be routed directly, without intermediate node routing, between APPN nodes attached to the SATF.

Border Node

Base APPN architecture does not allow two adjacent APPN network nodes to connect and establish CP-CP sessions when they do not have the same network ID. The border node is an optional feature of an APPN network node that overcomes this restriction.

A border node can connect to an APPN network node with a different network ID, establish CP-CP sessions with it, and allow session establishment between LUs in different network ID *subnetworks*. Topology information is not passed between the subnetworks. Similarly a border node can also connect to another border node. A particular type of border node that is mentioned in this publication is the extended border node.

Extended Border Node

The extended border node allows the connection of network nodes with different network IDs and session establishment between LUs in different network ID subnetworks that need not be adjacent.

An extended border node provides directory session setup and route selection services across the boundary between paired or cascaded nonnative network ID subnetworks. An extended border node can also partition a single network ID subnetwork into two or more clusters or topology subnetworks with the same network ID, thus isolating one from the topology of the other.

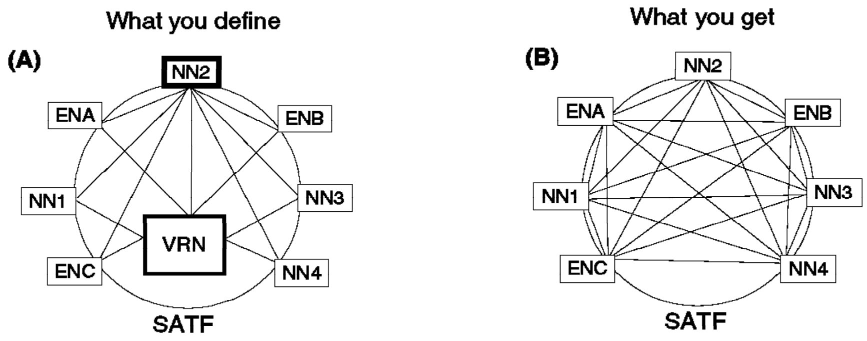

Figure 2-2. Virtual Routing Node

APPN Node Structure

This section briefly describes the structure and components of APPN network nodes, end nodes, and LEN nodes. An overview of each is presented in Figure 2-3 on page 2-11.

Node Operator

This component defines all information required by the node (for example, on links to adjacent nodes, and on LUs within its domain) and causes activation and deactivation of the node and its resources (for example, links). It may also query the status of a node's resources.

A node operator can be a *person* using a system-specific dialog manager, which converts the information entered by the individual into node operator commands that are passed to the *node operator facility*. Also, a node operator can be a *command file* whose records are read and converted into node operator commands that are passed to the node operator facility. Finally, the node operator can be a *transaction program*. In this case, a remote transaction program communicates with a local transaction program and this local program converts the information received into node operator commands.

All three types of node operators use a program within the system to interact with the node operator facility.

Node Operator Facility (NOF)

The function of this component is to allow communication between the node operator and the control point (CP), intermediate session routing (ISR), and LUs. NOF initializes the CP and ISR components when the node is started. It also performs functions such as the following when requested to do so by the node operator:

- Defining (creating) and deleting (destroying) LUs
- Activating and deactivating links
- Querying the CP and ISR for database and status information

Application Transaction Program (TP)

These programs communicate with other local or remote application transaction programs (TPs) to perform user-defined functions. Communication is accomplished by establishing conversations between TPs. Data is then exchanged between the TPs using an LU verb interface such as CPI-C. These verbs facilitate the task of writing a transaction program by isolating the programs from the protocol levels.

No matter what the verbs are, the tasks performed by them in a transaction program are the following:

- Allocate a conversation
- Accept a conversation
- Send data
- Receive data
- Grant permission to send
- Request a confirmation
- Grant or reject a confirmation
- End a conversation

Control Point (CP)

The function of the CP is to manage the resources of the node. It creates the path control (PC) and data link control (DLC) components. The CP also manages session resources and provides facilities such as directories and topology information. The CP is created by NOF when the node is started and is composed of several functions.

The following is a brief description of the CP components:

- **Configuration Services (CS)** manages the node's local resources such as links to adjacent nodes. It is responsible for the definition of ports, types of data link control (DLC), adjacent link stations, adjacent nodes and attached connection networks. In addition, the following functions are performed: link activation (including XID exchange), nonactivation XID exchange (SSCP takeover, for example), link deactivation, link queries and connection networks capabilities (not supported in LEN end node). The NOF initializes, starts, stops and queries the configuration services. As it manages the local resources, it interacts with the other components of the CP, described below, and with the path control (PC) and the data link control (DLC).

- **Topology and Routing Services (TRS)** on LEN end nodes and APPN end nodes, collects information on links and adjacent nodes. In APPN network nodes, additionally, TRS collects and exchanges information on other network nodes and links between them. Furthermore, it provides the optimum route for LU-LU sessions.

- **Directory services (DS)** is responsible for locating network resources throughout the APPN network. On APPN end nodes, it searches network resources in its local database first, and if the resource cannot be located, uses the services provided by the network node server that this APPN end node has a CP-CP session with. On LEN end nodes, it only searches resources in its local database, since it cannot keep CP-CP sessions with a network node.

- **Session Services (SS)** is responsible for activating and deactivating CP-CP sessions used by the CP components to exchange network information. Also, SS is responsible for assigning unique session identifiers to sessions, and to assist LUs in activating and deactivating sessions.

- **Address Space Manager (ASM)** administers address space information that is used by path control to identify each individual session on a given link. This address space is a set of binary numbers, each one being 17 bits long. These numbers are used to uniquely identify a session between two adjacent nodes and are called the local form session identifier (LFSID). The LFSID, for a session being established through a intermediate node, is assigned by the ASM and passed to the session connector manager (SCM) that is a component of the intermediate session routing (ISR) function described in the next item. Also, ASM interacts with LUs and ISR at BIND/RSP(BIND) and UNBIND/RSP(UNBIND).
- **Management Services (MS)** monitors and controls resources of a node. It can generate alerts to the network operator.

Intermediate Session Routing (ISR)

The intermediate session routing (ISR) component is present only in an APPN network node. The primary function of ISR is to route session traffic received from one node and destined to another node. So, its node does not contain a destination LU. ISR is created by NOF when the node is started.

The components of the ISR are:

- **Session Connector Manager (SCM)** interfaces with the ASM to obtain the LFSID and the transmission group in the direction of the destination LU. Perform intermediate BIND processing.
- **Session Connector (SC)** connects two stages of a session. The ISR must receive a frame and forward it to another node that can be the destination or another intermediate node. This is performed using a table that associates a source address in an incoming frame with an address on which the frame must be forwarded. These addresses are based on the LFSID of each session and are carried in the TH field of the SNA frame (specifically in the origin and destination address fields).

Logical Unit (LU)

The LU serves as a port into the network for one or more application transaction programs. It establishes sessions with other LUs. Conversations are allocated on these sessions that allow communication between TPs. Also, the session-level pacing is done by the SC.

Path Control (PC)

This component routes message units from LUs, ISR, and CP within the node to DLC for transmission to adjacent nodes. Messages received by path control from DLC are routed to the appropriate component (CP, LU, or ISR). PC also routes message units between LUs within the local node.

Additionally, it performs segment generation, handles transmission priorities and performs error checking on TH.

Data Link Control (DLC)

DLC provides the protocols necessary to ensure reliable delivery of messages between link stations in adjacent nodes attached to a common transmission medium. DLC also controls the node attachment to various types of transmission media.

Names

Resource naming is important because it allows end users to start a session without knowing the exact location of different resources within the network.

The Network Accessible Unit

In an APPN network, all components that can establish sessions with one another are called network-accessible units. Examples are CPs and LUs. The term NAU was previously used as an abbreviation for network addressable units. The terminology has changed with APPN, and now NAUs are represented by names rather than by addresses.

Within an APPN network, the names of network accessible units must be unique. To ensure the uniqueness of names within the network, a consistent naming convention is required. To make the administering of resource names easier, the network can be divided into partitions.

Network Identifiers

A partition of the network may be given a unique network identifier (net ID). Net IDs are 1 to 8 bytes long. The net ID is used throughout SNA, both in the subarea and the APPN part of a network. Because names of LUs and CPs have to be unique only within the scope of a net ID, they can be assigned and administered independently for each distinct partition of the network.

Network Names

A network name is an identifier of a network resource. Each CP, LU, link, and link station in an SNA network has a network name. The network names are assigned through system definition. In an APPN node, the system definition is done using the node operator facility (NOF).

Network-Qualified Names

A network-qualified name identifies both the resource and the network in which the resource is located. It is a concatenation of the network ID and the network name of the resource; for example, the names NETA.LUA and NETB.LUA refer to different entities.

Addresses

Network addresses uniquely identify a resource throughout the subarea network. Local addresses uniquely identify a session on a link. APPN uses local addresses. Traditional SNA subarea networking uses local *and* network addresses. Local addresses are used between peripheral nodes and the boundary functions of VTAM and NCP; network addresses are used when routing data between subarea nodes and bear no relation to specific sessions.

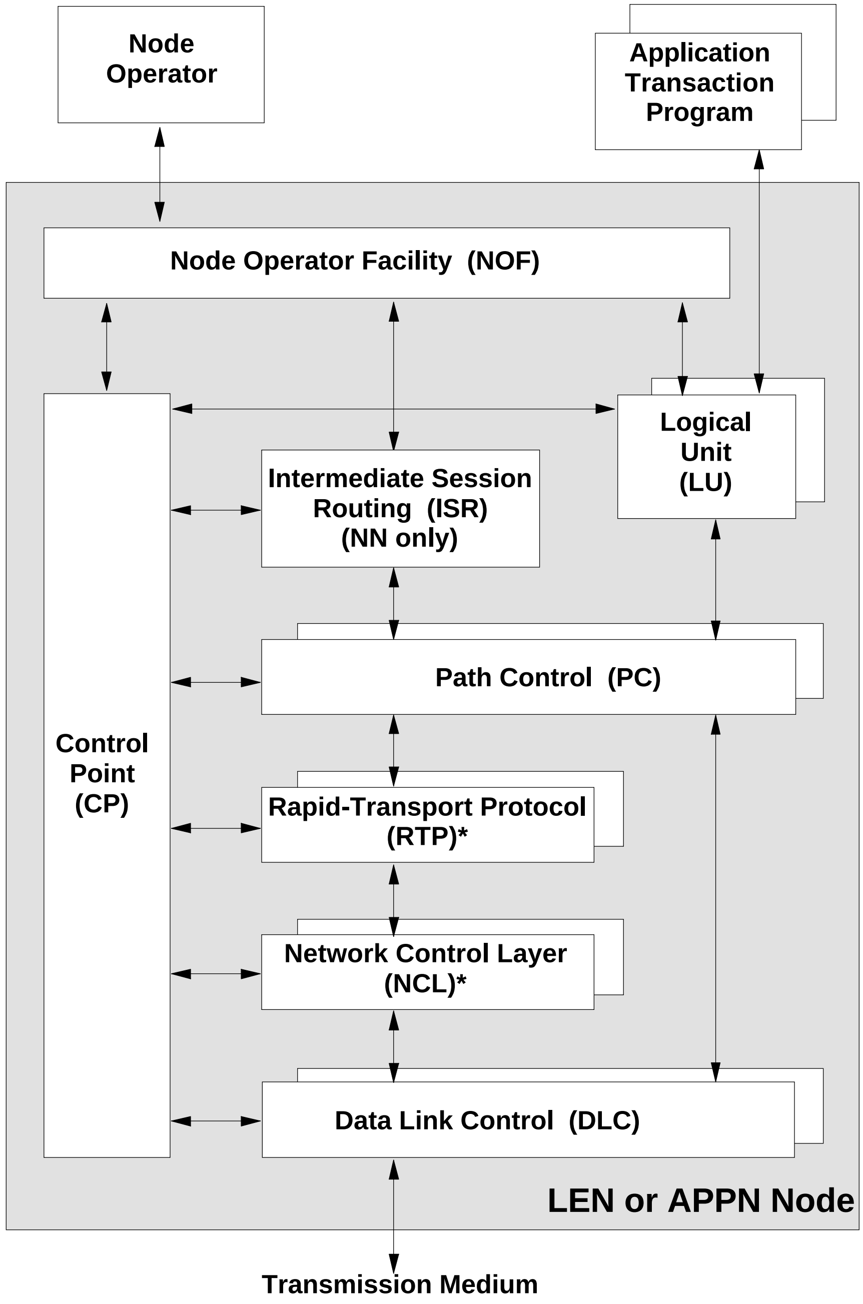
Node Operator
Application Transaction Program
Node Operator Facility (NOF)
Control Point (CP)
Logical Unit (LU)
Intermediate Session Routing (ISR) (NN only)
Path Control (PC)
Rapid-Transport Protocol (RTP)*
Network Control Layer (NCL)*
Data Link Control (DLC)
LEN or APPN Node
Transmission Medium

```
Legend:
 *  =  APPN Node only
```

Figure 2-3. Structure of an APPN or LEN Node

The address used in an APPN transmission header is an identifier unique on the given link for a particular session rather than the address of the NAU.

Addresses are used for routing. Routing in an SNA network uses a combination of two things:

- Information carried in the transmission header of the message
- Information stored in the intermediate node

In an APPN network, routing information is session oriented. The transmission header carries session identifiers that are locally defined for each pair of adjacent routing nodes and are only *temporarily* assigned. They are assigned at session initiation, and released when the session ends. The session initiation request (BIND) carries routing information about the full session path that determines the sequence of links used from origin to destination. The local session identifier stored in each intermediate node in a session path is contained in a session connector and kept for the life of the session only.

The session identifier is associated with:

- A particular session
- A transmission group (link) between two nodes

Figure 2-4 on page 2-13 shows a session between two LUs, LU1 and LU2, residing on two nonadjacent APPN end nodes. The session data is routed through two intermediate network nodes. The session can be thought of as a sequence of three session stages or “hops” with a distinct session identifier assigned to each session stage.

Session identifiers vary at different session stages, which is why they are called local form session identifiers (LFSID). The LFSID is set up during session establishment by the address space manager component of the CP and assigned for the “lifetime” of an LU-LU (or CP-CP) session.

Each session is uniquely identified by a network-unique identifier, the fully qualified procedure correlation ID (FQPCID).

Domains

A domain is an area of control. A domain in an APPN network consists of the control point in a node and the resources controlled by the control point. Consequently, all APPN networks are multidomain networks.

Although all APPN nodes are peers with respect to session initiations and do not rely on other nodes to control their resources, APPN end nodes and LEN end nodes use the services of network nodes. The domain of an APPN end node or LEN end node contains the node's own (local) resources. The domain of an APPN network node contains its local resources *and* the resources of those nodes that use the network node's services. Thus, the domains of the APPN end nodes and LEN end nodes are included in the domains of their respective network node servers.

Note: In traditional subarea networking, a domain is the part of the network managed by a VTAM system services control point (SSCP). Within this Redbook, when using the term domain, we refer to an APPN domain unless explicitly stated otherwise.

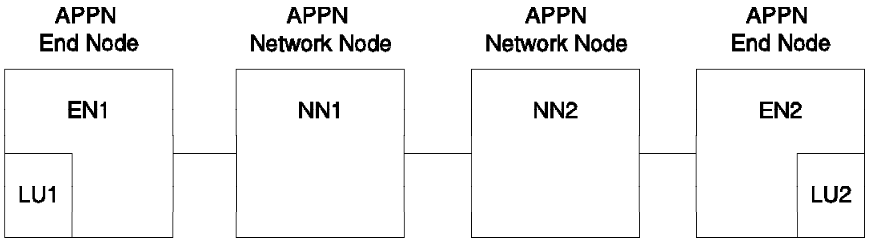

Figure 2-4. Session with Several Session Stages

High-Performance Routing (HPR)

High-performance routing (HPR) is an extension to the APPN architecture. It can be implemented on an APPN end node or an APPN network node and does not change the basic functions of the architecture. It is intended that HPR will be implemented by making software upgrades to existing APPN products, without needing to change the hardware platform.

HPR enhances the routing mechanisms used by APPN to provide the following functions:

- HPR improves the performance over existing APPN routing, especially when using high-speed links.
- It can nondisruptively route sessions around links or nodes that have failed.
- It provides a new mechanism for congestion control that can improve traffic throughput.
- It reduces the amount of storage required in APPN intermediate nodes.

In an HPR network, a new form of routing is used, which is called automatic network routing or ANR. ANR is a source-routing protocol, which means the sender of a packet provides the information about the physical path the packet will use through the network in the network header. As HPR provides the ability to do nondisruptive path switching, the HPR architecture handles the case where the route changes in mid-session.

ANR uses a new form of addressing to identify the route through an HPR network. However, unlike the APPN session-oriented addresses (LFSIDs), the addresses in ANR are based purely on the links which make up the route. The network header contains a list of ANR labels which identify the route through the network. Each ANR label describes a link which is to be taken to exit a node.

In addition to the ANR labels, there are still addresses which are associated with sessions in HPR. Each session will have a pair of unique session addresses, one for each direction. Unlike the LFSID which identifies each stage of the APPN session, the HPR session addresses are used only on an (HPR) end-to-end basis. These are known as enhanced session addresses.

The process of supporting the end-to-end sessions across the HPR network is known as rapid-transport protocol or RTP.

In a network that is supporting both existing APPN nodes and HPR nodes, both the APPN and the HPR methods of addressing are used.

HPR Base and Towers

In order to facilitate implementations across a wide range of products, the following (optional) portions of HPR have been identified:

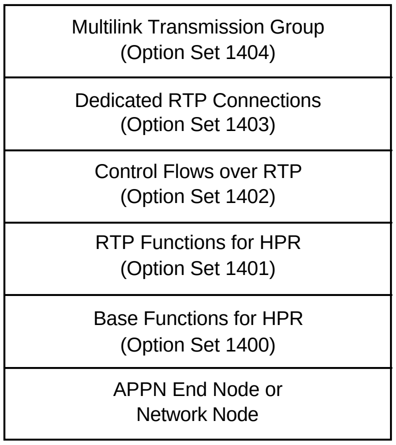

Figure 2-5. HPR Base and Towers

Base Functions (Option set 1400)

The primary function of the HPR base is to provide ANR routing. Products that only implement the base can participate as intermediate nodes for RTP connections. Nodes that do not support the RTP Tower cannot be the end points of RTP connections. The main function in the base is the *Intermediate ANR Routing*.

A new packet format, called a network layer packet (NLP), is used to transport data in the HPR subnet.

NLPs flowing over RTP connections may be efficiently routed through intermediate nodes using ANR routing. The CP-CP session traffic flows still use FID2 PIUs and APPN LU-LU session traffic not flowing over RTP connections will also use FID2 PIUs. Both FID2 PIUs, and NLPs may flow over a single link and are distinguished from one another by the first 4 bits in the packet.

Prior to establishing an RTP connection, a route setup protocol is done over the desired path. Link and node APPN topology update information indicates the appropriate level of HPR support by means of a new HPR control vector. Base-level nodes participate by adding the appropriate ANR information for the inbound and outbound links. When the route setup messages are exchanged between two nodes where one or both are base-level nodes, they flow in a FID2 PIU.

RTP Functions for HPR (Option set 1401)

Nodes that support the RTP functions for HPR Tower are able to transport LU-LU session traffic across HPR networks over RTP connections, thus enabling the use of HPR's high-speed ANR routing and non-disruptive path switch functions. An RTP connection can only be made between nodes that support the RTP Tower so it is essential that there be such nodes in the network. If all the HPR nodes in the network support only the base, there will be no advantages over APPN (in fact, pure APPN protocols will be used). All data flowing over an RTP connection is carried in a network layer packet (NLP). The following functions are included in the RTP Tower:

- Rapid Transport Protocol (RTP)

  This is the transport protocol used in HPR for transporting data across HPR subnets.

- Non-disruptive Path Switch

  If the current path being used by an RTP connection fails, the connection may be switched to a new path automatically. Sessions that are being transported by the RTP connection are not disrupted.

- APPN/HPR Boundary Function Support

  APPN (FID2) traffic is mapped to HPR (NLP) traffic and vice versa.

  When setting up a search for an LU, and the session to the LU over an RTP connection, the search reply will contain the ANR label of the network connection endpoint (NCE) at the end of the RTP associated with that LU.

  If the current path being used for an RTP connection fails, the RTP connection is switched to a new path (whenever possible). Sessions transported over the RTP connection are not disrupted.

  The APPN/HPR boundary function provides the mapping of APPN (FID2 PIU) traffic to HPR (network layer packet) traffic and vice versa.

Control Flows Over RTP Tower (Option set 1402)

Nodes supporting the HPR control flows over the RTP option use RTP connections (if both adjacent nodes support this option) for CP-CP sessions. When a link connecting two nodes that both support this option is activated, a long-lived RTP connection is established that is used to forward route setup messages.

Dedicated RTP Connections (Option set 1403)

Multilink Transmission Group (Option set 1404)

HPR Link Support

Since HPR is an enhancement of APPN it will operate over links supported by today's APPN, supporting both APPN and HPR traffic on the same link. That means that existing hardware adapters and DLCs currently being used for APPN communications can continue to be used for HPR.

During link activation, DLCs are used by HPR in the same manner as by APPN. That is, XID3s are exchanged and the appropriate set mode signals are sent when the exchange is complete. For HPR, a new control vector is included in the negotiation-proceeding XID3 that contains additional HPR-specific information. If both sides include the new control vector, the link is referred to as an HPR link. If both nodes indicate support for the HPR transport option, then HPR transport option protocols are used; otherwise, HPR base protocols are used.

CP-CP sessions are established in the same manner as in APPN when the link is activated.

Immediately after an HPR link is activated and both nodes support the HPR control flows over RTP option, a long-lived RTP connection is established across the link for route setup and CP-CP control flows. This RTP connection remains active as long as the link remains active (hence the “long-lived” RTP connection).

To avoid the necessity to segment the transport layer header, the minimum “maximum packet size” that has to be supported for an HPR link is 768 bytes. For performance reasons it might be advisable to support larger packet sizes.

Automatic Network Routing

Automatic network routing (ANR) mode is a low-level routing mechanism that minimizes cycles and storage requirements for routing packets through intermediate nodes. ANR routing is significantly faster than current APPN routing. No intermediate node storage is required (APPN requires 200-300 bytes per session) and no pre-committed buffers are necessary, which APPN recommends should be used.

HPR uses the ANR routing mode to route session traffic, including binds and unbinds, through an HPR network between nodes supporting the high-performance routing transport option.

HPR employs a route setup protocol in order to obtain ANR and RTP connection information of the selected path.

Each packet is routed through the network as a self-contained unit and is independent of all other packets. There is no table lookup or processing necessary at transit nodes such as the LFSID swapping procedure used by APPN. Any processing of packets required at the network connection and transport connection sublayers is the responsibility of the origin and destination endpoints of the packets. Endpoint processing includes flow control, segmentation and reassembly, and recovery of lost packets.

ANR is designed to be simple enough so high-performance switching can be accomplished. A major goal is to optimize the design for hardware implementation to get the appropriate performance level that is required by the new generation of high-speed networks.

Rapid Transport Protocol (RTP)

RTP is a connection-oriented, full-duplex protocol designed to transport data in a high-speed network. HPR uses RTP connections to transport LU-LU and optionally CP-CP session traffic.

Rapid transport protocol provides end-to-end error recovery with selective retransmission, non-disruptive path switch and adaptive rate-based (ARB) flow control. RTP may be implemented on network nodes or end nodes.

The physical path utilized by the RTP must satisfy the class of service (COS) associated with the session traffic it is carrying. Session traffic is carried over the RTP connection in such a way that intermediate nodes are not aware of the SNA sessions, or even the transport connection itself. Traffic from many sessions may be carried by a single RTP connection, provided they all use the same COS. An RTP connection provides two important advantages:

1. It transports data at very high speeds by using low-level intermediate routing and by minimizing the number of flows over the links for error recovery and flow control protocols. The flows are minimized by performing these functions at the RTP connection endpoints rather than at each hop (link) along the path. Data resequencing takes place at the RTP endpoints.
2. An RTP connection's path may automatically be switched to reroute data around a failed node or link without disrupting the sessions. The new path for the RTP connection is selected that best fits the same class of service as the failed connection. Higher layer protocols are not even notified of the rerouting.

The endpoints of an RTP connection must support the RTP transport option, whereas intermediate nodes need only support HPR base functions. APPN session endpoints can be in APPN or HPR nodes.

Class of Service

RTP connections are used to transport session data between HPR nodes operating within an HPR subnetwork. They provide a full-duplex logical connection or *pipe* between two HPR nodes over a specific path through the HPR subnetwork.

Each RTP connection transports session data between two RTP-capable endpoints for a single class of service (COS) as specified in the BIND. An RTP connection is not used for more than one COS, in order to simplify the route selection process during path switching. A node may activate multiple RTP connections (using different paths) to the same partner node for the same COS in order to distribute the traffic over multiple physical paths.

The same RTP connection that was activated by a node to carry sessions to the remote node may be used for the remote node to establish its sessions along. All traffic from an individual session flows over a single RTP connection, but many sessions may be multiplexed over a single RTP connection. All sessions requesting the same COS and following the same path through the HPR subnetwork are transported over a single RTP connection between the HPR nodes containing the session endpoints.

ARB Flow Control and Congestion Control

In HPR, sessions with the same COS are multiplexed over a single RTP connection. The ARB flow control mechanism, used by RTP, addresses the fairness issue among multiple RTP connections through the network. It does not address the fairness issue among multiple sessions using one RTP connection.

In order to provide this fairness, we not only use the ARB mechanism for RTP connections but also use the existing session pacing over an RTP connection to prevent one session from monopolizing buffers in the two RTP components. In an APPN sense, the RTP connection can be seen as one stage (hop) on the session path.

The APPN hop-by-hop Windows-based flow control protocol, known as adaptive session pacing, is inadequate for high-speed data routing. HPR uses a protocol suitable for high-speed routing called adaptive rate-based (ARB) flow/congestion control. It regulates the flow of data over an RTP connection by adaptively changing the sender's rate based on feedback on the receiver's rate. This protocol allows for high-link utilization and prevents congestion before it occurs.

The input traffic entering the network is regulated by the ARB algorithm based on the conditions in the network and the partner RTP endpoint. An increased delay and decreased throughput indicates that congestion is occurring, and so input traffic is reduced. When the capacity of the network or partner RTP endpoint increases, input traffic is increased.

The ARB algorithm is designed to ensure that maximum throughput is attained. The ARB algorithm has the following properties:

- It adapts to network conditions in such a way as to maximize throughput and minimize congestion. It therefore operates within the operating region.
- It smooths the input traffic into the network, avoiding bursts when the physical capacity of the access link to the network is larger than the allowed input rate. This prevents long queues from developing in the network and helps minimize oscillation in the network traffic patterns.
- It provides end-to-end flow control between the RTP endpoints so that one endpoint does not flood the other.
- It requires minimum overhead in both processor cycles and network bandwidth.
- It provides equal access, or fairness, to all RTP connections.

The ARB algorithm is implemented on the RTP endpoints of a connection. There are two components at the RTP endpoints, an ARB sender and ARB receiver. The intermediate nodes have no awareness of the ARB protocol and therefore do not participate in it.

The sender continually queries the receiver by sending a *rate request* along with the data in order to determine the state of the network and state of the receiver node. The sender may reduce or increase its send rate depending upon the information it gets in the *rate reply* received from the receiver. Fixed characteristics, such as speed of the slowest link in the path and transmission time, are factored into the algorithm when the RTP connection is set up. These path characteristics are communicated by using an ARB *setup* message. Either RTP endpoint may send a setup message as a result of a path switch.

RTP Alive Timer

The RTP *Alive* timer is used to make sure that both the partner endpoint of the RTP connection and the path between the endpoints are operational after a period of inactivity. When this timer expires and no packet has arrived from the partner since it was last set, a packet with a status indicator will be sent and the SHORT_REQ timer is started. Should a status segment be received from the partner, the SHORT_REQ timer is stopped. Otherwise, when the SHORT_REQ timer expires the status request is retransmitted. After a predetermined number of retries and no response, an attempt will be made by the sender to find a new path for the connection. If the partner is not operational or there is no suitable path to the partner, the sender will eventually terminate the connection.

The purpose of the Alive timer is as follows:

- Keep limited resource links active, where limited resource links are automatically deactivated in HPR when no traffic flows over the link for a specified period of time. So when there is no session traffic, RTP sends *liveliness* messages at intervals set by the Alive timer.

- If the partner RTP endpoint or a link on the path fails, and the RTP endpoint is idle awaiting session traffic from the partner, then the RTP connection is "hung." The Alive timer triggers a liveliness message that is used to detect this condition. Such a detection triggers a path switch.

Nondisruptive Path Switch

The HPR path switch function is used to automatically route data around a failed link or node. This function only operates within an HPR subnetwork and is supported by all HPR network nodes and end nodes. When a failure occurs and an alternate path exists that satisfies the class of service for the failed RTP connection, a new RSCV is calculated and the RTP connection is switched; session traffic will be rerouted over the new path without disrupting the existing sessions.

Multilink Transmission Groups

A multilink transmission group (MLTG) consists of multiple DLC-level connections between two nodes made to appear to higher layers as a single connection. An MLTG is available for service as long as one or more of its constituent links are available.

Multilink transmission groups are supported in traditional subarea SNA networks and in APPN HPR networks, but not in base APPN.

Although superficially similar to multilink transmission groups in subarea networks, MLTGs in APPN HPR networks are significantly different in operation. This section describes HPR MLTGs.

HPR MLTG Requirements

Multilink transmission groups (MLTGs) have advantages over single-link TGs and parallel TGs in a number of cases:

Where the traffic demand can exceed existing TG capacity

Traffic demand can exceed existing TG capacity when a single session reaches the point at which it needs more bandwidth than the TG can provide. Aggregate available bandwidth can be raised simply by the addition of more links dynamically. If the demand subsequently falls, the extra bandwidth can be taken back by deletion of the extra links, saving network charges. Parallel TGs cannot help in this circumstance.

The need may also arise because of varying loads placed on a TG by a collection of sessions, rather than any single session. In this instance, adding parallel TGs *might* be an alternative solution, or not, depending on class-of-service and route selection implementations. But a single session could not use more capacity than the link offers that carries this session.

Where multiple lower-speed links are less expensive than a single higher-speed link

There are cases where multilink transmission groups prove less expensive than single-link TGs. In certain countries circuit capacities of 64 Kbps and 2 Mbps are available, but nothing in between. If you live in one of these countries and have to provide 100 Kbps of bandwidth, for example, you may find it costs less to put two 64 Kbps links into a multilink transmission group than to have a single 2 Mbps link.

Where individual links are unreliable
Although HPR provides a fast nondisruptive path switch capability, not even this will be necessary if your TGs never fail. If you are considering MLTGs to avoid TG failures, however, you must plan for the potential effects of temporarily reduced TG capacity. When one of several active links in an MLTG fails, effective capacity will be reduced even though the TG does not itself fail.

Where you have a subarea network including multilink transmission groups
If you have grown used to having the multilink transmission group facility in subarea networks you may feel more comfortable about migration to APPN HPR, knowing a similar facility is there.

Additional design objectives of the MLTG architecture include:

- The need to support mixed link types within MLTGs

  All supported SNA link types are also supported in HPR MLTGs.

- The need to support mixed link speeds within MLTGs
- The need to minimize system definition

HPR MLTG Overview

The critical parameter determining whether two links belong to one MLTG or to two parallel TGs is TG number (given of course that the links connect the same pair of nodes). If the links share the same TG number, then they belong to an MLTG; if they have different TG numbers, then they belong to parallel TGs. In this regard, subarea SNA and HPR do not differ.

One of the architectural problems with subarea multilink transmission groups was the need for resequencing of packets. Higher layers required DLC to guarantee delivery of packets, hop-by-hop, and to guarantee FIFO order. This dictated, among other things, that SNA subarea nodes had to act as *store-and-forward* switches, being unable to make forward routing decisions until entire packets had been safely received. It could easily happen that two packets, transmitted on different links within a multilink transmission group, would reach this point in reverse order of their initial order. The receiving node would have to buffer the second packet, pending the arrival of the first. This TG resequencing function could impose large processing overheads, especially where there were widely varying line speeds, propagation delays, or packet lengths, or where there were significant line error rates. In today's high-speed networks, resequencing delays en route would be unacceptable.

HPR eliminates the need for TG resequencing and for hop-by-hop error recovery by shifting these functions to RTP endpoints. When a VR-based transmission group (VR-TG) crossing the subarea network includes a subarea multilink transmission group, resequencing is not done for HPR network layer packets transported over that subarea MLTG.

In the HPR MLTG architecture, error recovery on individual links is optional, and TG resequencing en route is absent. Because FID2 packets have to be transmitted reliably and in sequence, HPR MLTGs do not support any FID2 traffic. HPR MLTGs must carry ANR network layer packets exclusively. This means, in turn, that RTP connections must be used for CP-CP sessions and route setup flows.

Both nodes connected by an HPR MLTG must hence support the control flows over RTP option.

As regards routing and ANR labels, MLTGs are treated the same as single-link TGs. See “Automatic Network Routing” on page 2-16. An MLTG is assigned one ANR label for each direction.

MLTGs and single-link TGs are also considered alike by TRS when it comes to the generalities of topology databases, TDUs, and route calculations. Differences show up when an MLTG's characteristics change *in flight*; for instance, when a new link is added. Such circumstances cannot arise in single-link TGs. When MLTG characteristics do change, topology database records are modified and TDUs generated.

Some functions are not supported in HPR MLTG:

- Limited resource
- Connection networks
- Nonactivation XID

Much of the HPR MLTG architecture revolves around the handling of TG number and other characteristics governed by XID3 exchanges during link activation. In particular, it deals with the exceptions that can occur when differently defined links are put together.

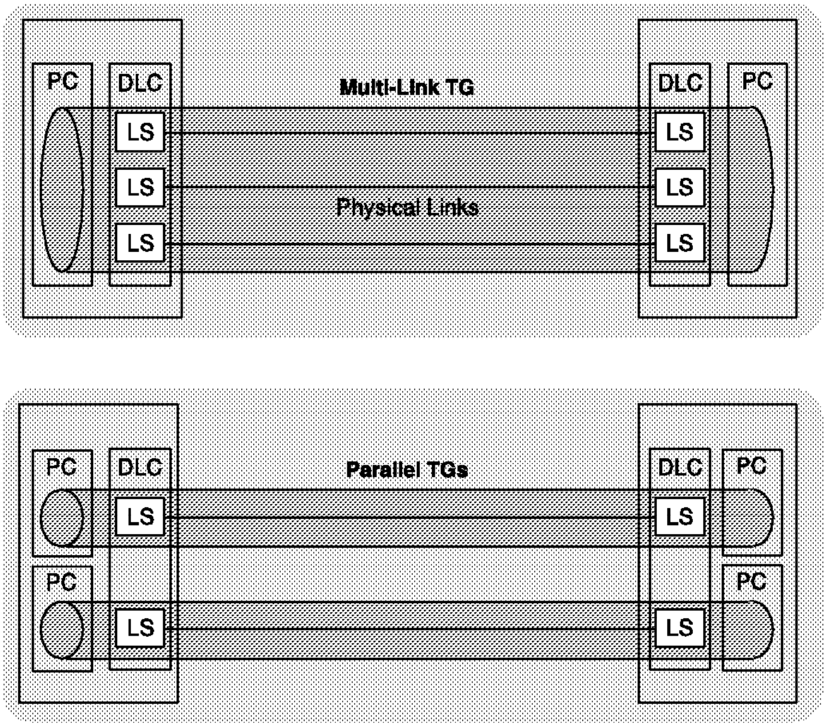

Figure 2-6. Multilink and Parallel TGs

Priority

HPR uses the same transmission priorities as the base APPN architecture (low, medium, high, and network). Transmission priority is associated with a class of service (COS). APPN has a set of architecturally defined COSs, each having a specified transmission priority. For example, the CPSVCMG COS, used by CP-CP sessions, has a transmission priority of network (the highest). COSs that are not architecturally defined may be used for LU-LU sessions; these COSs are associated with the priority of the session, either high, medium, or low.

APPN intermediate network nodes provide priority queues to allow higher-priority traffic to be routed before lower-priority traffic. The transmission priority function is also provided for packets flowing on RTP connections. Each RTP connection is associated with a COS and, therefore, has an assigned transmission priority.

Route Calculation

The APPN route calculation algorithm is used by HPR. HPR links are marked in the topology database, and the routes (paths) within HPR networks are specified with ANR labels instead of transmission group (TG) numbers and CP names. Nonetheless, the algorithm to compute these paths is unchanged.

The APPN architecture provides various means that could be used to make HPR links preferred to regular APPN links. Here are two examples:

- HPR links can be made to have better characteristics in terms of cost, delay, and so on, than APPN links even if the physical link characteristics are the same.
- Another possibility would be to use the user-defined fields in the COS tables and link characteristics to give HPR links smaller weights. The route calculation algorithm is then more likely to choose HPR links than APPN links.

However, as the routing decision is not a local optimization (within a node), but rather a global one (within the network), a small change in a link's characteristics may change the whole network's distribution of traffic. As a result, by artificially lowering the weight of an HPR link, the network could route all its traffic through that link causing a network collapse.

Therefore, when a node activates an HPR link, the link characteristics, which is broadcast as part of the topology function, should accurately characterize the link, and the node should not artificially modify the link characteristics.

To take full advantage of the HPR function:

- Adjacent APPN nodes should be upgraded. HPR-HPR-APPN routes are better than HPR-APPN-HPR routes.
- Links with the heaviest traffic, for example backbone links, should be added to the HPR network first.

In summary, this architecture allows seamless migration from an APPN network to an HPR network.

HPR-Only Route For Path Switch

When a time-out is detected while attempting to transmit data or Alive messages, RTP requests a new path from route selection services (RSS). If the new path between the two RTP endpoints contains one or more APPN links, that is, supporting only ISR function, the path switch will fail and the sessions traversing that RTP connection will be deactivated. This is not the case if there is a path between the two nodes that uses only HPR links. Because of the importance of session availability, a best-fit, HPR-only path is used, even if it does not match the required COS weight.

Timers

HPR uses timers to ensure that non-operational paths do not consume useless bandwidth to the detriment of other traffic. The timers are:

- Alive timer
- Short_req timer
- Path switch timer.

Each is described in more detail below.

Alive Timer

This user-defined timer ensures that both endpoints of an RTP connection and the path between them are still operational after a period of inactivity. When this timer expires and no packet has arrived from the partner since it was last set, an *“are_you_there”* packet is sent to the partner and a short_req timer is started. There are then two possible actions:

- If a response is received, the short_req timer is stopped.
- If the short_req timer expires, the *“are_you_there”* packet is resent.

If no response is received after a specified number of retries, a path switch is attempted.

Note: Both the Alive time and number of retries can be configured via CCM.

Short_req Timer

This timer is used for error recovery by RTP. Each RTP endpoint periodically “tags” a frame with a “request status” and waits for a response from the other endpoint. If no response is received with the time set in this timer, another frame is tagged and the wait starts again. This tag and wait cycle is repeated until a response is received or specified number of tags have been attempted.

The initial value of Short_req is 1 second (not configurable) and is computed regularly by RTP based on the round trip delay [old name: Round Trip Time (RTT)]:

```
Short_Req = (0.9 x Previous_Short_Req) + (1.8 x RTT).
```

Path Switch Timer

This timer is used to monitor the length of time that RTP should attempt a path switch for a connection before failing the patch switch.

Three path switch timers can be configured in the 3746 via CCM, one per transmission priority (high/medium/low). This improves the detection time of path switch failures for high-priority connections.

Migration

The following features of HPR ease migration from APPN:

- Interoperation with existing APPN nodes
- No configuration restrictions (*drop-in* migration)
- Use existing APPN Control Point protocols and algorithms
- Shared topology

Each is described in more detail below.

Interoperation with Existing APPN Nodes

Transforming APPN protocols into HPR and vice versa is done by the *APPN/HPR boundary function*. This function provides all the necessary transformations to allow APPN-level nodes and HPR-level nodes to interoperate seamlessly.

No Configuration Restrictions

New HPR nodes may be added and existing APPN nodes may be upgraded to HPR in any manner desired. There are no configuration restrictions whatsoever (*drop-in* migration). However, the benefits of HPR do not appear until contiguous clusters of HPR nodes, *HPR subnets*, are formed.

HPR Uses APPN Control Point Protocols and Algorithms

HPR uses the control point (CP) protocols of APPN (directory, topology, CP capabilities, and so on). CP-CP sessions are employed, just as in APPN, to transport these protocols. The APPN route selection algorithm, with a modification to calculate HPR-only paths for the path switch, is also used by HPR. Using existing APPN control flows significantly reduces the amount of code required for migration, that is, to implement the APPN/HPR boundary function.

Shared Topology

All nodes and links are reflected in every APPN and HPR network node's topology to facilitate migration and network management.

Migration Planning

The benefits appear when you have at least one HPR subnet operating, see “Route Calculation” on page 2-22. You could start by nominating for HPR operation either:

- The longest chain of HPR capability.
    1. List all HPR-capable stations and nodes.
    2. Starting from a VTAM host, work outwards to the furthest HPR-capable unit. That becomes your first HPR subnet, or chain of subnets.
    3. Using CCM, define all the required 3746 Nways Multiprotocol Controllers and stations as HPR-capable. (See Chapter 40, “CCM Worksheets for Controller Configuration Definitions” on page 40-1 for HPR options on APPN Parameters worksheets).
- The links with the heaviest traffic, either by current or predicted performance, then work outwards to adjacent APPN nodes.

Dependent LU Requester/Server

Figure 2-7 is used as reference in this section.

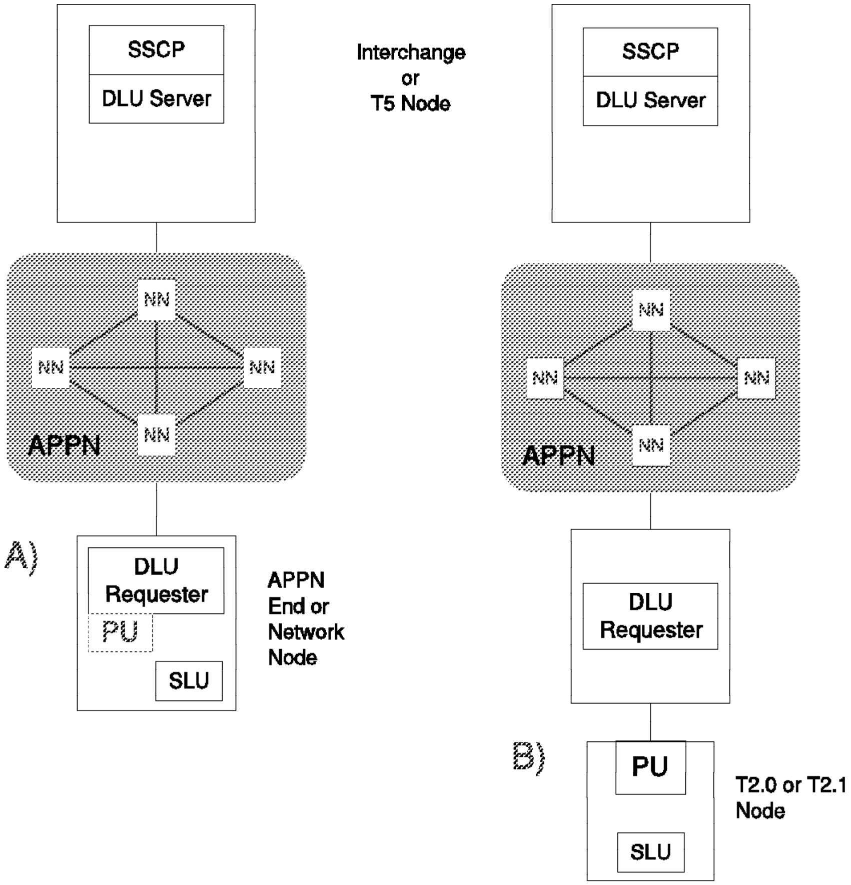

Figure 2-7. Dependent LU Requester/Server.

A) The dependent LU requester in the same node as the dependent LU(s).
B) The dependent LU requester in an APPN end node or network node directly connected to the PU T2.0 or APPN or LEN node containing the dependent LU(s).

The dependent LU requester (DLUR) and dependent LU server (DLUS) functions allow SSCP-PU and SSCP-LU data to flow through an APPN network. The session establishment data for dependent LUs flows on top of two APPN LU6.2 sessions. (For each one of DLUR and DLUS, there is one *contention winner* session and one *contention loser* session.) From the DLUR function to the dependent PU, the flow is the same of the NCP or VTAM boundary function to a PU (ACTPU, ACTLU, BIND, etc.).

The dependent LU server function (option set 1066) is a product feature of an interchange node or a T5 network node supporting session services extensions. This function provides server support for dependent LU requester clients in which SSCP-PU and SSCP-LU flow to a PU T2.0 or APPN or LEN node externally attached to the requester, or a PU T2.0 or APPN or LEN node image within the requester, are encapsulated within LU 6.2 sessions.

The dependent LU requester function (option set 1067) is an enhancement to an APPN end node or network node. This function is the client side of the dependent LU server function in which SSCP-PU and SSCP-LU flows to a PU T2.0 or APPN or LEN node attached to the requester are encapsulated within LU 6.2 sessions as mentioned previously.

The requester function provides a remote boundary function for dependent LUs. This option set relieves the restriction that PU T2.0 nodes be directly attached (or bridged, or data link switched, or frame relayed) to the VTAM or NCP boundary function. The dependent LU requester function may reside in the same node as the secondary LU or be provided by a node adjacent to and upstream from the secondary LU (see Figure 2-7 on page 2-25).

For a more detailed discussion of the DLUR function see *Inside APPN with HPR: The Essential Guide to New SNA* SG24-3669, and *Subarea to APPN Network Migration Experiences*, SG24-4656.

Note: : Bisynchronous (BSC) 3270 sessions are not supported over APPN links. If BSC sessions are required, you must maintain subarea paths to all potential session partners.

3746-9x0 APPN/HPR Network Node Implementation

This section describes details of the APPN features implemented by the 3746-9x0 APPN/HPR Network Node.

Terminology and Implementation Specifics

The 3746 NN is composed of a 3746 frame connected via a dedicated token-ring LAN to its service processor (SP) and network node processor (NNP). The token-ring LAN used for communication between the NNP and the service processor is referred to as the service LAN. The SP and NNP each contains a token-ring adapter attaching them to the same service LAN. MOSS-E traffic travels over the SP adapter onto the service LAN. Likewise APPN traffic session establishment traffic travels over the NNP adapter onto the service ring. Note that APPN user traffic does *not* flow on the service ring.

Figure 2-8 on page 2-27 depicts how the APPN functions are split up between the network node processor (NNP) and the adapters within the 3746 frame.

Note: By adapter we mean the CLP, TRP2, ESCP2, or CBSP2 processor and the associated line interface (LIC), token-ring interface (TIC3), Ethernet interface(TIC3 bridged), or ESCON (ESCC) couplers.

Node Operator Facility (NOF) functions (for example, port and link activation), APPN topology and routing services, and session establishment tasks are executed in the NNP, while intermediate session routing (user traffic) is done within the 3746 adapter.

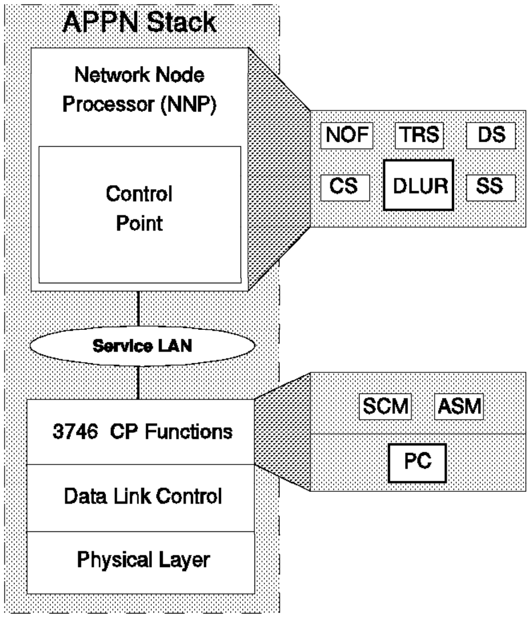

Figure 2-8. 3746 NN Structure. A full APPN stack is composed of functions performed on the network node processor (NNP) and within the 3746-9x0.

The APPN functions that run on the network node processor are:

- NOF - Node Operator Facility
- TRS - Topology and Routing Services
- DS - Directory Services
- CS - Configuration Services
- SS - Session Services
- DLUR - Dependent LU Requester

APPN functions performed within the 3746 adapters are:

- DLC - Data Link Control
- PC - Path Control
- ASM - Address Space Manager
- SCM - Session Connector

The following section details how these components interoperate during session establishment and routing for APPN (independent LU 6.2) sessions.

Session Establishment and Routing

During APPN session establishment, CP functions on the NNP participate in locating session partners and are responsible for APPN route calculation. Figure 2-9 on page 2-29 depicts how CP-CP session data flows between the NNP and the control points of adjacent nodes. Irrespective of the coupler the node is connected to, CP-CP data will always traverse:

- The adapter (coupler and processor) the APPN node is attached to
- The connectivity switch (CS)
- CBSP2
- Token-ring port 2080

- Service LAN

Note: In only two cases CP-CP session data will not traverse the connectivity switch:

1. When APPN nodes connect via token-ring port 2080

   However, with the introduction of the APPN NN functions the attachment of user equipment via the service LAN is no longer supported.

2. When using an internal APPN link between the 3746-900 NN and any of the CCUs of the attached 3745 Model A

The BIND, which is the first SNA request unit flowing on the newly calculated route between two session partners, will trigger the address space manager (ASM) function running on the 3746 processors to assign LFSIDs. In addition, a session connector (SC) will be generated to enable intermediate session routing on the 3746 NN. The SC can be an intra-processor (within the same 3746 processor), or an inter-processor (between two different processors connected via the 3746 connectivity switch (CS)). See Figure 2-10.

Figure 2-11 on page 2-30 illustrates the data flows during and after session establishment. End node A (EN A) is token-ring-connected to 3746 NN, while end node C (EN C) is SDLC-connected. In both cases the 3746 NN (NN B) is providing the network node server function, having CP-CP sessions with both end nodes.

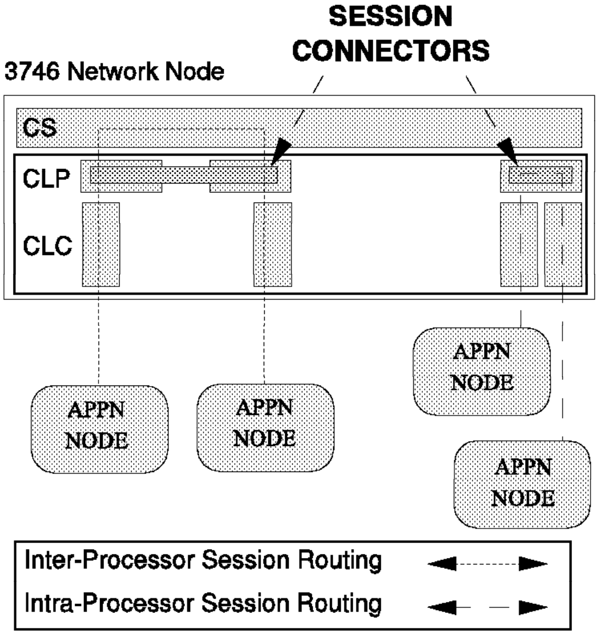

Figure 2-10. Intermediate Session Routing

To locate the session partner and calculate the best session path, APPN functions within the NNP are invoked. Initiated by the BIND, CP functions available on the 3746 processors will assign local form session identifiers (LFSIDs) for this session and generate a session connector (SC). Note that for this session an

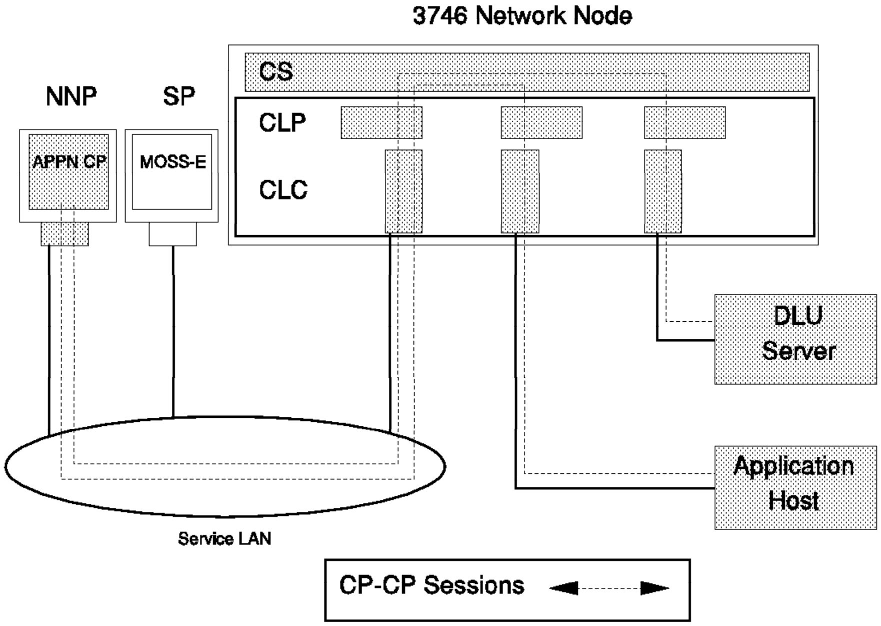

```
Legend:
CS   =  Connectivity Switch
CLP  =  Communications Line Processor
CLC  =  Communications Line Coupler
NNP  =  Network Node Processor
SP   =  Service Processor
```

Figure 2-9. CP-CP Sessions

inter-processor SC applies. If both EN A and EN C were connected to couplers controlled by the same processor, an intra-processor SC would result.

When the 3746 NN is performing intermediate session routing, the session connector manager (SCM) performs the LFSID swapping required to forward session data.

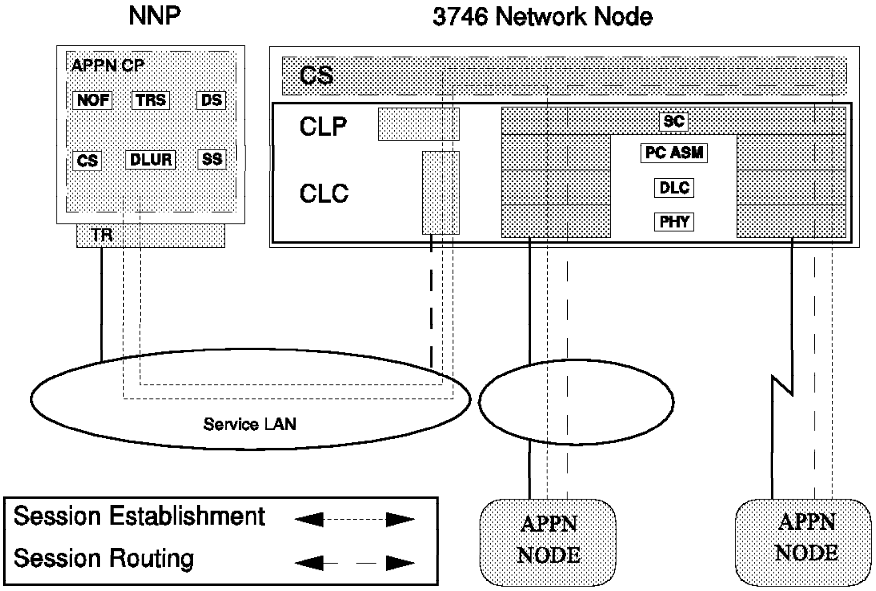

Figure 2-11. 3746 NN Intermediate Session Routing. All shaded components are involved in session setup. The dark-shaded components are also involved in intermediate session routing.

3746 Network Node Processor (NNP) Backup

In order to provide the network node function in the 3746-9X0, the network node processor feature is used. This provides the network node processor (NNP) hardware resources and the licensed internal code required to support the APPN network node functions.

The NN processor feature includes the APPN CP and the APPN NN configuration control and management software (CCM) along with the hardware and token-ring interface. A keyboard and a display are not required on the network node processor; access is provided from facilities available on the service processor (SP).

Functions running on the control point can be accessed from the service processor. To allow configuration and management of the APPN NN functions the configuration control and management (CCM) tool which runs on the SP is used.

To provide additional resilience a second network node processor can be installed. The backup control point also attaches to the service LAN and can take over the functions from the primary control point. This process is controlled from the service processor; the CP backup can be done either manually or automatically. In case of a malfunctioning primary control point no new sessions can be established. To allow new session establishment, either the primary CP needs to be restarted, or the backup CP must take over. Once the original failing NNP again becomes available, it will be regarded as the backup network node processor.

Maximum Connectivity of the 3746-9x0 APPN/HPR Network Node

The number of PUs, frame-relay DLCIs, and sessions available on the 3746-9x0 are given in the following tables.

Adapter Connectivity

Table 2-1 on page 2-32 gives the maximum number of PUs, frame-relay DLCIs, and APPN or dependent LU sessions that the various 3746-9x0 adapters can handle, assuming that the IP routing software is not loaded in these adapters.

For adapters providing IP routing, the maximum number of PUs and sessions controlled by the 3746 network node may be lower, due to the storage used by the IP routing software.

Depending on the storage available in the processors, the actual maximum number of 3746-controlled PUs and sessions may be different. The maximum number of ESCON logical link stations (16) and, in case of 3746 Model 900, the maximum number of NCP-controlled PUs (see column **NCP** in Table 2-1 on page 2-32) and total number of PUs (see column **Total** in Table 2-1 on page 2-32) are absolute maximum numbers which cannot be exceeded.

Table 2-1. Adapter Level Connectivity

Adapter	3746 Model 900				3746 Model 950	
	PUs[1]			Sessions[2] 3746 NN	PUs[1]	Sessions[2]
	3746 NN	NCP	Total			
ESCP	0	16	16	0	-	-
ESCP2	16[9]	16	16[9]	4900	16[9]	4900
ESCP3	16[9]	16	16[9]	14 000	16[9]	14 000
TRP	0	2000	2000	0	-	-
TRP2[10]	1000	2000	2000	4500	1000	4500
TRP3[10]	2000	2000	2000	13 500	2000	13 500
For CCU B[3]**:**						
TRP	0	500	500	0	-	-
TRP2[10]	1000	2000	2000	4000	-	-
TRP3[10]	2000	2000	2000	13 000	-	-
CBSP	-	500	500	-	-	-
CBSP2/CBSP3[4]**:**	-	500	500	-	-	-
CBSP2/CBSP3[5]**:**	0	0	0	0	0	0
CLP with:						
3000 DLCIs[4]	-	4000[6]	4000[6]	-	-	-
500 DLCIs[10]	1000[8]	2000[7]	2000[7]	3500	1000[8]	3500
CLP with:						
3000 DLCIs[4]	-	4000[6]	4000[6]	-	-	-
2000 DLCIs[10]	3000[12]	3000[11]	3000[11]	12 500	3000[12]	12 500

Legend:

CBSP2	Controller bus and service processor (type 2)
CBSP3	Controller bus and service processor (type 3)
CCU	Central control unit
CLP	Communication line processor
CLP3	Communication line processor (type 3)
DLCI	Data link connection identifier
ESCP2	ESCON processor (type 2)
ESCP3	ESCON processor (type 3)
LU	Logical unit
NN	Network node
PU	Physical unit
TRP2	Token-ring processor (type 2)
TRP3	Token-ring processor (type 3)

Notes related to table A-4:

1. These are adjacent PUs (or ESCON logical link stations), such as end nodes, network nodes, LEN nodes, dependent PUs, gateway downstream PUs, and X.25 virtual circuits. For the 3746-900, the total of NCP-controlled and 3746-controlled stations can not exceed the total that is in the Total column.
2. These are all the LU sessions (independent and dependent LUs) routed by the 3746 adapter, including LU-LU sessions involving non-adjacent nodes. HPR/ANR sessions between HPR/RTP nodes, that do not begin or end in the 3746, are not part of these numbers and can be any number. For the 3746-900, these numbers do not include the sessions routed by NCP. The quantity of NCP-routed sessions depends on the 3745 storage capacity.

   These figures apply only to processors that have a few PUs or ESCON stations.
3. This is the TRP, TRP2, or TRP3 used to connect the 3745 CCU-B to the 3746-900.
4. For a 3746-900, if neither 3746 APPN/HPR nor 3746 IP routing is used in any CLP/CLP3.
5. For any 3746-950, and any 3746-900 using the 3746 APPN/HPR network node or IP Routing support.
6. Up to 1000 SDLC PUs and any mix of up to 3000 frame-relay PUs, ISDN PUs, and X.25 virtual circuits (one PVC or SVC per PU).
7. Up to 1000 SDLC PUs and any mix of up to 1000 Frame-relay PUs, ISDN PUs, and X.25 virtual circuits (one PU per PVC or SVC).
8. Up to 1000 PUs over SDLC, frame-relay, and X.25 lines.
9. This includes any logical stations (TCP/IP) used by the 3746 IP Router.
10. Not all the maximum connection capabilities are possible simultaneously: for a given processor,the maximum number of resources in a category (3746-controlled PU s, NCP-controlled PUs, 3746-controlled sessions, SDLC links) depends on the number of active resources in other categories, on the presence of the IP Routing feature, and, in case of a CLP, on the mix of lines (SDLC, frame-relay, X.25).

    For example: TRP2s (without IP Routing feature) support simultaneously a total of 500 APPN/HPR PUs and 3000 data sessions, or 1000 dependent PUs and 1500 data sessions.
11. Up to 1000 SDLC PUs and any mix of up to 2000 frame-relay PUs, ISDN PUs, and X.25 virtual circuits (one PU per PVC or SVC).
12. Up to 1000 SDLC PUs and any mix of up to 2000 frame-relay PUs and X.25 virtual circuits (one PU per PVC or SVC).

Network Node Connectivity

Table 2-2 on page 2-34 gives the total number of PUs, APPN and dependent LU sessions, and lines that a 3746 network node can handle (no matter what (type 2) adapter configuration is used).

Table 2-2. Network Node-Level Connectivity

Connectivity		**Comments**
Type	**Number**	
PU	5000	End nodes, LEN nodes, network nodes, Dependent PUs.
Sessions	15000	All the LU-LU sessions using 3746 DLUR and APPN routing, including sessions involving non-adjacent nodes. HPR/ANR sessions between HPR nodes connected to the 3746 are in addition to this number of sessions and can be in any quantity.
CLP or Serial Lines	120	Frame-relay, SDLC, X.25, and PPP.

Note: For the 3746 Model 900, the resources beyond these network node quantities are controlled by NCP(s) either as part of a PU type 4 (SNA) node or part of an APPN composite network node (CNN).

DLC Support of Error Recovery

Network layer packets (NLPs) may be sent over an HPR link using link-level error recovery procedures or not. (Note that FID2 PIUs always use the link-level error recovery procedures, just as in APPN.) Whether or not error recovery is used on NLPs is determined during the initial exchange between adjacent nodes:

- Using link-level error recovery

  NLPs are transmitted in the same manner as in APPN where the LLC elements of procedure provide full error recovery. For HPR, this is recommended for links with high error rates.

- Not using link-level error recovery

  NLPs are transmitted in a manner such that the LLC will not perform any error recovery for it. This mode of operation is recommended for low error rate reliable links. The method for bypassing the link-level error recovery mechanisms varies depending on the link type (see descriptions of each link type below).

The following sections describe the link types currently supported by the 3746.

Token Ring: Both *link-level error recovery* and *no error recovery* are supported in the 3746, but you are recommended to use *no error recovery*. HPR provides the capability to bypass error recovery for NLPs on LANs using 802.2. From an architecture point of view, a separate service access point (SAP), different from the SAP currently used to transmit APPN traffic, could be used to transmit NLPs with no link level error recovery. NLPs requiring link-level error recovery use the same SAP as existing APPN traffic. However, in the 3746, HPR SAP (error recovery or no error recovery) is identical to the APPN SAP for this port. All traffic, both error recovery and non-error recovery ones, travel over the same physical path. And, if that physical path is switched (due to link-level switching) to a new one, the new path is then used for all traffic.

Frame Relay: Both *link-level error recovery* and *no error recovery* are supported in the 3746, but you are recommended to use *no error recovery* especially on good quality lines. In the 3746, HPR SAP (error recovery or no error recovery) is identical to the APPN SAP for this port).

SDLC: There are no changes required in order to run HPR over SDLC. SDLC is run exactly as it is run without HPR, that is, *only link-level error recovery* is supported in the 3746.

X.25: The X.25 protocol always provides error recovery between a DTE and a DCE using the LAPB procedure. HRP runs over X.25 using the same QLLC protocol as the regular APPN traffic.

Neither the X.25 standard protocol nor the IBM QLLC protocol have been modified to transport HPR traffic.

ESCON: *Only link-level error recovery* is supported in the 3746.

Dependent LU Requester Implementation

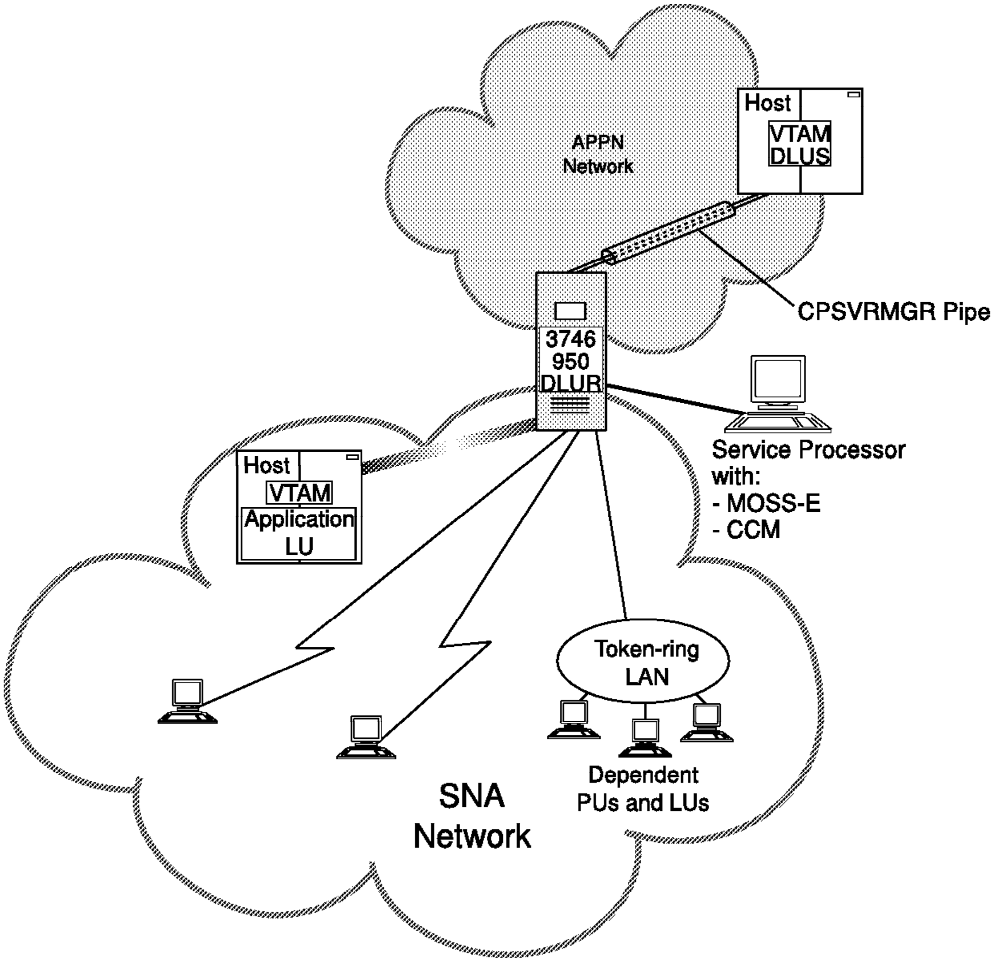

Figure 2-12. 3746 Nways Multiprotocol Controller DLUR Connectivity

The Dependent LU Requester (DLUR) function facilitates conversion from a subarea environment to an APPN environment, allowing you to maintain central management of remote dependent LUs while benefiting from APPN throughout a network.

Two LU 6.2 sessions (one inbound, one outbound) are established between a DLUR and a Dependent LU Server (DLUS). These LU 6.2 sessions are collectively known as the CPSVRMGR pipe. SSCP-PU and SSCP-LU session flows use the CPSVRMGR pipe. An SSCP-PU session is established between a VTAM network node and the PU that owns the dependent LU, and an SSCP-LU session is established between VTAM and the dependent LU. Session initiation flows for the dependent LU are sent over the SSCP-LU session, and VTAM can use a subarea or APPN path to initiate a session between the dependent LU and the primary LU (the application). BIND and session data are then routed directly between the primary LU and the dependent LU. To implement a DLUS/DLUR environment, consider the following:

- For the DLUR to initiate activation, via a call-in (to VTAM), no system definitions are required for existing token-ring LANs. Refer to page “For Token-Ring Lines” on page 16-13.

For SDLC lines, new definitions are needed in VTAM. Even if the SDLC line is leased, it must be referenced via a switched major node definition. Refer to Chapter 19, “3746 SDLC Support” on page 19-1.

The dynamic switched definition facility can be used to define the PU. For information about using the ISTEXCCS exit routine for dynamic definitions, refer to "Dynamically Defining Switched Resources" and related topics in the VTAM *Network Implementation Guide* that corresponds to your level of VTAM.

- For VTAM to initiate activation, via a call-out (from VTAM), of a PU, define the dependent LU requester by including the DLURNAME and DLCADDR operands on the PATH definition statement in the switched major node. DLURNAME specifies the CP name of the DLUR that owns the PU, and DLCADDR includes data link control (DLC) information used by the DLUR to locate the PU. Also include the MAXDLUR operand on the VBUILD definition statement to indicate the maximum number of unique DLURs defined for this switched major node.

  Note: If the DLURNAME is not fully qualified, the NETID of the DLUS is used.

- A DLUR may be served by multiple DLUS VTAMs simultaneously. It is possible for each of up to five downstream PUs on a DLUR to use a different DLUS, but this is not practical.

  Normally, a 3746 DLUR has a SSCP VTAM as its generic primary DLUS (and possibly, another VTAM as its generic back-up DLUS). One or more downstream PUs can grouped to use a different SSCP VTAM as its generic primary DLUS (with a different generic backup DLUS).

- Multiple DLURs may support the same PU, but only one at a time. The SSCP-PU session is through only one DLUR at any given time.

- Redial occurs as follows for DLUS supported PUs:

  – When a DLUS initiates the activation of a PU but receives a negative response, VTAM attempts to redial over each valid path statement for the PU until successful or until all valid paths have been tried. Redial does not occur if the negative response indicates that the PU is already active, or that the fully qualified procedure correlation identification (PCID) is not unique. If the fully qualified PCID is not unique, the DLUS attempts to redial the PU over the same path with a newly generated PCID.

  – Redial does not occur if PU activation was initiated by the DLUR.

  – When a protocol violation, topology database update (TDU) error, or CPSVRMGR session outage signal is received for a particular DLUR, VTAM attempts to redial every active or pending-active PU served by that DLUR for which a valid PATH statement is found. If the PU is already active, then VTAM performs the following:

    - If the PU was defined with ANS=CONT and the DLUR supports this function, giveback processing is performed prior to attempting redial.

    - If the PU was defined with ANS=STOP or the DLUR does not support ANS=CONT, the PU is deactivated prior to attempting redial.

- When a CP-CP session between a DLUS and DLUR fails, the CPSVRMGR session is inactivated, enabling reactivation of this CP-CP session.

- Information regarding DLUR-attached resources can be processed by the configuration services XID exit. This exit routine can be coded so that VTAM processes or denies requests for contact from known switched devices or so that VTAM processes or denies requests for PU activation from a DLUR. For details, refer to the *VTAM Customization*, LY43-0063. (Available to IBM licensed customers only).

VTAM DLCADDR Keyword in PATH Statement

The value of the DLCADDR keyword in the VTAM PATH definition statement is used when the DLUS attempts to establish a connection to a PU attached through a DLUR. As there is no direct connection to the DSPU, VTAM sends the information to the DLUR specified (DLURNAME), which then uses the information to establish a connection to the DSPU. You must code a DLCADDR keyword for each element of the information needed used to connect to a DSPU.

The format of the DLCADDR keyword is:

DLCADDR=(element,data_format,element_value)

- The first suboperand indicates the element you are defining. A value between 1 and 96 can be specified.
- The second suboperand indicates the format of the third suboperand value, when transmitted to DLUR:

  ```
  BCD:= Binary coded decimal
  C:  = Any EBCDIC character
  D:  = Decimal
  X:  = Hexadecimal
  I:  = Any ASCII character.
  ```

  Note: To be able to use the ASCII format, the following PTFs are required for:

  VTAM V4 R2, PTF UW28497
  VTAM V4 R3, PTF UW28498.

- The third suboperand defines the value of the element.

An example of DLCADDR encoding is given below:

```
CP900DPT PATH  PID=1,                                             X
               DLURNAME=NETNODE1,            CPname of DLUR       X
               DLCADDR=(1,C,TR),             Type of attachment   X
               DLCADDR=(2,X,504F525432313434), PORT2144           X
               DLCADDR=(3,D,04),             DSAP of downstream CM/2 X
               DLCADDR=(4,X,400052005160),   MAC  of downstream CM/2 X
               USE=YES                 INITIALLY ACTIVE
```

The CP name of the 3746 DLUR is *NETNODE1*, the DSPU is attached via the token-ring on port 2144 of the 3746. The DSAP and MAC address of the DSPU are specified in parameters 3 and 4.

Migration of SNI Connections to 3746 Models 900 and 950

SNA Network Interconnection (SNI) is a function of NCP which allows the interconnection of two SNA/subarea networks. Peripheral Border Node (PBN) and Extended Border Node (EBN) are functions that allow the interconnection of two APPN/HPR networks. They are provided today by VTAM and NCP using 3745/3746-900 Composite Network Nodes (CNNs):

- Peripheral Border Node (PBN) allows access to adjacent APPN/HPR networks
- Extended Border Node (EBN) allows access to non-adjacent APPN/HPR networks.

You cannot connect an SNA/subarea SNI node to an APPN/HPR node; once you migrate from a 3745/3746-900 to a 3746-900 network node or a 3746-950, the connections with the other nodes must use APPN/HPR flows.

Therefore, if your “partners” do not migrate to APPN but keep using SNI, you must keep a 3745/NCP that will handle your SNI traffic, while the coupled 3746-900 Network Node handles your APPN traffic. As shown in Figure 2-13, the SNI traffic may flow over 3746 adapters to the NCP.

One exception to the above statement is when an SNI network is connected to a 3745 which is part of a CNN. Figure 2-20 on page 2-47 shows such a configuration. The 3745/NCP handles the SNI flows and routes these flows via an internal APPN link to the 3746-900 network node (Figure 2-16 on page 2-43 shows an example of how an internal APPN link works). The TRP2 which handles the APPN link also contains the HPR/RTP function; RTP converts APPN (ISR) flows from the SNI connection to HPR flows, which can then be routed through the APPN network as any other HPR flows.

Use the following procedure to migrate SNI nodes to APPN nodes:

Step 1. Add APPN/HPR Network Node support to the 3746-900 (NN).

This does not affect the SNI connections on the 3746-900; they stay under control of the NCP running in the 3745 CCU. See Figure 2-13.

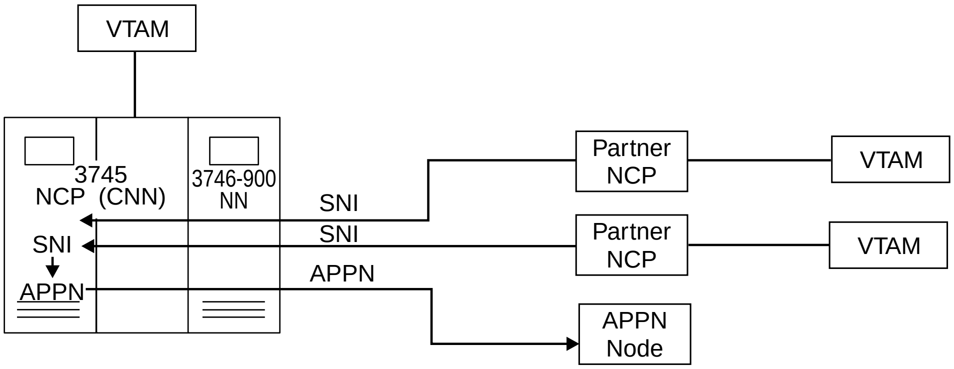

Figure 2-13. Example Network: Adding APPN Before SNI-to-APPN Migration

If there is a need to route traffic between an SNI port (or any SNA/subarea port of the 3745/3746-900) and an APPN port controlled by the 3746-900 network node, then the NCP must be defined as part of a Composite Network Node (CNN). NCP does the SNI-to-APPN and APPN-to-SNI conversions for the communication with the 3746-900 network node. This communication uses the coupling between the 3745 and the 3746-900, refer to “Internal APPN Connection between a 3745 and a 3746-900” on page 2-41.

Step 2. Progressively upgrade each partner network to support APPN connectivity to the 3746-900.

The minimum requirement is that the partner NCP becomes part of a Composite Network Node (VTAM plus one or more NCPs). This supports the APPN border node connection to the 3746-900 network node.

In the partner NCP, the connection is border node and in the 3746-900, the line will be defined in the 3746-900 network node instead of NCP. See Figure 2-14.

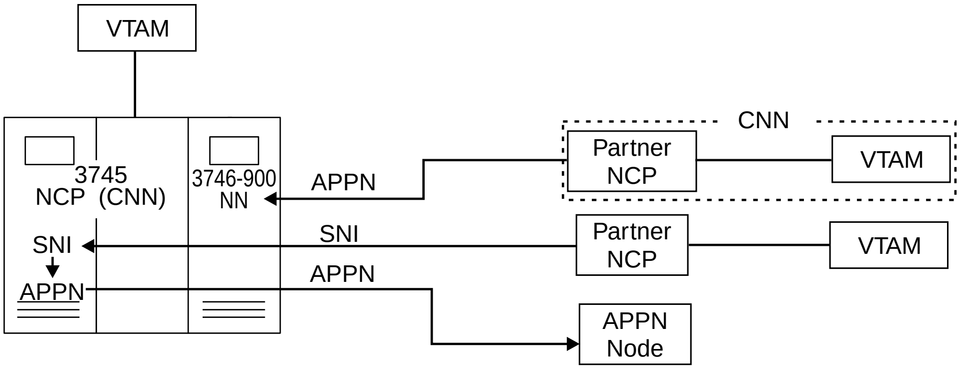

Figure 2-14. Example Network: SNI-to-APPN Migration

During this migration, the “partner” link stays physically attached to the same port of the 3746-900. The Controller Configuration and Management (CCM) is used to define this link in the 3746-900 network node.

Step 3. Upgrade the 3746-900 to a 3746-950.

This can be accomplished when there are no more SNI links or other need for NCP. See Figure 2-15 on page 2-41.

Important Note:

Until Session Services Extension (SSE) is supported by the 3746 NN, this configuration supports independent sessions across the network boundary between the 3746NN and a CNN providing the Extended Border Node (EBN) function. Dependent sessions across this boundary are not supported. The support of dependent sessions requires the 3746 NN directly adjacent to the CNN to support SSE. Currently, this is supported by the 3746 and VTAM (which has SSE support) acting together as a CNN.

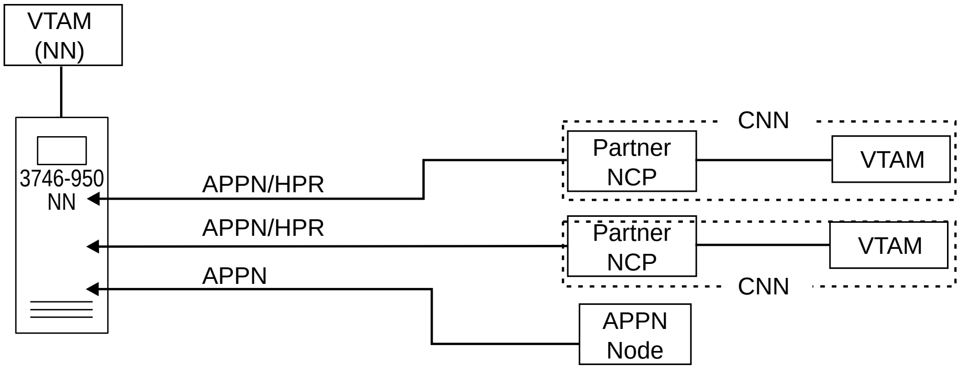

Figure 2-15. Example Network: Upgrading 3746-900 to a 3746-950

During this step, the 3745/NCP is removed from the controller configuration.

Internal APPN Connection between a 3745 and a 3746-900

On a 3745/3746-900, the controller bus coupler (CBC) provides an internal connection (an APPN transmission group) between the 3746-900 network node and the 3745 (NCP) operated as a composite network node (CNN). This internal APPN link is set up using a 3746-900 token-ring port. Any token-ring port, except the TIC3 port of the CBSP2 (port address 2080), can be used (see Figure 2-16 on page 2-43). The TRP2 for this TIC3 performs the APPN routing between the CBC link and the other ports of the 3746 Nways Multiprotocol Controller.

Although data transmitted between the 3746 Nways Multiprotocol Controller and the CNN does not actually pass through the token-ring port, a TIC3 must be present. To enable the activation of the token-ring port, a LAN must be connected to this TIC3. If this TIC3 port is used by the stations connected to the LAN, all the LAN traffic over this port must use a single point of control. For example, all the LAN stations that use this TIC3 as a destination address must be defined with the same destination SAP (for either NCP or the 3746 Network Node). However, this may not be practical as the TRP2 may already be loaded to its capacity by the internal APPN traffic.

Example of an Internal APPN Link

Figure 2-16 on page 2-43 shows an example of a 3745/3746-900 configuration using the internal APPN communication between the 3746 Nways Multiprotocol Controller and the CNN. The components are:

- User workstation (PU/node) with a dependent or independent logical unit (LU) communicating with an application running in the host.
- CNN consisting of a host with VTAM, and a 3745 with a single NCP (one CCU). NCP communicates with VTAM via a parallel channel.
- 3746 Nways Multiprotocol Controller.

The figure shows how the data flows between the user and the application.

3746-900 Configuration Requirements

The token-ring hardware required for the internal APPN link to the 3745 depends on the configuration of the 3745 and its mode of operation:

- 3745 with a single CCU (Models 17A, 21A, and 31A)

  One TIC3 on a TRP2 must be available.

- 3745 with two CCUs (Models 41A, and 61A):

  – If only **one** NCP communicates with the 3746 Nways Multiprotocol Controller at a time (for example, in the “3745 twin-standby mode”), one TIC3 on a TRP2 must be available, for example, on of the TRP2 connecting CCU B.

  – If **two** NCPs communicate at the same time with the 3746 Nways Multiprotocol Controller, two TIC3s must be available. For load distribution and backup reasons, they should preferably be on two different TRP2s. For example, the internal APPN link to CCU B could use a TIC3 on the TRP2 connecting CCU B and the APPN internal link CCU A could use a TIC3 on another TRP2.

Definitions Required for the Internal APPN Link

From a definition perspective, the internal link appears as a token-ring connection with the same token-ring address defined in the 3746 Nways Multiprotocol Controller and the CNN:

- **Important**

  The CNN SAP value must be different from the NCP SAP, which is X'04' by default.

- **3746 Nways Multiprotocol Controller (CCM Definitions)**

  One token-ring port definition is required for each CCU/3746 Nways Multiprotocol Controller link. Figure 2-16 on page 2-43 is a single link example.

  To allow simultaneous communication between the 3746 Nways Multiprotocol Controller and both CCU A and CCU B, two port definitions are required. Each port definition must use a different TIC3 address.

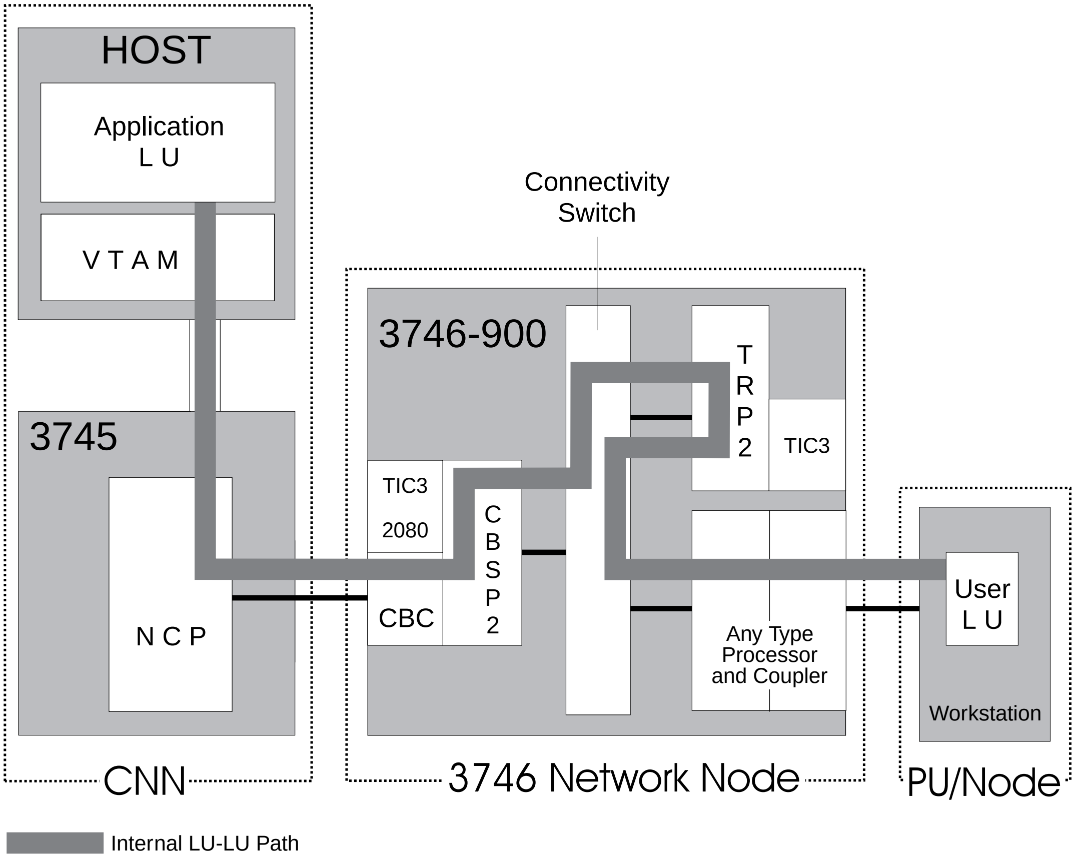

Figure 2-16. Internal APPN Link (3746-900)

The link station definition depends on who controls the TIC3:

1. If the TIC3 port connects LAN stations controlled *only by the 3746 Nways Multiprotocol Controller and these stations do not need to be defined in the 3746 Nways Multiprotocol Controller*, then, using the Controller Configuration and Management (CCM) application:

   – Define the TIC3 port with the option “Accept any incoming calls” set to YES.

   – Do not make a link station definition for this TIC3 port.

2. If the TIC3 port connects LAN users controlled *only by NCP in the CNN*, then, using CCM:

   – Define the TIC3 port with the option “Accept any incoming calls” set to NO.

   – For this TIC3 port define NCP as a link station:

     - The locally administered address (LAA) of the station must have the same value as the port LAA.

     - The destination service access point of the station must be set X‘04’ (the default value for the NCP service access point).

3. If the TIC3 port connects LAN users controlled *only by the 3746 Nways Multiprotocol Controller and defined in the 3746 Nways Multiprotocol Controller*, then, using CCM:

    - Define the TIC3 port with the option “Accept any incoming calls” set to NO.
    - For this TIC3 port define NCP as a link station:
        - The locally administered address (LAA) of the station must have the same value as the port LAA.
        - The destination service access point of the station must be set X‘04’ (the default value for the NCP service access point).

- **CNN (VTAM and NCP Definitions)**
    - Token-ring (port) definitions in NCP.

      The token-ring port address is the address of the TIC3 defined in the 3746 Nways Multiprotocol Controller (using the CCM). The locally administered address must be the same in the NCP keyword LOCADD as in CCM.

      When using a dual CCU Model 3745, this port can be defined in both NCPs, but only one NCP is able to activate the port at a time. This means that the internal APPN link can be active via either CCU A or CCU B, but not both at the same time.

      To allow simultaneous communication between the 3746 Nways Multiprotocol Controller and both CCU A and CCU B, two port definitions are required. Each port definition must use a different TIC3 address.
    - Link station definitions in VTAM.

      To represent the 3746-900 network node, a link station definition is required in VTAM. The local service access points value for the 3746 Nways Multiprotocol Controller must be used (the default is X‘08’).

      VTAM dynamics can be used to generate the appropriate switched major node definitions.

3746 Nways Multiprotocol Controllers in APPN/HPR Networks: Examples

This section contains three examples of 3746 Nways Multiprotocol Controllers in APPN/HPR networks as:

- ANR node, see Figure 2-17
- APPN/RTP Node with Boundary function, see Figure 2-18 on page 2-46
- DLUR/RTP Node with Boundary function, see Figure 2-19 on page 2-46.

Each example shows the active path, end-to-end, and the network layers involved in the 3746, assuming that the traffic routing involves two adapters (ADP#1 and ADP#2). The legend for the examples is:

ADP	=	adapter
ANR	=	automatic network routing
APPN	=	APPN
DLC	=	data link control
DLUS	=	dependent LU server
HPR	=	High Performance Routing
ISR	=	intermediate session routing
RTP	=	rapid transport protocol

Example 1: ANR Node

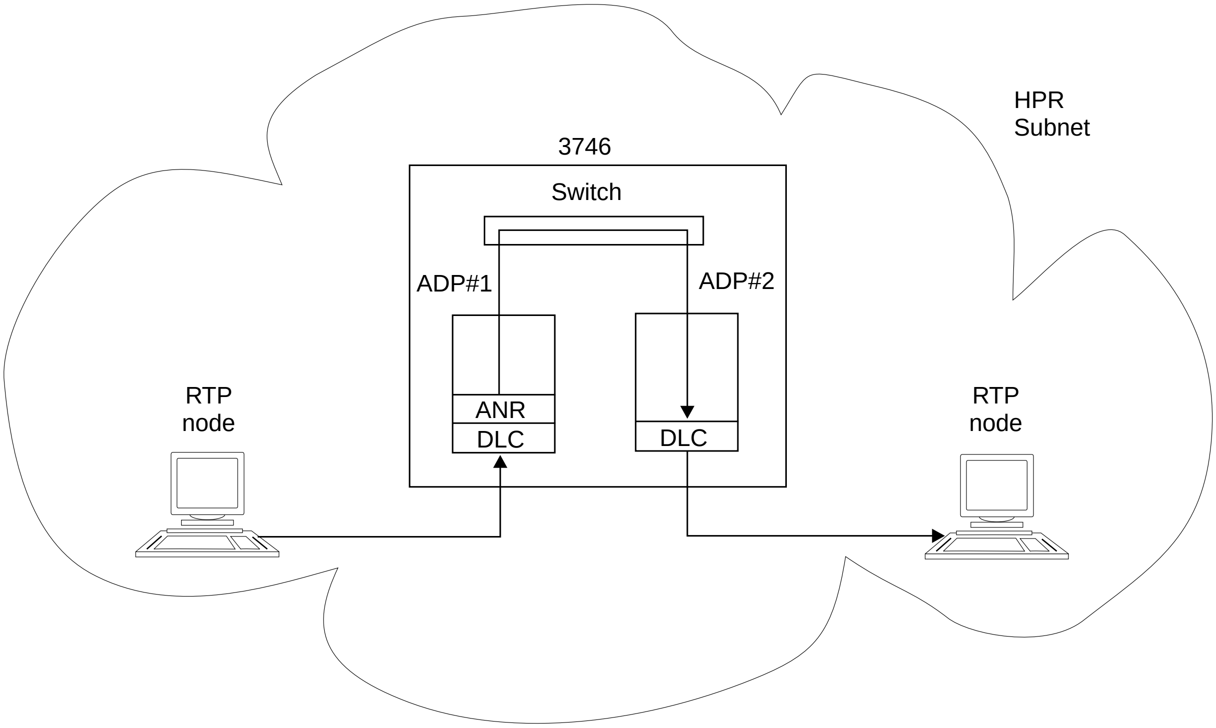

Figure 2-17. 3746 as an ANR Node

Example 2: APPN/RTP Node with Boundary Function

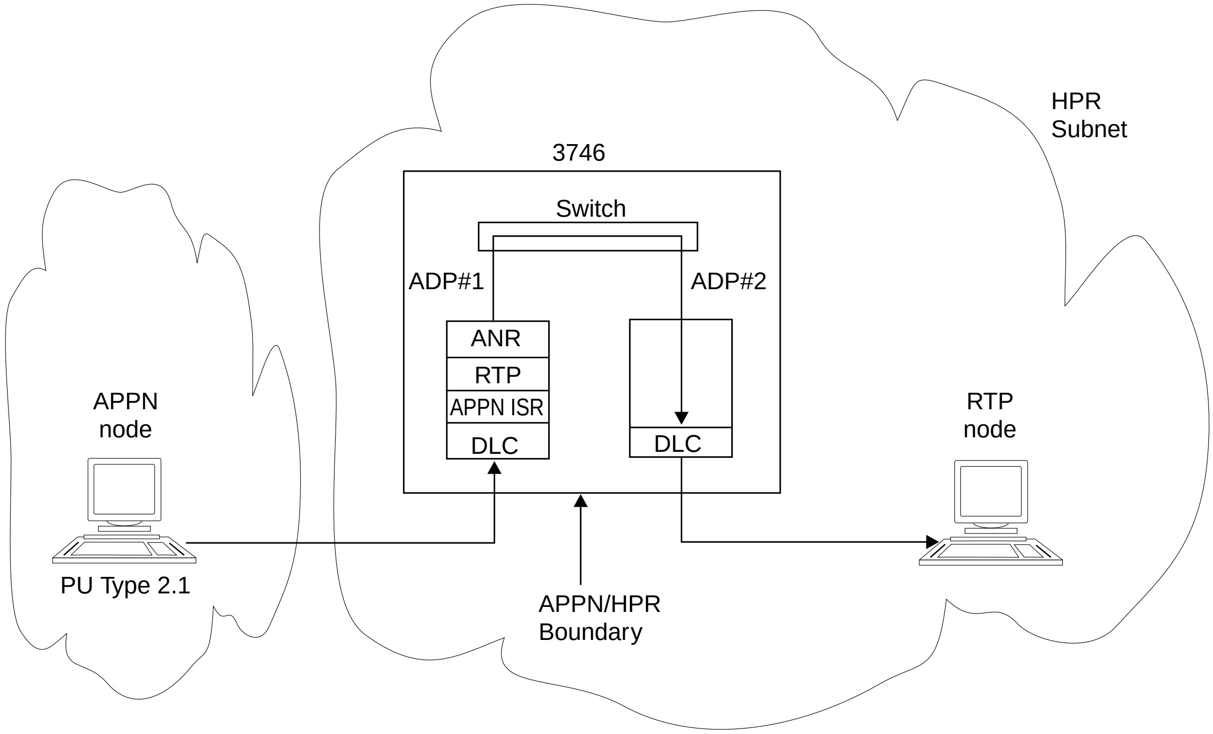

Figure 2-18. 3746 as an APPN/RTP Node (with Boundary Function)

Example 3: DLUR/RTP Node with Boundary Function

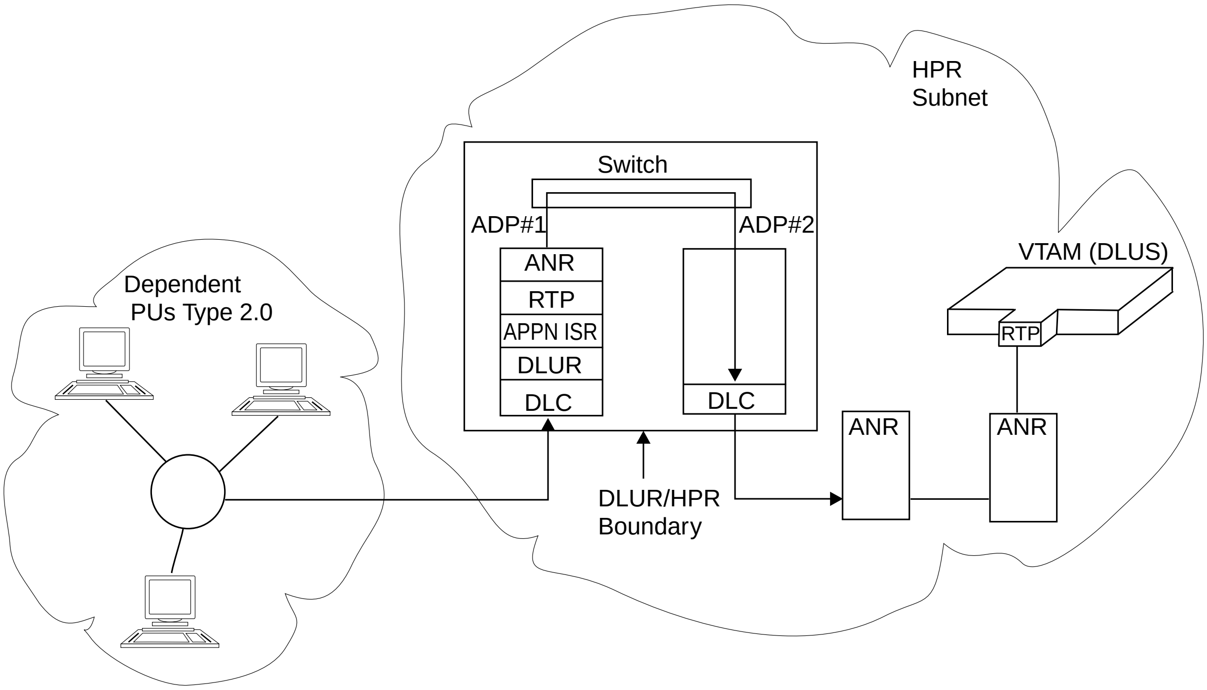

Figure 2-19. 3746 as a DLUR/RTP Node (with Boundary Function)

Example 4: APPN/RTP Node with SNI Network

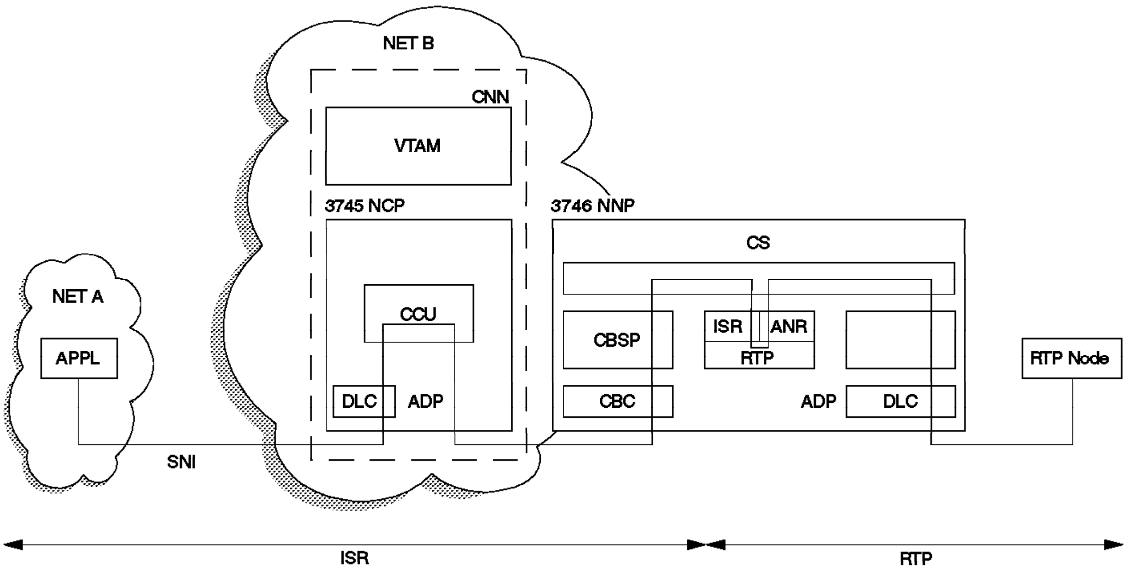

Figure 2-20. 3746 as an APPN/RTP Node with SNI Network

External APPN Connection between a 3745-Composite Network Node and a 3746-900

A simple way to interconnect a 3745-CNN and the attached 3746-900 Network Node is to use an external connection (token-ring or serial link). Compared to the internal APPN connection (see "Internal APPN Connection between a 3745 and a 3746-900" on page 2-41 form=pageonly), this connection does not support as much traffic between the 3745-CNN and the 3746-NN, for example traffic between nodes such as ENA and ENB.

For example, in Figure 2-21, the advantages of a connection network, duplicate MAC addresses, and optimal route computation are achieved by using an "external" token ring interconnecting the 3745-CNN and the attached 3746-900 NN.

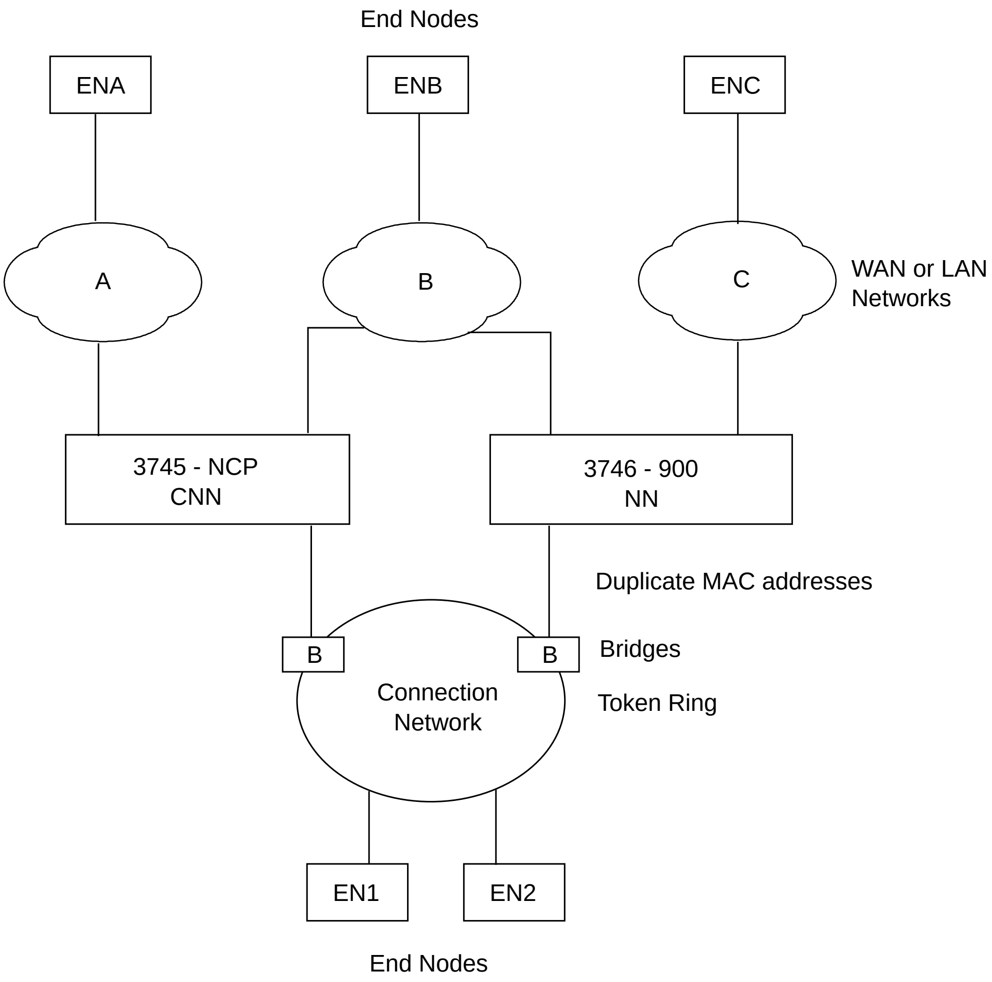

Figure 2-21. External Connection of 3745-CNN and 3746-NN

If ENA wants to establish a session with ENC, the route will be
ENA → CNN → token-ring LAN (Connection Network) → 3746 NN → ENC.

End nodes cannot have multiple network node servers at the same time, they can have *only* one.

By using duplicate MAC addresses between the 3745 TIC and the 3746 TIC, end nodes connected to the connection network LAN will use "at random" either the 3745 or the 3746 as a network node server.

Referring to Figure 2-21, suppose EN1 has CNN as a network node server:

Case 1: EN1 establishes an LU session with EN2
The CNN will send back the TG vector of the connection network and EN1 will establish the session with EN2 through the connection network. The 3746 NN is not involved, and the 3745 CNN is no longer involved once the session has been set up.

Case 2: EN1 establishes an LU session with ENA
As ENA is not reachable by the 3746 NN, the 3745 CNN will send back the RSCV containing the route:
EN1 → CNN → ENA.
Here, only CNN is used and the 3746 NN is not involved.

Case 3: EN1 establishes a LU session with ENB
In its route computation, the CNN (as a network node of EN1) will find two possible routes:
EN1 → CNN → ENB
OR
EN1 → 3746 NN → ENB.
This is because CNN uses the topology database with all nodes and TGs. CNN will choose the best route that maps the COS of the required session.

If the best route is:
EN1 → 3746 NN → ENB,
EN1 will send the BIND and all data directly to 3746 NN. The CNN is not involved.

Case 4: EN1 establishes a LU session with ENC
This is more or less identical to case 3. CNN finds one route that is:
EN1 → 3746 NN → ENC,
and thus EN1 will send the BIND and all data directly to 3746 NN. The CNN is then no longer involved.

HPR MLTGs in the 3746

HPR MLTGs on the 3746 can consist of any mix of CLP and TRP links (ESCON is not supported) running any mix of SDLC, frame relay, X25 and 802.2 LLC protocols.

HPR MLTGs are defined in CCM from the APPN parameters window for APPN stations. You can specify the name and TG number (1-20) of the MLTG here. Figure 2-22 shows the MLTG definition window.

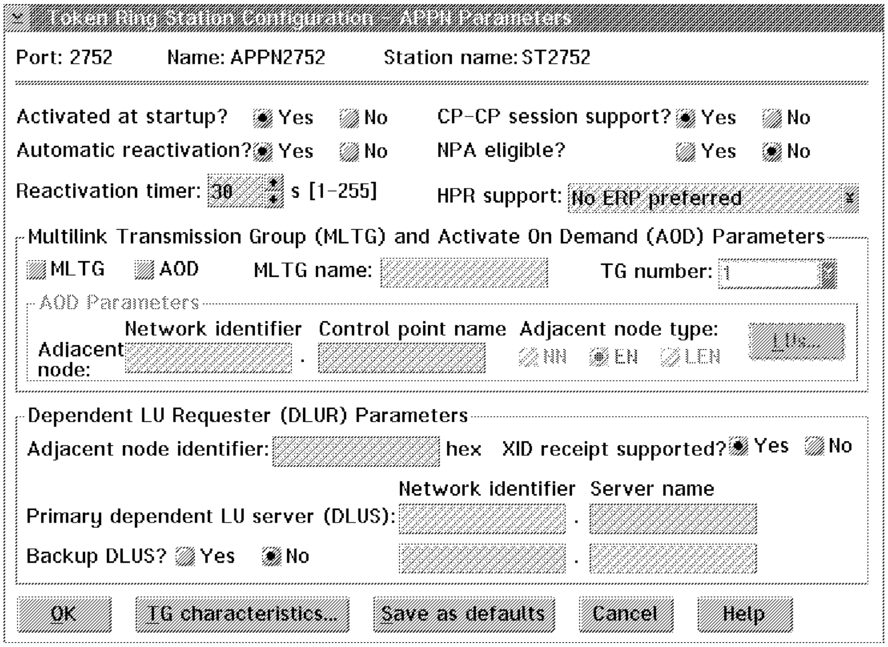

Token Ring Station Configuration - APPN Parameters

Port: 2752 Name: APPN2752 Station name: ST2752

Activated at startup? ● Yes ○ No CP-CP session support? ● Yes ○ No

Automatic reactivation? ● Yes ○ No NPA eligible? ○ Yes ● No

Reactivation timer: 30 s [1-255] HPR support: No ERP preferred

Multilink Transmission Group (MLTG) and Activate On Demand (AOD) Parameters

☐ MLTG ☐ AOD MLTG name: TG number: 1

AOD Parameters

Adjacent node: Network identifier . Control point name Adjacent node type: ○ NN ● EN ○ LEN

Dependent LU Requester (DLUR) Parameters

Adjacent node identifier: hex XID receipt supported? ● Yes ○ No

Primary dependent LU server (DLUS): Network identifier . Server name

Backup DLUS? ○ Yes ● No

Figure 2-22. Defining MLTGs

Other links that belong to the same TG must be defined with the same name and TG number. If during XID exchange, the 3746 APPN CP discovers that a new TG does not go to the same adjacent node as the already active TGs in that MLTG, then the TG activation will fail.

In addition, the CCM performs checking that an MLTG name is only associated with a single MLTG number across all MLTG definitions.

The previous example shows how to define MLTGs at each link station, there is also a central way of defining and checking all MLTG definitions in CCM. From the main CCM screen select Configuration -> APPN -> MLTG. This displays a screen where all defined MLTGs, and all defined links stations can be seen. From this screen it is possibble to change existing MLTGs and define new MLTGs.

3746 APPN Activate on Demand (AOD)

The AOD function gives a 3746 user more control over when a link used by APPN is activated. With AOD active, the link is activated when needed and not before. This can lead to savings when switched lines are used, and charges are calculated according to the time the link was in use. AOD is only used to activate a link, a different function (Limited Resource links for example) must be used to deactivate the link when it is no longer in use. If the link is deactivated, AOD will reactivate the link again when it is needed.

AOD gives the user a new option for controlling when an APPN link is activated. Previously the user could specify "Activate at startup", this caused the physical link to be activated after a 3746 IML, or the user could manually start a link.

The adjacent APPN node at the end of an AOD link can be either a network, end, or LEN node.

APPN Processing

The following sections decribe how AOD works.

Locating Partner LUs: Normal APPN processing before sessions can be started between two logical units (LUs), is to locate the partner LU. Although there are variations on this processing, a "Locate/Find, CD-Initiate" will arrive at the network node server (NNS) of an APPN end node, or at a network node (NN). The network node will give a positive answer if the LU resides on the NN, or search its attached end nodes, or if the EN has registered its resources with its NNS the NNS can respond for the EN. The NN can only search its EN, or know the resources on that EN if the link to the EN is active. In the case of an AOD link, the link will still be inactive as no sessions are using the link. Normally the NN would reply that the LU being searched for cannot be found. To resolve this problem for AOD attached nodes, the NETID, CP name, and LUs of AOD attached nodes must all be predefined on the NN.

In the case of predefined LUs, the NN will answer a Locate/Find with a Locate/Found and pass an RSCV which includes the AOD link back to the LU that issued the Locate/Find. The LU will then issue a BIND to start the LU-LU session. When this BIND arrives at the NN, the AOD link must be activated before the BIND can be passed to the EN. The link is automatically activated at this time, and when it becomes active the BIND is forwarded to the node at the end of the AOD link.

As long as an AOD link is inactive, only predefined LUs can be located which reside on an APPN node at the end of an AOD link. Once the AOD link becomes active, the network node can search the adjacent node and sessions can also be activated to undefined LUs.

APPN Topology: AOD links are repoted as being active in the APPN topology database, even when inactive. This means that APPN route calculation will use such links when calculating the optimal route for a session. If limited resource deactivates a link, this deactivation is not reported to APPN, this means that topology and route selection services will still calcualte routes that cross inactive AOD links. When the link is actually needed again, AOD will reactivate it.

Defining AOD in CCM

AOD links are defined in CCM in two stages:

1. Specify AOD for the link and predefine the NETID and CP name of the adjacent node
2. Predefine the LUs on the adjacent node

The following sections show this process in detail. The example shows an SDLC link with a single station.

Defining AOD Links: The first steps are to define the port and link stations on that port, this is shown in Figure 2-23 on page 2-52 and Figure 2-24 on page 2-52.

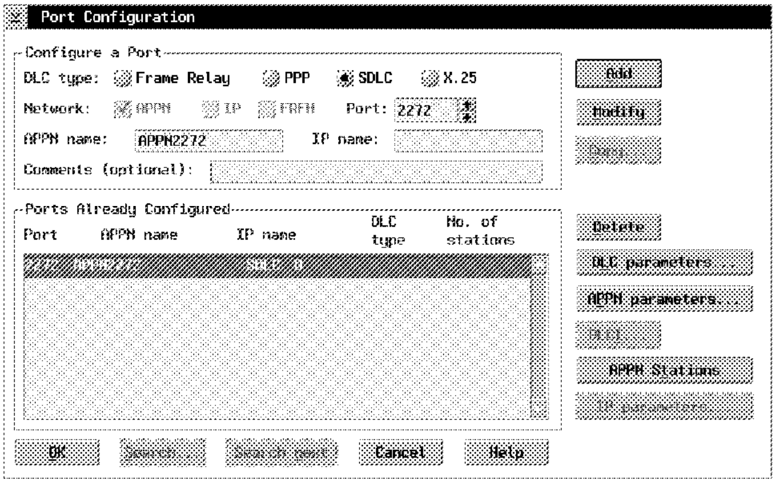

Port Configuration

Configure a Port

DLC type: ○ Frame Relay ○ PPP ● SDLC ○ X.25

Network: ☒ APPN ○ IP ○ FRFH Port: 2272

APPN name: APPN2272 IP name:

Comments (optional):

Ports Already Configured

Port	APPN name	IP name	DLC type	No. of stations
2272	APPN2272		SDLC	0

Figure 2-23. Port Configuration

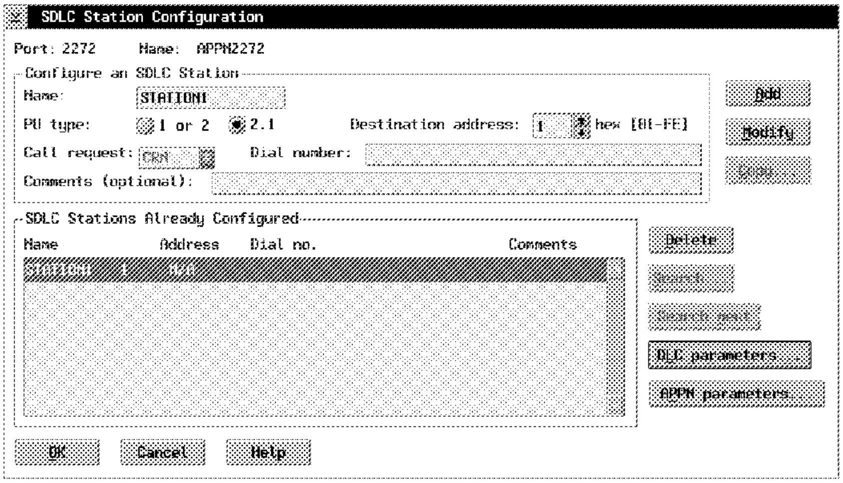

SDLC Station Configuration

Port: 2272 Name: APPN2272

Configure an SDLC Station

Name: STATION1

PU type: ○ 1 or 2 ● 2.1 Destination address: 1 hex [01-FE]

Call request: CRN Dial number:

Comments (optional):

SDLC Stations Already Configured

Name	Address	Dial no.	Comments
STATION1	1	N/A	

Figure 2-24. Station Configuration

For the station, we select *APPN parameters*, the resulting dialog is shown in Figure 2-25 on page 2-53 Here AOD can be activated for this link, and we must define the NETID, CP name, and APPN node type of the adjacent APPN node at the other end of the AOD link.

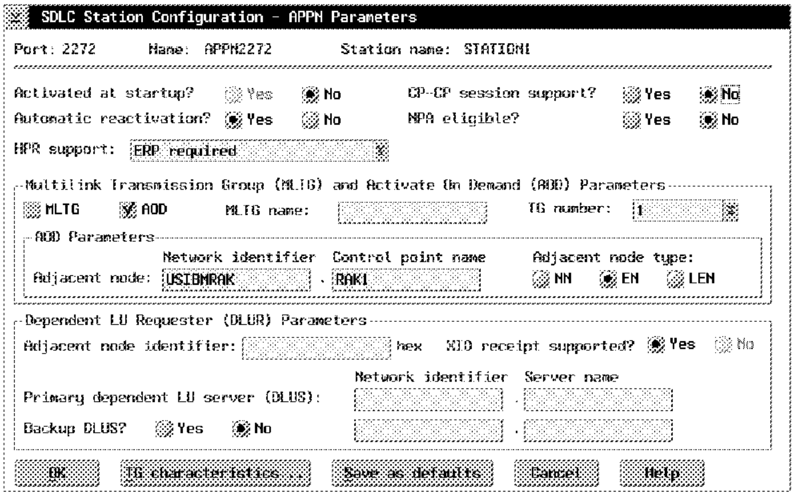

SDLC Station Configuration - APPN Parameters

Port: 2272 Name: APPN2272 Station name: STATION1

Activated at startup? ○ Yes ● No CP-CP session support? ○ Yes ● No

Automatic reactivation? ● Yes ○ No NPA eligible? ○ Yes ● No

HPR support: ERP required

Multilink Transmission Group (MLTG) and Activate On Demand (AOD) Parameters

☐ MLTG ☑ AOD MLTG name: TG number: 1

AOD Parameters

Adjacent node: Network identifier USIBMRAK . Control point name RAK1 Adjacent node type: ○ NN ● EN ○ LEN

Dependent LU Requester (DLUR) Parameters

Adjacent node identifier: hex XID receipt supported? ● Yes ○ No

Primary dependent LU server (DLUS): Network identifier . Server name

Backup DLUS? ○ Yes ● No

Figure 2-25. Station Configuration APPN Parameters

Predefining Logical Units: To predefine logical units, from the main CCM screen select *configuration*, then *APPN*, then *Adjacent Node remote LUs*. This is shown in Figure 2-26

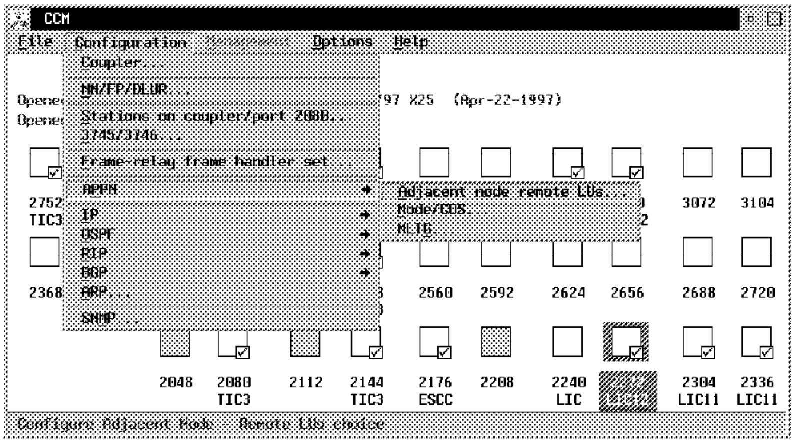

Figure 2-26. Adjacent Node Remote LUs

The following dialog allows the user to define the LUs on each adjacent node, shown in Figure 2-27 on page 2-54 In each case the NETID and CP name of the adjacent node, and the NETID and LU name must be specified. The example shows LUs defined with full, partial, and no wildcards, although the of definitions shown only one wildcard type is needed, here all three types were shown as an example.

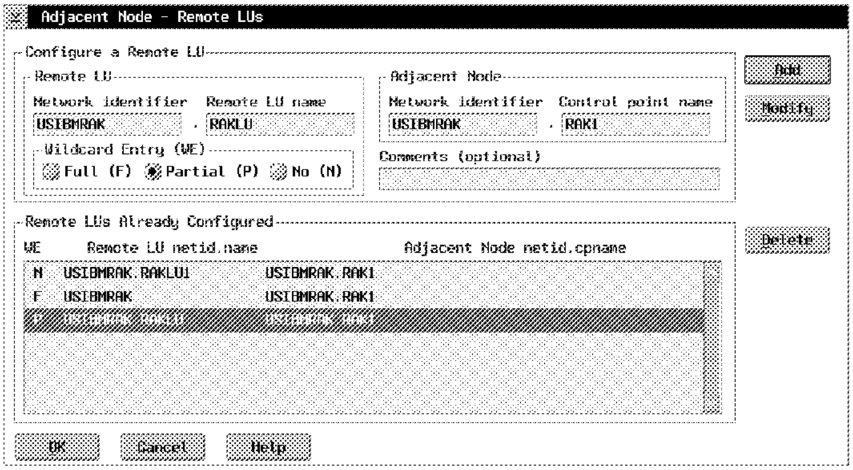

Configure a Remote LU

Remote LU

Network identifier: USIBMRAK . Remote LU name: RAKLU

Wildcard Entry (WE): ○ Full (F) ● Partial (P) ○ No (N)

Adjacent Node

Network identifier: USIBMRAK . Control point name: RAK1

Comments (optional):

Remote LUs Already Configured

WE	Remote LU netid.name	Adjacent Node netid.cpname
N	USIBMRAK.RAKLU1	USIBMRAK.RAK1
F	USIBMRAK	USIBMRAK.RAK1
P	USIBMRAK.RAKLU	USIBMRAK.RAK1

| Figure 2-27. Predefining Remote LUs

User Defined Parameters

The 3746-9x0 implements three "User Defined Parameters" (UDP1, UDP2 and UDP3). Their purpose is to allow the network operator to alter the TG weight between the 3746-9x0 and its adjacent Network Nodes and so favor some TGs for route selection over other TGs. APPN route selection will select the route with the lowest TG weight that satisfies all other parameters for the session. The TGs with a heavier TG weight being used in case of congestion or unavailability of the lower weight TG. There are three ways to make UDP definitions:

1. **Class Of Services and LOGMODE definitions**

   User Defined Parameters are enabled only for user defined Class Of Services (COS) and logmodes, changing standard logmodes such as #CONNECT or #INTER, would reduce UDP effectiveness.

In order to define your own COS and LOGMODE, go to the *CONFIGURATION* menu from CCM main panel, select **APPN** then select **MODE/COS** to get following screen:

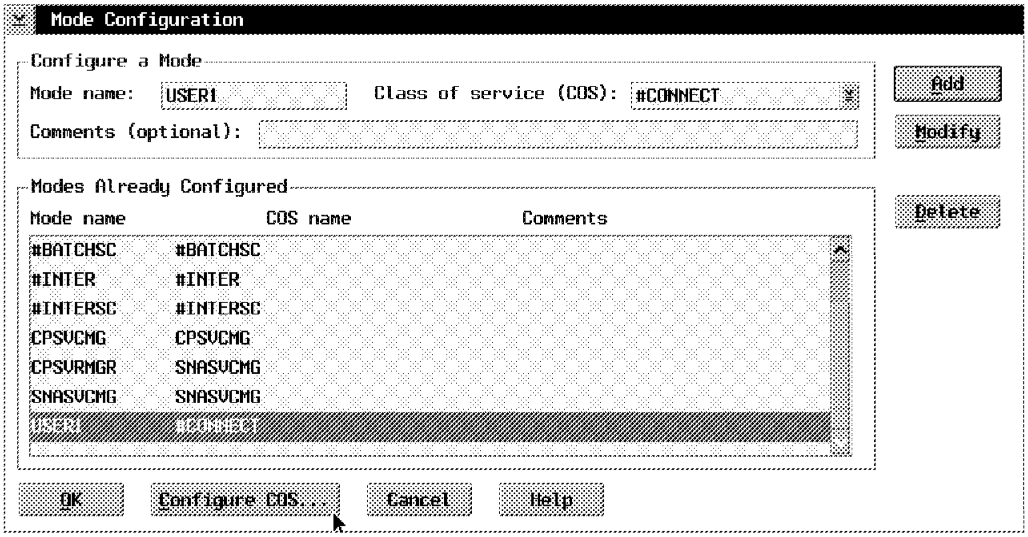

Mode Configuration

Configure a Mode

Mode name: USER1   Class of service (COS): #CONNECT

Comments (optional):

Modes Already Configured

Mode name	COS name	Comments
#BATCHSC	#BATCHSC	
#INTER	#INTER	
#INTERSC	#INTERSC	
CPSVCMG	CPSVCMG	
CPSVRMGR	SNASVCMG	
SNASVCMG	SNASVCMG	
USER1	#CONNECT	

Figure 2-28. Mode Configuration Screen

Select **CONFIGURE COS** to create your own Class Of Services Enter COS Name and click on **ADD** button. Select TG Row to specify the different TG weights and associated UDPs ranges. This gives the following screen:

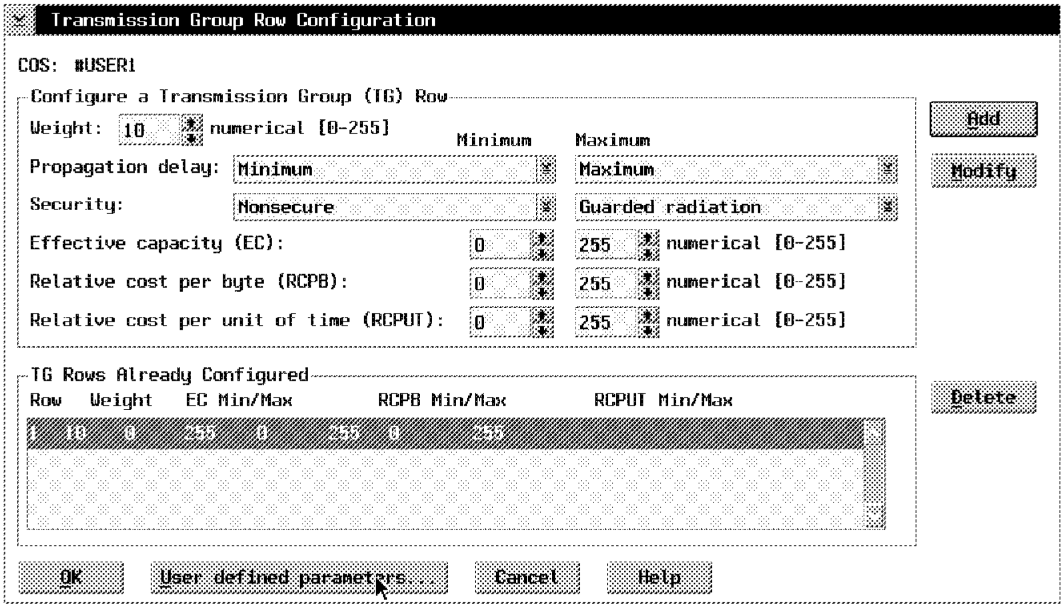

Transmission Group Row Configuration

COS: #USER1

Configure a Transmission Group (TG) Row

Weight: 10 numerical [0-255]

	Minimum	Maximum	
Propagation delay:	Minimum	Maximum	
Security:	Nonsecure	Guarded radiation	
Effective capacity (EC):	0	255	numerical [0-255]
Relative cost per byte (RCPB):	0	255	numerical [0-255]
Relative cost per unit of time (RCPUT):	0	255	numerical [0-255]

TG Rows Already Configured

Row	Weight	EC Min/Max		RCPB Min/Max		RCPUT Min/Max	
1	10	0	255	0	255	0	255

Figure 2-29. TG Row Configuration

Enter the TG weight you want to define, then click on **USER DEFINED PARAMETERS.** This gives the following screen:

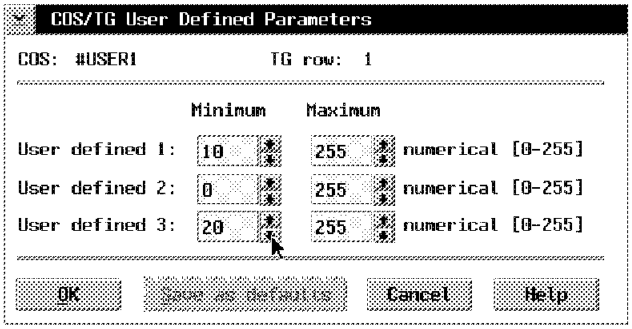
COS/TG User Defined Parameters

COS: #USER1 TG row: 1

	Minimum	Maximum	
User defined 1:	10	255	numerical [0-255]
User defined 2:	0	255	numerical [0-255]
User defined 3:	20	255	numerical [0-255]

Figure 2-30. COS/TG User Defined Parameter Screen

You can now specify the ranges for UDP1, UDP2 and UDP3 which will be associated to the TG weight previously defined. Click on OK to go back to TG Row configuration screen Repeat this operation for each different TG weight you want to define then click successively on OK to return to the CCM main panel.

2. **UDPs definitions at PORT level**

   You can specify specific UDPs values for each APPN Port you define from CCM. These values will be compared to the ranges defined in COS and LOGMODE tables and the corresponding TG weight will be associated to the PORT. According to UDPs values, different weights will be associated PORTs, thus allowing APPN to compute the best route.

3. **TGs characteristics modification at STATION level**

   By default, the PORT UDPs values are used for any STATION on that port. These UDPs can be specified for specific stations in order to force a different route weight.

   From station definition screen, under CCM, select **APPN PARAMETERS** then **TG CHARACTERISTICS** to get the following screen:

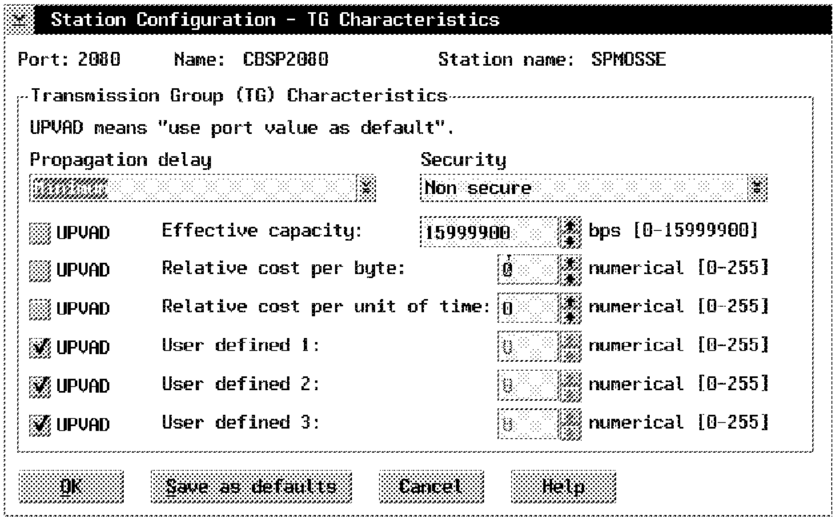
Station Configuration - TG Characteristics

Port: 2080 Name: CBSP2080 Station name: SPMOSSE

Transmission Group (TG) Characteristics

UPVAD means "use port value as default".

Propagation delay: [illegible]

Security: Non secure

□ UPVAD	Effective capacity:	15999900	bps [0-15999900]
□ UPVAD	Relative cost per byte:	0	numerical [0-255]
□ UPVAD	Relative cost per unit of time:	0	numerical [0-255]
☑ UPVAD	User defined 1:	0	numerical [0-255]
☑ UPVAD	User defined 2:	0	numerical [0-255]
☑ UPVAD	User defined 3:	0	numerical [0-255]

Figure 2-31. Station UDPs

Deselect the push button for the parameters you want to alter, then enter the new parameter value. Click on OK after having defined the new values. If

UPVAD (use port value as default) is selected, then the port parameters are used as default for this link.

3746 Nways Multiprotocol Controllers in SNA/APPN/HPR Networks: Operation

When operating in an APPN/HPR or mixed SNA and APPN/HPR environment, the 3746-900 and 3746-950 provide the following networking capabilities:

- APPN/HPR network node services
- Dependent logical unit requester (DLUR)
- Automatic Network Routing (ANR) and Rapid Transport Protocol (RTP) for HPR traffic

APPN/HPR Network Node

The 3746-9x0 supports APPN/HPR network node functions. They include the network node services for the APPN end nodes connected to the 3746 (adjacent end nodes). As a network node, the 3746 automatically and dynamically learns the full, up-to-date connection topology of the network. It dynamically locates any resource and computes routes within the network. The 3746-9x0 can also register the resources of adjacent end nodes to the APPN central directory server node such as VTAM.

The 3746 network node supports the following types of end-node connections (see Figure 2-32 on page 2-58):

- APPN (PU type 2.1, such as PS/2s and IBM 3174s)
- Non-APPN (including PU types 1.0 and 2.0, such as 3270-type devices)
- Low-entry networking (PU type 2.1, such as IBM System 36™ or nodes without APPN installed, for example, an IBM AS/400®, IBM 3174, IBM PS/2, or other PCs).

The 3746-9x0 APPN/HPR control point functions are performed by a dedicated processor (the network node processor). The routing of data is done by the adapters (either port-to-port within an adapter or adapter-to-adapter) without any control point intervention. This allows the 3746-9x0 to support high speed data transfer.

Dependent Logical Unit Requester (DLUR)

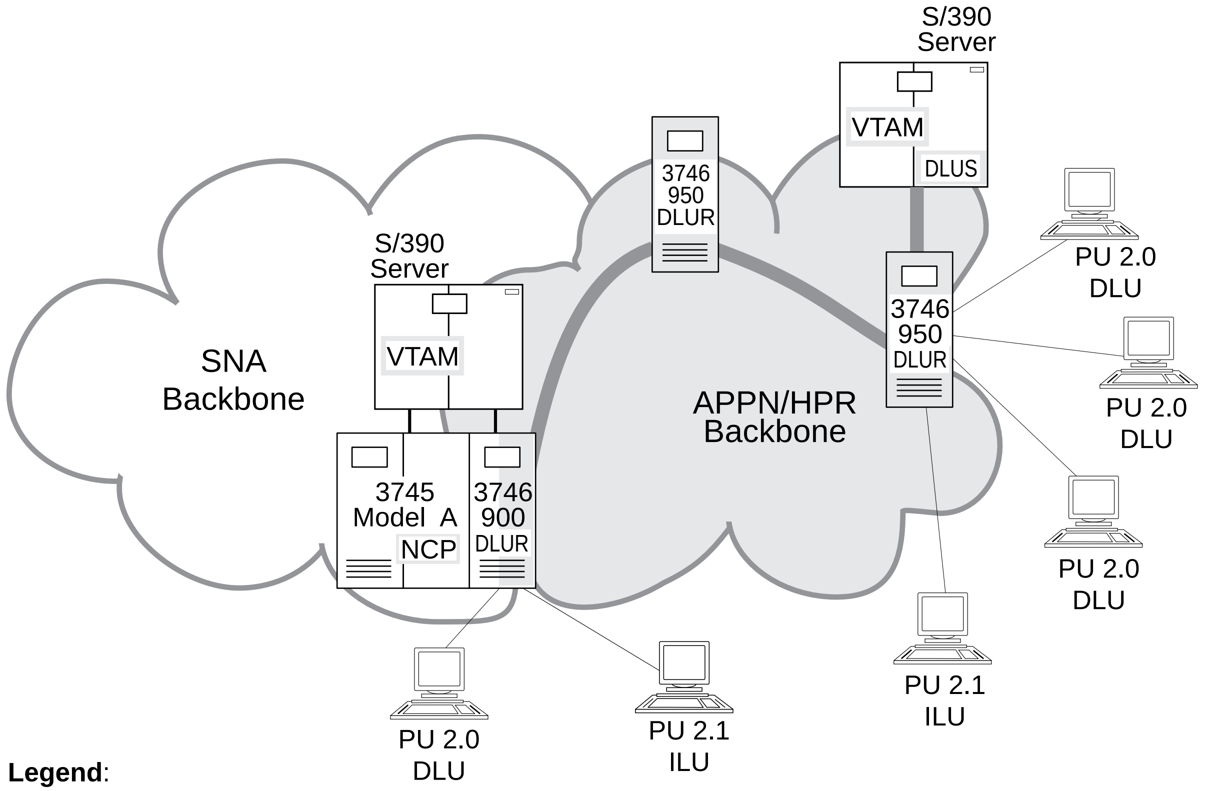

Legend:

DLU	Dependent LU (non-APPN)
DLUR	Dependent LU requester
DLUS	Dependent LU server
ILU	Independent LU (APPN)
▬	DLUR to DLUS path

Figure 2-32. Dependent LU Support

The 3746 network node allows existing host-dependent SNA devices to access S/390 applications across an APPN/HPR backbone network. For example (see Figure 2-32), a physical unit type 1.0 or type 2.0 attached to the rightmost controller can access applications in any of the two S/390 servers.

Host-dependent logical units (LUs) need a control session with their VTAM system services control point (SSCP). This SSCP-LU session allows the dependent LUs to request VTAM to set up LU-LU sessions for them. Once an LU-LU session has been established, the dependent LU (secondary LU) can exchange data with the application LU (primary LU).

In an APPN environment, the dependent LUs (DLU) must reside on, or be owned by an APPN node providing the DLUR function. The DLUR node requests the dependent logical unit server (DLUS) of a VTAM network node to provide the SSCP services for its dependent LUs. To support session establishment for the dependent LUs, the traditional SSCP-PU or SSCP-LU data flows through two LU 6.2 sessions between the DLUR node and DLUS node.

The 3746-9x0 provides the DLUR function, in conjunction with the DLUS support provided by VTAM Version 4 Release 2 or higher level. In a network with multiple VTAMs, only one VTAM with DLUS support is required.

The DLUS can be in any APPN/HPR network, provided that an APPN path exists between the DLUS and each DLUR.

High-Performance Routing (HPR)

The IBM 3746 Model 900 and 950 support the high-performance routing (HPR). HPR is an extension to the APPN architecture that takes advantage of fast links with low error rates. HPR enhances the routing mechanisms and provides the following benefits:

- Dynamic rerouting around failed nodes and links without session loss
- Better routing performance
- Enhanced congestion control which improves link efficiency
- Reduction of amount of storage required in intermediate nodes
- Very high data throughputs between the S/390 servers and the network
- Synergy with the parallel SYSPLEX processor implementation providing end-to-end non-disruptive path switching up to the applications. This includes the S/390 support of multi-node persistent session (product direction).

The HPR architecture is made of two layers:

- Automatic network routing (ANR), the HPR base, mainly in the intermediate nodes
- Rapid transport protocol (RTP), the HPR transport tower, in the edge nodes

ANR is responsible for data packets routing across an HPR network. ANR uses a new form of addressing to identify routes through an HPR network. This addressing is based on the links and nodes that make up the route. It consists of labels which are contained in the data packet header – each label describing the outbound link to be taken to exit an intermediate node, processing performed in each intermediate node is reduced. ANR provides several functions, for example:

- Source-independent routing
- Connectionless, stateless, fast routing without hop-by-hop error recovery procedure (non-ERP mode)
- Discarding of incoming packets in the event of congestion

RTP establishes end-to-end connections between edge nodes of the APPN/HPR network. Each RTP connection can carry traffic for multiple end-to-end user and control sessions. RTP relies on ANR to perform the packet forwarding across the HPR network and offers the following functions:

- Transport of APPN and SNA boundary traffic (DLUR)
- Selective transmission, based on class of service
- Flow control and network congestion avoidance (Automatic Rate Based - ARB)
- End-to-end error recovery and selective retransmission
- Non-disruptive rerouting around network failures

HPR Multilink Transmission Group (MLTG)

MLTG enables the 3746 Models 900 and 950 to use a variable bandwidth on a single logical Transmission Group (TG) composed of multiple physical links or LANs, most helpful in situations where single or multiple sessions require more bandwidth than a single physical link or LAN can provide.

The MLTG is defined with a single TG number. It is reported and seen as a single TG by Topology Management, and viewed as a single TG in the route calculation process. Based on link error rates, error recovery may be run on an individual link basis. Links may be dynamically removed from the MLTG when no longer needed, thus resulting in cost savings.

HPR MLTG is supported over SDLC, frame-relay and X.25 links, token-ring and Ethernet LANs.

3746 Model 900 as a Mixed SNA and APPN/HPR Node

In a mixed SNA and APPN/HPR network, as shown in Figure 2-32 on page 2-58, the 3746-900 NN can operate as an SNA node (PU type 4) for the network resources owned by the NCP running in the 3745, and as an APPN/HPR network node (PU type 2.1) for the resources owned by the 3746 APPN/HPR control point.

In Figure 2-32 on page 2-58, the 3745/3746-900 NN is channel-attached to a VTAM which operates as an interchange node (ICN). This ICN allows SNA devices connected to the SNA backbone to access applications over the APPN/HPR backbone, and SNA/APPN devices connected to the APPN/HPR backbone to access S/390 applications in the SNA backbone, therefore providing any-to-any networking.

X.25 Network Connectivity

The X.25 Support feature natively controls X.25 connections without using the support of NCP or NCP Packet Switching Interface (NPSI). It allows the 3746 to attach to private or public X.25 networks as DTE nodes for routing of APPN, SNA/DLUR, HPR, and IP traffic over those X.25 connections. Supported functionality includes:

- Support of Qualified Logical Link Control (QLLC) connections for APPN, SNA/DLUR, and HPR traffic
- Routing of mixed APPN, SNA/DLUR, HPR, IP, and SNA/NCP traffic on the same X.25 link
- PVC and SVC connections
- Up to 2.048 Mbps speed
- Direct DTE-to-DTE attachment
- Support of SDLC, frame relay, ISDN, X.25, X.25 NCP, and X.25 NPSI links attached to a single Communication Line Processor (CLP).

Frame-relay Networking

The IBM 3746 Models 900 and 950 provide frame-relay networking functions. They are similar to those provided by NCP for 3746-900, with the exception of frame-relay switching functions (see Chapter 7, “Frame Relay Overview” on page 7-1).

Physical Media

The 3746-900 network node and 3746-950 provide the same connectivity. They both support:

- Token-ring and Ethernet LANs
- Leased frame-relay links
- Switched and leased SDLC links
- ESCON channels
- Frame-relay network connections

- X.25 network connections

The 3746 network node supports RFC 1490, including:

- Frame relay BAN
- Frame relay BNN

More information about the adapters that support the network node connectivity can be found in “Evolution of the 3745 and 3746 Controller Family.”

3746-9x0 MAE APPN/HPR Implementation

The MAE, together with the Multi-Protocol Access Services V1R1 (MAS) software, has the The MAE has the capability of being an APPN network node (NN) with intermediate routing functions and provides network services to both APPN and LEN end nodes.

MAE Implementation Specifics

APPN over DLSw

The MAE supports APPN over DLSw for connectivity to nodes through a remote DLSw partner. An example is shown in Figure 2-33 on page 2-62. The MAE only supports remote DLSw.

Note: It is recommended that you use APPN over direct DLCs when available instead of APPN over DLSw.

When APPN is configured on the MAE to use a Data Link Switching (DLSw) port, DLSw is used to provide a connection-oriented interface (802.2 LLC Type 2) between the APPN component in the MAE and APPN nodes and LEN end nodes attached to a remote DLSw partner.

When configuring a DLSw port for APPN on the MAE, the network node itself is assigned a unique MAC and SAP address pair that enables it to communicate with DLSw. The MAC address for the network node is locally administered and must not correspond to any physical MAC address in the DLSw network.

Figure 2-33 on page 2-62 shows how TCP/IP and DLSw are used to transport APPN traffic over an IP network.

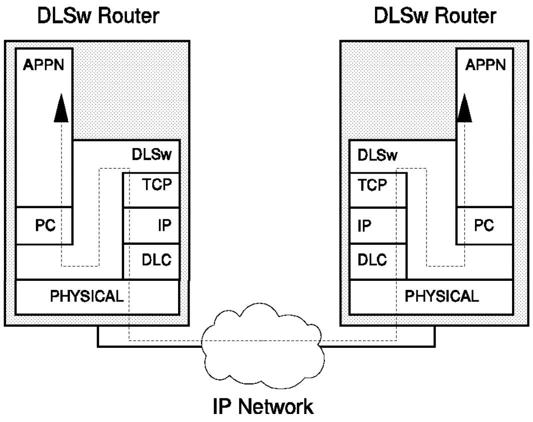

Figure 2-33. APPN over DLSw

Supported Traffic Types

APPN ISR uses the QLLC protocol for X.25 direct data link control, the IEEE 802.2 LLC Type 2 protocol for token-ring, Ethernet, PPP, and frame relay and SDLC protocol for the SDLC data link control. APPN HPR, which is supported on token-ring, Ethernet, PPP and frame relay, does not use LLC Type 2 protocol, but does use some functions of an APPN link station for XID and inactivity timeout. A single APPN link station is therefore used for ISR or HPR. Different mechanisms are used to distinguish between ISR and HPR traffic depending upon the DLC type:

For token-ring and Ethernet LAN ports

Each protocol that uses a port must have a unique SAP address, with the exception of DLSw (which may use the same SAP address as other protocols because DLSw frames will not be destined for the local MAC address, but rather a DLSw MAC address). A unique SAP address identifies the APPN link station for HPR traffic (local HPR SAP address parameter). If ISR traffic is destined for a link station, then a different SAP address (local APPN SAP address parameter) must be used. The ISR traffic uses LLC Type 2 LAN frames. The HPR traffic is handled in a similar fashion to LLC Type 1 LAN frames and must have a different SAP address. The default SAP address for HPR traffic is X'C8'. If X'C8' has already been used by another protocol on a port, the default must be overridden. Note that there is only one APPN link station even though APPN ISR and HPR traffic use different SAP addresses.

For frame relay ports

APPN ISR traffic and APPN HPR traffic transferred over a frame relay data link connection support both the RFC 1490 bridged frame format and the RFC 1490 routed frame format. RFC 1490 routed frame format APPN ISR traffic will be transferred over a frame relay data link connection using the connection-oriented multiprotocol encapsulation method defined in RFC 1490 using:

- NLPID = X'08' (Q.933 encoding)
- L2PID = X'4C80' (Layer 2 protocol identifier indicating 802.2 LLC)
- L3PID = X'7083' (Layer 3 protocol identifier indicating SNA-APPN/FID2)

APPN HPR traffic transferred over a frame-relay data link connection does not use IEEE 802.2 LLC. It uses a different multiprotocol encapsulation as defined in RFC 1490 using:

- NLPID = X'08' (Q.933 encoding)
- L2PID = X'5081' (Layer 2 protocol identifier for no layer 2 protocol)
- L3PID = X'7085' (Layer 3 protocol identifier indicating SNA-APPN/HPR)

APPN HPR does not use a SAP for traffic transferred using the RFC 1490 routed frame format because there is no layer 2 protocol. APPN HPR uses a SAP for traffic transferred using the RFC 1490 bridged frame format.

The 2210 and 2216 support both the routed and bridged frame formats, the 6611 supports only the routed frame format.

For PPP ports

APPN ISR traffic uses 802.2 LLC over the PPP connection. Since there is no layer 2 protocol used in HPR's RFC 1490 encapsulation (non-ERP), no SAP is used for HPR traffic.

Topology Safestore

Topology and Routing Services (TRS) can store the APPN topology database on the 3746 Multiaccess Enclosure's hard disk. In order to reduce the number of topology database updates (TDUs) transmitted over the network, the backup copy of the topology database is restored when the APPN topology database maintained in the MAE's memory is lost due to either a power loss or reboot. After the topology database is retrieved from the hard disk during startup, TRS advertises the last TDU sequence number received by the MAE. Only APPN network changes made after that sequence number will be broadcast. Without this feature, a complete set of TDU broadcasts are sent that significantly increases network traffic.

The 3746 Multiaccess Enclosure only saves the topology to its hard disk once a day during garbage collection.

Route Test

Currently, HPR's route test can only be invoked through SNMP. Two variations of route test exist. The first tests the wrap-around time of an established HPR connection. This route test is invoked specifying the NCL and RTP connection identifiers of an HPR connection. The NCL and RTP connection identifiers are unique identifiers used to identify HPR connections. These identifiers can be retrieved from the HPR RTP connection table MIB information.

The second tests the wrap-around time of an APPN selected route to a specified destination node. This route test is invoked specifying the APPN network and LU name of the destination node and the mode used for the session. APPN's route selection algorithm is used to calculate the best route to the destination. If the selected route is an HPR connection originating at the 3746 Multiaccess Enclosure, the route is then tested.

When a route test is initiated, a separate time-stamped message is sent to each node within the HPR connection. At the destination node, the message is returned to the originating node. The route test message is only processed by the NCL forwarders in the intermediate and destination nodes. Upon receipt of the returned route test message, the wrap-around time is recorded along with the corresponding node. The individual link rates can be calculated by comparing the round-trip time for each route test message.

User API

The MAE implementation of APPN does not provide an application program interface to support user-written LU 6.2 programs.

Limited Resource Link Stations

On the MAE, limited resource link stations are supported on the following links:

- Connection network links
- X.25 SVC links (Previewed for APPN)
- PPP links running over ISDN or V.25bis
- Frame relay links running over ISDN

Session-Level Security

A session-level security feature can be enabled for connections between the MAE network node and an adjacent node. Both partners require a matching hexadecimal key that enables each node to verify its partner before a connection is established.

Parallel TGs

Parallel TGs are not supported between two router network nodes using the same port on each router. However, parallel TGs are supported between two router network nodes using different ports on one or both routers. Also, parallel TGs are supported between a router network node and another non-router remote node over the same port using different remote SAP addresses, provided that the remote node has a mechanism to define or accept different local SAP addresses for APPN on the same port.

DLUR Restrictions

The DLUR option, as implemented on the router network node, has the following functional restrictions:

- Only secondary LUs (SLUs) can be supported by the DLUR function. An LU supported by DLUR cannot function as a primary LU (PLU). Therefore, the downstream physical unit (DSPU) should be configured as secondary.
- Because only SLUs are supported, Network Routing Facility (NRF) and Network Terminal Option (NTO) are not supported. Extended recovery facility (XRF) and XRF/CRYPTO are not supported.
- You must be able to establish an APPN-only or APPN/HPR-only session between DLUS and DLUR. The CPSVRMGR session cannot pass through a subarea network.

Connection Network Restrictions

The router APPN support has the following connection network restrictions:

- Connection networks defined on the router network node are only supported on token-ring and Ethernet LAN ports.
- The same connection network (VRN) can be defined on only one LAN. However, the same VRN can be defined on multiple ports having the same characteristics to the same LAN.
- The same connection network can be defined on a maximum of five ports to the same LAN on the router network node.
- There is only one connection network TG from a given port to a given connection network's VRN.
- The same connection network TG characteristics apply for each port on which a given connection network is defined on this router network node. The TG characteristics could be different on a different node.
- Because the VRN is not a real node, CP-CP sessions cannot be established with or through a VRN.
- When a connection network is defined on the router network node, a fully qualified name is specified for the connection network name parameter. Only connection networks with the same network ID as the router network node may be defined. The network ID of the VRN is then the same as the network ID of the router network node.

APPN over DLSw Restrictions

The following restrictions apply to APPN over DLSw:

- Connectivity through remote DLSw partners only
- Only 1 DLSw port per router
- Use of a locally administered MAC address
- HPR is not supported on DLSw ports
- DLSw ports cannot be members of connection networks
- Parallel TGs are not supported over more than one DLSw port

  A parallel TG may contain a single DLSw port, and any combination of other supported DLCs, but a parallel TG may never contain more than one DLSw port.

Summary of Implemented APPN Functions

Note: See the notes at the end of the tables (page 2-68) for abbreviations and explanation of terms used.

Table 2-3 (Page 1 of 3). Summary of Implemented Base Functions (APPN Version 2)

Function		VTAM	3746	MAS
No.	**Description**	**V4R3**	**V4R2**	**V1R1**
Configuration Services				
001	LEN-level XID3	B	N	N
002	All XID3 States	B	N	N
003	Link Station Role Negotiation	B	N	N
006	CP Name on XID3	B	N	N
007	TG Number Negotiation	B	N	N
008	Multiple TGs	B	N	N
010	Single-Link TG	B	N	N
1001	Secondary-Initiated Non-Activation XID	B	N	N
1004	Adjacent Node Name Change	B	N	N
Intermediate Session Routing				
011	LFSID Addressing	B	N	N
013	Priority Queuing for Transmission	2	N	N
Address Space Manager				
020	Extended BIND and UNBIND	B	N	N
021	Adaptive Pacing for Independent LU BINDs	3	N	N
023	Bind Segmenting and Reassembly	-	N	N
024	Adaptive Pacing for Dependent LU BINDs	-	N	N
Session Services				
030	CP-CP Sessions	B	N	N
031	CP-CP Capabilities Exchange	B	N	N
033	FQPCID Generation	B	N	N
034	CD-Initiate	B	N	N
035	Reconstruct CD-Initiate Reply	B	N	N
036	COS/TPF	B	N	N
037	BIND (ILU=PLU)	B	N	N
038	Limited Resource	B	N	N
039	BIND without RSCV from Any LEN or APPN Node	B	N	N
040	Propagate Unrecognized CVs	N	N	N
041	Session RU Segmenting and Reassembly	B	N	N
042	Interleaved Segments	B	N	N
1015	CP-CP Session Activation Enhancements	B	N	N
Directory Services				
050	Register EN Resources	B	N	N
051	Locate/Find/Found	B	N	N
052	Reconstruct GDS Variables for Locate Reply and CD-Initiate Reply	B	N	N
053	Participate in Network Searches	B	N	N
054	Send Wildcard Reply	-	N	N
055	Broadcast and Directed Searches	B	N	N
056	ENCP Search Control	B	N	N
057	Partial Directory Entries	-	N	N
059	Accept Unqualified LU Name	B	N	N
060	Locate Chains - Locate(keep)	B	N	N
061	Sending Locate to a Gateway	N	N	N
062	Cache Resource Locations	N	N	N
063	Favor Explicit Replies	N	N	N
064	Network-Qualified LU Names	B	N	N
065	Central Directory Client	N	N	N
066	Abbreviated Resource Hierarchy	N	N	N
068	Authentic Net ID Indicator	B	N	N
069	DS Support for Domain LEN Resources	6	N	N
1103	Retry Referred Search	N	N	N

Table 2-3 (Page 2 of 3). Summary of Implemented Base Functions (APPN Version 2)

Function		**VTAM**	**3746**	**MAS**
No.	**Description**	**V4R3**	**V4R2**	**V1R1**
1104	Topology-Based Directory Nonverify	5	N	N
1105	PCID Modifier	B	N	N
1109	Surrogate Owner	N	N	N
1117	Bypass of Directed Locate Not Allowed	5	-	-
Topology and Routing Services				
070	Process Local Resource Change	B	N	N
073	Initial Topology Exchange	B	N	N
074	Flow Reduction Sequence Numbers	B	N	N
075	Resource Sequence Numbers	B	N	N
076	Topology Broadcast	N	N	N
077	Garbage Collection	N	N	N
078	Topology Isolation at Net ID Boundaries	N	N	N
079	Build RSCV	N	N	N
080	Calculate Route Using Connection Networks	N	N	N
081	Class-of-Service Manager	B	N	N
082	Route Randomization	N	N	N
083	Member of Connection Network	B	N	N
084	Select One-Hop Routes	N	N	N
085	Select Network Routes	N	N	N
086	Topology Awareness of CP-CP Sessions	N	N	N
087	Garbage Collection Enhancements	N	-	-
088	TDU Flow Improvements	N	-	-
1202	Safe-Store of Topology DB	N	-	7
Node Operator Command Set				
090	Common Node Operator Command Set	B	N	-
091	Network Node Node Operator Command Set	N	N	-
Intermediate Session Routing				
100	Extended/Unextended BIND and UNBIND	B	N	N
101	Fixed Session-Level Pacing	B	N	N
102	Adaptive Session-Level Pacing	B	N	N
103	Intermediate Session Segmenting/Reassembly	-	N	N
104	Routing BIND and UNBIND	6	N	N
105	Intermediate Session Routing for Dependent LU Sessions	6	N	N
106	Intermediate Session Routing for Type 6.2 LU-LU Sessions	6	N	N
Management Services - Multiple-Domain Support				
150	MDS Common Base	4	N	N
151	MDS End Node Support	4	N	-
152	MDS Network Node Support	4	N	N
153	MDS High Performance Option	4	-	-
154	MDS Transport Confirmation Option	-	-	-
Management Services - MS Capabilities				
160	MS_CAPS Base End Node Support	4	N	-
161	MS_CAPS Have a Backup or Implicit FP	4	N	1
162	MS_CAPS Be a Sphere of Control End Node	4	-	-
163	MS_CAPS Base Network Node Support	4	N	N
164	MS_CAPS Have a Subarea FP	5	-	-
Management Services - Entry Point Alerts				
170	EP Alert Base Subset	4	N	N
171	Problem Diagnosis Data in Alert	-	N	N
174	Operator-Initiated Alert	4	-	-
175	Qualified Message Data in Alert	-	-	-
176	Self-Defining Message Text Subvector in Alert	-	-	-
177	LAN Alert	-	-	-
178	SDLC/LAN LLC Alert	-	-	-
179	X.21 Alert	-	-	-
180	Hybrid Alert	-	-	-

Table 2-3 (Page 3 of 3). Summary of Implemented Base Functions (APPN Version 2)

Function		VTAM	3746	MAS
No.	**Description**	**V4R3**	**V4R2**	**V1R1**
181 182 183 184	X.25 Alert Held Alert for CPMS Resolution Notification Support Operations Management Support in Alert	- - - -	- - - -	- N - -
Miscellaneous				
1013	Interoperability with Peripheral Border Node	N	N	N

notes:

E = End Node

B = End and Network Nodes

N = Network Node

- = Not Supported

1 Backup Focal Point only

2 Supported by NCP's BF only

3 VTAM will respond to a BIND pacing request received, but will never set the pacing request indicator on a BIND.

4 Function supported through NetView

5 VTAM does perform topology database lookup (to see if an unknown resource is a NN CP), but does not skip sending the APPN locate. This locate is then sent as a directed search to the NN. Because of this processing, VTAM has implemented option set 1117.

6 VTAM NN, and VTAM EN providing LEN attachment

7 MAS only

Table 2-4 (Page 1 of 2). Summary of Implemented Optional Functions (APPN Version 2)

Function		VTAM	3746	MRS MAS
No.	**Description**	**V4R3**	**V4R2**	**V1R1**
Configuration Services				
1002	Adjacent Link Station Name	B	N	N
1003	Short-Hold Mode	B	-	-
1006	Dynamic Name Change	B	-	-
1007	Parallel TGs	B	N	N
CP Capabilities				
1011	Multiple Local LUs	B	N	-
1012	LU Name = CP Name	-	N	N
1014	Peripheral Border Node	-	-	-
1016	Extended Border Node	N	-	-
1017	Gateway	-	-	-
1018	Delete EN Resources Before Registering	B	-	-
Dependent LU Support				
1060	Prerequisite for Session Services Extensions CP Support	B	-	-
1061	Prereqs. for SSE NNS Support	N	-	-
1062	Session Services Ext. CP Support	B	-	-
1063	Session Services Ext. NNS Support	N	-	-
1064	Session Services Ext. PLU Node Support	B	-	-
1065	Session Services Ext. CP(SLU) (SSCP) Support	B	-	-
1066	Dependent LU Server	N	-	-
1067	Dependent LU Requester	-	N	N
1071	Generalized ODAI Usage	-	N	N
Cryptography Support				
1070	Session Cryptography	B	-	-
Directory Services				
1100	Safe-Store of Directory Cache	N	-	-
1101	Preloaded Directory Cache	N	N	N
1102	EN Authorization	-	-	-
1106	Central Directory Server	N	-	-
1107	Central Resource Registration (of LUs)	B	N	N
1108	Nonverify	N	-	-
1116	DLUS-Served LU Registration NNS Support	-	N	N
1118	EN TG Vector Registration	B	-	-
Topology and Routing Services				
1200	Tree Caching and TG Caching	N	N	N
1201	Permanent Storage Medium	B	N	N
1203	Detection and Elimination of TDU Wars	-	-	-
Intermediate Session Routing				
1300	Tuning Values for ISR	B	N	-
1301	Nonpaced Intermediate Session Traffic	B	-	-
High Performance Routing				
1400	HPR Base (ANR)	1	N	N
1401	Rapid Transport Protocol	1	N	N
1402	Control Flows over RTP	-	N	N
1403	Dedicated RTP Connections	-	-	-
1404	Multilink TG (MLTG)	-	-	-
Management Services - File Services				
1500	File Services Support Base	-	-.	-
1501	Network Operator Support for File Services	-	-	-

Table 2-4 (Page 2 of 2). Summary of Implemented Optional Functions (APPN Version 2)

Function		VTAM	3746	MRS MAS
No.	**Description**	**V4R3**	**V4R2**	**V1R1**
Management Services - Change Management				
1510	CM Base	-	-	-
1511	CM Production Only Activate	-	-	-
1512	CM Execution Window Timing	-	-	-
1513	CM Activate Report	-	-	-
1514	CM Alter Active Install	-	-	-
1515	CM Object Disposition Install	-	-	-
1516	CM Initiate Command	-	-	-
1517	CM Cancel Command	-	-	-
1518	CM Activate Last	-	-	-
Management Services - Operations Management				
1520	Common Operations Services	-	-	-
1521	Operations Management	[1]	-	-

Notes:

E = End Node

B = End and Network Nodes

N = Network Node

- = Not Supported

[1] Within composite network nodes (CNNs), only ANR base (1400) is supported and only if the RTP path enters and exits the CNN through ANR-capable NCPs. In VTAM NNs (VTAMs with no subarea capability at all), both ANR base (1400) and RTP (1401) are supported. In VTAM ENs (with no subarea capability at all), only RTP (1401) is supported.

Summary of Supported DLCs and APPN Functions

The following table gives an overview of the DLC types supported by APPN capable hardware. The APPN function supported over these DLCs are listed.

Table 2-5. Summary of Supported DLCs and APPN Functions

DLC TYPE	MAE	3746
Token Ring	IHD	IHD
Ethernet	IHD	IHD
Twinax	n/a	n/a
Frame Relay BNN	IHD	IHD
Frame Relay BAN	IHD	IHD
Point-to-Point Protocol	IH	n/a
APPN over DLSw	ID	ID (see note)
HPR over IP	H	H (see note)
SDLC Leased Line	ID	IHD
X.25 PVC and SVC (Previewed)	ID	IHD
APPN over PPP over ISDN	IH	n/a
APPN over Frame Relay over ISDN	IHD	n/a
APPN over ATM Forum Compliant LAN Emulation	IHD	n/a
Native HPR over ATM	H	n/a
APPN over PPP over V.25BIS	IH	n/a
SNA over Asynch	n/a	n/a
ESCON	IH	IH
PARALLEL CHANNEL	n/a	n/a

Note: Need the MAE installed

Legend

I = Intermediate Session Routing

H = High Performance Routing

D = Dependent LU Requester, this refers to the port providing the connection to the downstream PU (DSPU)

Chapter 3. Internet Protocol (IP) Overview

This chapter gives an overview of the Internetwork Protocol (IP). *The Internetwork* is better known as *The Internet*. It is not just one protocol, but a group of protocols that make up an architecture designed for communication between and across diverse network platforms and network architectures. (The word *internet* used here refers to internetwork operations in general rather than the cooperative public domain network known as *The Internet*.)

This chapter describes:

- The need for a means of communicating across multiple network architectures
- How IP meets those needs
- The main features of IP

The Need for Transparency in Communication

The computer networks of large organizations have evolved through the merger, interconnection, and expansion of previously autonomous networks. Such interconnected networks support diverse protocols, multiple speed requirements, and incompatible topologies. For example, suppose there is a community of users, A, B, C, and D, see Figure 3-1, who need to reach each other across four networks.

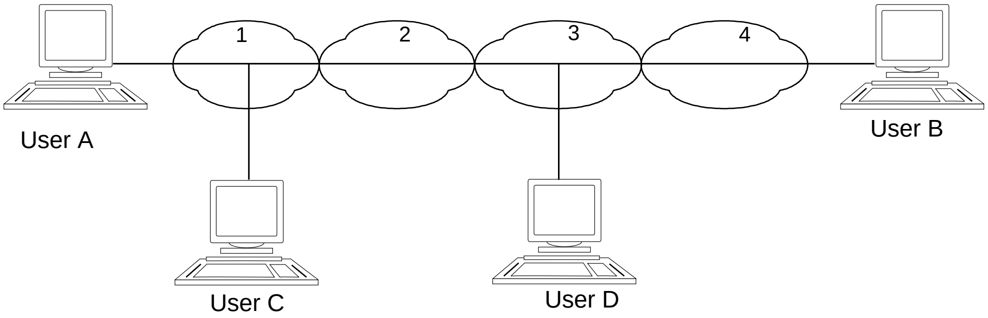

Figure 3-1. Communicating across Multiple Networks

They would each need to know about intervening networks:

- Users A and C only need to know about Network 1.
- Users B and D need to know about each other's networks (3 and 4).
- Users A and C, and User B, need to know about *all* networks between them before they can communicate (1, 2, 3, 4), including two (2 and 3) to which they are not themselves connected.

This is obviously an impractical and unsatisfactory situation.

A solution is one that overlays existing network architectures and provides a single unified view of all networks between any two users or applications.

The *Internet Protocol (IP)* is just such an architecture[1] . IP provides a unified appearance to users and applications independent of the actual networks connecting them by the simple expedient of a common addressing scheme. This relationship of IP to networks is shown in Figure 3-2 on page 3-2.

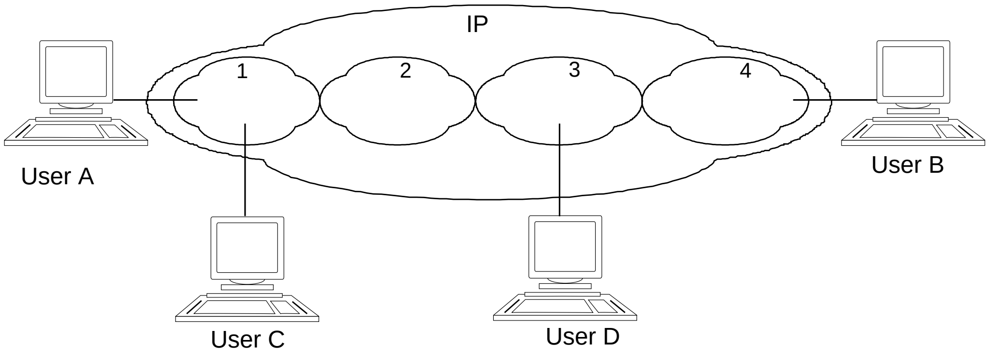

Figure 3-2. The Internet Protocol (IP) Overlays Networks

To communicate with each other now, users need only know the IP address of their target destination. Communication is handled transparently across all intervening networks. A 3746 Nways Multiprotocol Controller with the IP Feature (Feature number 5033) can be used to connect incompatible or independent networks. Its role is to route data from one network to the other, and when used in this context, is called an *IP router*, or *IP gateway*.

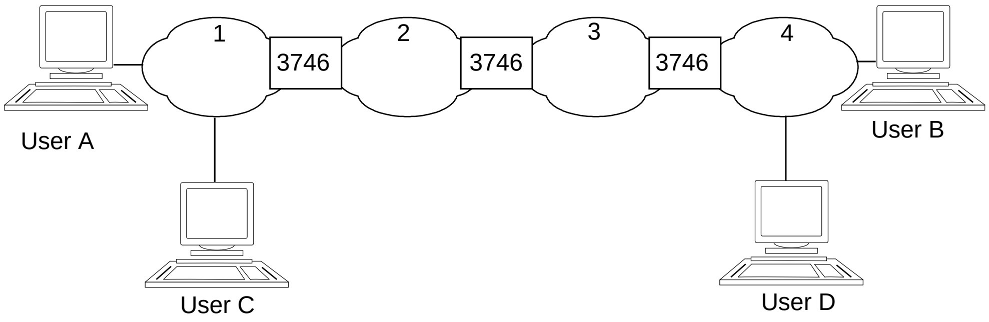

Figure 3-3. 3746 Nways Multiprotocol Controllers with the IP feature (Feature Number 5033). They link the networks and route data between the them.

Physically, users are connected into their local respective networks. The IP addressing scheme is known to the 3746 Nways Multiprotocol Controller IP routers that connect the networks, so they can forward data if the destination address is in a network different to the source address.

IP Addressing

The addressing scheme provides a network-transparent means of communicating. To send data, you only need to know the destination address, regardless of what types of network serve source and destination, and regardless of any intervening type of network.

[1] Internet also refers to an *interconnected network*, a network consisting of two or more otherwise autonomous networks.

The IP Address

IP addresses can be symbolic or numeric. The symbolic form is easier to read; for example:

```
myname@ibm.com.
```

The numeric form is a 32-bit unsigned binary value that is normally expressed in *dotted decimal format*. For example:

```
128.2.7.9
```

is a valid Internet address. Its binary form is:
B'10000000 00000010 00000111 00001001'
The numeric form is used by the IP software, which also provides the mapping from symbolic to numeric forms.

IP Address Classes

The address shown above is an example of one class of IP address but is not suitable for all network configurations. The proportion of networks to hosts can be vastly different from one enterprise to another. IP addresses are thus arranged in *classes* to suit different network characteristics. See Figure 3-4

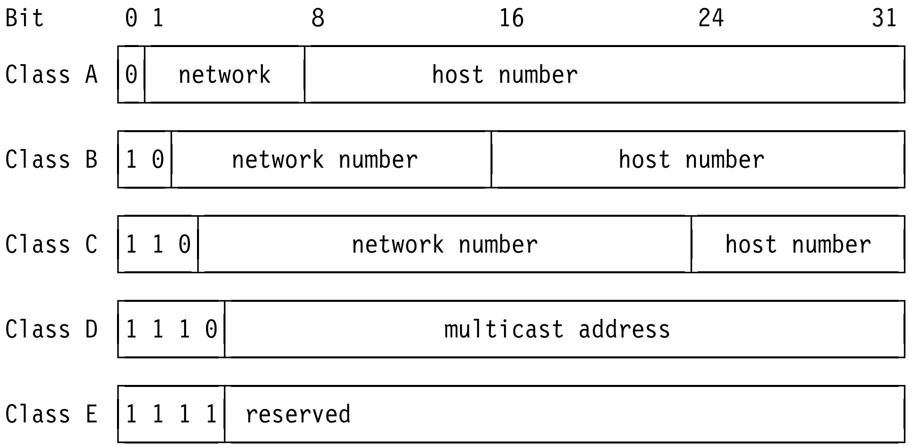
Bit 0 1 8 16 24 31
Class A 0 network host number
Class B 1 0 network number host number
Class C 1 1 0 network number host number
Class D 1 1 1 0 multicast address
Class E 1 1 1 1 reserved

Figure 3-4. Classes of IP Addresses

Each class, having a different proportion of network number to host number, can be assigned to different sizes of network. See Table 3-1.

Table 3-1. Use of IP Address Classes

Class	Maximum Number of Networks	Maximum Number of Hosts per Network
A	126	16 777 214
B	16382	65534
C	2 097 150	254
D	Not applicable	Not applicable
E	Not applicable	Not applicable

Host and Network Numbers

The IP address consists of a pair of numbers:

```
IP address = <network number><host number>
```

The network number is administered centrally by the Internet Network Information Center (the InterNIC) and is unique for each network. The host number (also known as host address or hostID), is assigned to each host within a network and belongs to the authority that owns the network. Table 3-2 shows the components of the example address given in “The IP Address” on page 3-3.

Table 3-2. IP Address Structure

IP address=	**<Network number>**	**<Host number>**
128.2.7.9	128.2	7.9

IP Addresses for Routers

A router appears on both networks it connects. It must have a network number and a host number to make a complete IP address. For example, suppose a router routes data between networks 128.2 and 131.5.

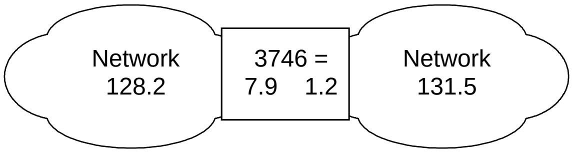

Figure 3-5. Routers Need Two Addresses

It must only have one address when seen from each network, but each network has its own network number. Thus, the router must have two addresses, one for each network it connects. Its resulting addresses are shown in Table 3-3.

Table 3-3. IP Addresses for a Router

IP address	**Network number**	**Host number**
128.2.7.9	128.2	7.9
131.5.1.2	131.5	1.2

Note: A *router* is not the same as a *bridge*, despite the apparent similarity of their positions between networks.

Subnets

With the massive growth of networks, the use of assigned IP addresses becomes too inflexible to allow easy changes to local network configurations. These changes can occur when:

- A new physical network is installed.
- Growth of the number of hosts requires splitting the local network into two or more separate networks.
- Reorganization of an enterprise requires new address structures.

To avoid having to request additional IP network addresses in these cases, the concept of *subnets* was introduced. The host number part of the IP address is subdivided again into a network number and a host number. This second network is termed a *subnetwork* or *subnet*.

Subdividing the Network

The main network now consists of a number of subnets and the IP address is interpreted as:

```
IP address=<network number> <<subnet number><host number>>
```

The combination of the subnet number and the host number is often termed the *local address* or the *local part*.

Subnetting is transparent to remote networks. A host within a network which has subnets is aware of the subnetting but a host in a different network is not. It still regards the local part of the IP address as a host number.

Administrators can assign subnet addresses to represent, for example:

- Groups of similar types of host
- Departmental clusters of hosts
- Hierarchical or security level of the host users

Subnets provide a customizable address system within the IP address scheme.

Note: They do not have to represent the physical network connections, as it is a purely *logical* addressing system.

The Subnet Mask

The division of the local part of the IP address into subnet number and host number parts can be chosen freely by the local administrator. Any bits in the local part can be used to form the subnet. The division is done using a *subnet mask* which is a 32-bit number:

- Ones (1) indicate bit positions assigned to the subnet number.
- Zero (0) bits in the subnet mask indicate bit positions assigned to the host number.

The bit positions in the subnet mask belonging to the network number are set to ones but are not used. Subnet masks are usually written in dotted decimal form, like IP addresses. The special treatment of *all bits zero* and *all bits one* applies to each of the three parts of a subnetted IP address just as it does to both parts of an IP address that has not been subnetted.

Example Subnet Mask

For example, a subnetted Class B network, which has a 16-bit local part, could use one of the following schemes:

- The first byte is the subnet number, the second the host number. This gives us 254 (256 minus 2 with the values 0 and 255 being reserved) possible subnets, each having up to 254 hosts. The subnet mask is `255.255.255.0`.

- The first 12 bits are used for the subnet number and the last four for the host number. This gives us 4094 possible subnets (4096 minus 2) but only 14 hosts per subnet (16 minus 2). The subnet mask is `255.255.255.240`.

There are many other possibilities.

Building Good Masks

While the administrator is completely free to assign the subnet part of the local address in any legal fashion, the objective is to assign a number of bits to the subnet number and the remainder to the local address. Therefore, it is normal to use a contiguous block of bits at the beginning of the local address part for the subnet number because this makes the addresses more readable (this is particularly true when the subnet occupies 8 or 16 bits). With this approach, either of the subnet masks above are good masks, but masks such as `255.255.252.252` and `255.255.255.15` are not.

Multiple Destinations

The point of the IP addressing scheme is to get messages, whatever their content, from one user to another. There are three ways this might be done:

- Unicasting. The majority of IP addresses refer to a single recipient. These are called *unicast addresses*.
- Broadcasting. Messages are sent out to every user on the network.
- Multicasting. Messages are broadcast to a select group of users.

There are two special types of IP address that are used for addressing multiple recipients:

- Broadcast addresses
- Multicast addresses.

These addresses are used for sending messages to multiple recipients. Each requires different handling from the point of view of IP addressing.

Broadcasting

There are a number of addresses that are used for IP broadcasting. All use the convention that all-bits-1 indicates all-addresses. Broadcast addresses are never valid as source addresses, only as destination addresses. The different types of broadcast address are listed here:

Limited broadcast address

The address 255.255.255.255 (1s in all bits of the IP address) or 0.0.0.0 (0s in all bits of the IP address) is used on networks that support broadcasting, such as token-rings, and it refers to all hosts on the subnet. It does not require the host to know any IP configuration information at all. All hosts on the local network will recognize the address, but routers will never forward it.

Network-directed broadcast address

If the following conditions are true:

- The network number is a valid network number
- The network is not subnetted
- The host number is all ones or all zeros

then the address, for example, `128.2.255.255`, refers to all hosts on the specified network `128.2`. Routers should forward these broadcast messages unless configured otherwise.

Subnet-directed broadcast address

If the following conditions are true:

- The network number is a valid network number
- The subnet number is a valid subnet number
- The host number is all ones or all zeros

then the address refers to all hosts on the specified subnet. Since the sender's subnet and the target subnet may have different subnet mask, the sender must somehow find out the subnet mask in use at the target. The actual broadcast is performed by the router which receives the datagram into the subnet.

Multicasting

Broadcasting has a major disadvantage - its lack of selectivity. If an IP datagram is broadcast to a subnet, every host on the subnet will receive it, and have to process it to determine whether the target protocol is active. If it is not, the IP datagram is discarded. Multicasting avoids this overhead by using groups of IP addresses. Each group is represented by a 28-bit number, which is included in a Class D address. See Figure 3-4 on page 3-3.

Multicast group addresses are IP addresses in the range 224.0.0.0 to 239.255.255.255. For each multicast address there is a set of zero or more hosts which are listening to it. This set is called the *host group*. There is no requirement for any host to be a member of a group to send to that group.

There are two kinds of host groups:

Permanent

The IP address is permanently assigned. The membership of a host group is not permanent; a host may leave or join the group at will. Some IP addresses are assigned permanently to host groups. Important ones are:

- 224.0.0.0 - Reserved base address
- 224.0.0.1 - All systems on this subnet
- 224.0.0.2 - All routers on this subnet

A permanent group exists even if it has no members.

Transient

Any group which is not permanent is transient and is available for dynamic assignment as needed. Transient groups cease to exist when their membership drops to zero.

Domains

The system of symbolic names (see “The IP Address” on page 3-3,) is more convenient to use than the IP numeric address structure. However, this requires:

- The mapping of all host names to their numeric addresses
- The maintenance of a database in all hosts of such a mapping
- Considerable work in sending changes in the database to all possible hosts

In all but the smallest networks this is an almost impossible task. An alternative method is required.

A new concept was developed to handle this problem. The whole internet community was divided into *domains*, and the mapping of symbolic names to IP addresses called the *Domain Name System*.

The Domain Name System allows a program running on a host to perform the mapping of a high-level symbolic name to an IP address for any other host. It does, however, require a formalized use of symbolic names if mapping is to be successful.

The Hierarchical Name Space

Consider the internal structure of a large organization. As the chief executive cannot do everything, the organization will probably be partitioned into divisions, each of them having autonomy within certain limits. Specifically, the executive in charge of a division has authority to make direct decisions, without permission from the chief executive. Domain names are formed in a similar way, and will often reflect the hierarchical delegation of authority used to assign them.

For example, consider the name `lcs.mit.edu`. The lowest-level domain name, `lcs`, is a subdomain of `mit`, which again is a subdomain of `edu` (education) which is called a top-level domain. We can also represent this naming concept by a hierarchical tree. See Figure 3-6.

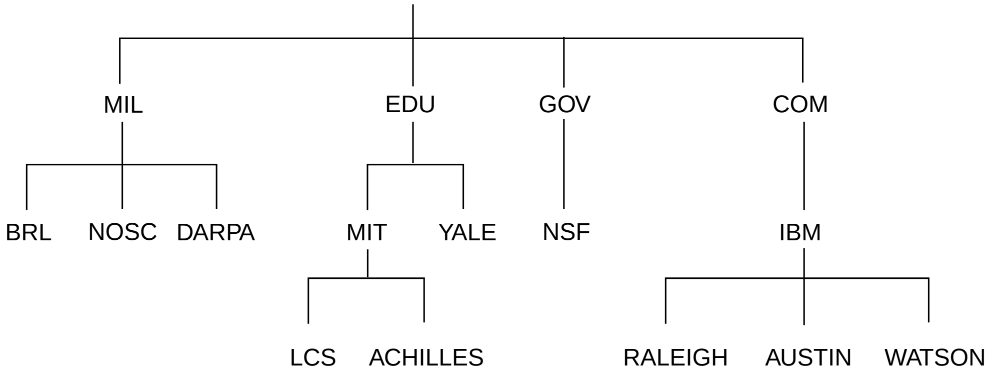

Figure 3-6. Hierarchical Structure of Domains

Figure 3-6 shows the chain of authority assigning domain names. This tree is only a tiny fraction of real namespace. Table 3-4 on page 3-9 shows some of the top-level domains. The single domain above the top-level domains has no name and is referred to as the root domain.

Fully Qualified Domain Names (FQDNs)

When using the Domain Name System, it is common to work with only a part of the domain hierarchy, for example the `ral.ibm.com` domain. The Domain Name System provides a simple method of minimizing the typing necessary in this circumstance. If a domain name ends in a dot, for example, `wtscpok.itsc.pok.ibm.com.` it is assumed to be complete. This is termed a Fully Qualified Domain Name (FQDN) or an absolute domain name. If, however, it does not end in a dot, for example `wtscpok.itsc` it is incomplete and will be completed by appending a suffix such as `.pok.ibm.com` to the domain name. The rules for doing this are implementation dependent and locally configurable.

Generic Domains

The three-character top-level names are called the *generic domains* or the *organizational* domains.

Table 3-4. The Generic Top-Level Domains

Domain Name	Meaning
edu	Educational institutions
gov	Government institutions
com	Commercial organizations
mil	Military groups
net	Networks
int	International organizations
org	Other organizations

The organization of the hierarchical namespace initially had only US organizations at the top of the hierarchy, and it is still largely true that the generic part of the namespace contains US organizations. However, only the .gov and .mil domains are restricted to the US.

Country Domains

There are also top-level domains named for each of the ISO 3166 international two-character country codes, for example .uk for the United Kingdom and .nl for the Netherlands. These are called the *country domains* or the *geographical domains*.

Many countries have their own second-level domains which parallel the generic top-level domains. For example, in the United Kingdom, the domains equivalent to the generic domains .com and .edu are .co.uk and .ac.uk (ac is an abbreviation for academic).

3746-9x0 Router Node Implementation

The 3746 Nways Controller Models 900 and 950 provide the following set of IP functions:

- IP routing providing connectivity over:
  - ESCON channels:
    To IBM and other equipment manufacturers (OEM) processors that adhere to the ESCON architecture, possibly via one or two cascaded ESCON directors (IBM 9032 or 9033).
  - Token-ring LANs
  - Ethernet LANs:
    Ethernet access requires the 3746-9x0 Ethernet feature. See Chapter 17, “Ethernet Adapters” on page 17-1 and “Ethernet Port Specifications (Features 5631 and 5632).” on page 44-43 for detailed information on this feature.
  - Leased lines:
    Leased lines using either the PPP or frame relay protocol.

  - Frame relay and X.25 networks.

- Dynamic routing protocol support, using:
  - RIP Version 1.
  - OSPF Version 2.
  - BGP Version 4.
- ARP and proxy ARP.
- BOOTP relay.

If you use only one type of adapter, the maximum connectivity is up to:

- 16 ESCON channel couplers
- 32 token-ring couplers
- 32 line interface couplers, with 120 lines (of any permissible speed).

Overview of 3746 IP Routing

First, it is necessary to understand the model on which IP is based. Each station has to be attached to a least one network. Most often this is a local area network attachment, but stations might be attached via a serial line (PPP, or frame relay), or, for that matter, be ESCON attached.

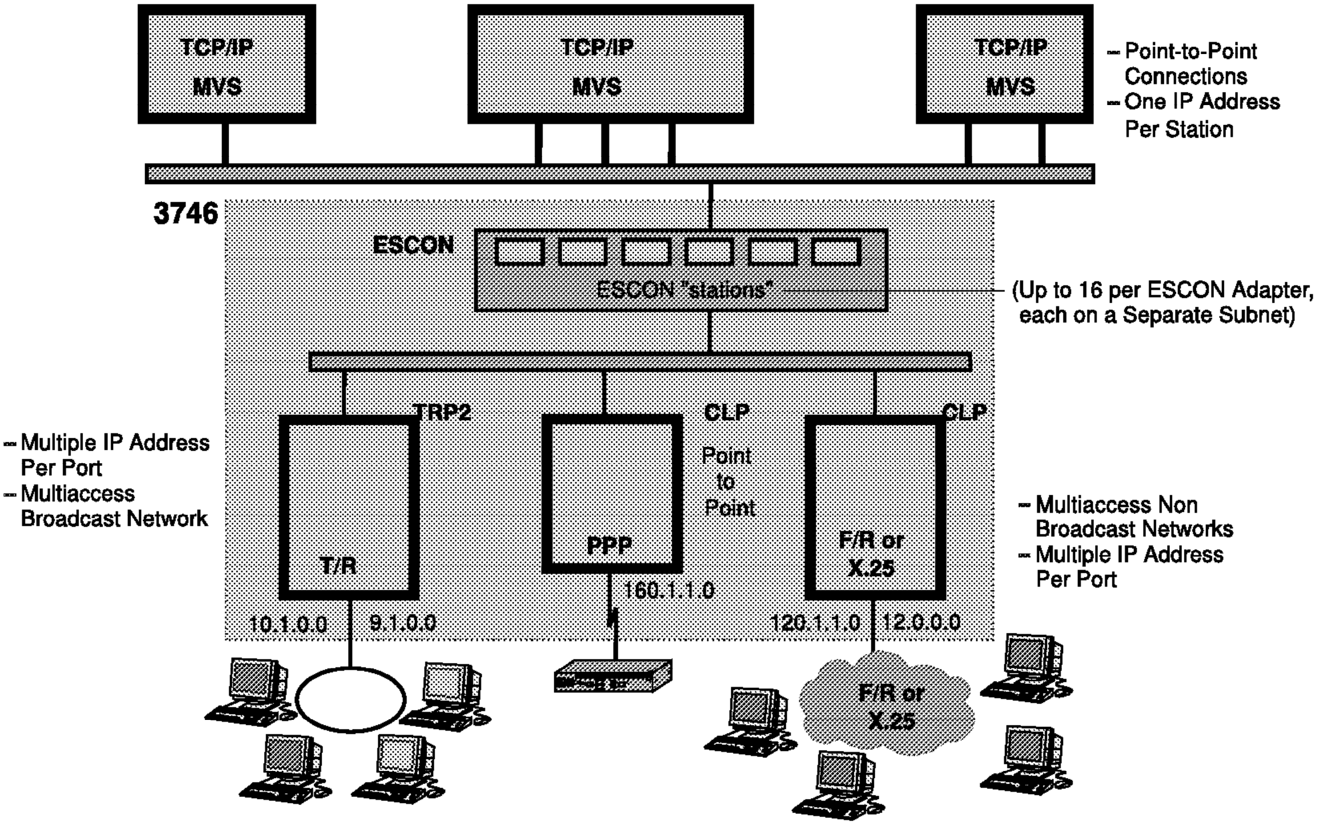
TCP/IP
MVS
TCP/IP
MVS
TCP/IP
MVS
– Point-to-Point Connections
– One IP Address Per Station
3746
ESCON
ESCON "stations"
(Up to 16 per ESCON Adapter, each on a Separate Subnet)
TRP2
CLP
CLP
– Multiple IP Address Per Port
– Multiaccess Broadcast Network
T/R
Point to Point
PPP
F/R or X.25
– Multiaccess Non Broadcast Networks
– Multiple IP Address Per Port
160.1.1.0
10.1.0.0
9.1.0.0
120.1.1.0
12.0.0.0
F/R or X.25

Figure 3-7. IP Address Support Example

Stations can send datagrams to any other system on its own network.

Networks can be:

- Multiaccess broadcast networks

  Multiaccess broadcast networks allow the attachment of multiple stations. To send data to a destination node attached to the same network, the physical address of the node must be learned, and direct communication can take place. Broadcasting allows a node to send a datagram to multiple stations at the same time.

  Examples of multiaccess broadcast networks are token ring, Ethernet, and FFDI[2].

- Multiaccess nonbroadcast networks

  Multiaccess nonbroadcast networks also allow the attachment of multiple stations. To send data to a destination node attached to the same network, the physical address of the node must be learned, and direct communication can take place. In general, unless special provisions are taken, broadcasting is not possible. Examples of multiaccess nonbroadcast networks are frame relay, and X.25.

- Point-to-point (PtP) networks

  Point-to-point networks allow direct communication between two adjacent nodes. No broadcast is required on PtP connections From an IP routing perspective, PtP connection comprise a separate network, although only two stations connect. Instead of the term PtP network, the term PtP connection is more commonly used. Examples of point-to-point networks are PPP, SLIP[2], SNAlink, and ESCON connections

IP General Functions

Broadcasting

The 3746 IP Router is capable of using one the following types of broadcast:

- Limited or local-wire broadcast.

  Uses either all zeros (0.0.0.0), or all ones (255.255.255.255) as the broadcast address.

- Subnet-directed broadcast.

  Uses a broadcast address comprised of the subnetted network identifier and a host identifier of either all zeros or all ones. (For example, 9.132.56.0 or 9.132.56.255).

During customization either local-wire or subnet-directed broadcasting is selected, using either all ones or all zeros

[2] Not supported by the 3746.

The 3746 IP Router will accept the following types of broadcasts:

- Limited or local-wire broadcast.

  Destination IP address 255.255.255.255, and source IP address 0.0.0.0 are accepted.

- Subnet-directed broadcast.

  Destination IP address containing the subnetted network identifier and a host identifier of either all zeros or all ones, will be accepted, if the source IP address belongs to the network interface on which it has been received.

- Network-directed broadcast.

  Destination IP address containing the (non-subnetted) network number and a host number of either all zeros or all ones, will be accepted, if the source IP address belongs to the network interface on which it has been received (for example, 9.255.255.255 or 9.0.0.0).

The 3746 IP Router never forwards limited broadcast. Directed broadcasts will be forwarded, if not received from the network to which it is directed, and only if the packet was not received as a link-level (for example token-ring LAN) broadcast. Directed broadcasts are forwarded as a link-level broadcast. The 3746 contains an option to disable the forwarding of directed broadcasts.

Supernetting (Route Aggregation)

There is a major problem with the use of a range of Class C address instead of a single Class B addresses: each network must be routed separately. Standard IP routing understands only the class A, B and C network classes. Within each of these types of network, subnetting can be used to provide better granularity of the address space within each network, but there is no way to specify that multiple class C networks are actually related. The result of this is termed *the routing table explosion problem:* a Class B network of 3000 hosts requires one routing table entry at each backbone router, but if the same network is addressed as a range of Class C networks, it requires 16 entries.

The solution to this problem is a scheme called *Classless Inter-Domain Routing (CIDR).* CIDR is a proposed standard protocol with a status of elective. CIDR does not route according to the class of the network number (hence the term classless) but solely according to the high order bits of the IP address, which are termed the *IP prefix.* Each CIDR routing entry contains a 32 bit IP address and a 32-bit network mask, which together give the length and value of the IP prefix. This can be represented as <IP_address network_mask>. For example <194.0.0.0 254.0.0.0> represents the 7 bit IP prefix B '1100001'. CIDR handles the routing for a group of networks with a common prefix with a single routing entry. This is the reason why multiple Class C network numbers assigned to a single organization have a common prefix. This process of combining multiple networks into a single entry is termed *address aggregation* or *address summarization*. It is also called *supernetting* because routing is based upon network masks that are shorter than the natural network mask of the IP address, in contrast to subnetting where the network masks are longer than the natural mask.

Unlike subnet masks, which are normally contiguous but may have a non-contiguous local part, supernet masks are always contiguous. The 3746 IP Router uses CIDR in its BGP-4 implementation.

Multiple IP Addresses for a Network Interface

The 3746 IP Router supports class A, B, C, and D addresses. Class A, B, and C addresses are unicast addresses. Class D addresses are for multicast services, for example by the OSPF routing protocol.

One or up to 16 IP addresses can be assigned to a network interface for token-ring and frame relay. A subnet mask must be assigned together with each IP address. The IP addresses must be unicast addresses. Two reasons to use multiple addresses per interface are:

1. Consolidation of two routers into one with no change to the host definitions.

   Figure 3-8 shows a situation that requires three routers to handle communication between three IP networks.

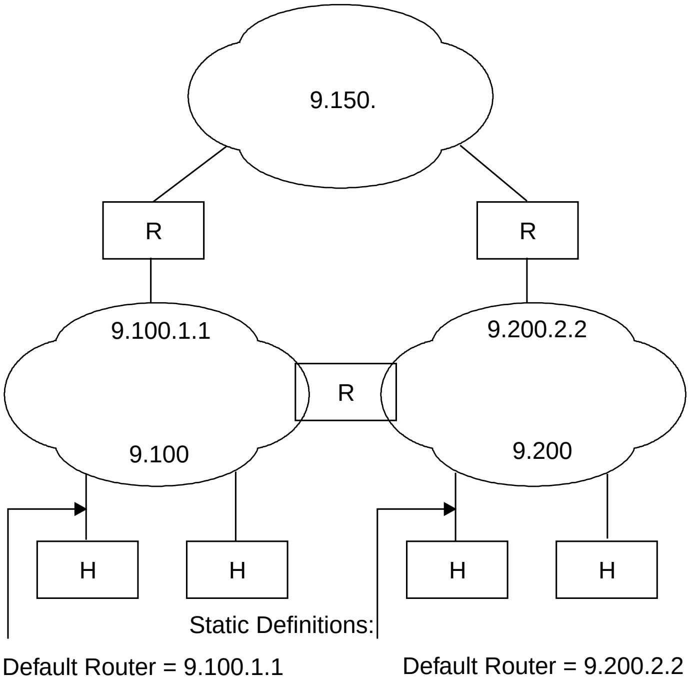

Figure 3-8. Three Networks and Three Routers Before Consolidation of the Routers

Figure 3-9 shows the situation after consolidation into one 3746 IP Router. Users see the network addresses as before and require no definition changes.

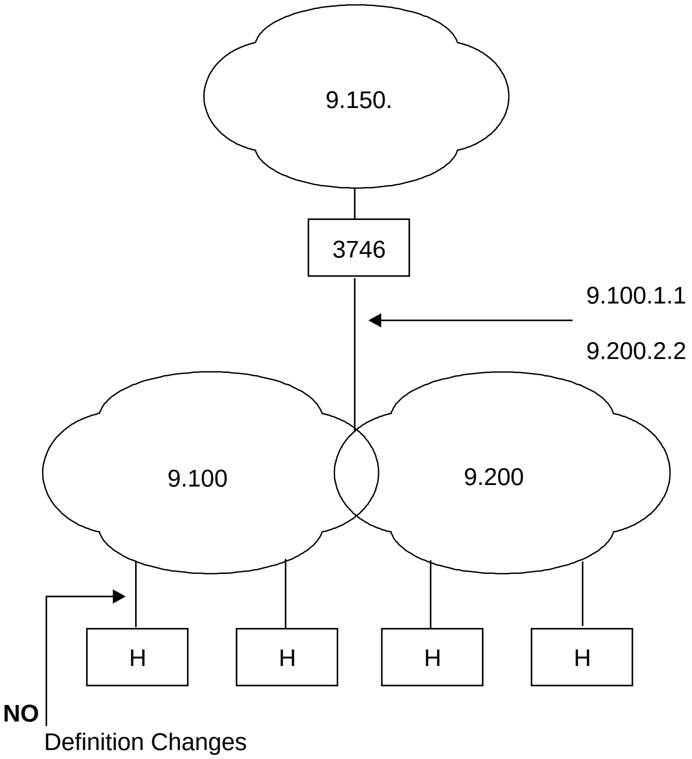

Figure 3-9. Three Networks and One 3746 IP Router After Consolidation

2. Route aggregation of multiple Class C addresses.

   Suppose an enterprise requests 1000 new IP host addresses for its internal network. They are given *four* contiguous ranges of class C addresses, but the four networks are connected to a single port on the 3746 IP Router. Figure 3-10 shows that the single port can be assigned four IP addresses, one for each IP network. To each network, the 3746 IP Router appears as a single IP address.

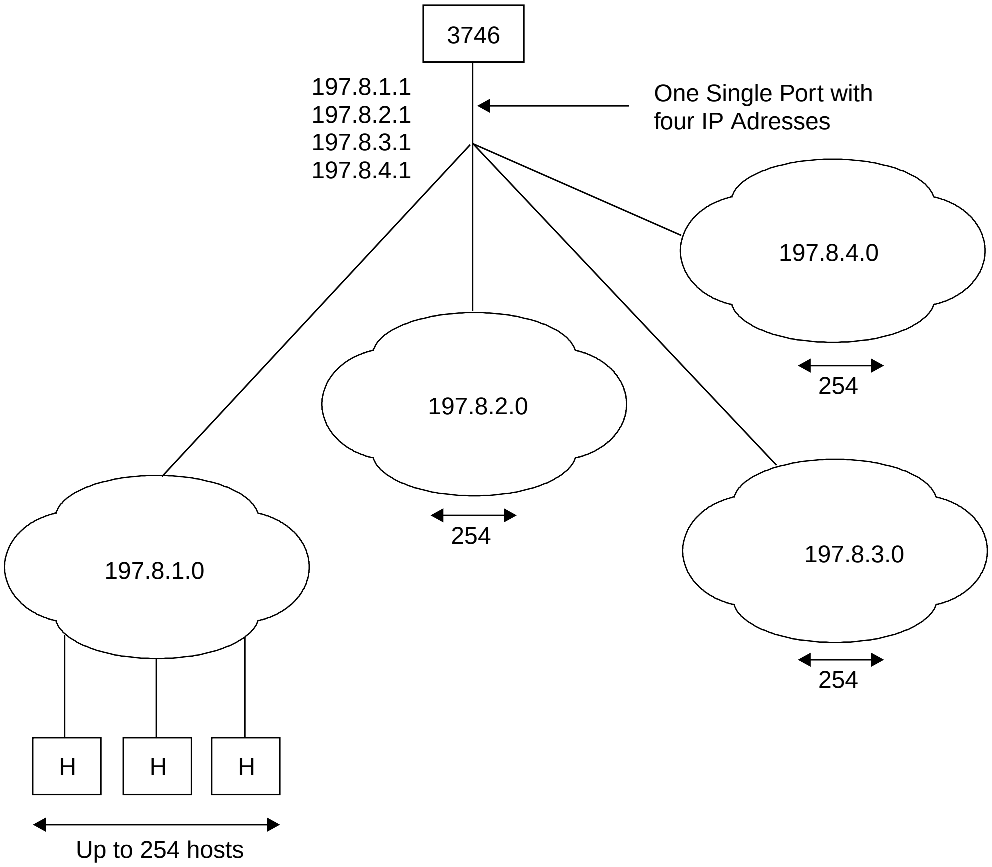

Figure 3-10. A Single 3746 IP Router Serving Four IP Networks via a Single Port

It is not necessary to assign an IP address to the 3746 PPP interface, but frame relay and X.25 interfaces do need at least one IP address. However, unnumbered interfaces (that is, without IP address) cannot be addressed (for example, Ping and Traceroute cannot be directed to the interface) and require OSPF as a interior routing protocol.

In addition to the interface addresses a *router_ID* must be assigned. Rather than specifying a specific interface, this IP address specifies the router as a whole. The router_ID is used for Ping, Traceroute, and OSPF messages originating from the 3746 IP Router. To avoid routing problems, it is recommended to set the router_ID to one of the 3746 interface IP addresses.

User Datagram Protocol (UDP)

The 3746 IP Router uses UDP for RIP and/or SNMP data transport. If a UDP packet with an unknown (unregistered) destination port is received, this event is logged. For destination ports other than 9 (RFC863, *Discard protocol*) an ICMP destination unreachable message (port unreachable) is returned, unless the packet was a link-level broadcast or the source IP address was a broadcast or multicast IP address.

Transmission Control Protocol (TCP)

The 3746 IP Router uses TCP for its Telnet, SNMP, and BGP functions. Note that OSPF uses IP directly, it does not pass through TCP (nor UDP).

Internet Control Message Protocol (ICMP)

The 3746 IP Router recognizes the following ICMP types:

0 Echo Reply

When an Echo request is received, an Echo reply is returned containing the data from the Echo request.

3 Destination Unreachable

The 3746 may generate the following ICMP destination unreachable messages:

- Network unreachable.

  When there is no active route to the destination address, the destination network is zero, or the destination address is not a valid class A, B, C, or D address.

- Fragmentation needed [the Don't Fragment (DF) bit is set].

  An outgoing packet needs fragmentation but the DF bit in its IP header is set, or fragmentation results in an invalid fragment offset.

- Protocol unreachable.

  An IP datagram destined for the 3746 IP Router contains an unsupported protocol number.

- Port unreachable.

  An UDP datagram destined for the 3746 IP Router contains an unsupported UDP port number.

- Source route failed.

  The next hop IP address in a datagram strict source routing is not on a directly connected network.

4 Source quench

Source quench messages are logged but no action is taken. The 3746 never generates a source quench message.

5 Redirect

Redirect messages are logged but no action is taken. The 3746 generates a redirect message (redirect datagram for the host) when a packet is forwarded on the same interface as received, the network is not point-to-point, and the subnet or network number of the source IP address matches that of the next hop IP address.

8 Echo Request

Echo request messages are responded with the echo reply message. ICMP Echo request are generated by the Ping application (see “Ping” on page 3-24) at the user's request.

11 Time exceeded

Time exceeded messages are logged. The 3746 generates a time exceeded message (time to live exceeded in transit) if the time to live (TTL) field in a datagram that the 3746 IP Router is forwarding is 0 or 1.

12 Parameter problem

Parameter problem messages are logged. The 3746 generates a parameter problem message (pointer indicates the error), if:

- Options field in IP header indicates that the options extend beyond the end of the IP header.
- Pointer is less than 4 in loose source route, strict source route, or record route option. Pointer is less than 5 in a timestamp option.
- Pointer is less than the length in a record route or timestamp option, but insufficient room for the IP router to insert its IP address and/or timestamp.
- In a record route option, the option is full (no data can be added).

13 Timestamp

Timestamp messages are logged. The 3746 never generates or responds to timestamp messages.

14 Timestamp reply

Timestamp reply messages are logged. The 3746 never generates or responds to timestamp reply messages.

15 Information request

Information request messages are logged. The 3746 never generates or responds to information request messages.

16 Information reply

Information reply messages are logged. The 3746 never generates information reply messages.

17 Address mask request

Address mask request messages are dropped if:

- The destination IP address matches the 3746's internal IP address.
- The source IP address is 0.0.0.0 and the receiving interface has more than one IP address assigned to it.
- The source IP address is 0.0.0.0 and the receiving network is not broadcast or point-to-point. The 3746 never generates address mask request messages.

18 Address mask reply

Address mask reply messages are returned in response to valid address mask request messages. An address mask reply message is returned with the destination IP address set to 255.255.255.255 if the (request) source IP address is 0.0.0.0; otherwise, the (reply) destination IP address is set to the (request) source IP address.

When the 3746 IP Router receives an ICMP message with a Type value it does not recognize, it logs the event. If the Type is greater than 18, it increments the count of erroneous received ICMP messages. The 3746 keeps transmit and receive counts per-Type, for values less than 19.

The TTL value in the IP header of the ICMP messages generated by the 3746 IP Router is equal to the value it generally sets for all IP datagram, *60*. The precedence and Type of Service (TOS) fields are set to zero. No options are included. No mechanism exists to limit the rate at which ICMP messages are sent.

Address Resolution Protocol (ARP)

The 3746 IP Router supports ARP requests and responses. An ARP request is sent when the requested address is not found in the ARP cache table. Requests are broadcast on the local physical network. An ARP request is also sent when attempting to refresh an existing entry. In this case, the ARP request is directed to the current entry's hardware address.

A received ARP request contains the requester's protocol and hardware address, and the destination IP address. ARP responds by providing the 3746 IP Router's IP address of the interface on which the ARP request is received. The requester's address translation is also gleaned from the received ARP request and an entry made in the ARP table cache.

An ARP response is sent when a valid ARP request is received. An unsolicited ARP response is broadcast when a new IP address pair (resulting from new local hardware) is registered, or a previously registered pair is changed. When an ARP response is received, the address translation information is cached in the ARP cache.

ARP Caching

To improve performance ARP entries are cached in the ARP table. Each 3746 CLP maintains independent ARP tables. The table contains dynamic and permanent entries. Dynamic entries are learned, while permanent entries are operator defined.

Each ARP entry contains a 6-byte token-ring LAN address and protocol-specific data referred to a *side-car.* Side-car information maintained for IP over token-ring is the routing information field (RIF). On the 3746 IP Router, the support for source-route bridging (SRB) is always active.

ARP table entries that have not been recently used or updated (refreshed) are discarded. Permanent entries are never discarded. The default time for an entry to be retained before being discarded is five minutes. User commands exist to delete all (dynamic), or specific ARP (dynamic and permanent) entries. Dynamic ARP table entries are also removed when a network interface is considered down.

The 3746 IP Router can be configured to attempt to refresh an unused entry before it ages out. This is referred to as the *auto-refresh* function. If configured, ARP will send out a request to the hardware address in the table to verify the address. The request is sent out approximately 30 seconds before an entry ages out. The default is to disable this function. The ARP table sizes are limited by available memory. ARP will continue to request memory for available entries until 3746 processor memory is depleted, at which time the IP routing function is restarted.

Care should be taken in setting the aging timers, as this could impact the ARP table size. Special care should also be taken in disabling this timer.

Proxy-ARP

The 3746 IP Router supports proxy-ARP. By user configuration, either RFC 925 or RFC 1027 can be selected. RFC 925 proxy-ARP is referred to as *ARP net routing*. RFC 1027 proxy-ARP is referred to as *ARP subnet routing*. The proxy-ARP support is optional; it can be enabled or disabled for the whole 3746 IP, but it cannot be enabled or disabled on an interface basis. It is disabled by default. Proxy-ARP is an appendage to ARP. When enabled and an ARP request is received in which the destination address does not match any of the router's IP addresses, proxy-ARP is called. Proxy-ARP looks up the destination protocol address in the routing table, and, in general, if it finds an active route to the destination, it tells ARP to send an ARP reply. No ARP reply is returned, if:

- There is no route to the destination IP address in the ARP request.
- The source IP address in the ARP request is not in a network directly attached to the 3746 IP Router.
- Source and destination IP address do not have the same network identifier.

  Note: An option exists to disable this option.
- The physical transmit interface of the route to the destination IP address is the same interface from which the ARP request is received.
- The route found for the destination protocol address is a filter.

Routing Protocols

The IBM 3746-9x0 supports the following IP routing protocols:

- Routing Information Protocol, Version 1 (RIP)
- Open Shortest Path First Protocol, Version 2 (OSPF)
- Border Gateway Protocol Version 4 (BGP)

Each is explained in more detail below.

Routing Information Protocol (RIP)

The IBM 3746-9x0 implementation of RIP Version 1 includes support for *split horizon* and *poison reverse*. The two are mutually exclusive. One of the two is always in effect. To speed network convergence, triggered updates are sent.

Route Acceptance Policy: Per interface the user can configure, whether or not:

- RIP will listen.
- Network, subnet, and/or host routes are learned.
- Default and/or static routes can be overridden with a learned route.

Note: Default routes are indicated by a destination and mask of 0.0.0.0.

Route Advertisement Policy: Per interface the user can configure, whether or not:

- RIP will be advertised.
- Network, subnet, and/or host routes are advertised.
- Default and/or static route are advertised.

Subnets will be advertised if they are part of the same natural network (that is, class A, B, or C) as the interface's IP address, and the route subnet mask must be the same as the interface's subnet mask. Host routes (mask 255.255.255.255) are exempt from these rules.

Open Shortest Path First Protocol (OSPF)

The 3746 IP Router supports the following features of OSPF, Version 2, which have implementation-specific behavior:

- Area border (AB) router support
- Support for stub areas
- Autonomous system border (ASB) support
- OSPF interface support
- Equal-cost multipath routing
- Simple authentication
- OSPF routing policies
- Multicast OSPF (MOSPF) support.

There is no support for Type of Service (TOS) based routing; that is, TOS 0 is the only supported TOS.

Area Border (AB) Router Support: 3746 IP supports attachment to multiple areas and summarization of routing information between areas. Area border routers must attach to the backbone (0.0.0.0) and at least one other area. Summarization information from one area will manifest itself as type 3 and 4 link state advertisements (LSAs) in other areas.

Support For Stub Areas: 3746 IP supports attachment to stub areas. OSPF autonomous system external (ASE) LSAs will not be advertised into stub areas. Rather a type 3 network summary LSA for the default route (destination/mask 0.0.0.0) is generated.

Autonomous System Border (ASB) Support: 3746 IP can be configured as an autonomous system border (ASB) router. Non-OSPF routes can be imported into OSPF as OSPF autonomous system externals (ASEs). This implies that the 3746 IP will generate type 5 LSAs.

OSPF Interface Support: 3746 IP supports the following types of OSPF interface:

- Numbered point-to-point (PtP)

  Numbered PtP connections are links to which an IP address has been assigned. Examples are ESCON and PPP connections.

- Unnumbered point-to-point (PtP)

  Unnumbered PtP connections are links to which no IP address has been assigned. OSPF packets will be sent using the 3746 IP Router-ID as source address.

- Broadcast

  A broadcast is, for example, a token-ring interface. Link-level multicast is used to broadcast OSPF frames to all attached OSPF routers.

- Nonbroadcast multiaccess (NBMA).

  Viewing a network, for example a frame relay network, as an NBMA network can be used when the network is fully meshed (meaning virtual circuits exist between any pair of routers). In routers that are eligible to become designated routesr, neighbors must be configured whether or not the neighbor is eligible to become designated router. The configurable poll interval defines the interval at which the designated router will attempt to contact the neighboring routers to establish an adjacency. NMBA connections are supported for frame relay networks. If the network is partially meshed, it is more useful to view the network as a point-to-multipoint network.

- Point-to-multipoint (PtM)

  This interface type is used to allow BMBA topologies to be non-fully meshed. Rather than electing a designated router for the network and having that router generate network LSAs for the network, each router includes its neighbors in its router LSAs. When the route table is calculated, the network topology will appear as multiple point-to-point links rather than a single cloud. On one side of the frame relay virtual circuit, the neighboring router must be configured to allow the two routers to form an OSPF adjacency. Point-to-multipoint connections are supported for frame relay networks.

- Virtual link

  Virtual links are supported to extend the backbone area's connectivity through a transit area. The two endpoints of the virtual link are area border routers.

Equal-Cost MultiPath Routing: 3746 IP supports up to four equal-cost next hops for a route. When multiple next hops exist, the traffic to the destination is spread over the next hops round-robin.

Simple Password Authentication: 3746 IP supports both simple password and no authentication. When simple password authentication is used, an 8-byte password is included in each OSPF packet. Upon reception, this password is validated with packets failing validation being dropped. The authentication type (simple or none) is configured on the area level, while the authentication key is configured for each interfaces in areas with simple password authentication enabled.

OSPF Routing Policies: The 3746 IP OSPF policy can be explained in terms of rules for:

- Advertisements of OSPF routes
- Import of external routes into OSPF as OSPF AS external (ASE) routes
- Generation of the default route and import as an OSPF ASE LSA
- OSPF route policy when multiple routes to a destination exist

OSPF Advertisement of Routes: OSPF allows filtering of LSAs on area boundaries only. All routers within an area have the same view of the area topology. To limit the number of LSAs advertised outside the area, one can configure the network ranges associated with an area at the area boundary. This, in effect, will aggregate a number of networks into a single advertisement. The cost associated with the network range will be the lowest for any of the component networks. Additionally, one can define a network range that will not be visible. Normally, OSPF ASE LSAs are flooded throughout the entire routing domain. One can prevent this for a given area, by defining the area as a stub area. Within the stub area, the area border router will advertise single default route (destination/mask 0.0.0.0).

Importing Routes into OSPF: The 3746 IP Router configuration determines whether or not non-OSPF routes can be imported as OSPF AS external routes, and advertised as OSPF ASE LSAs throughout the OSPF routing domain. Imported routes can be imported as ASE type 1 (router) or ASE type 2 (network) routes. ASE type 1 route always override type 2 routes. It has a single metric that is the sum of the path cost to the AS border router, and the AS border router metric. Conversely, the metric for a type 2 ASE has internal and external components, namely the path cost to the AS border router and the route's external metric. When comparing ASE type 2 routes to the same destination, the one with the lower external metric will be always be preferred, independent of the internal metric. The metric used for both the OSPF ASE type 1 route and the external component of the type 2 route is the metric from the protocol from which it is imported.

Default Route: 3746 IP allows generation of a default routes and advertisement of the routes through an AS external (ASE) LSA. For details see “Interoperability of Routing Protocols.”

Border Gateway Protocol Version 4 (BGP)

The IBM 3746-9x0 supports a full implementation of BGP Version 4 with null authentication. Only BPG Version 4 support is provided; earlier BPG versions are not supported. Configuration options exist to:

- Enable BGP, and specify the local autonomous system number.
- Define BGP neighbors.
- Define (exclude) ASs from which no routing information will be accepted.
- Define send, receive, and originate policies.
- Define aggregate routes. The 3746 IP Router requires that aggregated routes are preconfigured. When defining aggregated routes, make sure that the individual routes that make up the aggregated route are not exported (that is, define a send policy).

Interoperability of Routing Protocols

When combinations of routing protocols are used on the IBM 3746-9x0 it is necessary to specify how to pass routes learned by one protocol to another. Examples of situations where this is required include:

- Advertising static/default route
- Passing route information between RIP and OSPF gateway protocols
- Passing route information between the interior gateway protocol(s) (RIP/OSPF) and BGP

The principle behind route export is simple but there are a number of implementation details that must be understood for the IBM 3746-9x0.

RIP Specifics: The following rules apply when there is both a RIP route and a route from another protocol to the same destination:

- The RIP route will override a BGP route.
- The RIP route will be overridden by an OSPF internal route.
- The RIP route will override a static route if the RIP route's metric is less.
- The RIP route will override a default route if the RIP route's metric is less.
- The RIP route will override an existing OSPF autonomous system external (ASE) type 1 route if the OSPF external comparison switch is set to type 1, and the RIP metric is lower than the OSPF metric.
- The RIP route will override an existing OSPF autonomous system external (ASE) type 2 route if the OSPF external comparison switch is set to type 1, or the OSPF external comparison switch is set to type 2 and the RIP metric is lower than the OSPF metric.

The 3746-9x0 RIP implementation does not allow the import of OSPF or BGP learned routes. The 3746 IP Router may advertise itself as the default router within the RIP routing domain depending on the values configured within the Originate Default Route field. You have the option:

- To always originate the default route.
- To originate the default route when a BGP route for a specific destination has been received. The default route will be generated if the route corresponds with a specific AS. AS 0 indicates that the AS number should not be checked.
- To originate the default route when OSPF routes are present in the routing table.

OSPF Specifics

The 3746-9x0 has a configuration option to enable route imports. The type of routes that are imported can be specified. Options exist to import static routes, direct routes, subnet routes, RIP routes, and BGP routes.

All routes are imported as OSPF AS external routes and advertised into the OSPF network using AS external link state advertisements.

When OSPF imports routes, it is necessary to specify additional parameters to define how the external route will be advertised into the OSPF network. These parameters are the metric type (1 or 2) and the external route tag.

Metric type 2
defines the metric to the destination to be larger than any internal path.

Metric type 1
means that the metric is comparable to internal paths.

It is recommended that type 2 be selected to ensure that external routes have a larger metric than internal ones. Type 1 should be selected only if the operational effect of using it is fully understood.

The external route tag provides additional information in the AS external link state advertisement, although it is currently only defined for use with BGP. The tag can normally therefore be set to any value.

The 3746 IP Router may advertise itself as the default router within the OSPF routing domain depending on the values configured within the Originate Default Route field. You have the option:

- To always originate the default route.
- To originate the default route when a BGP route for a specific destination has been received. The default route will be generated if the route corresponds with a specific AS. AS 0 indicates that the AS number should not be checked.

The forwarding address indicates to which router default packets are routed. It must be a router interface address that is adjacent, that is, part of the same (sub)net, to the 3746-9x0. 3746 IP routing policies when multiple routes to the same destination are available, are used in the following order:

- Interior gateway protocol (RIP, OSPF) routes are preferred over BGP routes.
- OSPF routes are always preferred over other type of routes (other than routes to directly attached networks).
  - OSPF intra-area routes are preferred to inter-area routes.

- RIP routes and static routes can override OSPF ASE routes. It can be configured if external routes will be compared to OSPF ASEs as type 1 or type 2 routes. Given the preceding definitions of the configurable OSPF ASE route, static and RIP routes may override OSPF ASE routes in the following situations:
  - The OSPF comparison is configured as type 1, and the OSPF ASE route is an OSPF ASE type 2 route.
  - The OSPF comparison is configured as type 1, and the OSPF ASE route is an OSPF ASE type 1 route but has a higher cost than the RIP or static route.
  - The OSPF comparison is configured as type 2, and the OSPF ASE route is an OSPF ASE type 2 route but has a higher cost than the RIP or static route.

BGP Specifics: The routing information that is sent by BGP to its BGP neighbor depends on the routing information within the IP routing table (that is, locally defined static routes and routes learned from RIP and/or OSPF), and routes received from BGP neighbors. This process is controlled by defining originate and send policies, and by defining aggregated routes.

Static/Default Routes: The default referred to in this subsection is the default route that is generated by BGP if the generate default route option is enabled. Export default cannot be used to export an intra-AS default route without BGP being configured. If this is a requirement it must be met by the use of the static route (0.0.0.0) exported to the interior gateway protocol within the AS. Default should be exported to the interior gateway protocol being used within the AS. It will become advertised when and only when BGP communications are established with a neighbor that is configured to generate a default route.

Internal Applications

Two important applications implemented on the 3746 IP Router that rely on ICMP messages are:

- Ping
- Traceroute

Both applications are used in diagnosing connectivity problems. Each application can be invoked by using either the CCM or a TELNET command line mode.

Ping

Ping is the simplest of all TCP/IP applications. It sends one or more messages to a specified destination host requesting a reply and measures the round trip time. The word *ping,* which is used as a noun and a verb, is taken from the sonar operation to locate an underwater object. It is also an abbreviation for *Packet InterNet Groper*.

Traditionally, if you could ping a host then other applications such as Telnet or FTP could reach that host. With the advent of security measures on the Internet which control access to networks by application protocol and/or port number, this is no longer strictly true. Nonetheless, the first test to reach a host is still to attempt to ping it.

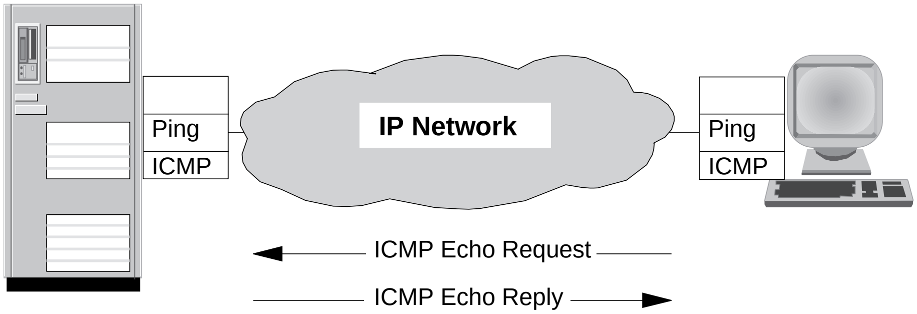

Figure 3-11. Packet InterNet Groper (PING)

The syntax that is used in different implementations of ping varies from platform to platform. The following syntax is for the 3746 IP implementation:

```
ping host
```

Where host is the target IP address. On the 3746 IP Router the target IP address specified must be in dotted decimal notation. Host names are not supported.

Ping uses the ICMP Echo and Echo Reply messages. Since ICMP is required in every TCP/IP implementation, hosts do not require a separate server to respond to pings. Ping is useful for verifying a TCP/IP installation. Consider the following forms of the command; each requires the operation of an additional part of the TCP/IP installation.

ping loopback	Verifies the operation of the base TCP/IP software.
ping my-IP-address	Verifies whether the physical network device can be addressed.
ping a-remote-IP-address	Verifies whether the network can be accessed.

Be aware that IP is a connectionless protocol and that a Ping request and its reply may traverse different routes through the network. Being able to ping a node does not necessarily mean that the reply can be returned as well. Ping requests contain a destination address that is set to the target host indicated in the Ping command. The 3746 IP Routers sets the source IP address to its router_ID. Therefore, to enable the target host to return the Ping reply, an active route between this host and the 3746 Router ID is required.

Note: For Pings to be successful, make sure that the Ping ICMP messages are not filtered along the route between the 3746 IP Router and the destination host.

In normal operation Ping builds and sends an ICMP echo request. Ping waits up to 1 second to receive an ICMP echo reply to calculate the round-trip time, and display the size of the echo request/reply packet (not including IP header), its source IP address, and the round-trip in milliseconds. 3746 IP echo request/reply messages are always 64 bytes long (excluding IP header).

Ping keeps count of the echo requests it sends and the echo replies it receives. It keeps a sum of the round-trip times, and calculates maximum and minimum round-trip times. In addition, it keeps count of the ICMP destination unreachable (network or host unreachable). 3746 IP Ping stops when the user presses any key. Ping statistics are then displayed which include:

- The number of Echo request sent.
- The number of Echo replies received.
- The percentage of packets lost (no replies received).

- The number of destination unreachable messages received.
- The number of Echo request not sent due to lack of buffers.

Traceroute

The Ping command described in the previous section can be used to verify connectivity between two hosts. In addition, for example when the Ping command fails, the Traceroute command can be used to determine the route that IP datagrams follow from host to host.

Traceroute relies on ICMP time exceeded and destination unreachable messages. It sends a UDP datagram with a TTL of 1 to the destination host. The first router to see the datagram will decrement the TTL to 0 and return an ICMP Time Exceeded message as well as discarding the datagram. In this way, the first router in the path is identified. This process can be repeated with successively larger TTL values in order to identify the series of routers in the path to the destination host. Traceroute sends UDP datagrams to the destination host that reference a port number outside the normally used range. This enables Traceroute to determine when the destination host has been reached, that is, when an ICMP destination (Port) unreachable message is received.

For each TTL value, the UDP datagram is sent three times, waiting up to three seconds between transmissions to receive a response. Traceroute continues this process until either a TTL value of 33 has been reached or an ICMP destination unreachable message received.

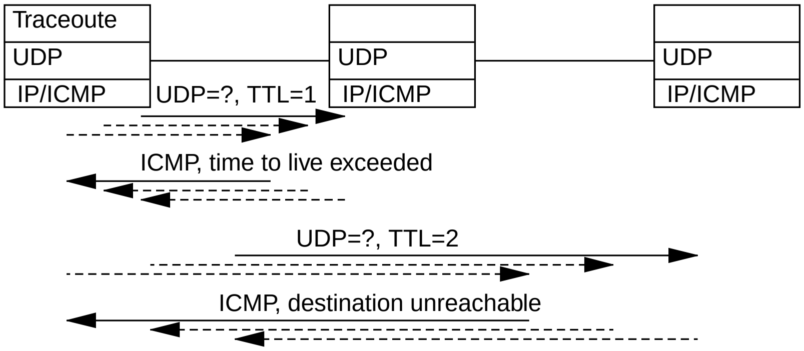

Figure 3-12. TraceRoute

When a route exist between the 3746 IP Router and a destination node, each intermediate router will return an ICMP time exceeded message. However, a situation might exist where somewhere along the route the Traceoute frame cannot be forwarded.

Note: For Traceroute to be successful, make sure that the Traceroute ICMP messages are not filtered along the route between the 3746 IP Router and the destination host. For each TTL value, Traceroute displays a line with the following information:

- `BF` each time the 3746 IP is unable to send the UDP frame due to lack of buffers.
- * each time no response is received.
- The source IP address of the response, if different from the previous response received.
- The round-trip time in milliseconds.

- !P if the response is a destination unreachable (protocol unreachable).
- !H if the response is a destination unreachable (host unreachable).

Bootstrap Protocol (BOOTP)

One restriction with the scheme described in the previous section is related to the use of the limited broadcast address for BOOTP requests; it requires that the server is on the same subnet as the recipient.

The 3746 IP Router, however, has a mechanism to forward BOOTP requests to one or multiple preconfigured BOOTP servers. If the BOOTP request's router address is zero, the 3746 BOOTP relay agent sets it to the IP address of its receiving interface.

The 3746 IP BOOTP relay agent discards the BOOTP request message when:

- BOOTP request's “hops” value is greater than the value configured by the user.
- BOOTP request's “hops” value is greater than or equal to 16.
- BOOTP request's “seconds (retry)” value is greater than the value configured by the user.

  This allows BOOTP servers on the local network a chance to respond before the BOOTP relay agent has forwarded to remote BOOTP servers.

The BOOTP relay agent forwards a BOOTP reply message to the interface whose IP address matches the BOOTP reply's router IP address. Because the BOOTP client may not know its IP address and therefore may not be able to respond to an ARP request, the 3746 IP BOOTP relay agent installs an entry in the outgoing interface ARP cache table using the client IP and hardware address specified in the BOOTP reply message.

Once a BOOTP server has responded and the BOOTP client has processed the reply, it may proceed with the transfer of the boot file and execute the full boot process. The transfer of the boot file is usually done, using the trivial file transfer protocol (TFTP). The full boot process will replace the minimum IP protocol stack used by BOOTP and TFTP by a normal IP protocol stack transferred as part of the boot file and containing the correct customization for the client.

Security

The 3746 IP Router offer two methods to control its IP forwarding functions, namely by defining:

- Filters
- Access controls

Both are described in more detail below.

3746 IP Filters

The 3746 IP filters are essentially user-defined static routes. Any packet received or originated by the 3746 IP Router in which the destination IP address matches a filter is discarded. No error message will be generated. Each filter, like a static route, consists of an IP address and a mask. When an IP datagram is forwarded, the destination address is ANDed with the filter mask. If the result equals the filter address, then the datagram is discarded.

```
destination_address & filter_mask = filter_address = discard!
```

The 3746 IP Router installs filters in its routing table; therefore the numbers of filters is limited by the size of the routing table. Routes that match a filter are not advertised when using a dynamic routing protocol.

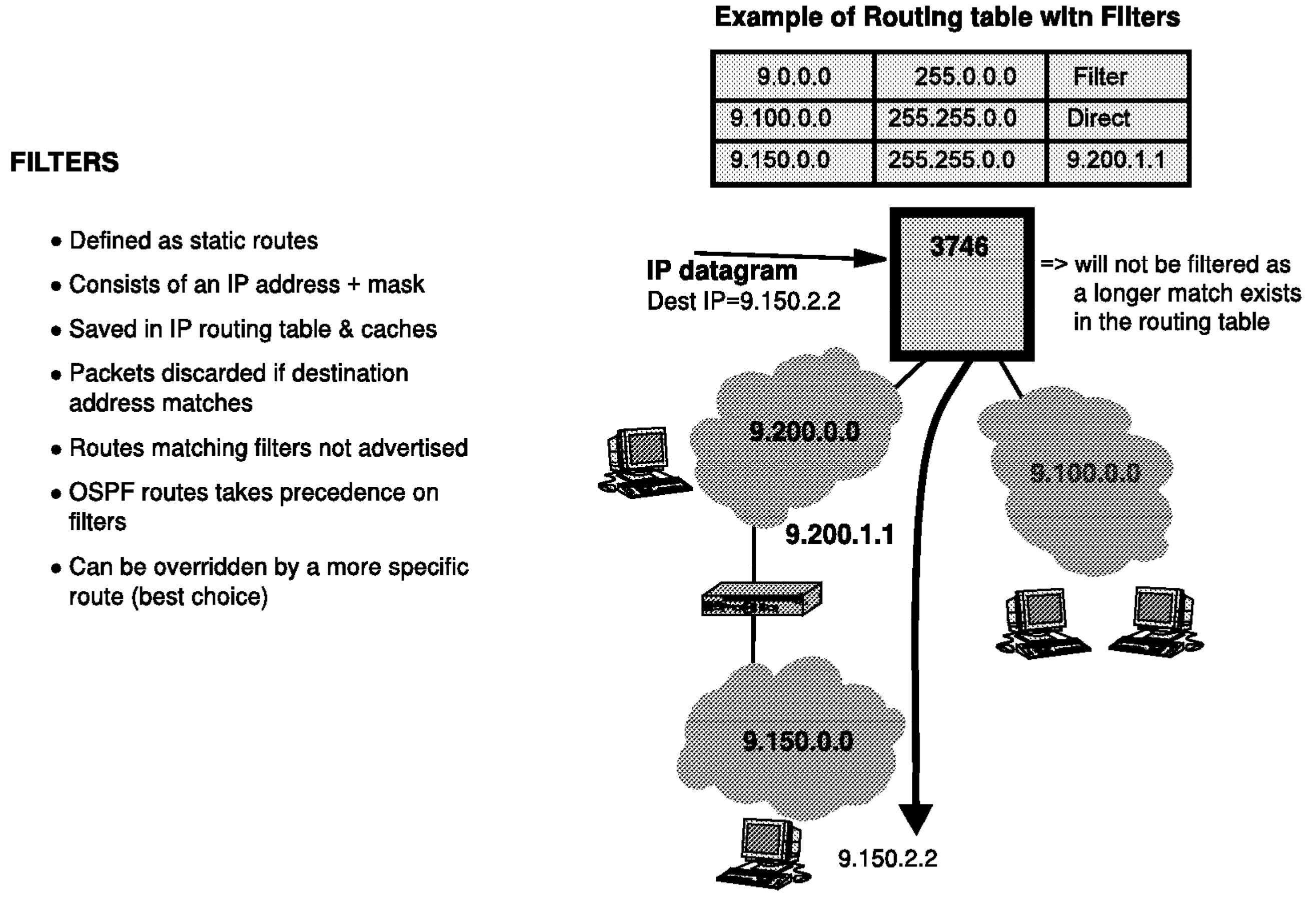

FILTERS

- Defined as static routes
- Consists of an IP address + mask
- Saved in IP routing table & caches
- Packets discarded if destination address matches
- Routes matching filters not advertised
- OSPF routes takes precedence on filters
- Can be overridden by a more specific route (best choice)

Example of Routlng table wltn Fllters

9.0.0.0	255.0.0.0	Filter
9.100.0.0	255.255.0.0	Direct
9.150.0.0	255.255.0.0	9.200.1.1

=> will not be filtered as a longer match exists in the routing table

Figure 3-13. IP Filters

3746 Access Control

Access control provides a way for the user to control which IP packets are forwarded by the 3746 IP Router based on the packet's source IP address, destination IP address, IP protocol number, and TCP or UDP port number. It is applied to all IP packets that are received or originated by the 3746 IP Router. Access control is an optional function that by default is disabled.

Note: The 3746 implements access control for the whole of the 3746, *not at the port (interface) level*. You cannot select individual ports for access control.

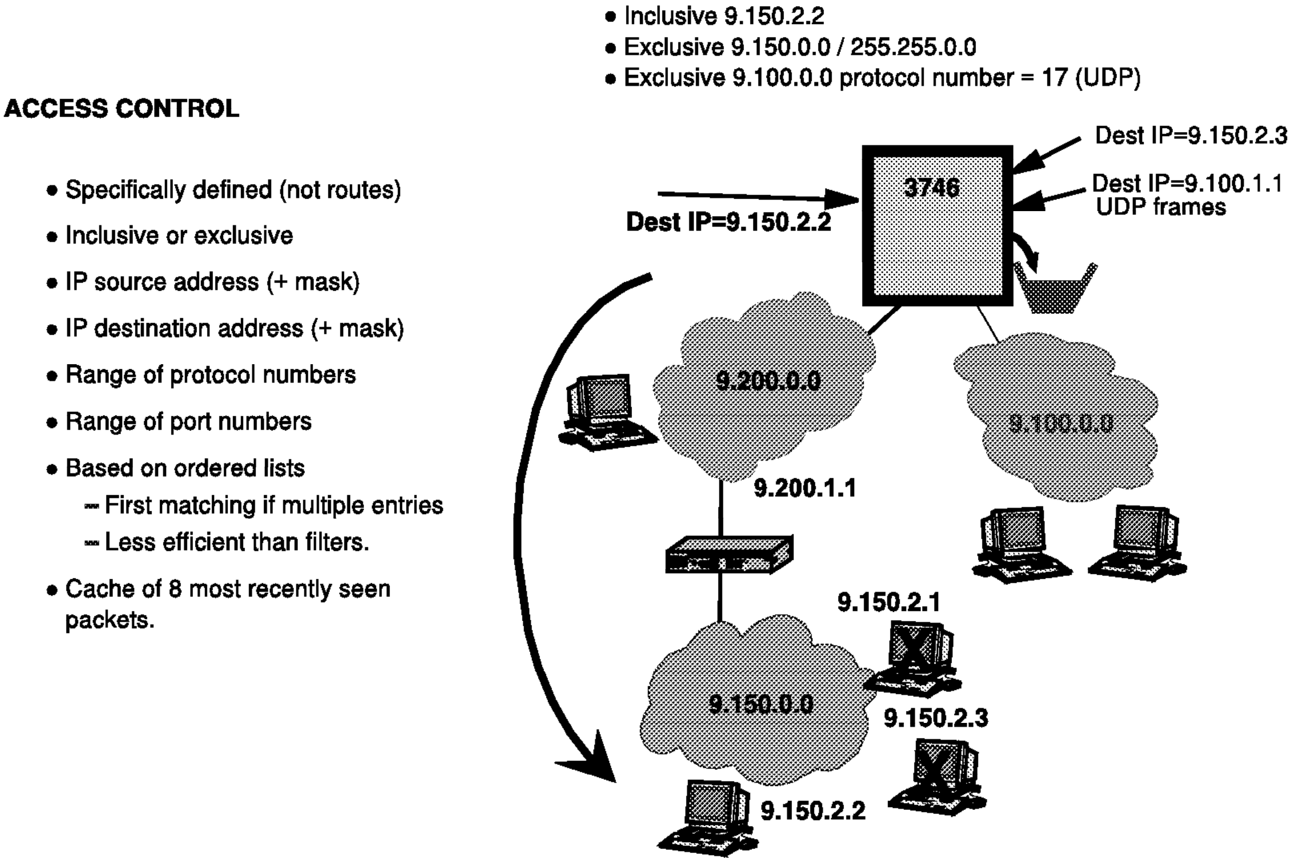

ACCESS CONTROL

- Specifically defined (not routes)
- Inclusive or exclusive
- IP source address (+ mask)
- IP destination address (+ mask)
- Range of protocol numbers
- Range of port numbers
- Based on ordered lists
  - First matching if multiple entries
  - Less efficient than filters.
- Cache of 8 most recently seen packets.

- Inclusive 9.150.2.2
- Exclusive 9.150.0.0 / 255.255.0.0
- Exclusive 9.100.0.0 protocol number = 17 (UDP)

Figure 3-14. IP Access Control

Multiple access entries can be specified, each having the parameters shown in Table 3-5.

Table 3-5 (Page 1 of 2). IP Multiple Access Parameters

Parameter	Description
Access control type	The access control type can be inclusive or exclusive indicating that if this entry is matched the datagram will be forwarded, or discarded, respectively.
Source IP address	Value compared to the datagrams source IP address.
Source IP address mask	Value ANDed with the datagrams source IP address before comparing it with the access control entry's source IP address value.
Destination IP address	Value compared to the datagrams destination IP address.

Table 3-5 (Page 2 of 2). IP Multiple Access Parameters

Parameter	Description
Destination IP address mask	Value ANDed with the datagrams destination IP address before comparing it with the access control entry's source IP address value.
First protocol number	The lower bound of a range of IP protocol numbers; a datagram matches the access control entry only if its IP protocol number is greater than or equal to this value.
Last protocol number	The upper bound of a range of IP protocol numbers; a datagram matches the access control entry only if its IP protocol number is less than or equal to this value.
First port number	The lower bound of a range of TCP/UDP port number; a packet in which the IP protocol number indicates TCP or UDP matches the access control entry only if its TCP or UDP port number is greater or equal to this value.
Last port number	The upper bound of a range of TCP/UDP port numbers; a packet in which the IP protocol number indicates TCP or UDP matches the access control entry only if its TCP or UDP port number is less or equal.

The access control entries are organized into a single list which is used for all interfaces. The access list is ordered, meaning that if a packet matches multiple entries in the list, the first matching entry in the list is the one that takes effect. If no matching access control entry is found, the 3746 discards the datagram.

Note: CCM always adds an all-inclusive entry at the end of the list, meaning that datagrams are forwarded unless discarded because of other entries in the list. Use the TELNET operator interface to insure that this entry is on that list.

Warning: PTR Z237641 needs to be examined to insure correct IP addressing on the Service LAN. The *exclusive* and *inclusive* statements should be coded with care according to the particular installation.

Filters Versus Access Controls

Filters

A filter on a destination IP address is added to the IP routing table with indication *filter*. The position of a route in the routing table depends on its IP address. Filters are saved into the main routing table maintained on the CBSP2 and the routing table cache on each of the 3746 processors.

A filter is a route like a static route. The primary difference is that filters are not advertised by RIP, or BGP, and if a filter is replaced by another type of route (for example, OSPF) that later goes away, the filter is not automatically restored. OSPF internal routes (intra-area and inter-area) take preference over filter routes. This is to prevent unexpected reachability problems in OSPF networks. Because it is a link state protocol, OSPF does not take filters into account in its advertisements, so filters in an OSPF network can cause reachability problems. To be sure to filter the traffic, access control must be used instead.

Filters have another shortcoming; they can be partially overridden by a more specific route. Assume a filter for 9.0.0.0 with mask 255.0.0.0, and a route for

9.100.0.0 with mask 255.255.0.0 has been installed. Because the routing table search looks for the most specific match, datagrams to 9.100.x.x will follow the route instead of the filter.

Access Control: Access control entries are searched sequentially for each datagram until a matching entry is found. This may lead to considerable overhead when many access control entries have been specified. There is no special search of a routing table for a filter. The routing table is searched once for the best matching route, and if that match turns out to be a filter, the datagram is dropped.

Recommendation: In general, the use of filters is advised rather than using access control. However, you should be aware of the shortcomings of filters mentioned above.

IP over 3746 Networks

Figure 3-15 depicts how TCP/IP functions have been implemented on the 3746-9x0 components identified.

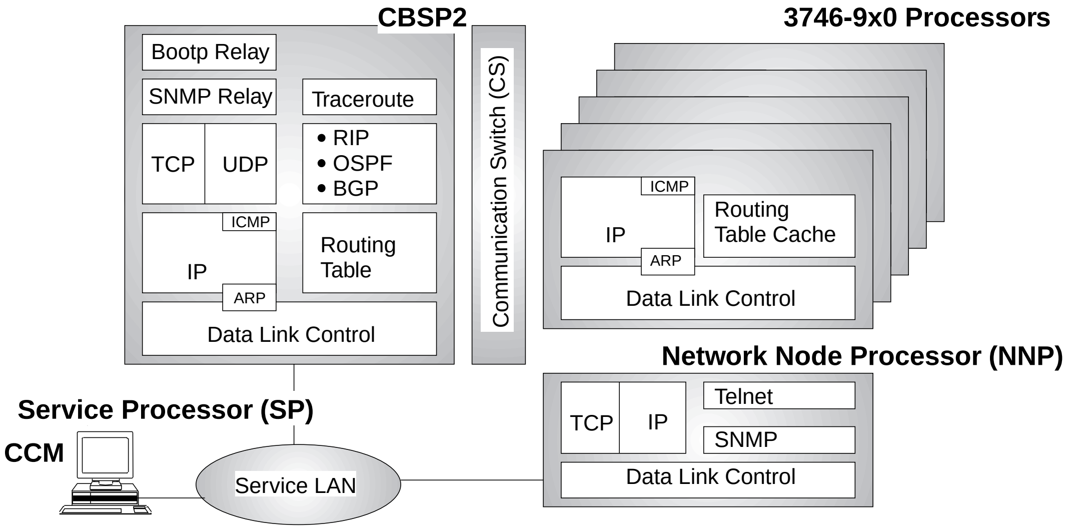

Figure 3-15. 3746-9X0 IP Base

Each 3746-9x0 processor, including the CBSP2, provides IP routing and data link control (DLC) functions. The DLC functions enable data transport to adjacent nodes. The DLC functions are shared between protocol stacks (IP and APPN/HPR). The IP functions allow IP transport between each of the 3746-9x0 adapters, and to and from externally attached IP equipment. To perform the IP routing, each adapter maintains a routing table cache. The code that performs the IP routing functions is present on all processors. However, it will only be activated on processors on which an IP interface has been configured. The routing cache is built during the transport of IP data, using routing information learned from the CBSP2.

In addition to the IP forwarding and DLC functions, more functions are required. All dynamic IP routing protocols (RIP, OSPF, and BGP) run on the CBSP. The CSBP2 is responsible for processing routing (RIP/OSPF/BGP) packets received from, and generating routing packets destined for, external equipment. The adapters are instrumental in forwarding these packets between the external equipment and the CBSP2, and vice versa.

The CSBP2 builds and maintains a routing table using static routes and IP interfaces defined on the 3746-9x0, and routing information learned from the dynamic routing protocols. To minimize the routing information maintained on each of the adapters, each processor maintains a routing table cache that contains a subset of the information stored in the CSBP2 routing table. When, during IP forwarding, a processor is not capable of making a routing decision based on the information maintained in its routing table cache, the appropriate routing information is retrieved by querying the CBSP2. Routes learned from the CSBP2 are stored within the routing table cache, and used to route successive IP datagrams. The Routing Table Caches contain one entry per *destination host* and *not* per *destination network*. This is to prevent less specific routes from masking more specific routes.

For example, if the most specific route that matches a particular destination IP address is the default route, if the default route is saved in the cache, then everything matches it. Thus everything is forwarded on the *default route* and more specific routes are ignored. Caching everything as host routes avoids the complications of dealing with more specific routes.

All cached entries are subject to an aging mechanism, and will be periodically removed.

Attached to the service LAN is the network processor (NNP). The main functions of the NNP are the 3746 IP operator and SNMP functions. In addition, the NNP contains the 3746 IP configuration file. When the 3746 IP code is restarted, each adapter retrieves a copy of the 3746 IP configuration file and configures itself. For details of SNMP functions, see “Managing Your IP Resources with SNMP” on page 29-21.

Also attached to the service LAN is the service processor (SP). The SP provides a repository of the microcode running on the 3746-9x0 adapters. Each 3746-9x0 adapters loads its microcode from the SP during an initial microcode load (IML). The 3746-9x0 controller configuration and management (CCM) tool can be accessed from the SP.

IP Interfaces

The 3746 IP Router supports attachment of IP equipment using ESCON, token-ring, Ethernet, or serial lines. For serial line-attached equipment frame-relay, X.25, and Point-to-Point Protocol (PPP) connections can be used (see Figure 3-16 on page 3-33). Both switched, for example, using a public switched telephone network (PSTN), and leased-line PPP and frame relay connections are supported.

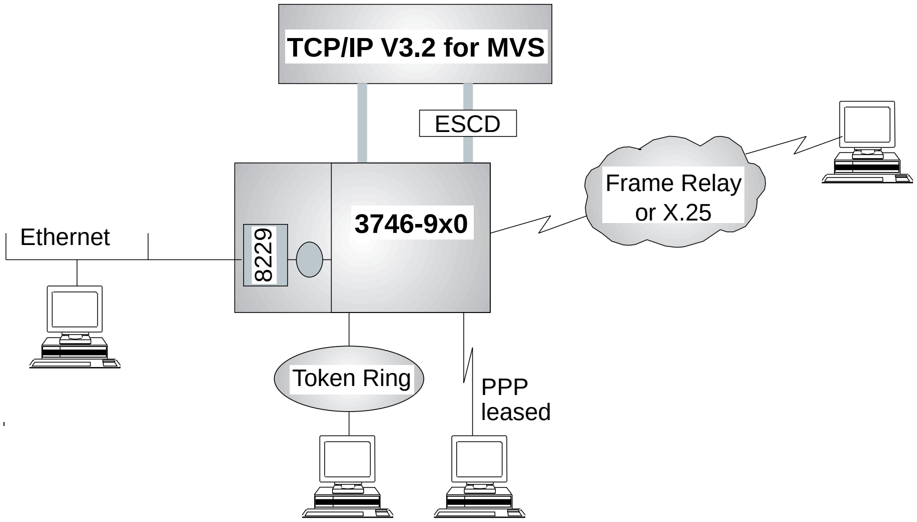

Figure 3-16. IP Interfaces

Each of these interfaces will be discussed in more detail in the following section.

ESCON

ESCON attachments can be used to provide native IP transport, using the channel data link control (CDLC), between the 3746 IP and host system(s) running TCP/IP for MVS (at least Version 3, Release 1). See Figure 3-17. The host systems can be directly attached to the 3746-9x0, or connected via an ESCON director (ESCD).

The ESCON implementation on the 3746-9x0 is based on the concept of host links and link stations. A host link is a logical connection between a host system and the 3746-9x0. A link station represents a point-to-point link between the 3746-9x0 and host programs such as TCP/IP for MVS, VTAM, and/or TPF. Host link and link station definitions need to be entered before IP (or, for that matter, APPN/HPR or SNA subarea) communication is possible.

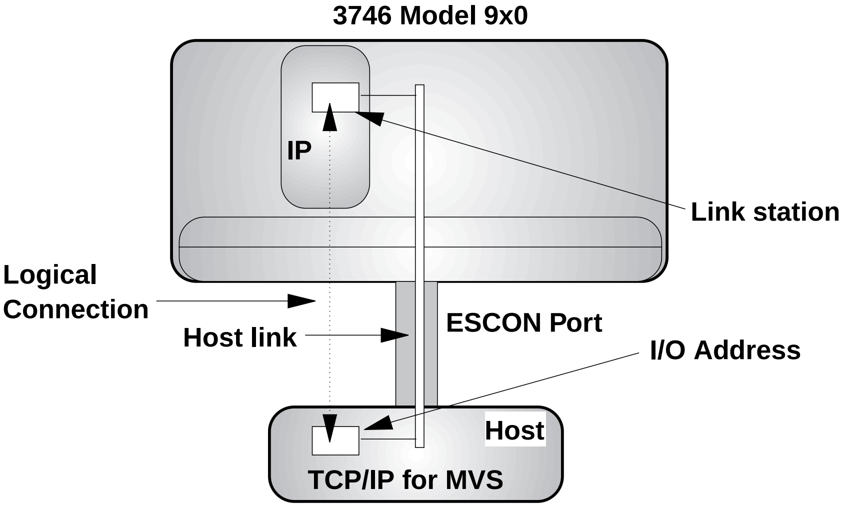

Figure 3-17. ESCON Port, Host Link, and Link Station

On each 3746-9x0 you can install up to 16 ESCON adapters (15 on a 3746-900 attached to a dual CCU Model 3745). Each ESCON adapter controls a single

ESCON coupler. An ESCON coupler provides a single physical connection to a single host system, unless you are using an ESCON director (ESCD). An ESCD provides a physical connection to multiple host systems (See Figure 3-18 on page 3-34).

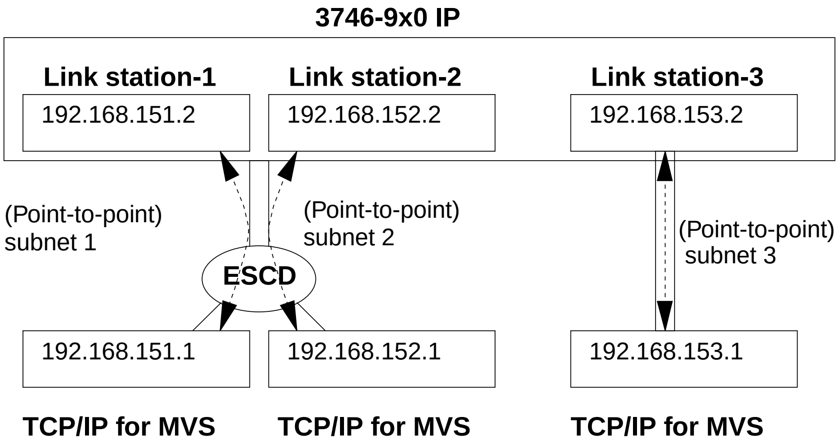

Figure 3-18. Multiple 3746 IP ESCON Links

Single IP addressing is possible on multiple ESCON stations to provide easier coding for multiple host access over the ESCON adapter.

On each ESCON adapter you can define up to 16 host links. For each host system you can define only one host link, unless the host system uses the ESCON Multiple Image Facility (EMIF), in which case you can define a separate host link (up to 16) to each of the logical partitions (LPARS). On each host link you can define one or multiple link stations; however, the maximum number of link stations that can be defined per ESCON adapter is sixteen.

Note: Using the above figures a single 3746-9x0 can provide ESCON attachment for up to 256 (16x16) host systems, and communicate with up to 256 host programs.

The ESCON configuration process consists of two phases:

1. A general definition part in which host links and link stations are defined, and the link stations are assigned to the appropriate protocol stack (3746 IP for IP routing).
2. The actual IP port configuration.

ESCON: General Configuration

To define an ESCON link, an ESCON port, a host link, and a link station definition are required. A host link, comprising a connection between a S/390® operating system and the 3746-9x0, is the logical equivalent of a port. On a single port, multiple host links can be defined, with multiple link stations on each host link. Figure 3-19 on page 3-35 depicts in configurations 1 and 2 an ESCON port on which a single host link and a single link station has been defined.

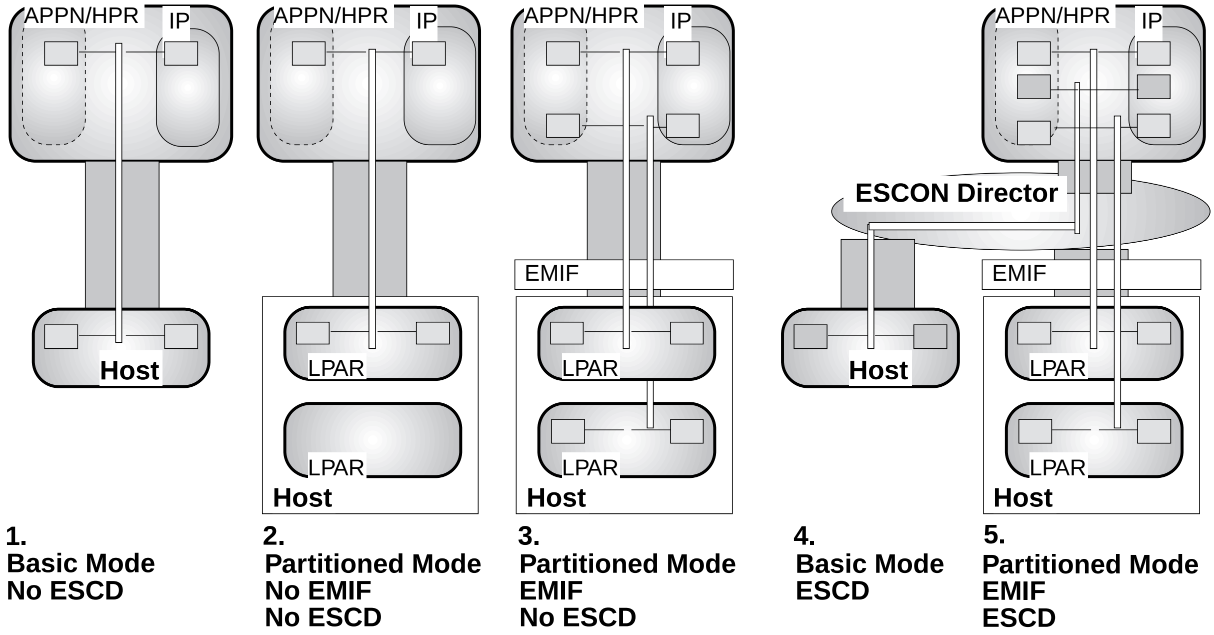

Figure 3-19. ESCON Host Links and Link Stations

The maximum number of host links that can be defined per ESCON adapter is 16. The actual number depends on the type of physical attachment between your host and the 3746-9x0:

Configurations 1 and 2
If your host is directly attached and running in basic (non-partitioned) mode only, or in the partitioned mode without the IBM ESCON Multple Image Facility (EMIF), then a single host link can be defined.

Configuration 3
If your host is directly attached, is running in partitioned mode, and uses the ESCON multiple image facility (EMIF), then a single host link can be defined to each LPAR.

Configuration 4
If your host is connected via an ESCON director (ESCD), then a single host link is possible to each host running in basic (non-partitioned) mode.

Configuration 5
If your host is connected via an ESCON director (ESCD), then a single host link is also possible to each LPAR on the host running in partitioned mode and using the ESCON multiple image facility (EMIF).

Link stations map one-to-one with I/O addresses on your host. By using a specific I/O address to, for example, TCP/IP for MVS and assigning the corresponding link station to the 3746-9x0 IP functions, native IP communication between the two is possible.

ESCON: Port Sharing

As discussed in the previous section up to 16 link stations can be defined per ESCON adapter. Each link station provides a point-to-point connection. See Figure 3-20.

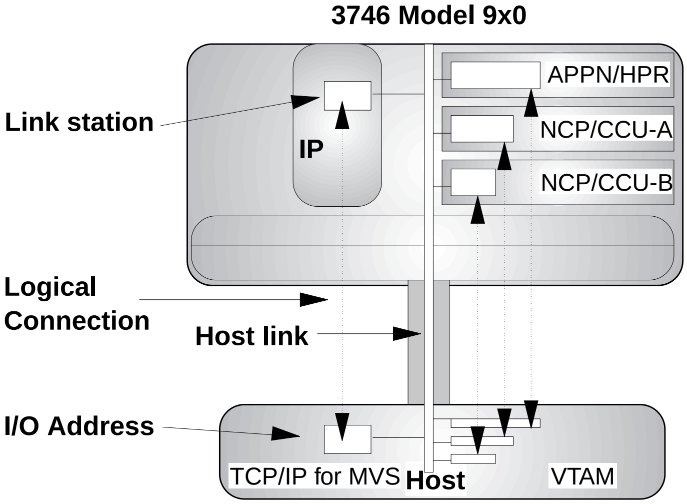

Figure 3-20. ESCON Port Sharing

ESCON ports can be defined and simultaneously activated from all protocol stacks available (IP, APPN/HPR, and for 3746-900 only; NCPs running on CCU-A and CCU-B on the attached 3745). On a 3746-900, each link station must be assigned to either 3746 IP, NCP/CCU-A, NCP/CCU-B, or the 3746 NN function; on a 3746-950, a link station must be assigned to either the 3746 IP, NCP or the 3746 NN function.

Link stations map one-to-one on I/O addresses defined on attached host systems. The I/O addresses are used by host programs to communicate with the corresponding protocol stack. The 3746-9x0 IP protocol stack can only communicate with TCP/IP for MVS (Version 3, Release 1 or 2), while the 3746 APPN/HPR and/or NCP protocol stacks can communicate with VTAM and/or TPF.

Link stations are independently assigned, and, therefore, an ESCON port and each host link can be shared between up to four (two on the 3746-950) 3746 protocol stacks. The total number of link stations per physical port cannot exceed 16.

ESCON: IP Addresses and Subnet Addresses Rules

The following section explains the rules for IP address definitions for ESCON channels.

The following general rules apply:

- One (physical) host can run multiple IP address spaces. Each IP address space is a full TCP/IP MVS stack (partition).
- Each IP address space is independent of the others.
- Each IP address space can run one or multiple CDLC instances.
- Each CDLC instance of a given IP address space must run on a separate ESCON fiber.

The following rules are defined by TCP/IP for MVS:

- CDLC instances of a *single* IP address space cannot have the same next-hop IP address.
- CDLC instances of *different* IP address spaces can have the same next-hop IP address, as they are in independent IP stacks.

The following rules are defined by the 3746:

- For hosts with only 1 IP address space and 1 CDLC in the IP address space.

  Users can define a single *subnet* between the hosts and the 3746 stations.

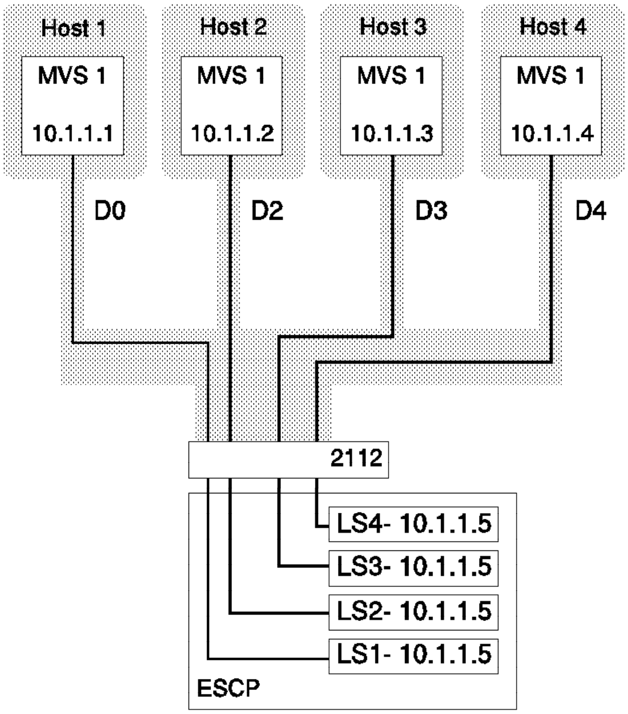

Figure 3-21. Four ESCP Stations with the Same IP Addresses

- For hosts with many IP address spaces and 1 CDLC per IP address space.

  Users can define a single *subnet* between the hosts and the 3746 stations.

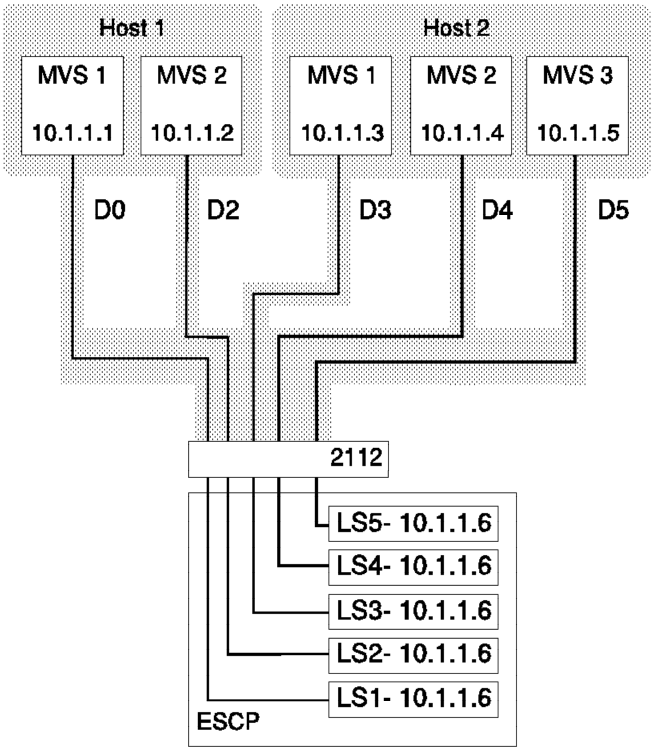

Figure 3-22. Five ESCP Stations with the Same IP Addresses

- For hosts with many IP address spaces and many CDLC per IP address space.

  Users *must* define multiple *subnets*, the basic rule being that two CDLCs of the same MVS IP address space cannot be in the same subnet.

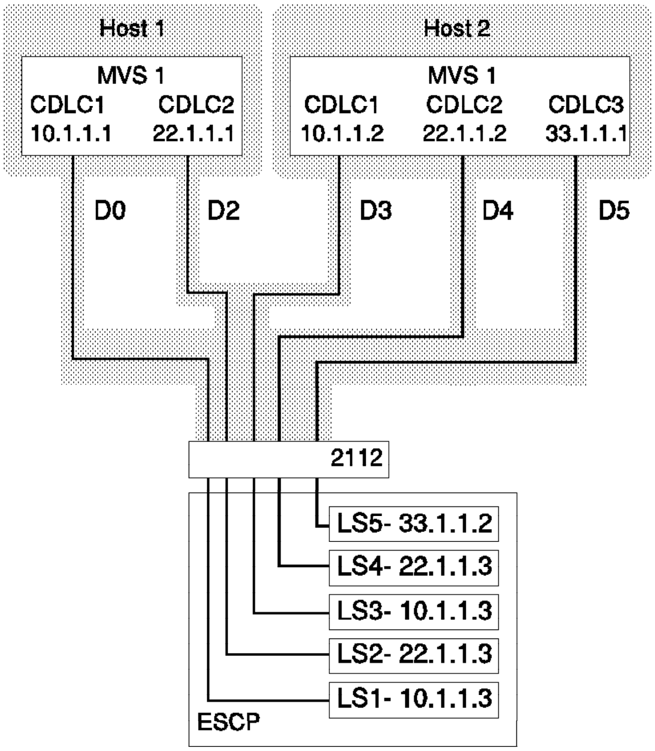

```
Subnets:
- 10.0.0.0 used for HOST1/CDLC1, HOST2/CDLC1 and ESCP LS1 and LS3
- 22.0.0.0 used for HOST1/CDLC2, HOST2/CDLC2 and ESCP LS2 and LS4
- 33.0.0.0 used for HOST2/CDLC3 and ESCP LS5
```

Figure 3-23. ESCP Stations on Separate Subnets

ESCON: IP Configuration

The IP configuration for ESCON link stations consists of:

- Assigning an IP address and subnet mask to the link station. Each link station represents a separate point-to-point CDLC link, therefore, for each link station an IP address must be assigned that is part of a different subnet.
- Enabling dynamic routing protocol on this interface (or define static routes). Assigning an IP address (refer to “IP Addressing” on page 3-2) and subnet mask (refer to “The Subnet Mask” on page 3-5) is mandatory. In addition, the channel adapter slowdown time (CASDL), the attention timer (TIMEOUT), and the delay time (DELAY) timer have to be defined.

Enabling dynamic routing protocols is optional. It is only required when the TCP/IP for MVS system has multiple IP connections, and static routing does not suffice. The only dynamic routing protocol currently supported by TCP/IP for MVS (Version 3, Release 1 and 2) is RIP.

Token Ring

Token-ring attachments can be used to provide IP transport between the 3746 IP and IP stations or routers running TCP/IP.

For the IP communication over token-ring, the 3746-9x0 uses (connectionless) IEEE 802.2 logical link control (LLC). IEEE 802.2 LLC requires the use of a local and a remote service access point (SAP). For IP communication both local and remote SAP should be equal to X'AA', indicating subnetwork access protocol (SNAP) encapsulation. See Figure 3-24. The (optional) *routing information field (RIF)* in the MAC header indicates the support for source route bridging, which is by default available.

SD	**AC**	**FC**	**DMAC**	**SMAC**	**RIF**	**DSAP**	**SSAP**	**03**	**OUI PID**	**User Data • • • • •**	**FCS**	**ED**	**FS**

SD	- Starting Delimiter
AC	- Access Control Field
FC	- Frame Control Field
DMAC	- Destination MAC address
SMAC	- Source MAC address
RIF	- Routing Information Field
DSAP	- Destination SAP (AA for IP)
SSAP	- Source SAP (AA for IP)
OUI PID	- Organization/Protocol Identifier
OUI	- 00 00 00
PID	- 0800,0806 (IP, ARP)
FCS	- Frame Check Sequence
ED	- End Delimiter
FS	- Frame Status Field

Figure 3-24. Token-Ring Encapsulation

On each 3746-9x0, you can install up to 16 token-ring adapters. Each token-ring adapter can control up to two token-ring interface couplers (TIC3). Each token-ring coupler provides a single physical connection to a token-ring LAN.

Note: The 3746-9x0 supports the attachment of up to 32 token-ring LANs (31 on a 3746-900 attached to a dual CCU Model 3745).

The token-ring configuration process consists of two phases:

1. A general definition part in which token-ring port specifics are defined
2. The actual IP port configuration.

Token-Ring: General Configuration

To define a token-ring port you have to specify its IP and APPN/HPR (if relevant) names, the speed of the token-ring LAN, and the local MAC address. The APPN and HPR SAPs are only relevant if the port is also used for APPN traffic. Make sure the IP maximum transmission unit (MTU) is high enough to prevent fragmentation of IP datagrams on the 3746 IP Router. However, in a bridged environment, make sure that the MTU is smaller than the lowest value configured on any on the LAN bridges. Select IP parameters to enter the IP over token-ring configuration. See Chapter 40, “CCM Worksheets for Controller Configuration Definitions” on page 40-1. For IP transport no station needs to be defined.

Token-Ring: Port Sharing

Token-ring ports can be defined on all protocol stacks. However, simultaneous activation of the port is limited to the IP protocol stack, APPN/HPR protocol, and (3746-900 only) either the NCP running on CCU-A or the NCP running on CCU-B of the attached 3745. Therefore, on a 3746-900, a token-ring port can be shared between the 3746 IP, a single NCP and the 3746 NN function; on a 3746-950, the token-ring port can be shared between the 3746 IP and the 3746 NN function. The 3746-9x0 data link control layer differentiates traffic from/to any of these protocol stacks using the service access points (SAPs) within the token-ring MAC header.

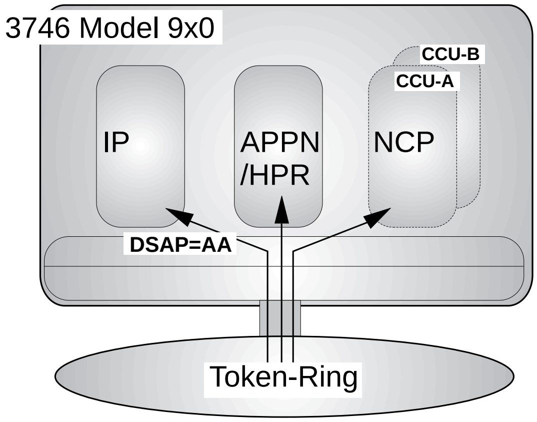

Figure 3-25. Token-Ring Port Sharing

When receiving data the 3746-9x0 will accept the following source SAPs:

- X'AA' - data to/from 3746 IP stack.
- X'04' - SNA data to/from NCP protocol stack (3746-900 only).
- X'C8' - HPR non-ERP data to/from NCP protocol stack (3746-900 only).
- X'xx' - APPN data to/from 3746 NN protocol stack *xx* is configured in the APPN local SAP field during token-ring port configuration.
- X'yy' - HPR non-ERP data to/from 3746 NN protocol stack *yy* is configured in the APPN/HPR local SAP field during token-ring port configuration.

Token-Ring: IP Configuration

The IP configuration for token-ring ports consists of:

- Assigning an IP address and subnet mask to the token-ring port.

  Note that the 3746 IP Router allows you to specify multiple IP addresses and subnet masks to the same physical interface. When defining multiple IP addresses, the 3746 IP Router connects to multiple subnets. The 3746-9x0 is capable of routing between subnets on the same physical interface, although in general this is not recommended.

- Enabling dynamic routing protocol for each IP address (or define static routes).

  The 3746 IP dynamic routing protocols operate independently for each IP address defined on a token-ring interface; therefore, each IP address is defined in the configuration.

Assigning an IP address and subnet mask is mandatory. At least one but optionally multiple (up to 16) IP addresses can be defined. If multiple IP addresses are defined, make sure that the addresses are part of different subnets.

The RIF timer indicates how many seconds will elapse between the last reference to an ARP entry and when it will be deleted. Increasing the timer may result in more storage being required to cache ARP entries. A zero value will disable the ARP caching function.

Enabling dynamic routing protocols is optional. It is only required when any of the token-ring attached devices have multiple IP interfaces, and static routing does not suffice. Dynamic routing protocols supported by the 3746-9x0 are RIP, OSPF, and BGP.

External IP Connection Between 3745 and 3746-900

With the introduction of NCP V7R6 it is now possible to link the NCP IP router and the 3746 IP router together via the CBC between the 3745 and 3746. Previously, an external link between the 3745 and 3746 was need to connect the two routers. To support this function from the 3746 side, the optional feature code #5800 is needed. This is supported by microcode level D46130I (ECA 170) and later versions.

The internal connection or connections between the two routers are seen as token-ring point-to-point (PtP) IP connections, these connections must be defined in the NCP and 3746 CCM. Each connection uses a separate TIC3 on the 3746, this TIC3 may also be used to drive a token ring, but beaconing on that token ring or other problems may interfere with the IP routing function. For that reason it is recommended, but not mandatory, that each TIC3 used for an internal IP connection is used only for that function.

Figure 3-26 on page 3-43 shows a representation of the internal point-to-point IP connection.

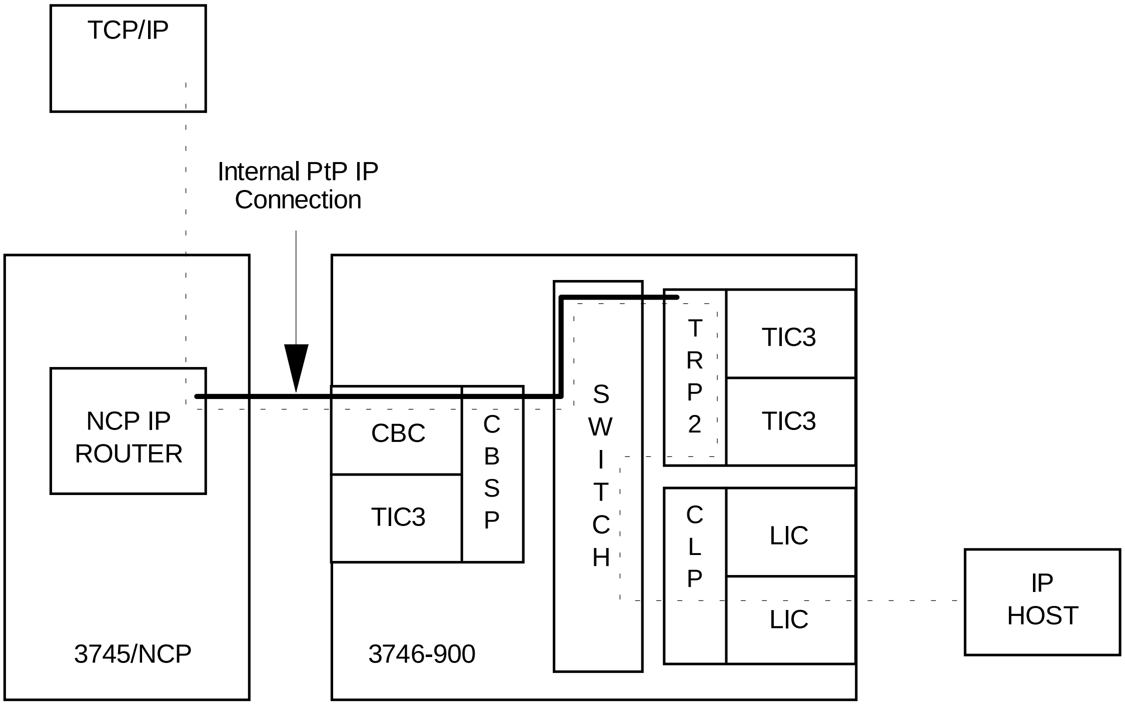

Figure 3-26. Internal IP PtP Connection

Internal IP Connection Definitions

The following sections explain how to define an internal point-to-point IP connection. Figure 3-27 on page 3-44 shows the configuration which is being defined. The defintions shown here include all the information needed for this part of the NCP generation process, the actual PtP link is defined in the LN2880 LINE (1), and IP2880 PU (2) statements, and the IPLOCAL statement for LADDR=10.04.00.99 (5). The NETWORK must be IP (3), and PUTYPE must be 1 (4).

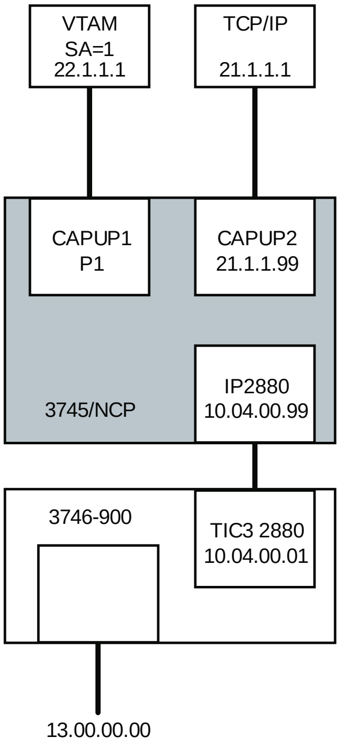

Figure 3-27. Example Configuration

NCP Definitions: The following examples show the NCP definitions needed to configure the internal IP PtP link from the 3745 side.

```
***********************************************************************
* VTAM HOST DEFIINTIONS
***********************************************************************
*
HOSTA01  HOST  BFRPAD=0,DELAY=0.0,INBFRS=6,MAXBFRU=32,
               SUBAREA=1,
               STATMOD=YES,UNITSZ=1024
*
**********************************************************************
*  TOKEN RING LINE 2880
**********************************************************************
*
ODLCGRP1 GROUP ECLTYPE=(PHY,ANY),ADAPTER=TIC3,
               DIAL=NO,LNCTL=SDLC,TYPE=NCP
LN2880   LINE  ADDRESS=(2880,FULL)  1 ,
               LOCADD=400000002880,
               MAXPU=2,
               MAXTSL=16732,
               NPACOLL=YES,
               PORTADD=2,
               SPEED=9600,
               TRSPEED=16
IP2880   PU    ADDR=02,
               INTFACE=(TR2880,2048)  2 ,
               NPACOLL=YES,
               NETWORK=IP  3 ,
               PUTYPE=1  4 ,
               ARPTAB=(1)

***********************************************************************
*              CHANNEL ADAPTERS
***********************************************************************
*
CAGRP1   GROUP LNCTL=CA,CA=TYPE7,NCPCA=ACTIVE,MONLINK=CONT,
               TIMEOUT=240.0,ISTATUS=ACTIVE,CASDL=420.0,DELAY=0.0
*
***********************************************************************
*  VTAM HOST CHANNEL ATTACHMENT
***********************************************************************
*
CALNP1   LINE  ADDRESS=P1
CAPUP1   PU    PUTYPE=5,TGN=1
*
***********************************************************************
*  CHANNELS FOR IP CHANNEL ATTACHED ROUTERS
***********************************************************************
*
CALNP2   LINE  ADDRESS=P2,CASDL=420,DELAY=0.0,MONLINK=NO
CAPUP2   PU    PUTYPE=1,INTFACE=CPU2,ARPTAB=(10,,NOTCANON)
```

```
*********************************************************************
*       IP ROUTING DEFINITIONS
*********************************************************************
*
        IPOWNER HOSTADDR=21.1.1.1,NUMROUTE=(25,25,25),UDPPORT=580,
               INTFACE=(CPU2),MAXHELLO=6,NUMDRIF=10

        IPLOCAL LADDR=21.1.1.99,INTFACE=CPU2,METRIC=1,
               P2PDEST=21.1.1.1,PROTOCOL=RIP
*
        IPLOCAL LADDR=10.04.00.99 5 , INTFACE=TR2880, METRIC=1,
               P2PDEST=10.04.00.01 6 , PROTOCOL=RIP, SNETMASK=255.255.0.0
*
        IPROUTE DESTADDR=13.0.0.0 7 ,
               NEXTADDR=10.4.00.01 8 ,
               INTFACE=TR2880,
               METRIC=1,DISP=PERM,HOSTRT=NO
*
GENEND  GENEND
        END
```

For reference, the following is a short description of some of the parameters used:

IPOWNER Identifies the TCP/IP MVS/VM host that manages the routing tables (NCPROUTE)

IPLOCAL Defines an interface to the NCP IP router.

IPROUTE Defines an entry in the NCP IP routing table.

HOSTADDR Defines the IP address of the owning IP host.

INTFACE Defines the name of the IP interface to NCPROUTE.

LADDR Defines the IP address of the interface.

P2PDEST Defines IP address of the destination IP host.

SNETMASK Defines subnet mask for the interface.

3746 CCM Definitions:

1. The token ring port 2880 must be configured first, and IP must be activated on that port.

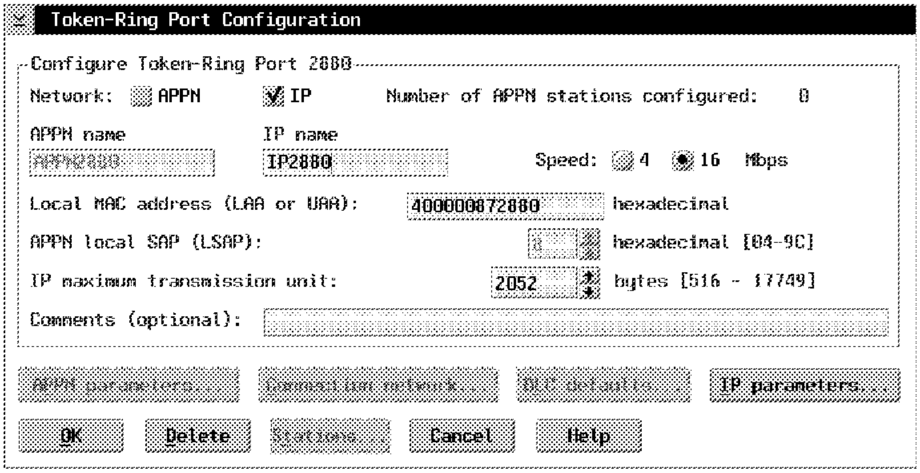

Token-Ring Port Configuration

Configure Token-Ring Port 2880

Network: □ APPN ☑ IP Number of APPN stations configured: 0

APPN name: APPN2880

IP name: IP2880

Speed: ○ 4 ● 16 Mbps

Local MAC address (LAA or UAA): 400000872880 hexadecimal

APPN local SAP (LSAP): 8 hexadecimal [04-9C]

IP maximum transmission unit: 2052 bytes [516 - 17749]

Comments (optional):

Figure 3-28. Port Configuration

2. On port 2880, an IP address must be defined, this is the IP address of the 3746 end of the PtP IP connection, 10.4.0.1. This must match the IP address specified on the P2PDEST operandon the IPLOCAL statement in NCP.

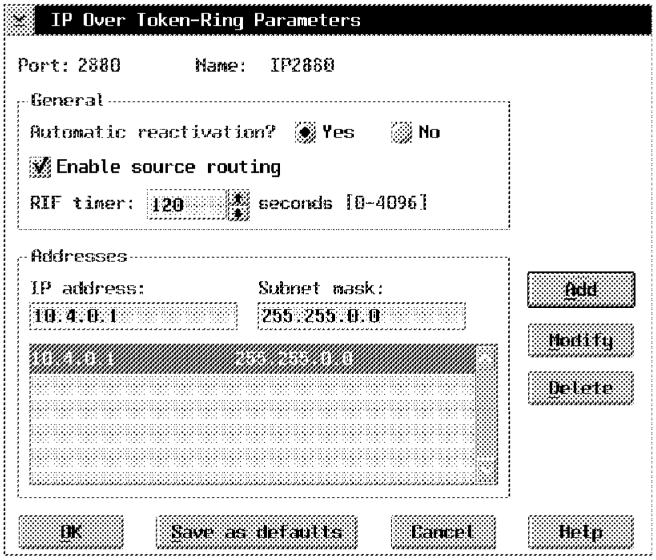

```
IP Over Token-Ring Parameters

Port: 2880        Name:  IP2880

General
Automatic reactivation?  ● Yes   ○ No
☑ Enable source routing
RIF timer:  120   seconds [0-4096]

Addresses
IP address:            Subnet mask:
10.4.0.1               255.255.0.0

10.4.0.1               255.255.0.0
```

Figure 3-29. IP over Token Ring Parameters

3. RIP should be activated as shown on port 2880.

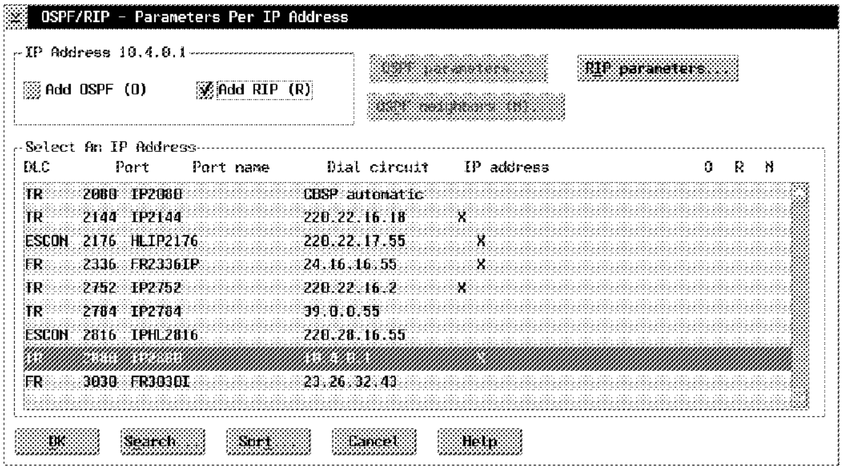

```
OSPF/RIP - Parameters Per IP Address

IP Address 10.4.0.1
□ Add OSPF (O)        ☑ Add RIP (R)

Select An IP Address
DLC      Port     Port name      Dial circuit    IP address          O  R  N
TR      2080  IP2080             CBSP automatic
TR      2144  IP2144             220.22.16.18      X
ESCON   2176  HLIP2176           220.22.17.55        X
FR      2336  FR2336IP           24.16.16.55         X
TR      2752  IP2752             220.22.16.2       X
TR      2784  IP2784             39.0.0.55
ESCON   2816  IPHL2816           220.28.16.55
TR      2880  IP2880             10.4.0.1            X
FR      3030  FR3030I            23.26.32.43
```

Figure 3-30. OSPF/RIP Parameters per IP Address

4. The following RIP parameters should be activated for IP address 10.4.0.1.

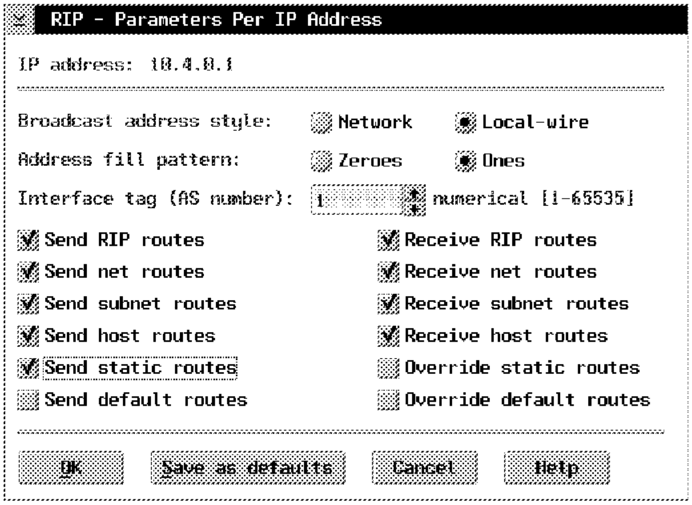

RIP - Parameters Per IP Address

IP address: 10.4.0.1

Broadcast address style: ○ Network ● Local-wire

Address fill pattern: ○ Zeroes ● Ones

Interface tag (AS number): 1 numerical [1-65535]

☑ Send RIP routes ☑ Receive RIP routes
☑ Send net routes ☑ Receive net routes
☑ Send subnet routes ☑ Receive subnet routes
☑ Send host routes ☑ Receive host routes
☑ Send static routes □ Override static routes
□ Send default routes □ Override default routes

Figure 3-31. RIP Parameters per IP Address

To use the internal PtP link for IP routing, to reach the destination network 13.0.0.0 (DESTADDR=13.0.0.0 7) from NCP, the next hop (NEXTADDR 8) points to the IP address of the 3746 end of the PtP connection 10.4.0.1.

Internal IP Connection Between 3745 and 3746-900

In the previous sections it has been explained that on a 3746-900, NCP and 3746 IP functions can coexist. As NCP is also providing IP routing functions, it is important that a distinction is made between the 3746 IP (stand-alone) and the NCP-controlled IP (NCP-IP) routing functions.

3746 IP provides IP routing for ESCON channel, token-ring (TIC3), Ethernet, and serial line (frame-relay, X.25, and PPP) attached equipment. Connectivity is provided through 3746-900 attachments only. NCP-IP provides IP routing for ESCON channel, parallel channel, token-ring (TIC1/TIC2), Ethernet, and serial line (frame relay) attached equipment. However, with the exception of ESCON, NCP-IP connectivity is provided through base attachments only, that is, for equipment not attached to the 3746-900.

3746 IP and NCP-IP perform their IP routing functions in a fully independent fashion. IP connectivity between both IP routers requires an external connection. Internal (via the CBC) IP connections is not supported. Figure 3-32 depicts how direct IP connectivity between NCP-IP and 3746 IP can be realized via token-ring, Ethernet, and frame relay. In all cases, a port (token-ring, Ethernet, or frame relay respectively) is required on both the 3746-900 and the 3745 base frame (or a 3746 Model L1x).

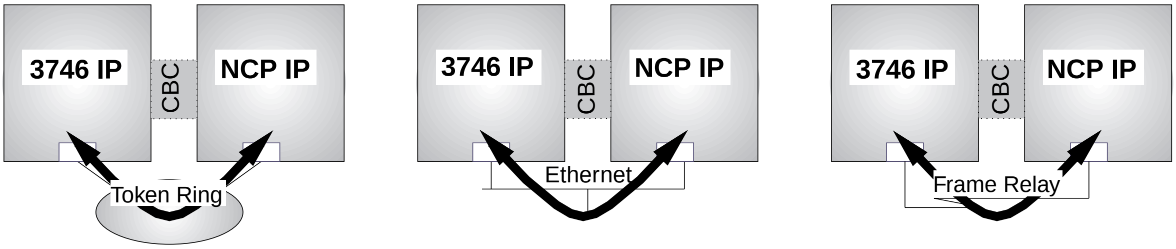

Figure 3-32. Direct NCP-IP 3746 Routing

Figure 3-33 on page 3-49 shows an alternative for the direct 3746 IP to NCP-IP attachment. A single ESCON channel can be used to connect TCP/IP for MVS (Version 3, Release 1 or 2) to both 3746 IP and NCP-IP. The result, however, of IP connectivity between 3746 IP and NCP-IP is that all IP datagrams that are sent between 3746 IP and NCP-IP will be routed via TCP/IP for MVS.

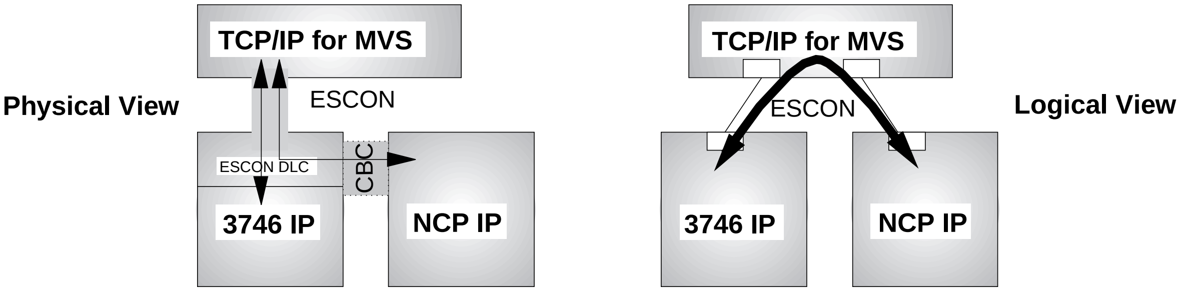

Figure 3-33. NCP-IP 3746 Routing Over ESCON

Telnet Operations via the 3746 Nways Multiprotocol Controller

The Telnet protocol provides a standardized interface through which a program on one host (the Telnet client) may access the resources of another host (the Telnet server) via a 3746 Nways Multiprotocol Controller as though the client were a local terminal connected to the server (see Figure 3-34). For example, a user on a workstation on a LAN may connect to a host as though the workstation were a terminal attached directly to the host.

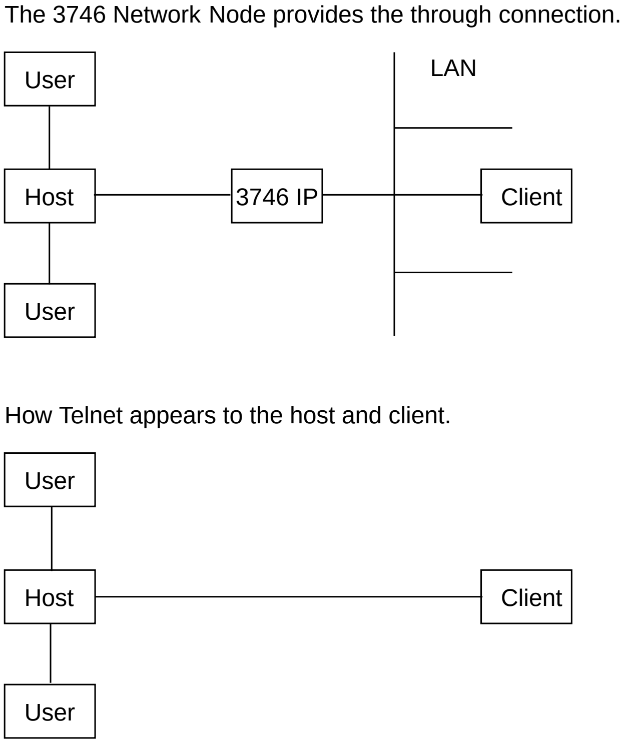
The 3746 Network Node provides the through connection.
User
LAN
Host
3746 IP
Client
User
How Telnet appears to the host and client.
User
Host
Client
User

Figure 3-34. Telnet Connection via a 3746 Nways Multiprotocol Controller

Telnet allows the LAN-attached user to log in the same way as the local terminal user.

Note: The 3746 Nways Multiprotocol Controller provides the connection but does not provide the actual host/client function. Refer to the documentation of the respective host and client systems for the functions available on those hosts/clients. This chapter gives an overview of Telnet possibilities.

Basic Operation

The Telnet protocol is based on three ideas:

- Network Virtual Terminal (NVT) concept.

  An NVT is an imaginary device having a basic structure common to a wide range of real terminals. Each host maps its own terminal characteristics to those of an NVT, and assumes that every other host will do the same.

- A symmetric view of terminals and processes.
- Negotiation of terminal options.

  The principle of negotiated options is used by the Telnet protocol, because many hosts wish to provide additional services, beyond those available with the NVT. Various options may be negotiated. After this minimum understanding is achieved, they can negotiate additional options to extend the capabilities of the NVT to reflect more accurately the capabilities of the real hardware in use. Because of the symmetric model used by Telnet, both the host and the client may propose additional options to be used.

Telnet operations, enabling and disabling, can be protected by passwords.

Network Virtual Terminal

The NVT has a printer (or display) and a keyboard. The keyboard produces outgoing data, which is sent over the Telnet connection. The printer receives the incoming data. The basic characteristics of an NVT, unless they are modified by mutually agreed options are:

- The data representation is 7-bit ASCII transmitted in 8-bit bytes.
- The NVT is a half-duplex device operating in a line-buffered mode.
- The NVT provides a local echo function.

All of these may be negotiated by the two hosts.

NVT Printer

An NVT Printer has an unspecified carriage width and page length. It can handle printable ASCII characters (ASCII code 32 to 126) and understands some ASCII control characters.

Full-Screen Capability

Full-screen Telnet is possible provided the client and server have compatible full-screen capabilities. For example, VM and MVS provide a TN3270-capable server. To use this facility, a Telnet client must support TN3270.

Command Structure

The communication between client and server is handled with internal commands, which are not accessible by users. All internal Telnet commands consist of 2 or 3-byte sequences, depending on the command type. An Interpret As Command (IAC) character is followed by a command code. This command deals with option negotiation, and has a third byte to show the code for the referenced option.

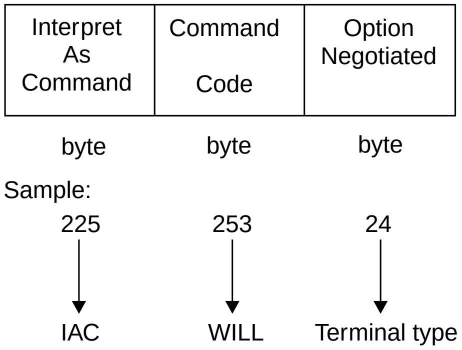

Figure 3-35. Telnet Command Structure

Host/Client Implementations

You can only use Telnet if it is implemented on ***both*** client ***and*** host machines, refer to your documentation for both. The respective documentation should also tell you the full range of Telnet functions available for those systems.

Access Control Mandatory Entry

The following entry in the access control list is mandatory and may *not* be removed.

```
PERMIT   0.0.0.0   0.0.0.0   0.0.0.0   0.0.0.0
```

If this entry is removed, NetView/6000 cannot gain access to the SNMP agent that resides on the NNP. No SNMP management of the 3746 will be possible.

Chapter 4. 3746 ATM Support

The Multiaccess Enclosure supports 155 Mbps ATM multimode and singlemode fiber interfaces (see “1-port ATM 155-Mbps Multimode Fiber Adapter - LIC284 (FC#3284)” on page 21-4 and “1-port ATM 155-Mbps Single-Mode Fiber Adapter - LIC293 (FC#3293)” on page 21-6 for details of the physical adapters). for details of the physical adapters). Using these interfaces, you can connect software that supports *Legacy LAN* (Ethernet or Token-ring) networks, and IP applications across an ATM network.

The two types of ATM connection methods supported are:

- LAN Emulation (LANE)
- Routing over ATM (Classical IP, or IPX over ATM)

The Multiaccess Enclosure implements the *LANE over ATM Version 1.0 Specification* (which is widely accepted as the industry standard for multivendor multiprotocol interoperability), and RFC1577 (Classical IP with ARP). The Multiaccess Enclosure uses ATMARP and InATMARP to get the ATM address of an IP node from an ARP server, but is not an ARP server itself. In addition, the Multiaccess Enclosure supports *Multiprotocol Encapsulation over AAL5* (RFC1483). This describes two encapsulation methods for carrying routed and bridged protocol data units (PDUs) over an ATM network. This includes:

LLC Encapsulation
Enables multiplexing of multiple protocols over a single ATM virtual circuit.

VC-Based Multiplexing
Each protocol is carried over a separate ATM virtual circuit.

This section gives a description of the ATM functions. If you need more detailed information about the ATM features of the Multiaccess Enclosure, please refer to *IBM 2210/2216 Nways Multiprotocol Router APPN/ATM Function Description and Configuration Scenarios,* SG24-4956.

ATM Addressing

ATM uses 20 byte addresses as shown below.

The first 13 octets of an ATM address are the network prefix. There are three kinds of this network prefix format type: ISO DCC, ISO ICD, ITU-T(CCITT) E.164. Each frame type is discriminated by the *AFI* (Authority and Format Identifier) field, which is above the network prefix. These AFI numbers are fixed:

ISO DDC ='0x39'

ISO ICD ='0x47'

ITU-T E.164 ='0x45'

Each switch in your ATM network must have a unique network prefix. ATM switches use the Network Prefix to route VCC setup requests to the destination ATM switch. End systems, such as this router, retrieve their network prefix from their ATM switch when they activate.

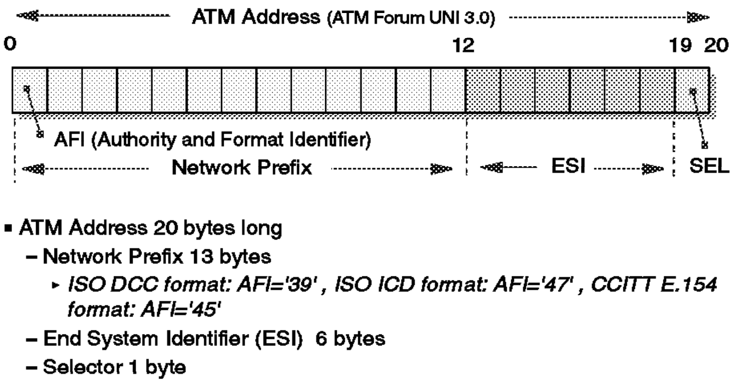
ATM Address (ATM Forum UNI 3.0)
0
12
19 20
AFI (Authority and Format Identifier)
Network Prefix
ESI
SEL

- ATM Address 20 bytes long
  - Network Prefix 13 bytes
    - *ISO DCC format: AFI='39' , ISO ICD format: AFI='47' , CCITT E.154 format: AFI='45'*
  - End System Identifier (ESI) 6 bytes
  - Selector 1 byte

Figure 4-1. ATM Address Format Overview

The next 6 octets (octets 14-19) of an ATM address are the End System Identifier (ESI). Each end system that is attached to the same switch must use a unique identification number.

When an end system activates, it attempts to register its ESI with its ATM switch by using the Interim Local Management Interface (ILMI). ILMI defines a set of SNMP-based procedures used to manage the interface between an end system and an ATM switch. End systems use ILMI to:

- Obtain the network prefix from the switch.
- Register their ESIs with the switch.
- Dynamically determine the UNI version of the ATM switch.
- Get a list of LECS addresses from the switch.

The octet 20 of an ATM address is the selector. End stations obtain their network prefix from the switch and from their own addresses by appending an ESI and selector. These addresses must then be registered with the switch, which rejects the registration if the ATM address is not unique.

Emulated and Virtual LANs

An ideal environment would only have ATM as the network transport layer and all applications would connect directly to this layer. Current network users have invested heavily in their physical networking infrastructure and applications that are built to use this infrastructure. The investment in existing networks means that a migration path must be offered to enable a gradual migration to ATM networks. In order to enable such *legacy* applications to use an ATM network without changes, it was necessary to define a method whereby the applications can run unchanged, and the networking hardware can gradually be migrated to ATM. The ATM Forum LAN Emulation over ATM Specification V1.0 allows such a migration.

LAN emulation supports all layer 3 protocols transparently, while RFC 1577, or Classical IP and ARP over ATM, is designed only to support IP as the layer 3 protocol.

LAN Emulation Version 1.0

LAN emulation enables the implementation of emulated LANs over an ATM network. An emulated LAN provides communication of user data frames across all of its users, similar to a physical LAN. One or more emulated LANs could run on the same ATM network. However, each of the emulated LANs is independent of the other and users cannot communicate directly across emulated boundaries. This is exactly the same as physical LANs. Inter-LANE communication will only be possible through routers or bridges.

Figure 4-2 shows a physical view of an emulated LAN. All LAN Emulation components are connected to the same switch. The software that is running above the LAN emulation software layers sees a logical view of the network shown in Figure 4-3. All LANE components connected to the switch, and using the same LANE server, appear to be in the same LAN.

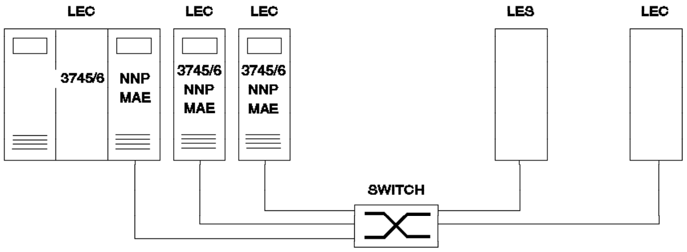

Figure 4-2. LAN Emulation Physical View

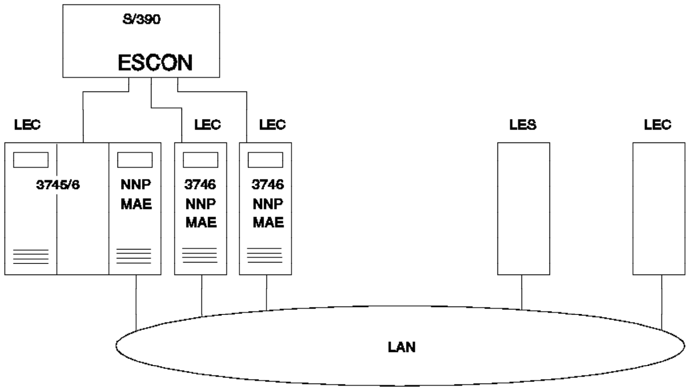

Figure 4-3. LAN Emulation Logical View

Each emulated LAN has only one type; either Ethernet/IEEE 802.3 or Token-Ring/IEEE 802.5. An emulated LAN is composed of several LAN emulation clients (LEC) and a single LAN emulation service entity (LE Service). The LE Service consists of an LE configuration server (LECS), an LE server (LES) and a broadcast and unknown server (BUS). The LE client resides in an ATM end station.

It represents a set of users, identified by their MAC address. The LE Service may be part of an end station or a switch. It can be centralized or distributed over a number of stations.

Communication between LE clients, and between LE clients and the LE Service, is performed over ATM virtual channel connections (VCCs). Each LE client must communicate with the LE Service over control and data VCCs. Emulated LANs operate in any of the possible environments:

- Switched virtual circuit (SVC)
- Permanent virtual circuit (PVC)
- Mixed SVC and PVC

In a PVC-only LAN, there are no call setup and shutdown procedures. Instead of that, layer management is used to set up and clear the connections. In this PVC environment, the layer management is responsible for both setting up and clearing connections and has the responsibility that the emulated LAN works correctly.

LAN Emulation Protocol Stack

LAN emulation over ATM operates within the data link layer of the OSI reference model. In Figure 4-4 on page 4-5, the position of the LAN emulation layer is shown.

The following list explains the flows depicted by the arrows in the figure:

1. LE to Higher Layer Services

   These services apply only to the LAN emulation client. The higher layer could be logical link control, or equivalent, or a bridging relay function. The services provide the capability to exchange user data frames over the LAN emulation service. Service definitions are compatible with ISO 10039 service architecture and ISO 10038 MAC bridging standard.

2. LAN Emulation to AAL Services

   These services apply to the LAN emulation clients and the LAN emulation service. These services provide the capabilities to transfer frames between peer LAN emulation layers. This specification assumes a null Service Specific Convergence Sublayer (SSCS), that is, the SSCS provides for the mapping of the equivalent primitives of the AAL and the Common Part Convergence Sublayer (CPCS). The common part of AAL5 makes use of the services provided by the underlaying ATM layer.

   A LAN emulation entity includes the following AAL service interfaces, each identified by a distinct SAP-ID. Each LAN emulation client includes the following SAPs:

   a. One or two control SAPs that handle initialization, registration and address resolution
   b. Two or more data forwarding SAPs
   c. Zero or one control SAP that handle configuration

3. Connection Management Services

   These services apply to the LAN emulation clients and the LAN emulation service.

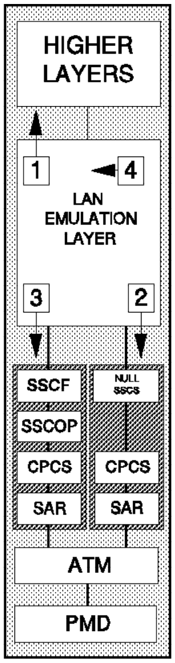

Legend:

- Higher Layers (802.2, NetBIOS, IPX,...)
- Control Plane (SAAL)
  - SSCF Service Specific Coordination Function
  - SSCOP Service Specific Connection Oriented Protocol
  - CPCS Common Part Convergence Sublayer
  - SAR Segmentation and reassembly Sublayer
- User Plane (AAL5)
  - SSCS Service Specific Convergence Sublayer
  - CPCS Common Part Convergence Sublayer
  - SAR Segmentation and reassembly Sublayer
- ATM ATM Layer
- PMD Physical Medium Dependent Layer

Figure 4-4. LANE Emulation Layers

The conceptual model assumed by the LAN emulation layer is shown in Figure 4-4 on page 4-5. The connection management services may use either PVCs or SVCs and provides the following primitives:

- Setup

  This service provides initial call establishment. It receives an ATM address and establishes a virtual connection identified by a SAP-ID.

- Release

  This service is used to request the network to clear an end-to-end connection identified by a SAP-ID.

- Add Party

  This service provides the capability to add a party to an existing connection.

  This service is used to drop or clear a party from an existing point-to-multipoint connection.

4. LAN Emulation to Layer Management Services

   These services enable initialization and control of the LAN emulation entities. These services differ between LAN emulation clients and the LAN emulation service.

LAN Emulation Components

An emulated LAN consists of several components. Clients, for example, ATM workstations and ATM bridges, each have at least one LE client entity. The LE Service consists of several components, including the LE server, the broadcast and unknown server, and the LAN emulation configuration server.

LAN Emulation Client (LEC)

The LAN emulation client is the entity in the end systems that performs data forwarding, address resolution, and other control functions. This provides MAC level emulated Ethernet/IEEE 802.3 or IEEE 802.5 service interface to higher-level software and implements the LUNI interface when communicating with other entities within the emulated LAN.

LE Server (LES)

The LE server implements the control coordination for the emulated LAN. The LE server provides a facility for registering and resolving MAC addresses and/or descriptors to ATM addresses. Clients may register the LAN destinations they represent with the LE server. A client will also query the LE server when the client wishes to resolve a MAC address and/or route descriptor to an ATM address. The LE server will either respond directly to the client or forward the query to other clients so they may respond. This is implemented by the Multiaccess Enclosure ATM adapter.

Broadcast and Unknown Server (BUS)

The broadcast and unknown server (BUS) handles data sent by an LE client to the broadcast MAC address ('FFFFFFFFFFFF'). All multicast and initial unicast frames that are sent by a LAN emulation client before the data direct VC target ATM address has been resolved (before a data direct VCC has been established) are handled by the broadcast and unknown server.

The multicast function provided in the BUS may be implemented by an underlying ATM multicast service. The BUS multicast function must be consistent with the ITU-T Recommendation X.6 Multicast Service Definition.

A LAN emulation client sends data frames to the BUS, which serializes the frames and retransmits them to a group of attached LAN emulation clients. Serialization is required to prevent AAL-5 frames from different sources from being interleaved.

In an SVC environment, the BUS needs to participate in the LE address resolution protocol (LE_ARP) to enable a LAN emulation client to locate the BUS.

The BUS must always exist in the emulated LAN and all LAN emulation clients must join its distribution group.

LE Configuration Server (LECS)

An ATM network can consist of several emulated LANs. The LE configuration server assigns LE clients to an emulated LAN based on its configuration database, its own policies and the information it receives from the respective LE clients. It assigns any client that requests configuration information to a particular emulated LAN service by giving the client the LES's ATM address. This method supports the ability to assign a client to an emulated LAN based on either the physical location (ATM address) or the identity of a LAN destination that it is representing.

It is optional for the LAN emulation client to obtain information from the LECS using the configuration protocol. The LECS allows the LAN emulation client to be automatically configured.

LAN Emulation Connections

VCCs are used for the connections of LAN emulation clients and other entities of LAN emulation such as the LECS, LES and BUS. For each different connection a separate VCC exists. In Figure 4-4 on page 4-5, the connections are shown with the stronger lines designating control connections.

Control Connections

Several control VCCs exist. They link the LEC to the LECS, or link the LEC to the LES, and carry LE_ARP traffic and control frames. A control VCC never carries data frames. Building control VCCs is a part of the LAN emulation client initialization process.

Three different control VCCs exist:

Configuration Direct VCC

The configuration direct VCC is a bidirectional VCC that may be set up by the LAN emulation client, or other entity, as part of the LECS connect phase. It is used for obtaining configuration information, including the address of the LES. The entity may maintain this VCC while participating in the emulated LAN. It may continue to keep it open for further queries to the LECS while participating in the emulated LAN. The configuration direct VCC may be used to inquire about an LE client other than the one to which the configuration direct VCC is attached. This connection is signaled using B-LLI to indicate it carries LE control packet formats.

Control Direct VCC

The LAN emulation client sets up a bidirectional point-to-point VCC to the LES for sending control traffic. This is set up by the LAN emulation client as part of the initialization phase. The LES has the option to use the return path for sending control data to the LAN emulation client. The LAN emulation client must thus accept control data from this VCC.

The LEC and LES must maintain this VCC as long as they are members of the emulated LAN.

Control Distribute VCC

The LES may optionally set up a unidirectional point-to-point or point-to-multipoint VCC to the LAN emulation client for distributing control traffic. This VCC may be set up by the LES as part of the initialization phase. If this VCC is set up, the LAN emulation client must accept the control distribute VCC.

The LEC and LES must maintain this VCC as long as they are members of the emulated LAN.

Data Connections

Data connections are used to connect LECs to LECs, and LECs to the broadcast and unknown server. The VCCs carry IEEE 802.3 or IEEE 802.5 data frames as well as flush messages. A flush message is the only time that a data connection will have control traffic. A flush message is generated by the flush protocol in order to ensure that the frames will remain in sequence when two data paths exist. This possibility exists when

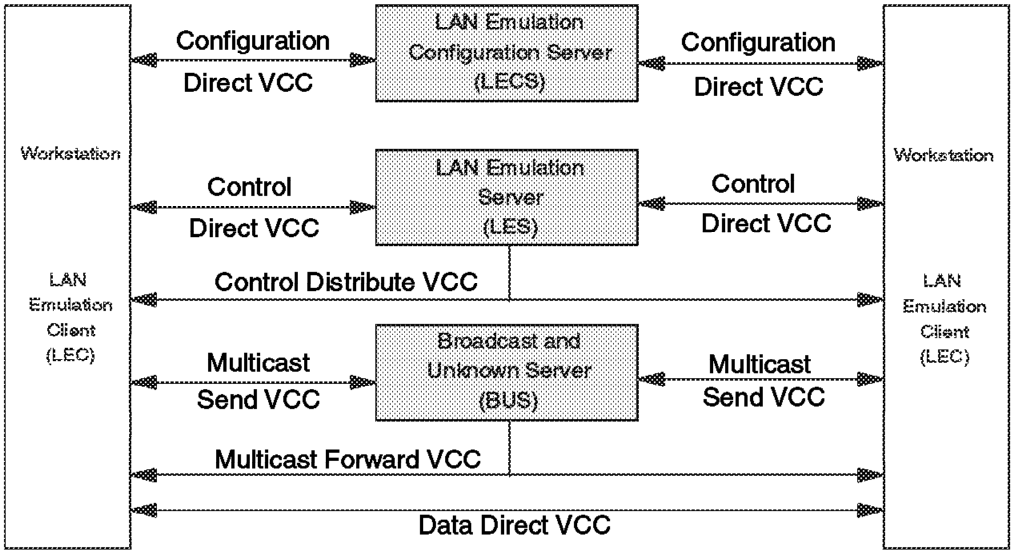

Figure 4-5. LANE Components

data is sent both via the broadcast and unknown server and via a direct VCC.

The following VCCs are defined:

Data Direct VCC

A data direct VCC is a bidirectional point-to-point VCC between LECs that wants to exchange unicast data traffic.

When a LAN emulation client has a packet to send and the ATM address for the destination MAC address is unknown, the LEC generates an LE_ARP request to resolve the ATM address for the destination. Once the LEC receives a reply to the LE_ARP, it sets up a point-to-point VCC, if not already established, over which to send all subsequent data to that destination.

The LEC that issues an LE_ARP request and receives an LE_ARP response is responsible for initiating the signalling to establish a bidirectional data direct VCC with the LEC sought in the LE_ARP request.

Multicast Send VCC

This VCC is used for sending multicast data to the BUS and for sending initial unicast data. The BUS may use the return path on this VCC to send data back to the LEC. This requires that the LEC accept traffic on this VCC.

A LAN emulation client sets up a bidirectional point-to-point multicast send VCC to the BUS. This VCC has the same setup as the data direct VCC. The LEC first sends an LE_ARP request and, when it receives the LE_ARP response, initiates signalling to establish the multicast send VCC to the BUS.

The LAN emulation client must maintain this VCC as long as it is a part of the emulated LAN.

Multicast Forward VCC

The multicast forward VCC is initiated by the broadcast and unknown server. This is done after the LAN emulation client has set up the multicast send VCC. The multicast forward VCC is used for distributing data from the broadcast and unknown server. It can be either a point-to-multipoint VCC or a unidirectional point-to-point VCC. The LEC emulation client must accept the multicast forward VCC regardless of type. A multicast forward VCC from the broadcast and unknown server must be established before a LAN emulation client can participate in the emulated LAN.

The LAN emulation client must attempt to maintain this VCC as long as it is a member of the emulated LAN.

The broadcast and unknown server may forward frames to a LAN emulation client on either the multicast send VCC or the multicast forward VCC. A LAN emulation client will not receive duplicate frames forwarded from the broadcast and unknown server on both the VCCs, but must be able to accept frames on either VCC.

LAN Emulation User-to-Network Interface

LE clients and the LE Service must interact in a well-defined manner. This is accomplished using PDUs and well-defined protocols. Four steps can be distinguished:

Initialization

- Obtaining the ATM address(es) of the LE Services that are available in an ATM network
- Joining or leaving a particular LAN specified by the ATM address of the LE Service
- Declaring whether this LE client wants to receive address resolution request for all the frames with unregistered destinations

Registration

Informing the LE Service of the following:

- The list of individual MAC addresses that the LEC client represents
- The list of source route descriptors (for example, segment/bridge pairs) that the LE client represents for source route bridging

Address resolution

Obtaining the ATM address representing the LE client with a particular MAC address (unicast, broadcast or segment/bridge pair).

Data transfer:

Moving the data from the source to the destination by:

- Encapsulation of the LE-SDU (service data unit) in an AAL-5 frame and transmission by the LE client
- Forwarding of the AAL-5 frame by the LE Service (if applicable)
- Reception and header removal of the AAL-5 frame by the LE client

LAN Emulation Functions

LAN emulation service is divided into seven functions. Some of the functions can be divided into several subfunctions. This chapter describes the functions and subfunctions as defined by the ATM Forum.

Initialization

Several steps are taken before the Initialization function is completed:

Initial State
In the initial state, there are parameters (for example, addresses, emulated LAN, maximum frame size, etc.) that are known to the LE server and the LE clients about themselves before they participate in the configuration and join phase functions.

LAN Emulation Configuration Server Connect Phase
In the LECS connect phase, the LE client establishes a configuration direct VCC to the LE configuration server.

Join Phase
In the join phase of ATM LAN emulation initialization, the LAN emulation client establishes its control connections to the LAN emulation server. The join phase can have two outcomes: success or failure.

Once the join phase has successfully completed, the LAN emulation client has been assigned a unique LAN emulation client identifier (LECID). It now knows the emulated LAN's maximum frame size and its LAN type. It also has established the control VCC(s) with the LAN emulation server.

Initial Registration
After joining, a LAN emulation client may register any number of MAC addresses and/or route descriptors. This is in addition to the single MAC address that can be registered as part of the join phase. Initial registration allows a LAN emulation client to verify the uniqueness of its local addresses before completing initialization and becoming operational.

Connecting to the Broadcast and Unknown Server
In order to establish a connection to the broadcast and unknown server, the LAN emulation client LE_ARPs for the broadcast MAC address and proceeds to set up the connection. The broadcast and unknown server then establishes the multicast forward VCC to the LAN emulation client.

Initialization Phases, Recovery and Termination

Registration

The address registration function is the mechanism by which clients provide address information to the LAN emulation server. An intelligent LAN emulation server may respond to address resolution requests if LAN emulation clients register their LAN destinations, defined as MAC addresses or, for source routing IEEE 802.5 LANs only, route descriptors, with the LAN emulation server. The LAN destinations may also be unregistered as the state of the client changes. A client must either register all LAN destinations for which it is responsible or join as a proxy.

Address Resolution

Address resolution is the procedure by which a client associates a LAN destination with the ATM address of another client or the broadcast and unknown server. Address resolution allows clients to set up data direct VCCs to carry frames.

When a LAN emulation client is presented with a frame for transmission whose LAN destination is unknown to that client, it must issue a LAN emulation address resolution protocol (LE_ARP) request frame to the LAN emulation server over its control point-to-point VCC.

The LAN emulation server may either:

1. Forward this LE_ARP to the appropriate client(s) using the control distribute VCC or one or more control direct VCCs. Different LAN emulation server implementations may use different distribution algorithms. If a client responds to a forwarded LE_ARP request with a LE_ARP reply, that reply is also sent and forwarded over the control VCCs to the original requester.
2. Or instead of forwarding the LE_ARP, the LAN emulation server may issue an LE_ARP reply on behalf of a client that has registered the requested LAN destination with the LAN emulation server.

A LAN emulation client must respond to an LE_ARP that it receives, asking for a LAN destination it has registered with the LAN emulation server or for which it is a proxy.

Each LAN emulation client maintains a cache of LE_ARP replies and uses a two-period time out mechanism to age entries in this cache. The aging time period is used for all entries learned from LE_ARP responses whose remote address flag was zero. That is, responses for registered LAN destinations are always timed out with the aging time. For aging entries learned from LE_ARP replies with the remote address FLAGS bit set to 1 and for entries learned from observing source addresses on data VCCs, which timeout to use is determined by the state of the LAN emulation client's topology change flag. When this flag is set, such entries are aged using the aging time parameter. The state of this flag may be altered either by local management or by reception of the LE_TOPOLOGY_REQUEST messages.

Connection Management

In switched virtual connection environments, the LAN emulation entities set up connections between each other using UNI signalling. The connections use best-effort quality of service as the minimum level.

Call establishment

When a call is being set up, the destination must not send its CONNECT message until it is ready to receive frames on the new VCC. The originator should expect that it can transmit frames after it has received the CONNECT message from the destination.

The CONNECT_ACK message is received by the destination and can be generated by its local switch. This message can reach the destination before the CONNECT message reaches the originator. The originator can only start to initialize its VCC after it receives the CONNECT message from the destination. Therefore, there is no guarantee for the destination that its initial data will be received by the originator until it receives some end-to-end indication from the originator.

The originator must send a READY_IND message as soon as it is ready to receive frames on the newly established VCC. At that point, the originator considers call establishment to be complete. The originator may also send data as soon as it is ready to receive frames on the newly established VCC. Data may be sent before or after sending the READY_IND.

It is possible that the READY_IND message can get lost. To recover it, the destination is responsible for timing the arrival of the READY_IND message. If the timer expires, the destination sends data or a READY_QUERY message on the VCC. Either party should always respond to receipt of a READY_QUERY message on an active VCC by transmitting a READY_IND message.

Tear down and timeout of VCCs

If a control direct VCC or control distribute VCC is ever released, a LAN emulation client must always return to the LAN emulation configuration server connect phase of initialization. If the broadcast and unknown server VCC is lost while the LAN emulation client is participating in a emulated LAN, the LAN emulation client may return to the broadcast and unknown server connect phase or go to the termination phase and restart.

Data Transfer

There are two different connections used for data transfer:

1. Data direct VCCs between individual LAN emulation clients
2. Multicast send and multicast forward VCCs that connect clients to the broadcast and unknown server

Unicast Frames

When a LAN emulation client has established, via the address resolution mechanism, that a certain LAN destination corresponds to a certain ATM address, and when that client knows it has a data direct VCC to that ATM address, then a frame addressed to that LAN destination must be forwarded via that data direct VCC.

If a LAN emulation client does not know which data direct VCC to use for a given unicast LAN destination, or if that data direct VCC has not yet been established, it may elect to transmit the frame over the multicast send VCC to the broadcast and unknown server. The broadcast and unknown server, in turn, forwards the frame to at least the client for which it is destined. If the LAN destination is unregistered, then the frame must be forwarded to at least all proxy clients and may be forwarded to all clients.

On an emulated LAN, the case can arise where a frame can only reach its destination through an IEEE 802.1D transparent bridge, and that bridge does not know the whereabouts of that destination. The only way such a frame can be assured of reaching its destination is for the frame to be transmitted to all of the IEEE 802.1D transparent bridges via the broadcast and unknown server so that they, in turn, can flood that frame to all of their other bridge ports, or at least the ones enabled by the spanning tree protocol. A LAN emulation client that chooses not to forward frames to the broadcast and unknown server, therefore, may not be able to reach destinations via transparent bridges, or perhaps other proxy agents.

Multicast Frames

LAN emulation clients may wish to send frames to a multicast MAC address, and/or they may wish to receive frames addressed to a given multicast MAC address. In order to send frames to a multicast MAC address, a LAN emulation client must send the frames to the broadcast and unknown server. The address resolution mechanism is used during the initialization process to provide the ATM address of the broadcast and unknown server for multicast and broadcast traffic, and connection management will provide a point-to-point multicast send VCC over which to send such frames.

All that is required in order for the LAN emulation client to receive frames addressed to a given multicast MAC address is for the LAN emulation client to connect to the broadcast and unknown server, after which the broadcast and unknown server will try to set up a return path for all broadcast and multicast traffic. When a client connects to the broadcast and unknown server, the broadcast and unknown server will try to establish a multicast forward VCC to that client. It is expected that multicast forward VCCs will be unidirectional point-to-multipoint VCCs, but they may be implemented as point-to-point VCCs. This decision is left to the LAN emulation service, not to the client.

A LAN emulation client will receive all flooded unicast frames and all broadcast and multicast frames over either its multicast send VCC or its multicast forward VCC. Which VCC the broadcast and unknown server uses to forward frames to the LAN emulation client is at the discretion of the broadcast and unknown server. A LAN emulation client will not, however, receive duplicate frames.

The LAN emulation header of any data frame sent from a client to the broadcast and unknown server must either contain the value 0 or the unique LECID value assigned to that client. The broadcast and unknown server is required to preserve the LAN emulation header of a relayed frame. Thus, a client can identify and filter frames that it sent by comparing the LECID field to its own LECID value. A transparent bridge LAN emulation client cannot reliably use the source MAC address to identify its own broadcast and unknown server traffic.

Token-ring functional addresses are treated just as any other multicast MAC address.

Frame Ordering

There may be two paths for unicast frames between a sending LAN emulation client and a receiving client: one via the broadcast and unknown server and one via a data direct VCC between them. For a given LAN destination, a sending client is expected to use only one path at a time, but the choice of paths may change over time. Switching between those paths introduces the possibility that frames may be delivered to the receiving client out of order. Out-of-order delivery between two LAN endpoints is uncharacteristic of LANs and undesirable in an ATM emulated LAN. The flush protocol is provided to ensure the correct order of delivery of unicast data frames.

Source Route Considerations

Source route bridging is the predominant bridging technology used within IEEE 802.5 token-ring networks. The use of source routing does not preclude transparent bridging in these networks. A token-ring end station will typically use a combination of source-routed and nonsource-routed frames. This allows a LAN emulation client to operate with both source routing and transparent bridging.

In addition to the Destination Address (DA) field and Source Address (SA) field, a source-routed frame contains a Routing Information (RI) field. The RI field contains a control field and a list of route descriptors (RD) that indicate the frame's path through the network. Therefore, the information in the RO field determines which SR bridges will forward the frame. The LAN emulation client determines if the frame is to be forwarded by an SR bridge or if the LAN destination is a station on the emulated LAN.

The LAN emulation client determines if the frame is to be forwarded by an SR bridge, or if the LAN destination is a station on the local emulated LAN by examining the RI field. If the LAN destination is accessible through an SR bridge, the LAN destination is the Next Route Descriptor (Next_RD); otherwise, the LAN destination is the frame's destination address.

Frames with specifically routed source routing information (an SRF frame) and unicast destination MAC address are sent down data direct VCCs following the usual LE_ARP and VCC setup process. Other source-routing frames are sent through the broadcast and unknown server.

Classical IP Overview

IP was the first network operating system that made use of ATM in a multivendor environment. Classical IP, or IETF RFC 1577, was the first standard available. Enhancements are made or proposed on all kinds of levels. This chapter gives an overview of work that is done or is under development and has IP as a base.

Classical IP over ATM (RFC 1577)

Since January 1993, the Internet Engineering Task Force (IETF) has had a formal recommendation on how to transport IP traffic over ATM. RFC 1577 describes the flows and mechanisms of Classical IP and ARP over ATM. RFC 1577 is the initial deployment of ATM within *classical* IP networks as a direct replacement for local area networks and for IP links that interconnect routers, either within or between administrative domains. The *classical* model refers to the treatment of the ATM host adapters as a networking interface to the IP protocol stack operating in a LAN-based paradigm.

Characteristics of the classical model are:

- The same maximum transmission unit (MTU) size is used for all VCs in a LIS.
- Default LLC/SNAP encapsulation of IP packets.
- End-to-end IP routing architecture stays the same.
- IP addresses are resolved to ATM addresses by use of an ATMARP service within the LIS; ATMARPs stay within the LIS. From a client's perspective, the ATMARP architecture stays faithful to the basic ARP model.
- One IP subnet is used for many hosts and routers. Each VC directly connects two IP members within the same LIS.

IP Subnetwork Configuration

In the LIS scenario, each separate administrative entity configures its hosts and routers within a closed logical IP subnetwork. Each LIS operates and communicates independently of other LISs on the same ATM network. Hosts connected to ATM communicate directly to hosts within the same LIS. Communication to hosts outside of the local LIS is provided via an IP router. This router is an ATM endpoint attached to the ATM network that is configured as a member of one or more LISs. This configuration may result in a number of disjoint LISs operating over the same ATM network. Hosts of differing OP subnets must communicate via an intermediate IP router even though it may be possible to open a direct VC between the two OP members over the ATM network.

The requirements for IP members (hosts, routers) operating in an ATM LIS configuration are:

- All members have the same IP network/subnet number and address mask.
- All members within a LIS are directly connected to the ATM network.
- All members outside of the LIS are accessed via a router.
- All members of a LIS must have a mechanism for resolving IP addresses to ATM addresses via ATMARP (based on RFC 826) and vice versa via InATMARP (based on RFC 1293) when using SVCs.
- All members of a LIS must have a mechanism for resolving VCs to IP addresses via InATMARP (based on RFC 1293) when using PVCs.

- All members within a LIS must be able to communicate via ATM with all other members in the same LIS; that is, the virtual connection topology underlying the intercommunication among the members is fully meshed.

The following list identifies a set of ATM-specific parameters that must be implemented in each IP station connected to the ATM network:

- The ATM hardware address (atm$ha) is the ATM address of the individual IP station.
- The ATMARP request address (atm$arp-req) is the ATM address of an individual ATMARP server located within the LIS. In an SVC environment, ATMARP requests are sent to this address for the resolution of target protocol addresses to target ATM addresses. That server must have authoritative responsibility for resolving ATMARP requests of all IP members within the LIS.

  Note: If the LIS is operating with PVCs only, then this parameter may be set to null, and the IP station is not required to send ATMARP requests to the ATMARP server.

Permanent Virtual Connections

An IP station *must* have a mechanism (for example, manual configuration) for determining what PVCs it has and, in particular, which PVCs are being used with LLC/SNAP encapsulation.

All IP members supporting PVCs are required to use the Inverse ATM Address Resolution Protocol (InATMARP) (refer to RFC 1293) on those VCs using LLC/SNAP encapsulation. In a strict PVC environment, the receiver will infer the relevant VC from the VC on which the InATMARP request (InARP_REQUEST) or response (InARP_REPLY) was received. When the ATM source and/or target address is unknown, the corresponding ATM address length in the InATMARP packet *must* be set to zero (0) indicating a null length, otherwise the appropriate address field should be filled in and the corresponding length set appropriately.

Switched Virtual Connections

SVCs require support for ATMARP in the nonbroadcast, nonmulticast environment that ATM networks currently provide. To meet this need, a single ATMARP server must be located within the LIS. This server must have authoritative responsibility for resolving the ATMARP requests of all IP members within the LIS.

The server itself does not actively establish connections. It depends on the clients in the LIS to initiate the ATMARP registration procedure. An individual client connects to the ATMARP server using a point-to-point VC. The server, upon the completion of an ATM call/connection of a new VC specifying LLC/SNAP encapsulation, will transmit an InATMARP request to determine the IP address of the client. The InATMARP reply from the client contains the information necessary for the ATMARP server to build its ATMARP table cache. This information is used to generate replies to the ATMARP requests it receives.

The ATMARP server mechanism requires that each client be administratively configured with the ATM address of the ATMARP server (atm$arp-req) as defined earlier in this chapter. There is to be one and only one ATMARP server operational per logical IP subnet. It is recommended that the ATMARP server also be an IP station. This station must be administratively configured to operate and recognize itself as the ATMARP server for a LIS. The ATMARP server must be configured with an IP address for each logical IP subnet it is serving to support InATMARP requests.

Enhancing RFC 1577

In RFC 1577, it was not possible to have more than one ATMARP server within a LIS. In the future there will be an environment with either a single or multiple synchronized servers. To make an ATMARP server capable of supporting server-to-server neighbor synchronization protocol and operations, several extensions must be made to the single ATMARP server model. These changes are still under consideration and are not discussed further here.

The MTU size is now negotiable. The same maximum transmission unit (MTU) is the default for all VCs in a LIS. However, on a VC-by-VC point-to-point basis, the MTU size may be negotiated during connection startup using Path MTU Discovery to better suit the needs of the cooperating pair of IP members or the attributes of the communications path. The Path MTU Discovery mechanism is Internet Standard RFC 1191 and is an important mechanism for reducing IP fragmentation in the Internet. This mechanism is particularly important because the new subnet ATM uses a default MTU size significantly different from older subnet technologies, such as Ethernet and FDDI.

In order to ensure good performance through the Internet, and also to permit IP to take full advantage of the potentially larger IP datagram sizes supported by ATM, all router implementations that comply or conform with this specification must also implement the IP Path MTU Discovery mechanism as defined in RFC 1191 and clarified by RFC 1435. Host implementations should implement the IP Path MTU Discovery mechanisms as defined in RFC 1191.

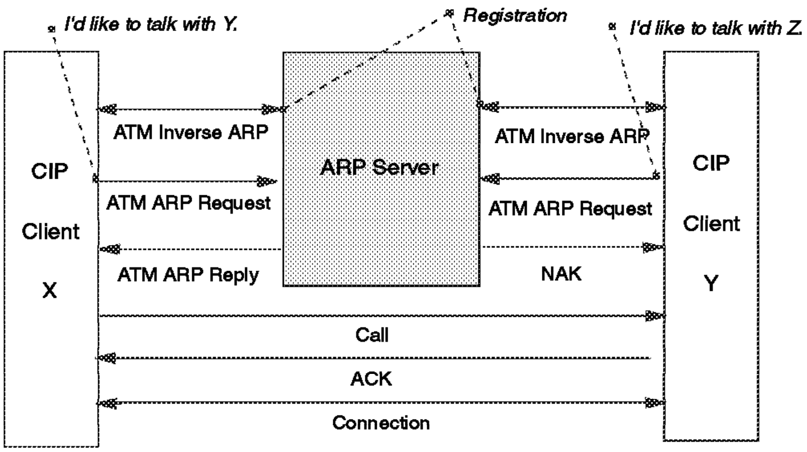

Figure 4-6. Classical IP Connection Overview

Differences between Classical IP and IPX over ATM

Configuring IPX over ATM (using RFC 1483) is similar to configuring Classical IP (RFC 1577). In the Multiaccess Enclosuresave configuration process, once you enter IPX as the protocol, some subsequent questions are different from those for the IP protocol. Since IPX over ATM does not use ARP servers, questions relating to ARP servers are not asked.

Also, IPX over ATM requires fewer parameters to be configured than CIP. The IPX network number and the IPX host number (IPX ATM-ARP-client) are the only required parameters for IPX over ATM. If you need to open a connection to a remote IPX router, you must additionally configure the desired channels (VCCs).

Chapter 5. Token-Ring/802.5

The IEEE 802.5 standard describes the token-ring medium access protocol and its physical attachments.

In a token-ring network the stations on the LAN are physically connected to a wiring concentrator usually in a star-wired ring topology. Logically, stations are connected in a pure ring topology. Each station has driver/transmitter as well as receiver circuitry; see Figure 5-1.

Differential Manchester code is used to convert binary data into signal elements, which are transmitted at 1, 4, or 16 Mbps (IEEE standard speeds). The standard does not prescribe the type of cabling to be used. In IBM's token-ring network implementation, shielded twisted pair cabling is recommended, although UTP may now be used.

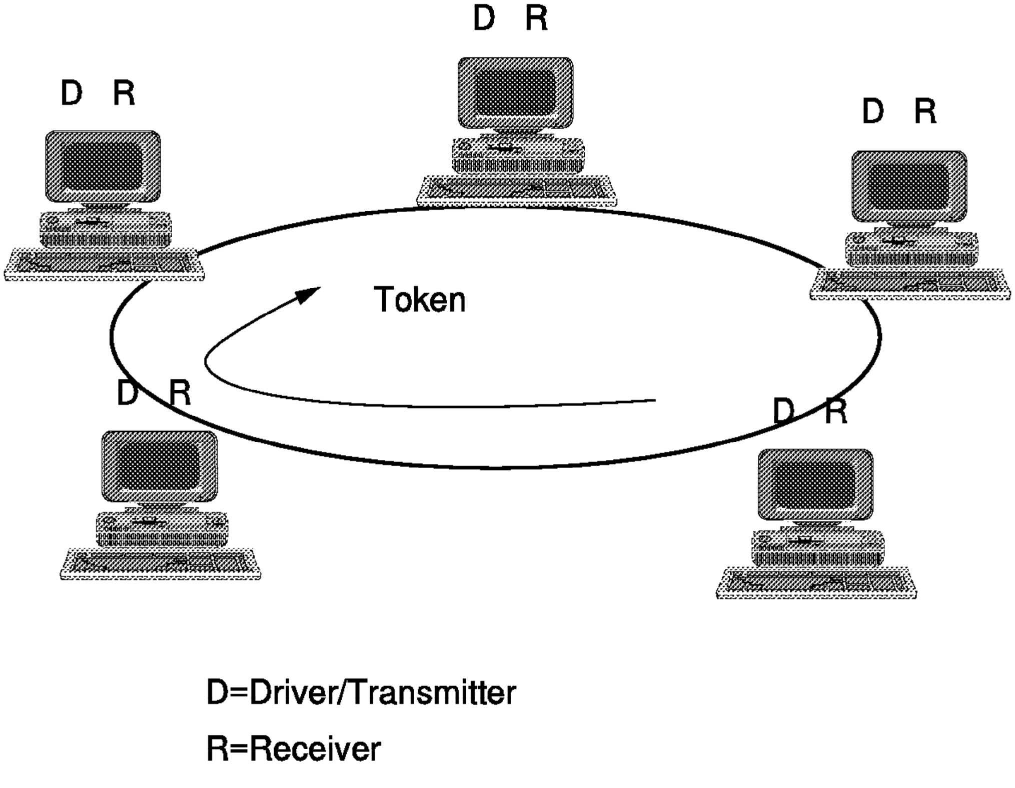

Figure 5-1. Sample Ring Configuration

Access to the ring is controlled by a circulating token. A station with data to transmit waits for a free token to arrive. When a token arrives, the station changes the token into a frame, appends data to it and transmits the frame. If the destination station is active, it will copy the frame and set the frame copied and address recognized bits, providing MAC level acknowledgment to the transmitting station. The sending station must strip the frame from the ring and release a new token onto the ring.

An option in the architecture allows the sending station to release a token immediately after transmitting the frame trailer, whether or not the frame header information has already returned. This is called early token release and tends to reduce the amount of idle time in 16 Mbps token-passing rings.

The token-passing protocol provides an extensive set of inherent fault isolation and error recovery functions, for implementation in every attaching device. The adapter network management functions include:

- Power-on and ring insertion diagnostics
- Lobe-insertion testing and online lobe fault detection
- Signal loss detection, beacon support for automatic test and removal
- Active and standby monitor functions
- Ring transmission errors detection and reporting
- Failing components isolation for automatic or manual recovery

The token-passing ring medium access protocol will be described in the following sections.

In summary, the token-passing ring protocol has the following elements:

- Active monitor
  - Ensures proper ring delay
  - Triggers neighbor notification
  - Monitors token and frame transmission
  - Detects lost tokens and frames
  - Purges circulating tokens or frames from the ring
  - Performs auto-removal in case of multiple active monitors
- Standby monitor (any other ring station)

  Detects failures in the active monitor and disruptions on the ring.

- Token claiming process

  A new active monitor is elected when the current active monitor fails. This process may be initiated by the current active monitor or by a standby monitor.

Figure 5-2 on page 5-3 shows the format of an 802.5 standard MAC frame as well as the token format and the format of the abort delimiter.

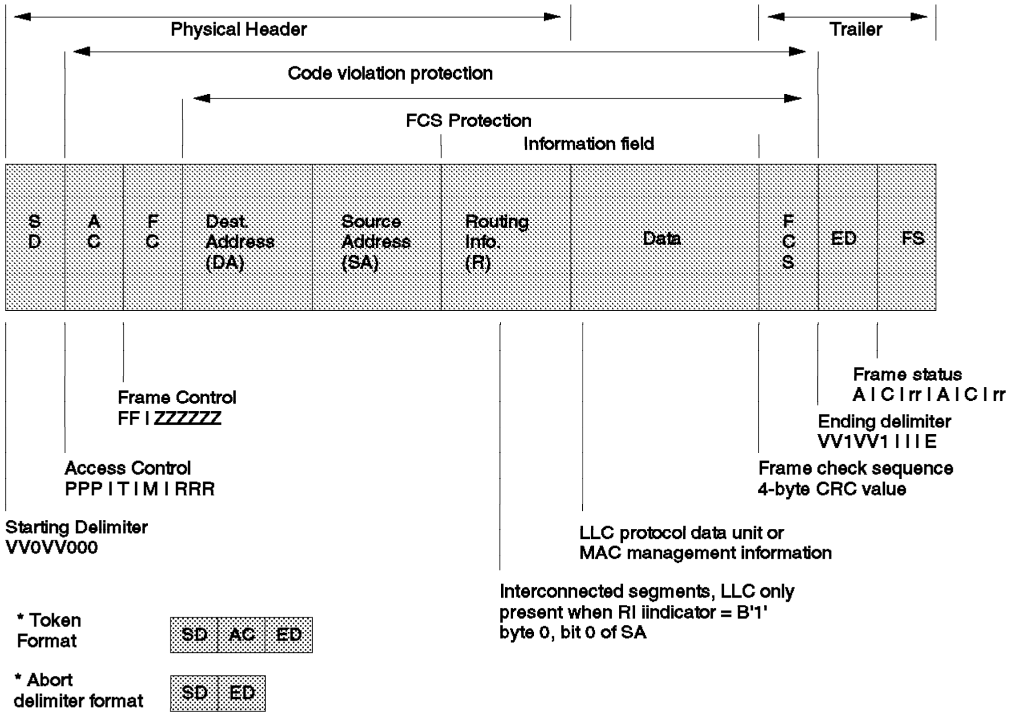

Figure 5-2. 802.5 Standard MAC Frame

The architecture describes 28 different MAC control frames, each identified by a unique major vector identifier (MVID). The main ones will be described in the following sections and are referred to as:

- Active Monitor Present MAC frame
- Ring Purge MAC frame
- Standby Monitor Present MAC frame
- Claim Token MAC frame
- Lobe Media Test MAC frame
- Duplicate Address Test MAC frame
- Request Initialization MAC frame
- Beacon MAC frame
- Soft Error Report MAC frame

Token-Ring Concepts

A token-ring network consists of the attaching medium and ring stations (devices able to attach to the ring and to use the link access protocols). A token-ring network uses one of several twisted pair media specifications, each having its own price/performance ratio, and all suitable to carry most other data communications signals. A token-ring network may also use optical fiber media.

A token-ring LAN installation is illustrated in Figure 5-3 on page 5-4.

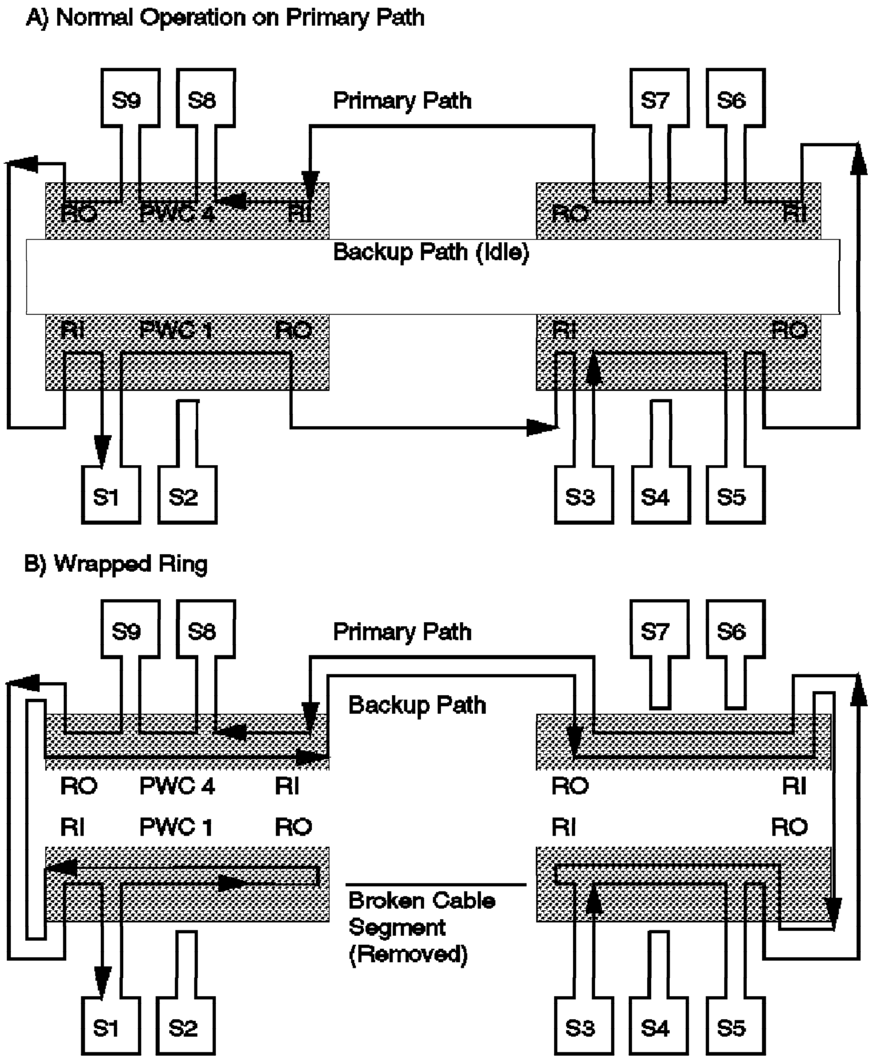

Figure 5-3. Normal and Wrapped Token-Rings

This figure shows four passive wiring concentrators (PWCs) and nine physically attached nodes. The power from the attached nodes when transmitted to the concentrator activates relays in the concentrator to allow the station to send signals across the LAN to other stations. In Part A of Figure 5-3, adapters (nodes S2 and S4) have not powered their respective PWC relays and, therefore, their lobe wires are internally bypassed. In part B of Figure 5-3, four adapters (nodes S2, S4, S6 and S7) are not actively inserted into the ring (their lobe wires are internally bypassed) and the primary path is wrapped to the backup path in PWCs 1 and 2.

When a cable segment between PWCs fails, manual removal from the appropriate ring-in and ring-out connectors causes automatic wrapping of the primary path to the backup path. How such a permanent wire fault is reported for LAN management is explained in the discussion of beaconing later on in this section. Recovery from the same error is automatic when using the IBM 8230 Controlled Access Unit, which is an active, or powered, wiring concentrator.

A ring station transfers data to the ring, in a data transmission unit called a frame. Frames are sent sequentially from one station to the next station physically active on the ring. This station is called the downstream neighbor. Each ring station repeats the frame. While doing so it performs error checking on the bit stream and it will copy the data if its own address, either its MAC or any of its functional addresses, are identified as a destination station in the frame. Upon return of the frame to the originating station, the latter will remove the data from the ring. In a token-passing protocol, a ring station can only transfer data to the ring while it is holding a token. The token is a specific bit sequence (24 bits) circulating around the ring at a rated speed of 100 Mbps according to the Fiber Distributed Data Interface specification. (4 Mbps or 16 Mbps in current implementations). Distributed Data Interface specification. Because of the high transmission speed with respect to the total ring length, a short ring might contain only a few bits at any point in time. Only one token may exist on a ring segment at any given point in time. Therefore, a delay equivalent to the time it takes for a token to circulate the ring is required to ensure that no overrun occurs which would result in a station receiving a token that it is transmitting and thinking that a second token exists on the ring. For a 24-bit token this means a minimum 24-bit delay. In addition to this delay, an additional elastic buffer is introduced to support the token protocols and speed.

In order to establish communication between any two ring stations, addressing mechanisms are needed. At the same time the integrity of the transmitted frames between ring stations must be preserved. Therefore data checking capabilities are required at the medium access control level of a ring station.

MAC Addressing

All ring stations are identified by a unique individual address. This address can be universally administered, assigned by the IEEE organization. Because it is set in read-only memory (ROM) on a token-ring adapter card, the universally administered address is also called a burned-in address.

Some manufacturers have been assigned universal addresses that contain an organizationally unique identifier. For instance, IBM has an identifier of x'10005A'. All IBM token-ring cards that use IBM token-ring chip sets have the first 6 digits of their address begin with those characters. Other identifiers are x'000143' for IEEE 802, and x'1000D4' for DEC. IEEE universal addresses, whether for token-ring or 802.3 stations, are all allocated out of the same common pool, but uniqueness is guaranteed.

A ring station's individual address can also be locally administered, that is set at adapter-open time and typically defined by a network administrator. A number of destination ring stations can be identified by a group MAC address. Some standard group addresses have been defined. These are listed in Table 5-1.

Table 5-1 (Page 1 of 3). Standardized Group Addresses

Bridge	X'8002 4300 0000'
Bridge management	X'8001 4300 0008'
Load server	X'8001 4300 0088'
Loadable device	X'8001 4300 0048'
ISO 10589 level 1 1 intermediate stations	X'8001 4300 0028'

Table 5-1 (Page 2 of 3). Standardized Group Addresses

ISO 10589 level 2 1 intermediate stations	X'8001 4300 00A8'
FDDI RMT directed beacon	X'8001 4300 8000'
FDDI status report frame	X'8001 4300 8008'
OSI network layer end stations	X'9000 D400 00A0'
OSI NL intermediate stations	X'9000 D400 0020'
Reserved for transparent bridging	X'8001 4300 000x'
All LANs bridge mgt group address (802.1D)	X'8001 4300 0008'
All cons end systems (ISO 10030)	X'8001 4300 0068'
All cons snares (ISO 10030)	X'8001 4300 00E8'
FDDI all root concentrator MACs (ANSI X3T9.5)	X'8001 4300 1004'
Reserved for FDDI	X'8001 4300 10X0'
Loopback assistance	X'F300 0000 0000'
AppleTalk support	X'9000 E000 0000'
AppleTalk highest address within range except broadcast	X'9000 E000 003F to X'9000 E0FF FFFF'
Novell IPX	X'9000 7200 0040'
Hewlett Packard probe	X'9000 9000 0080'
HP DTC	X'9000 9000 0020'
Apollo domain	X'9000 7800 0000'
Vitalink diagnostics	X'9000 3C40 00A0'
Vitalink gateway	X'9000 3CA0 0080'
LANtastic	X'FFFF 0006 0020'
LANtastic	X'FFFF 0002 0080'
LANtastic	X'FFFF 8007 0020'
Concord DTQNA	X'0000 9640 XXXX'
DEC DNA Dump/load assistance (MOP)	X'D500 0080 0000'
DEC DNA remote console (MOP)	X'D500 0040 0000'
DNA level 1 routing layer	X'D500 00C0 0000'
DNA routing layer end nodes	X'D500 0020 0000'
Customer use	X'D500 2000 XXXX'
System Communication Architecture	X'D500 2080 XXXX'
VAXELN	X'D500 D400 0040'
LAN traffic monitor	X'D500 D400 00C0'
CSMA/CD encryption	X'9000 D400 0060'
NetBIOS emulator (PSCG)	X'9000 D400 00E0'
Local area transport (LAT)	X'9000 D400 00F0'
All bridges	X'9000 D480 0000'
All local bridges	X'9000 D480 0080'

Table 5-1 (Page 3 of 3). Standardized Group Addresses

DNA level 2 routing layer routers	X'9000 D440 0000'
DNA naming service advertisement	X'9000 D440 8000'
DNA naming service solicitation	X'9000 D440 8080'
LAT directory service solicit (to slave)	X'9000 D440 8020'
FDDI ring purger advertisement	X'9000 D440 80A0'
LAT directory service solicit - X service class	X'9000 D440 80D0'
Local area system transport (LAST)	X'9000 D420 XXXX'
UNA prototype	X'5500 C000 XXXX'
Prom 23-365A1-00	X'5500 C080 XXXX'
Misc.	X'5500 C040 XXXX'
H400 - TA Ethernet transceiver tester	X'5500 C040 0000'
NI20 products	X'5500 C0C0 XXXX'
DECnet phase IV station addresses	X'5500 2000 XXXX'
Prom 23-365A1-00	X'1000 D40X XXXX'
Prom 23-365A1-00	X'1000 D48X XXXX'
Bridge mgt.	X'1000 D444 0000'
Prom 23-365A1-00	X'1000 D4C4 XXXX through X'1000 D4CX XXXX'
Shadow for prom 23-365A1-00	X'1000 D42X XXXX'
Shadow for prom 23-365A1-00	X'1000 D4AX XXXX'
Shadow for prom 23-365A1-00	X'1000 D4B6 XXXX' through X'1000 D4EX XXXX'
VAXft 3000 fault-tolerant LAN addresses	X'1000 D407 XXXX'
VAXft 3000 fault-tolerant LAN addresses	X'1000 D40F XXXX'

A token-ring LAN also provides a special case of a locally administered group address called functional addresses. Each (bit-significant) functional address represents a well-identified server function within the access protocol. Of 31 possible functional addresses, 22 have been defined while the remaining ones are reserved for future use or may be user-defined. They are listed in Table 5-2 on page 5-8.

Table 5-2. New and Current IEEE and IBM Functional Addresses

Active monitor	X'C000 0000 0001'
Ring parameter server	X'C000 0000 0002'
Network server heartbeat	X'C000 0000 0004'
Ring error monitor	X'C000 0000 0008'
Configuration report server	X'C000 0000 0010'
Synchronous bandwidth mgr	X'C000 0000 0020'
Locate - directory server	X'C000 0000 0040'
NetBIOS	X'C000 0000 0080'
Bridge	X'C000 0000 0100'
IMPL server	X'C000 0000 0200'
Ring authorization server	X'C000 0000 0400'
LAN gateway	X'C000 0000 0800'
Ring wiring concentrator	X'C000 0000 1000'
LAN manager	X'C000 0000 2000'
User-defined	X'C000 0000 8000' through X'C000 4000 0000'
ISO OSI ALL ES	X'C000 0000 4000'
ISO OSI ALL IS	X'C000 0000 8000'
IBM discovery non-server	X'C000 0001 0000'
IBM resource manager	X'C000 0002 0000'
TCP/IP	X'C000 0004 0000'
6611-DECnet	X'C000 2000 0000'
LAN Network Mgr & 6611	X'C000 40000 0000'

The most relevant protocol server functions will be described in greater detail in the Bridge and LAN management sections of *Local Area Network Concepts and Products: Routers and Gateways,* SG24-4755.

In addition, two special destination address values have been defined. The all-stations broadcast group address X'FFFFFFFFFFFF' identifies all ring stations as destination stations. A frame carrying the individual null address X'000000000000' as its destination MAC address is not addressed to any ring station; therefore, it can be sent but not received.

IEEE allows vendors to implement either 16-bit or 48-bit MAC addresses. The actual address field formats are shown in Figure 5-4 on page 5-9.

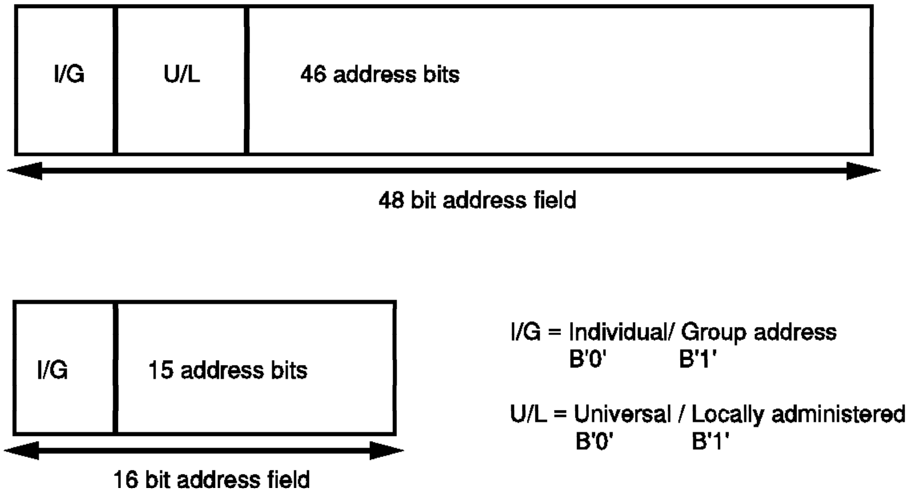

Figure 5-4. IEEE LANs - MAC Address Format

For the IBM LAN implementations, 48-bit addressing has been selected. The implementation format is shown in Figure 5-5.

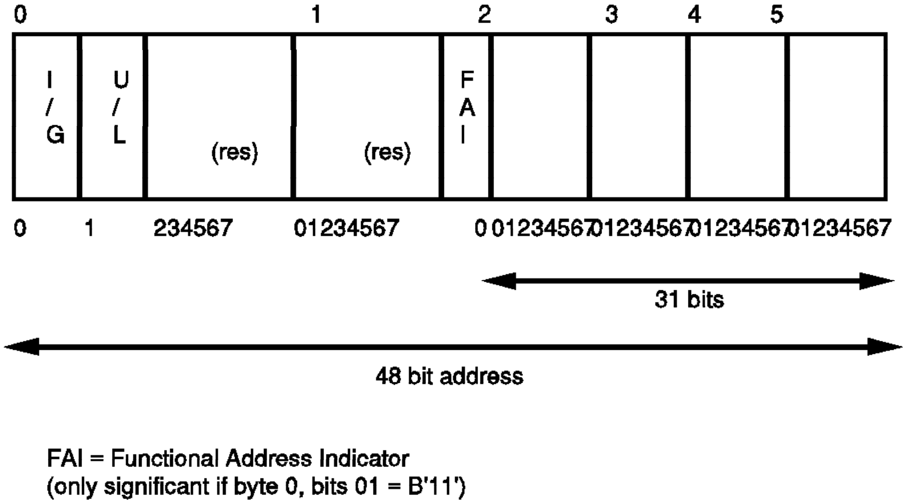

Figure 5-5. IBM Token-Ring Network - MAC Address Format

- The reserved bits are set to B'0' for locally administered addresses.
- Functional address indicator = B'0' indicates a functional address if I/G = B'1' (indicating a group address).

- For individual locally administered addresses, FAI must be B'0' by convention. This is an addressing anomaly.

These rules yield valid address ranges as described in Figure 5-6 for any IBM Token-Ring Network adapter.

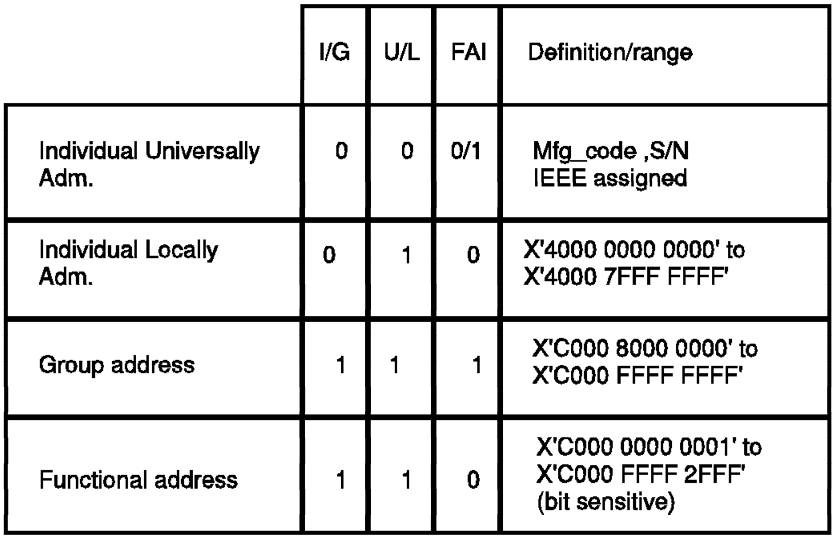

	I/G	U/L	FAI	Definition/range
Individual Universally Adm.	0	0	0/1	Mfg_code ,S/N IEEE assigned
Individual Locally Adm.	0	1	0	X'4000 0000 0000' to X'4000 7FFF FFFF'
Group address	1	1	1	X'C000 8000 0000' to X'C000 FFFF FFFF'
Functional address	1	1	0	X'C000 0000 0001' to X'C000 FFFF 2FFF' (bit sensitive)

Figure 5-6. Valid Address Ranges

Data Transmission

The transmission technique used in token-passing rings is baseband transmission. In a token-ring LAN, high-order bytes/bits are transmitted first; that is, byte 0 is transmitted before byte 1 and high-order bit 0 within a byte (of 8 bits) is transmitted first. This transmission order can be different for other types of LAN segments using different access protocols, for example, CSMA/CD. Opposite transmission order may be a diagnostic consideration when evaluating trace information from LAN segments of a different nature because of the possible need to reorder the bits. The ability to reorder the bits without significant performance degradation may also be a functional requirement of the bridge products being considered for a LAN segment interconnection. Figure 5-7 on page 5-11 shows the format of the information to be transmitted on a token-passing ring.

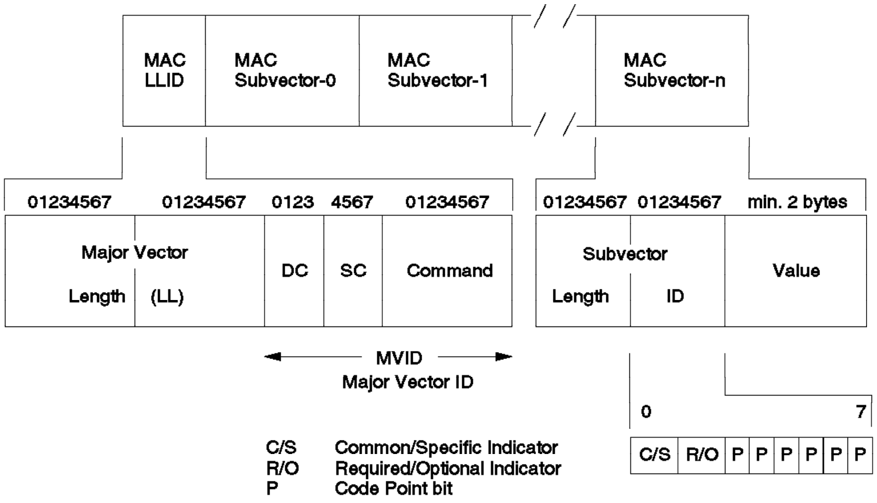

Figure 5-7. Token-Ring MAC Frame - Data Field Format

Examples of MVID code points are X'05' to indicate an Active Monitor Present MAC frame, X'02' for a Beacon MAC frame, etc. Unique MVID values for 28 different MAC frames have been defined. A complete description can be found in *IBM Token-Ring Network Architecture Reference,* SC30-3374. Subvectors provide additional information depending on the specific major vector identifier. MAC frames are processed according to destination and source function classes, including:

- **Ring Station**

  These are the functions necessary for connecting to the LAN and for operating with the token-ring protocols. A ring station is an instance of a MAC sublayer in a node attached to a ring.

- **DLC_LAN_MGR**

  The manager function of the data link control component activates and deactivates ring stations and link stations on command from the physical device. It also manages information transfer between data link control and the physical device.

- **Configuration Report Server (CRS)**

  The CRS function can reside on each ring in a multisegment environment in which configuration is being monitored. This function receives notifications about station insertion and removal, and notifications about active monitor failures.

- **Ring Parameter Server (RPS)**

  The RPS function can reside on every segment in a multisegment ring environment on which operational parameters are centrally managed. It may provide operational values to attaching stations upon request. For example, a ring station will request such parameters as ring number upon insertion into the ring.

- **Ring Error Monitor (REM)**

The REM server function is present on segments for which errors are to be monitored or analyzed. It collects error information from LAN stations attached to the local ring, analyzes soft error reports and possibly forwards error reports to a LAN Manager.

- **RPL Server**

  The RPL server function and its RPL functional address are involved during the power-on process of a LAN station equipped with the remote program load feature. Such a station will insert into the ring to find a control program server on the ring from which to download its control program and complete its initialization processes.

- **The Token-Passing Ring Protocol**

  To transmit data on the LAN medium, a ring station captures a token and sets the token bit in the access control field to identify that the data being transmitted is a frame. To this header, the transmitting station appends destination and source MAC addresses, data, a newly calculated frame check sequence field, and the ending delimiter and frame status fields.

  Any subsequent station will receive and retransmit the frame while performing a CRC check. Such a station is said to be in normal repeat mode.

  In general, a ring station in normal repeat mode checks the data in the tokens and frames it receives and sets the error-detected, address recognized or frame copied bits in a frame (bits E, A, or C) as appropriate while repeating the signal. A destination station will copy the data (frame copied) and pass the frame on. While processing the frame trailer, the destination station marks the A and C bits. Upon return to the originating ring station, the frame is removed from the ring and the A and C bits in the frame trailer's FS field are checked to see if the frame was recognized and read by the destination station or a bridge (this occurs for MAC frames only). When the frame header is received by the originating station the originating ring station must release a new token, possibly at a different priority level, for another ring station to capture and proceed with data transmission. The priority reservation bits in the access control field of the returned frame together with stored priority levels in the originating station determine the priority of the new token. See the description of access priority later in this chapter for details.

  This protocol is called a single token protocol, since only one token can circulate on the ring at any time.

Early Token Release

If an originating station releases a new token only when the frame header has circulated around the ring back to the source and the frame transmission time is shorter than the ring transmit time, then the originating station must generate idles until a header is received.

Token-passing ring protocols define the length of a token to be 2 bits and the shortest possible MAC frame to be 200 bits long. On a 4 Mbps token-ring LAN where the length of 1 bit is roughly 50 meters, a complete token is 1,200 meters long while the shortest frame length would be 10,000 meters. Therefore at 4 Mbps, the percentage of potential bandwidth which remains idle can be extremely small (that is high bandwidth utilization can be maintained at higher traffic levels).

If, however, we consider a 16 Mbps token-ring LAN, where 1 bit is 12.5 meters long, along with a complete token and the shortest possible MAC frame, both become four times smaller (300 and 2,500 meters respectively). Now we may wish to optimize the utilization of the medium by reducing the idle time required by waiting for a header. Obviously, when moving to even higher transmission speeds (for example, a 100 Mbps FDDI LAN), the token-passing protocol must be adjusted to achieve better utilization of the potential bandwidth.

The architecture provides an option called *early token release.* With this option a transmitting station will release the token after completing the transmission of the data frame before the receipt of the header of the transmitted frame; they thereby eliminate the idle time while waiting for the header to reappear. When such early release has occurred, an adapter indicator is set to prevent the adapter from releasing another token upon return of the frame header. This allows multiple frames but still only one token on the LAN.

The early token release option is enabled by default on the 16 Mbps IBM Token-Ring Network. It is an option for each station, and it is not required that all stations implement the option but is recommended.

Token Monitoring

Token-passing protocols provide relatively greater control and management at the medium access control (MAC) level than that provided by CSMA/CD protocols. The token-passing ring protocol concepts, described in the following sections, are implemented in the adapters themselves. They contribute to the availability, performance and manageability of a token-ring LAN. At any point in time, one and only one station per segment performs an active monitor function. Any ring station can act as the active monitor. Only one will have this function enabled. Active monitor tasks support the monitoring of the token and other ring management functions such as the following:

- Detection and recovery of a lost token or frame, including initiation of a token when a ring is started
- Detection and recovery of a circulating priority token or frame
- Detection and recovery of multiple tokens on the ring
- Detection and recovery of multiple active monitors
- Timing control to ensure accurate transmission regardless of the ring length

All other ring stations are said to be *standby monitors*, prepared to take over the active monitor function should it fail. The following description summarizes how the active monitor performs its ring management tasks:

- In every transmitted token or frame, the monitor bit (M) in the access control field is initially set to B'0'. As the active monitor repeats a frame or non-zero priority token, the M-bit is set to B'1'. If the M-bit had already been set to B'1', the active monitor assumes that the frame or token has already circled the ring once and that the originating station has not been able to remove the frame or priority token. The active monitor will purge the ring and generate a new token.
- To ensure that a complete (24-bit) token can be transmitted before the token returns to the originating ring station, the active monitor introduces a 24-bit ring delay.

- The active monitor periodically broadcasts (every seven seconds) an Active Monitor Present MAC frame. This forces each station on its ring to acquire the address of its nearest active upstream neighbor (NAUN) and to initiate a number of control timers within each station. This information is used when isolating errors on a segment.
- Loss of a token or frame is detected by expiration of an any-token timer whose timeout value exceeds the time required for the longest possible frame to circle the ring. The active monitor restarts this timer each time it transmits a starting delimiter. Upon expiration of this timer, the active monitor assumes a lost token or frame, purges the ring and originates a new token. The any-token timer value is defined as the sum of the physical trailer transmission delay plus the delay to transmit the longest frame. The IEEE 802.5 name for this timer is Valid_Transmission_Timer.

Ring Purge

To purge the ring, the active monitor broadcasts a Ring Purge MAC frame (indicated by X'04' in the frame control field) before originating a new token. Return of the Ring Purge MAC frame indicates proper signal propagation around the ring. The Ring Purge frame resets the ring stations to normal repeat mode, canceling or restarting all the appropriate timers. The active monitor starts a Ring-Purge timer when sending the purge frame. This timer will expire if the frame can't circulate and the monitor will enter a recovery process called a Claim Token.

Neighbor Notification

The neighbor notification process begins when the active monitor transmits an Active Monitor Present MAC frame to all stations on the ring (single ring broadcast). The first ring station that receives the Active Monitor Present MAC frame copies it (if possible) and sets the address-recognized (A) and frame-copied (C) bits to B'1'. It then saves the source address from the copied frame as its NAUN address (the address of the active monitor) and starts a timer called the Notification-Response timer.

All other active stations on the ring repeat, but do not otherwise process the Active Monitor Present MAC frame because the frame's A and C bits have already been set.

When the Notification-Response timer of the first station downstream from the active monitor expires, it broadcasts a Standby Monitor Present MAC frame.

The next ring station downstream copies its NAUN address from the source address field of the Standby Monitor Present frame, sets the A and C bits to B'1', and starts its own Notification-Response timer. When this timer expires, this station in turn transmits its Standby Monitor Present MAC frame.

In this way, neighbor notification proceeds around the ring, with each ring station receiving and transmitting Standby Monitor Present MAC frames until the active monitor copies its NAUN address from a Standby Monitor Present MAC frame. The active monitor then sets the Neighbor - Notification Complete flag to B'1', indicating that the process has been successfully completed. Neighbor notification thus enables a ring station to learn its NAUN address, and to provide its address to its downstream neighbor.

Standby Monitor

Any ring station that is not performing the active monitor function acts as a standby monitor. Its purpose is to detect a failing active monitor and disruptions on the ring.

Each time a token or frame is repeated, a standby monitor restarts its good-token timer to verify the presence of an active monitor.

A second timer, the Receive-Notification timer, is restarted by a standby monitor each time it copies an Active Monitor Present MAC frame. If any of these two timers expires, the standby monitor station will initiate the token-claiming process.

The Token-Claiming Process

This process, also called the monitor-contention process, is the procedure by which ring stations elect a new active monitor. This process is started upon any of the following conditions by:

- The active monitor detects the following:
  - Loss of signal.
  - The Active Monitor Present MAC frame doesn't return (Receive-Notification timer expires).
  - Failure of Ring-Purge MAC frames to return completely (Ring Purge timer expires).
- A standby monitor detects the following:
  - Loss of signal
  - Absence of active monitor's token management functions (good_token timer expires).
  - Missing Neighbor_Notification process (Receive-Notification timer expires).
- A ring station attaches to the ring and does not detect an active monitor (for example, when it is the first station on the ring).

The ring station detecting one of these conditions enters Claim-Token-Transmit mode by broadcasting a Claim Token MAC frame and repeating it at a defined interval. Each participating ring station compares the address in the Claim Token MAC frame's source address field to its own.

- If the source address is greater than the ring station's address, the station enters Claim Token Repeat operating mode.
- If the source address is less than the ring station's address, the station transmits its own Claim Token MAC frames.
- If the source address is the same as the ring station's address, it continues broadcasting until it has received three of its own Claim Token MAC frames. This indicates that the ring is viable and the station has won token-claiming.

  The station then adds the token delay to the ring, purges the ring, starts its active monitor timers, and releases a new token. It is now the new active monitor.

Ring Station Insertion

This process is executed by any ring station when entering the ring. It is also known as the five-phase insertion process.

- Phase 0: Lobe testing

  A series of Lobe Media Test MAC frames are sent on the lobe wire to the multistation access unit. The signal is wrapped at the entry into the multistation access unit causing the frames to return to the station. Then the receive logic is tested. If the tests are successful, a 5 volt DC current (also called phantom current) is sent to open the relay and attach to the ring.

- Phase 1: Monitor check

  The attaching station starts its Insert timer, and watches for an Active_Monitor_Present, Standby_Monitor_Present or Ring Purge MAC frame before this timer expires. If the timer expires, token-claiming is initiated. When it is on the first station on ring, the attaching station will become the active monitor.

- Phase 2: Duplicate address check

  The station sends a Duplicate Address Test MAC frame (destination address = source address = station's unique address). If a duplicate address is found (A-bit = B'1'), the station detaches from the ring.

- Phase 3: Participation in neighbor notification

  The station learns its nearest active upstream neighbor (NAUN) and reports its own address to its nearest active downstream neighbor.

- Phase 4: Request initialization

  A Request Initialization MAC frame is sent to the ring parameter server, if present (if not, default values will be used). The ring parameter server responds with an Initialize_Ring_Station MAC frame. Parameters which can be set are physical location, soft error report timer value, ring number and ring authorization level. In this way the last three parameters may be set to the same values for all stations on the ring.

Hard-Error Detection and Reporting

A hard error is a permanent fault that stops normal traffic on the ring. It is usually detected first at the receive side of the ring station downstream from the fault. A change in ring configuration is required to bypass such a fault and to restore normal operation. Reconfiguration may be the result of automatic recovery or, if this process fails to bypass the error, it may require manual intervention.

When a ring station detects a hard error, it starts transmitting beacon MAC frames at a specified time interval until its input signal is restored or until it removes itself from the ring. The detecting station also starts a Beacon timer. All other stations enter beacon repeat mode when they receive a beacon MAC frame.

A beacon frame identifies the address of the nearest active upstream neighbor of the beaconing station as well as error type information. When the beaconing station's NAUN has copied a number of these beacon frames, the NAUN will go offline and perform microcode and lobe tests. If the tests are successful, the station reattaches to the ring immediately. If the tests fail, the station stays offline.

When the beacon timer expires in the detecting (beaconing) station, and normal traffic has not been restored, the station assumes that its NAUN went offline, found no errors and came back online. It will now go through the same process as its NAUN. If the tests fail, the beaconing station remains detached. If successful, the station reattaches immediately. In the latter case, normal traffic may not have been restored during automatic recovery. Network management will be informed and manual intervention will be required. While reporting a permanent hard error, a set of adapter addresses is provided to identify the faulty part of the ring as a small fault domain.

Soft-Error Detection and Reporting

Intermittent faults that temporarily disrupt normal operation of the ring are called soft errors. They are usually tolerated by error recovery procedures but they may impair normal ring operation if excessive or non-random. The most critical soft errors are monitored in each ring station by a set of counters. Every two seconds the values of the soft error counters are sent as a Soft Error Report MAC frame to the Ring Error Monitor functional address (typically residing in a bridge or LAN Manager station), where the values for each counter are accumulated. If a soft-error counter exceeds a predefined threshold, a LAN Manager will be informed through its link with the LAN reporting mechanism. The LAN Manager may reconfigure the ring to bypass a faulty node, if the fault can be located.

Soft errors are said to be *isolating* if a fault domain can be specified. If not, they are called *non-isolating* soft errors.

Access Priority

The following discussion on access priority applies both to 4 Mbps and 16 Mbps token-ring LANs and is an integral part of the token-passing ring protocol. This access priority architecture is not applicable to the FDDI protocol where access priority is based upon timers rather than the contents of an access control field.

As stated earlier, access priority in a token or frame is indicated by the first three bits (PPP) of the access control field (AC). Any reservation of a priority level is indicated in the last three bits (RRR) of the AC field by a station requiring higher transmission priority.

A ring station wishing to transmit a frame at a given priority can use any available token with a priority level equal to or less than the priority of the frame to be transmitted. If such a matching token is not available, the ring station may reserve a token of the required priority in a passing token or frame according to the following rules:

- If the passing token or frame already contains a priority reservation higher than the desired one, the ring station must leave the RRR bits unchanged.
- If the RRR bits have not yet been set (RRR = B'000'), or indicate a lower priority than the desired one, the ring station will set the reservation bits to its required priority.

Upon removal of a frame by its originating station, the reservation bits in the header are checked. If they show a non-zero value, the station must release a non-zero priority token. The actual priority of the new token is based on the priority used by the ring station for the recently transmitted frame, the reservation received upon return of the frame and any stored priority.

A ring station originating a token of higher priority enters priority-hold state, (also called a stacking station in the IEEE 802.5 token-passing ring standards).

Table 5-3 lists the priority definitions as provided by the IBM Token-Ring Network architecture.

This protocol option however, impacts the priority handling mechanism, since a new token may be transmitted by the originating station before it is able to verify the access control field in its returned frame.

If the frame header was already received, the token will be issued according to the priority and reservation requested in the AC field of the frame and the resulting priority levels stored in the station.

If the frame header has not yet been completely received by the originating station, the token will be released with the same priority and reserved priority as the transmitted frame.

Table 5-3. Token-Passing Ring Protocol - Priority Allocation Table

B'000'	Normal user priority-MAC frames that need no token response type MAC frames
B'001'	Normal user priority
B'010'	Normal user priority
B'011'	Normal user priority - MAC frames that need token
B'100'	Bridge
B'101'	Reserved for IBM
B'110'	Reserved for IBM
B'111'	Specialized station management

To prevent a high-priority station from monopolizing the LAN medium and to make sure the priority eventually can come down again, the protocol provides fairness within each priority level.

Additional Token-Ring Considerations

Using an average frame size of 1000 bits to simulate the performance of a 4 Mbps token-passing ring with 100 active LAN devices results in a maximum throughput of about 3.6 Mbps. The token-passing protocol appears to be particularly stable and most efficient even under high load conditions.

The impact of increased transmission speeds, increased numbers of attached stations, or increased transmission distances on a token-passing LAN is significantly less than similar changes on a CSMA/CD LAN. Because each station regenerates the signal, increased distances are easier to support, while transmission speed is primarily limited by the choice of media. The use of bridges to provide additional device capacity and/or distance is an attractive growth option because the absence of collisions simplifies the processing requirements of bridges and maintains the deterministic characteristics of the protocols.

In a token-passing ring, fairness in the access protocol and high priority utilization by the bridge helps avoid frame loss. Even when a frame is rejected due to bridge congestion, successful recovery is simplified by the protocol.

Token-Ring Summary

The token-passing protocol provides for efficient use of the media under both light and heavy traffic loads. It guarantees fair access to all participating stations. This fairness is enhanced by an eight-level priority mechanism, based on priority reservations made in a passing token or frame. A key benefit of the token-passing ring protocol is its ability to handle increased traffic loads or peaks, making it an ideal protocol for larger and/or more heavily used LANs (including backbone rings). This also makes it a good base LAN for connection to even higher bandwidth LANs such as FDDI.

Chapter 6. Ethernet Overview

This chapter provides background information about IEEE 802.3/Ethernet, token-ring and FDDI local area networks. It is intended to provide the readers with a brief introduction to the various protocols and topologies used in designing today's local area networks.

Ethernet and IEEE 802.3

Ethernet (802.3) is currently the most widely used LAN protocol in the world. Since its introduction to the marketplace in the 1970s it has been established among a wide range of users.

Invented by Xerox in the early 1970s and brought to the marketplace as Ethernet V.1, the protocol was then developed by a consortium of DEC, Intel and Xerox. This consortium brought out a new version of Ethernet in 1980 called Ethernet (DIX) V2. They also published the architecture and took it to the Institute of Electrical and Electronics Engineers (IEEE) to have it accepted as an international standard. The IEEE ratified the Ethernet DIX V2 standards with some slight modifications as IEEE 802.3. The 802.3 standard has since been approved by a number of other organizations, including the American National Standards Institute (ANSI) and the International Organization for Standardization (ISO 8802-3).

Today both Ethernet and 802.3 LANs are widely implemented across all areas of the marketplace. It has not, as was widely predicted, been replaced by token-ring.

This is largely due to the fact that although the protocol used by Ethernet/802.3 LANs has not changed, the physical topology over which they can be implemented has changed significantly. This has enabled users to have access to some of the benefits (such as manageability) offered by other topologies such as token-ring while still enjoying the perceived advantages of Ethernet/802.3, which include:

1. Wide choice of equipment
2. Low cost of equipment

Though Ethernet and 802.3 are not identical, the term *Ethernet* is widely used to describe LANs that use either protocol. As most of the information in this chapter applies equally to both Ethernet and 802.3 LANS, the term Ethernet (802.3) will be used throughout this chapter. However, where there are differences, they will be indicated by using the appropriate terminology.

Note

Both Ethernet and the 802.3 protocol can be used on the same physical network simultaneously. However, stations using one protocol cannot interoperate with stations using the other protocol. This is due to the differences that will be explained later in this chapter.

Please note that this document will cover Baseband Ethernet only.

CSMA/CD

Carrier Sense Multiple Access/Collision Detection (CSMA/CD) is the name of the protocol used on the Ethernet (802.3) bus to control the operation of the network.

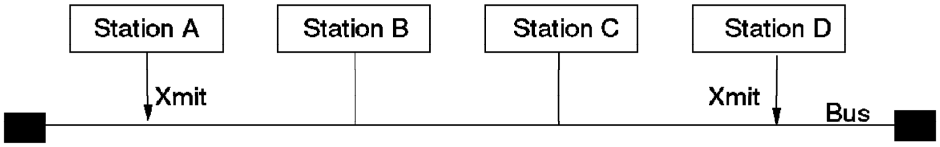

Figure 6-1. Ethernet CSMA/CD Bus

In a CSMA/CD bus, when a station wants to transmit data on the network bus, it first listens to see if the bus is free (that is no other station is transmitting). If the bus is available, the station starts transmitting data immediately. If the bus is not available (that is another station is transmitting), the station waits until the activity on the bus stops and a predetermined period of inactivity follows before it starts transmitting.

If there is a *collision* after transmission (that is another station starts to transmit at the same time), the stations will stop transmitting data immediately after the collision is detected, but they continue to transmit a jamming signal to inform all active stations about the collision.

In response to this signal, each transmitting station stops transmitting and uses a binary exponential backoff algorithm to wait before attempting to transmit again. This causes each station to wait for a random amount of time before starting the whole process again, beginning with the process of carrier sensing. If a station's subsequent attempt results in another collision, its wait time will be doubled.

This process may be repeated up to 16 times, after which the station, if still unsuccessful, reports a transmission error to the higher layer protocols.

The process of *collision detection* varies according to the type of media used in the LAN. This process is described in “CSMA/CD.”

The probability of a collision occurring is proportional to the number of stations, the frequency of transmissions, size of frames and length of the LAN. Therefore, care must be exercised in designing LANs with an excessive number of stations that transmit large packets at frequent intervals. Also, you must ensure that the length of individual *segments* and total length of the LAN does not exceed a certain length as defined by the 802.3 standards. These limitations are discussed later in this topic.

According to Ethernet and the 802.3 standard, to be able to detect collisions, a transmitting station should monitor the network for a period of time called a *slot time*. Slot time is the time during which a collision may occur and is the maximum delay for a transmission to reach the far end of the network and for a collision to propagate back. Slot time is defined to be 51.2 microseconds (512 bit times in a 10M Mps LAN). This time imposes a maximum length on the size of the network. It also imposes a minimum (64 bytes, excluding preamble and FCS) on the size of the frames transmitted by each station.

Frame Formats

The frame formats for Ethernet and IEEE 802.3 are not the same. However, both protocols use the same medium and access method. This means that, while LAN stations running these protocols could share a common bus, they could not communicate with each other.

Ethernet Frame Format

The layout of an Ethernet frame is as follows:

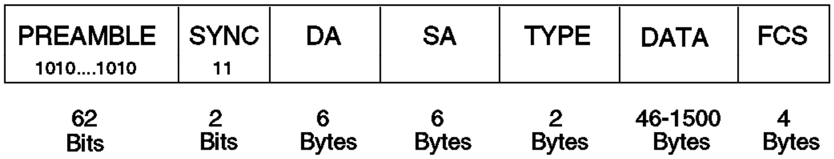

PREAMBLE 1010....1010	SYNC 11	DA	SA	TYPE	DATA	FCS
62 Bits	2 Bits	6 Bytes	6 Bytes	2 Bytes	46-1500 Bytes	4 Bytes

Figure 6-2. Ethernet Frame Format

- PREAMBLE - 62 bits, allows the Physical Layer Signalling (PLS) circuitry to synchronize with the receive frame timing circuitry.
- SYNC (Synchronize) - 2 bits, indicates that the data portion of the frame will follow.
- DA (Destination Address) and SA (Source Address) - 48-bits Media Access Control (MAC) address. Three types of destination addressing are supported:
  - Individual: The DA contains the unique address of one node on the network.

- Multicast: If the first bit of the DA is set, it denotes that a *group* address is being used. The *group* that is being addressed will be determined by a higher layer function.
- Broadcast: When the DA field is set to all 1s, it indicates a *broadcast*. A broadcast is a special form of multicast. All nodes on the network must be capable of receiving a broadcast.

- TYPE (Type Field) - 16 bits long, this field identifies the higher layer protocol that is used. Vendors must register their protocols with the Ethernet standards body if they wish to use Ethernet Version 2.0 transport. Each registered protocol is given a unique 2-byte *type* identifier. As this field is used as the *length* field by the 802.3 frames, the value assigned to the *type* field in Ethernet is always higher than the maximum value in the *length* field for the 802.3. This is to ensure that both protocols can coexist on the same network.
- Data field - This contains the actual data being transmitted and is 46-1500 bytes in length. Ethernet assumes that the upper layers will ensure that the minimum data field size (46 bytes) is met prior to passing the data to the MAC layer. The existence of any padding character is unknown to the MAC layer.
- FCS - 32 bits long, the result of a cyclic redundancy check algorithm (specific polynomial executed against the contents of DA, SA, length, information and pad fields). This field is calculated by the transmitting station and is appended as the last four bytes of the frame. The same algorithm is used by the receiving station to perform the same calculation and the results are compared with the contents of the FCS field in the received frame to ensure that transmission was error free.

IEEE 802.3 Frame Format

The layout of the IEEE 802.3 frame format is as follows:

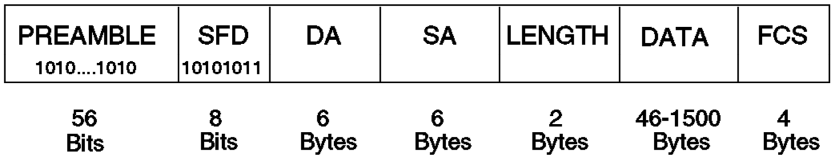

PREAMBLE 1010....1010	SFD 10101011	DA	SA	LENGTH	DATA	FCS
56 Bits	8 Bits	6 Bytes	6 Bytes	2 Bytes	46-1500 Bytes	4 Bytes

Figure 6-3. 802.3 Frame Format

- PREAMBLE - 56 bits long, allows the Physical Layer Signalling (PLS) circuitry to synchronize with the receive frame timing circuitry.

- SFD (Start Frame Delimiter) - 8 bits long, indicates that the data portion of the frame will follow.
- DA (Destination Address), SA (Source Address) - 48-bit Media Access Control (MAC) address. Three types of destination addressing are supported:
  - Individual - The DA contains the unique address of a node on the network.
  - Multicast - If the first bit of the DA is set, it denotes that a group address is being used. The *group* that is being addressed will be determined by a higher layer function.
  - Broadcast - When the DA field is set to all 1s, it indicates a *broadcast*. A broadcast is a special form of multicast. All nodes on the network must be capable of receiving a broadcast.
- LF (Length Field) - 16 bits long, indicates the number of *data* bytes (excluding the PAD) that are in the data field.
- DATA and PAD field - IEEE 802.3 (and Ethernet) specify a minimum packet size (header plus data) of 64 bytes. However, 802.3 permits the *data field* to be less than the 46 bytes required to ensure that the whole packet meets this minimum. In order to ensure that the minimum packet size requirement is met, 802.3 requires the MAC layer to add *pad* characters to the LLC data field before sending the data over the network.
- FCS - 32 bits long, the results of a cyclic redundancy check algorithm (specific polynomial executed against the contents of DA, SA, length, information and pad fields). This field is calculated by the transmitting station and is appended as the last four bytes of the frame. The same algorithm is used by the receiving station to perform the same calculation and the results are compared with the contents of the FCS field in the received frame to ensure that transmission was error free.

Chapter 7. Frame Relay Overview

This chapter examines the frame relay functions offered by the 3745 and 3746. It supplies the reader with an overview of the frame relay functions supported by NCP and the 3746 NNP, explains how IBM 3745 and 3746 Controllers can be used to build frame relay networks, and how to connect to frame relay networks built using 3745s and 3746 Model 900s. It also discusses the implementation of frame relay in the 3746-950 and 3746-900 under NNP control.

Frame relay is supported on all leased line attachments on the low-speed scanners, high-speed scanners, and CLP (3746 Models 900 and 950) serial ports with line speeds up to 2 Mbps.

Frame relay functions implemented in NCP and IBM 3745/3746 Models 900 and 950 comply with the ITU-T and ANSI Frame Relay standards, and to RFC1490 *Multiprotocol interconnect over Frame Relay* and the SNA extensions. Fragmenting as described within RFC 1490 is not supported; SNA segmenting however is supported.

The ability to manage a network is an important consideration. 3745/3746 frame relay is supported by the IBM network management products NetView, and NPM (NetView Performance Monitor), and provides a variety of well-known tools (Generic Alerts, PDSTATs, RECMS, NTune, etc.).

Note: For more information refer to the following publications:

- *IBM 3746-900 and NCP Version 7 Release 2*, GG24-4464
- *IBM 3746 Nways Controller Models 900 and 950: APPN Implementation Guide*, SG24-2536
- *IBM Frame Relay Guide*, SG24-4463
- *3746 Nways Controller 950 and 900 IP Implementation*, SG24-4845

ACF/NCP Frame Relay Support (3745, 3746-900)

The following sections give an overview of frame relay support of ACF/NCP for the 3745/3746-900.

ACF/NCP Version 6 Release 1

NCP provided its first frame relay support in NCP V6R1, which was available August 1992.

NCP V6R1 runs on all models of the 3745 Communications Controller family, providing a wide range of capacities and throughput. This frame relay support is implemented strictly in the NCP software without requiring any new hardware. It runs on the existing 3745 hardware and line adapters, including the T1/E1 High Speed Scanner. No host application software changes are required to use NCP frame relay.

Functionally, NCP V6R1 provides the ability to interconnect NCPs to each other through a frame relay network, providing NCP-to-NCP subarea traffic support. NCP acts as a frame relay end station, called Frame Relay Terminal Equipment (FRTE). This is referred to as NCP's INN FRTE function.

Note: A Frame Relay End Station is known as an FRTE but is sometimes referred to as an FR DTE (Data Terminal Equipment) in analogy with X.25 terminology.

ACF/NCP Version 6 Release 2

NCP V6R2 (June 1993) provided a general purpose frame relay switching capability for multiprotocol, SNA and non-SNA, traffic. The switching capability is referred to as the Frame Relay Frame Handling (FRFH) function.

This allows connection of FRTEs to a network of 3745s and provides end-to-end connections (PVCs) for FRTE pairs. Like the initial release, NCP V6R2 runs on all models of the 3745 Communications Controller and on existing line adapters, including the High Speed Scanner (HSS). NCP frame relay allows attachment using leased lines at speeds from 600 bits per second up to 2 Mbps. FRFH and INN FRTE functions can be shared on a single frame relay line.

Note: The FRFH function is often referred to as FR DCE (Data Communications Equipment) in analogy with X.25 terminology. In NCP publications the term Frame Relay Switching Equipment (FRSE) is used as well.

In addition, LMI support was added in this release.

ACF/NCP Version 7 Release 1

NCP V7R1 (January 1994) provides SNA peripheral device connectivity via frame relay This is referred to as NCP's BNN FRTE function, and uses the RFC1490 routed frame format. BNN FRTE, FRFH and INN FRTE functions can be shared on a single frame relay line. IBM networking equipment such as the IBM AS/400, IBM 3174, IBM 3172, and RouteXpander/2. can be used to directly access the SNA network, making use of the boundary function, and also transfer other frame relay traffic over different virtual circuits using the NCP FRFH support.

The frame relay BNN capability of ACF/NCP V7R1 supports multiple stations per DLCI for 3745 adapters (not 3746-900 attached).

ACF/NCP Version 7 Release 2

NCP V7R2 (October 1994) provides support for frame relay connections on the 3746 Model 900 that use the LIC11 and LIC12 couplers. The frame relay support on the 3746 Model 900 is similar to the support in the base frame. This support includes:

FRTE support

- Subarea frame relay connections to other NCPs
- Frame relay virtual circuits from peripheral devices to the boundary function in NCP

FRFH support

Switching of frames from a data link connection identifier (DLCI) of one port (called a subport) to another DLCI of another port. The switching paths are downloaded from NCP to the CLP at activation time. The frame switching between 3746-900 subports is done via the 3746-900 connectivity switch (or within the CLP) without going through the CCU. This improves the switching capacity of the 3745.

The same hardware as for SDLC and X.25 is used. SDLC, X.25, and frame relay lines can be attached to the same CLP.

Data Link Control for SNA support in the 3746-900

The RFC1490 routed format is used to identify the type of SNA support:

- Boundary Network Node (BNN)
- Intermediate Network Node (INN)

For SNA over frame relay 802.2 is used to provide a reliable transport. Both RFC1490 and 802.2 LLC run in the CLPs of the 3746-900.

Mixed media multilink transmission groups

Mixed media multilink transmission groups (MLTGs) are supported over an FRTE SA subport of the 3746-9x0.

3746 LMI Support

LMI support was added for 3746 Model 900 connected lines. The LMI support in the 900 will free up 3745 CCU cycles.

Communication Rate Support

NCP V7R2 introduces communications rate (CR) support for the 3746-900, which allows users to allocate a minimum bandwidth to each virtual circuit, depending on the traffic needs of the corresponding end stations. This guarantees that traffic will flow on a given virtual circuit at least at its communications rate. At the same time, any unused bandwidth is made available for use in active virtual circuits.

Comrate

NCP V7R3 extends the comrate support to the 3745 lines. The actual bandwidth assigned to a virtual circuit is defined either by COMRATE or CIR. These two parameters are mutually exclusive; either one or the other must be defined.

CIR is supported only by 3746 lines which are controlled by the NNP or IP components of the 3746. There is no NCP support for CIR (see “Committed Information Rate (CIR) and Burst Sizes” on page 7-24 for details on CIR). The 3746 NNP and IP lines also support COMRATE. See “Communication Rate (CR) and Committed Information Rate (CIR)” on page 7-26 for details on COMRATE.

Thus frame relay lines that are shared between NCP (V7R5), NNP and 3746 IP functions only support COMRATE. For details on COMRATE, See “Communication Rate (CR) and Committed Information Rate (CIR)” on page 7-26

Note: The communications rate must be defined per PVC segment. An end-to-end communications rate requires consistent definitions on each of the PVC segments that comprise a virtual circuit.

ACF/NCP Version 7 Release 3

NCP V7R3 (March, 1995) had the following frame relay support enhancements:

Multiple BNN stations per DLCI for 3746 lines

Support of multiple BNN stations per DLCI on 3746 connected lines was added.

Data Link Control for SNA support in the 3746-900

The RFC1490 routed format is used to identify the type of SNA support

- Advanced Peer-to-Peer Networking (APPN)
- Boundary Access Node (BAN)

The RFC 1490 bridged format, in conjunction with NCP V7R3 and above, can be used to carry traffic from downstream physical units (PUs) attached to the IBM 2210 and 6611 routers. This support is called *boundary access node* (BAN).

Frame Relay over Token-Ring

Frame relay over the TIC2 token-ring in the 3745 gives ACF/NCP Version 7 Release 3 the capability to support frame relay frame handler functions (FRFHs) between NCPs over token-ring (IEEE 802.5) physical connections. This will allow customers who interconnect NCPs with token-ring to provide a private frame relay network over these token-ring connections.

Frame relay over token-ring requires ACF/VTAM Version 4 Release 2 with the appropriate PTF.

IP over Frame Relay

NCP only supports IP on frame relay lines in the base 3745. IP over frame relay is an enhancement to NCP's frame relay function that allows IP frames to be transmitted over and received from a frame relay network without being encapsulated in SNA frames. This is termed native IP routing. Previously IP traffic routed from one 3745/NCP to another 3745/NCP over a communications link required SNA encapsulation.

With the implementation of IP over frame relay, NCP Version 7 Release 3 also supports dynamic reconfiguration (DR) of IP frame relay interfaces. It is possible to permanently define a frame relay IP PU on a frame relay physical link that does not already have IP resources defined. This requires the use of permanent dynamic reconfiguration function available with VTAM Version 4 Release 3.

The IP over frame relay capability is RFC 1490 compliant. This is for IP that is under NCP control.

Frame Relay BAN Support

Frame relay Boundary Access Node (BAN), which uses the RFC1490 bridged frame format, provides an extension to the previously announced frame relay Boundary Network Node (BNN) capability. BAN supports APPN and BNN traffic on 3745 and 3746 connected lines.

Increased DLCI Range

The DLCI range is increased from 16-215 to 16-991 for all lines. This allows greater flexibility in specifying DLCI numbers to match those assigned by frame relay providers when attaching to a frame relay network.

DLCI Sharing

NCP V7R3 supports the following DLCI sharing options:

- One INN station and IP per DLCI (3745 only)
- One INN station per DLCI (3746 only)
- Multiple BNN/APPN stations and IP (3745 only)
- Multiple BNN/APPN stations per DLCI (3746 only)

ACF/NCP Version 7 Release 4

NCP V7R4 (March 1996) had the following frame relay support enhancements:

Frame Relay Internal Frame Switching Support

This function will provide customers the ability to couple the 3745/NCP frame relay frame handler support function (FHFH) with the 3746 connected lines. For those customers using an external frame relay line to switch traffic between 3745 and 3746, this function will allow them to eliminate this external line by defining an internal PVC segment between a 3745 base line (either a frame relay physical line or NTRI frame handler logical line) and a 3746 line so that traffic can switched internally. This allows users to frame switch between a 3746 Model 900 frame relay line and either a 3745 frame relay line or a 3745 Token-Ring Interface Coupler (TIC1 or TIC2) supporting frame relay over token ring.

Improvements in performance, usability, and the elimination of the cost of external connections are advantages that can be achieved by frame relay users who employ this capability.

ACF/NCP Version 7 Release 5

NCP V7R5 (November 1996) had the following frame relay support enhancements:

Frame Relay Port Sharing

NCP V7R5 together with 3746-900 microcode allows 3746 frame relay ports to be shared by NCP, 3746 NNP and 3746 IP functions.

Boundary Access Node (BAN) for Subarea Links

Boundary Access Node (Bridged format) support over subarea (PU4) connections is provided as an additional frame relay connectivity option. This support allows 3746 Model 900 connections to remote token-ring capable physical units (PUs) through a remote frame relay BAN router such as an IBM 6611 or IBM 2210. The 3745 does not support BAN for subarea traffic.

Frame Relay Inoperative “error count” management upgrades

Previously, if NCP received 64 consecutive frames in error, the frame relay physical link and all its associated resources were brought down. This change allows the user to configure the allowed number of consecutive errored seconds before the physical link resources are brought down. This allows a frame relay physical line to survive a large burst of errors occurring in a very short time. With this change, NCP only counts one error in each 100 ms period for which an error occurs, thus allowing greater network control for the system manager.

3746 DLCI Sharing

For FRTES, 3746-900 microcode together with NCP V7R5 supports the following sharing options per 3746 DLCI:

- One INN station (routed or bridged frame format) per DLCI.
- On IP traffic (3746) and multiple BNN/APPN stations (routed or bridged frame format) each being controlled by NCP or NNP.

Enhanced dynamic windowing algorithm

Currently, NCP uses the IEEE 802.2 dynamic window algorithm in place of a Committed Information Rate (CIR) implementation to regulate the amount of bandwidth each logical SNA station is offered. NCP V7R5 provides two changes to support an enhanced dynamic windowing algorithm for 3745 Frame Relay logical lines:

1. To more closely approximate a CIR for Permanent Virtual Circuits (PVCs) that are assigned a low CIR on a high-speed link, NCP enables the transmission rate for a logical SNA station to be slowed even further than the dynamic windowing algorithm allows. NCP V7R5 introduces a variable time delay between I-frame transmissions when network congestion continues while the station's dynamic working window is 1. In order to keep the delays within reason, NCP V7R5 adds a new configuration parameter, DYNWIND, to specify the upper boundary of this delay.
2. To prevent NCP from overreacting to mild congestion, NCP ignores subsequent Backward Explicit Congestion Notification (BECN) for 100ms after an initial BECN is received and the dynamic working window is adjusted. NCP V7R5 adds a new configuration parameter to adjust this timer-DYNWIND.
3. NCP lines in 3745 and the base 3746-900 frame do not support CIR. In the case that CIR lines in the 950 talk to NCP 3745 lines, a fix must be applied to that NCP. It is APAR IR 35193 applied to the NCP V7.5 in the base box. (This is not a recommended configuration. Both partners should be configured the same.)

Frame Relay Overview

Frame relay is a high-speed packet switching technology based on permanent virtual circuits (PVC) for interconnecting data terminal equipment (DTE). The frame-relay standard defines only the physical DTE-DCE interface, layer 1, and the data link connection (DLC) link level interface, layer 2, for accurate exchange of data. Multiplexing on the physical layer is accomplished at the DLC layer by data link connection identifiers (DLCI).

A frame-relay DTE is called Frame Relay Terminating Equipment (FRTE).

The FRTEs communicate with each other via the frame-relay network using the I.233 frame format.

The frame-relay network nodes provide a frame switching function called Frame Relay Frame Handler (FRFH). Figure 7-1 on page 7-7 shows three PVCs: FRTE1 to FRTE2, FRTE1 to FRTE3, and FRTE2 to FRTE3, all going through the FRFH switch.

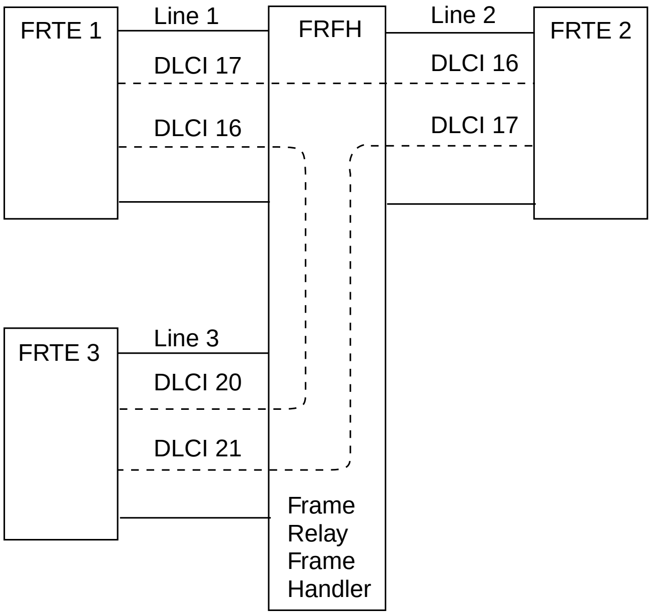

Figure 7-1. Frame-Relay Terminology

The frame-relay network does not provide guaranteed delivery of sent packets but does guarantee:

- The integrity of transported data by a frame check sequence that exists in every frame.
- That packets will not be duplicated across the network.
- That packets will be received in the order that they are sent.

Frame relay allows packets to be discarded during transport because of network congestion. It is the responsibility of the FRTE to ensure end to end delivery of packets.

Higher layer protocols (for example, IEEE 802.2 Logical Link Control (LLC), SNA/Subarea, APPN, HPR, or TCP/IP) are transparent to the frame relay network. It simply transports the higher protocols as user data. It is recommended that frame relay be used on *good quality lines* because recovery from transmission errors is done end-to-end instead of hop-by-hop.

Frame Relay Network: Example

The diagram below illustrates a simple frame relay network. Note that in some case a device functions solely as a FRTE, while at other times it is implemented with the FRFH function in addition.

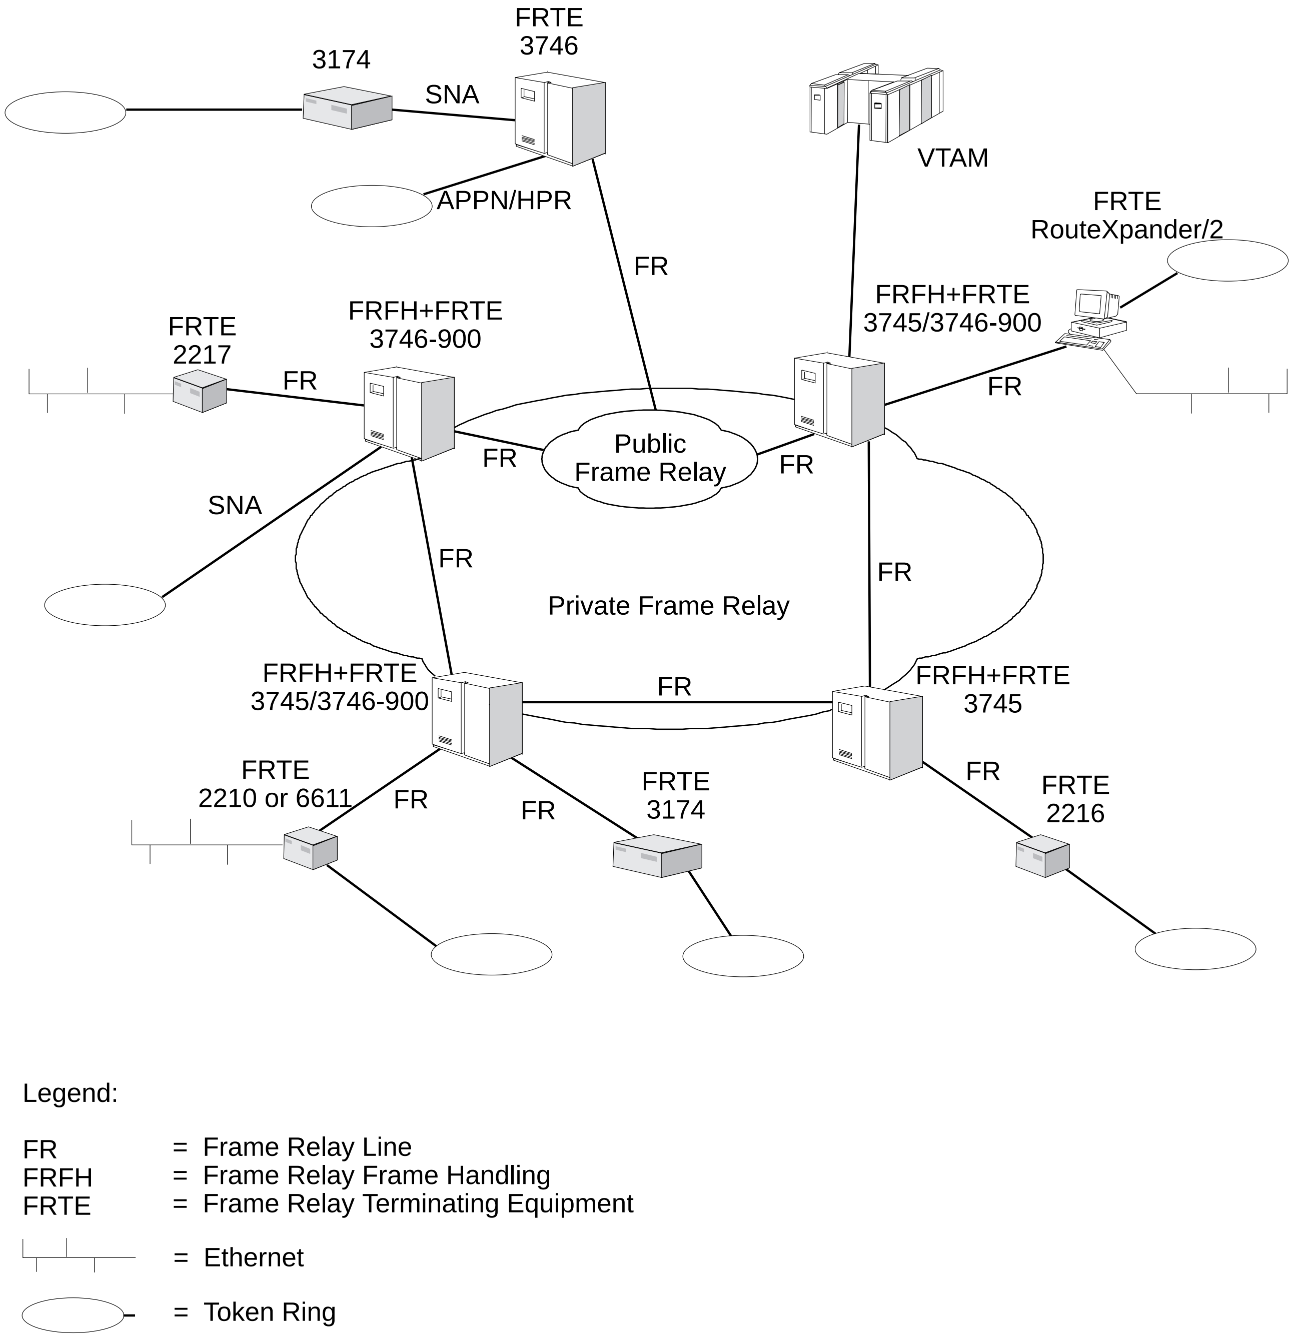

Figure 7-2. Example of a Frame Relay Network with 3745s and 3746s

Frame Relay Formats

The FRTE connections are supported by the 3746 in the same way as token-ring connections. The FRTE connections are defined as switched logical link stations. The RFC 1490 is used to identify the type of connection. Two formats are supported:

- **The RFC 1490 routed format** is used with equipment such as the IBM 3174 Establishment Controller. Multiple SNA stations connected to the 3174 can access the 3746 NN over a single DLCI and are identified by a different *Service Access Point* (SAP). The RFC 1490 routed format is used to provide the frame relay *boundary network node* (BNN) function. This format is also used between adjacent SNA or APPN node. IP traffic is routed with NLPID='CC'. See Figure 7-3 for the frame header for SNA traffic using the RFC 1490 routed format.

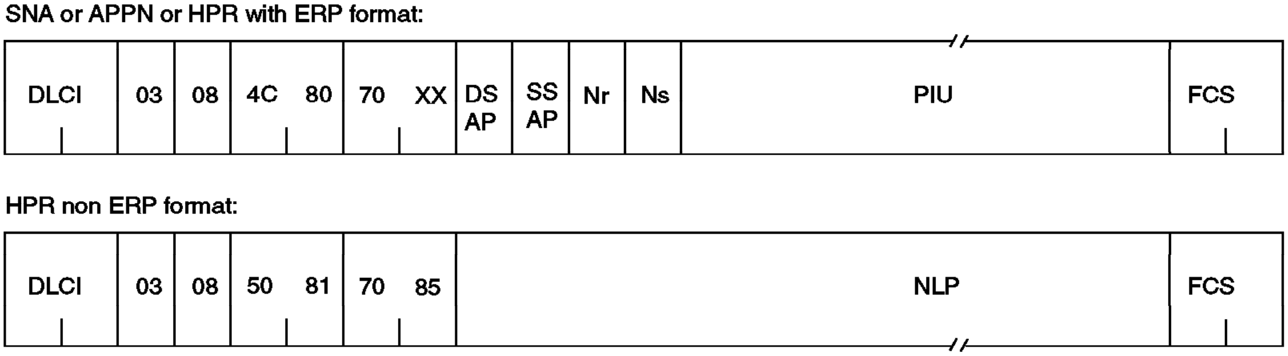

Legend:

DLCI = The Q.922 address on two bytes: B'dddddDxx ddddxxxx', where dddddd dddd is the data link control identifier (DLCI) number.

X'03' = Unnumbered Information (UI) frame (one byte)

X'08' = Network Layer Protocol Identifier (NLPID), which indicates that the IUT-T Recommendation Q.933 is used.

X'4C80' = Layer 2 is an IEEE 802.2 type 2 protocol (indicates error recovery)

X'70XX' = Layer 3, where XX can be:

- X'81' for SNA subarea, PU4-PU4 (FID4)
- X'82' for SNA subarea peripheral PU4-PU2 (FID2)
- X'83' for APPN (FID2)
- X'85' for HPR (NLP).

DSAP = Destination service access point

SSAP = Session service access point

Nr = Receive sequence number

Ns = Send sequence number

PIU = Path Information Unit (the transported SNA or APPN data)

FCS = Frame check sequence

Figure 7-3. RFC 1490 Routed Format for SNA Traffic

For additional information, refer to the ANSI T1.617 Annex F or the ITU-T Q.933 Annex E.

- **The RFC 1490 bridged 802.5 format** is used with equipment such as the IBM 2210 Nways Multiprotocol Router or IBM Nways Multiaccess Controller. This provides the frame relay *boundary access node* (BAN) function. IP over bridged traffic is *not* supported. See Figure 7-4 for the frame header for SNA BAN traffic using the RFC 1490 bridged format.

Frames sent and received without LAN FCS:

DLCI	03	00	80	00	80	C2	00	09	00	40	Token Ring Frames	FCS
					OUI		PID					

Frames received with LAN FCS:

DLCI	03	00	80	00	80	C2	00	03	00	40	Token Ring Frames	LAN FCS	FCS
					OUI		PID						

Legend:

DLCI = The Q.922 address on two bytes: B'dddddddxx ddddxxxx', where dddddd dddd is the data link control identifier (DLCI) number.

X'03' = Unnumbered Information (UI) frame (one byte)

X'00' = Unused byte (pad)

X'80' = Network Layer Protocol Identifier (NLPID), which indicates that the Subnetwork Access Protocol (SNAP) header is used

X'0080C2' = Organizationally unique identifier (OUI) for the IEEE 802.1 protocol (three bytes)

X'00XX' = Product identifier (PID), where XX can be:

- X'03' for the IEEE 802.5 protocol with the LAN FCS preserved (sent and received frames)
- X'09' for the IEEE 802.5 protocol without a LAN FCS (received frames only).

X'00' = Unused byte (pad)

X'40' = Flow control (FC) field

Token-ring frames = IEEE 802.5 frame that contains the SNA, APPN, or HPR transported data

LAN FCS = Frame check sequence used on the LAN

FCS = Frame check sequence

Figure 7-4. RFC 1490 Bridged Format for SNA Traffic

For additional information, refer to the RFC 1490.

The multiple stations attached to the 2210 or 2216 are identified by their MAC addresses, which allows an unlimited number of stations to use the same DLCI number. The number of stations using the same DLCI is only limited by the bandwidth of the frame-relay link between the 2210 or 2216 and the 3746 NN.

For both RFC 1490 formats, the integrity of the data is ensured by the 802.2 Logical Link Control (LLC) protocol that is connection oriented and provides reliable data link services. The 802.2 is the same LLC as the LLC used for SNA transport over a LAN.

3746-9x0 Frame Relay Implementation

Frame Relay DLCI Numbers

For FRTE connections, the DLCI number is assigned by:

- DLUS or CCM for connections initiated by the 3746 NN (call-out).
- The attached SNA, APPN/HPR, or IP device for connection initiated by the remote equipment (call-in).
- The FR network for IP connections initiated via the "orphanage" function discussed below.

Port Sharing

Frame relay is the only type of serial interface that is supported by all 3745 and 3746 protocol stacks. A single frame relay interface (port) can be activated simultaneously by each protocol stack, 3746 NN/DLUR and IP, and 3745 NCP. Frame relay, together with the 3745 and 3746, is ideally suited for building a multiprotocol network backbone.

Notes:

1. When using a dual-CCU 3746-900, a frame relay line can only be activated from either CCU-A (NCPA) or CCU-B (NCPB), not both at the same time.
2. NCP 3746-900 frame relay support is limited to SNA only. NCP does not support IP over frame relay on 3746-900 attachments.

CLP Sharing

A CLP line is dedicated, at any given time, to one of the following types of data link control (DLC):

- SDLC (NNP or NCP driven)
- X.25 (NNP, NPSI, or NCP driven)
- Frame relay (NNP or NCP driven)
- ISDN (NCP driven: LIC16, for Europe only)
- PPP (NNP driven)

SDLC, X.25, frame-relay, PPP, and ISDN lines can be mixed on the same CLP.

A frame-relay line is a leased point-to-point circuit that can be attached to a LIC11 or LIC12.

On any one of the frame relay ports, any mix of protocols is supported (for example: BAN and BNN, for SNA and APPN/HPR, and IP can be supported on a port). Any of the above protocols can be used on any DLCI of the same port.

PVC Sharing

To accomplish the IP, APPN, and NCP/subarea connectivity one can use either separate PVCs for each of the individual protocol stacks or use the 3746-9x0 PVC sharing facilities.

The 3746-9x0 PVC sharing facilities enable the 3746-9x0 to distinguish between and share a PVC for 3746 IP, 3746 NN/DLUR (routed and bridged frame format), and NCP/Boundary (routed and bridged frame format) traffic. The sharing of a 3746 controlled frame relay line with NCP requires NCP V7R5 or higher.

Additionally, devices such as the 2210 router support mixing IP (routed frame format) and BAN traffic (bridged frame format) on the same PVC. The 3746 supports this mixing of frame formats on a single PVC.Se "Frame Relay Formats" on page 7-9 for details on traffic encapsulation.

The following diagrams show how the 3745 and 3746 differentiate between the different traffic types arriving at either 3745 or 3746 adapters.

Internal routing of incoming frames is done as follows:

1. Check DLCI number (this may be a FRFH DLCI).
2. Routed frame format frames
   - NLPID=X'CC' this is IP traffic
   - NLPID=X'08' Q.933 Encapsulation
   - Layer 2 field indicates ERP or non-ERP.
   - Layer 3 field indicates INN, BNN or APPN, or HPR.
3. Bridged frame format frames

   NLPID=X'80', SNAP encoding

IP traffic is routed by the IP router of the 3746. INN traffic always goes to the NCP, BNN/APPN traffic is routed by the DSAP. HPR traffic is routed by ANR label. Therefore, in the case of HPR traffic controlled by NCP running as a CNN, HPR packets can be switched at the adapter level in the 3746 without being routed to software components.

Frame Relay Frame Handler Functions

The frame relay frame handler functions take the traffic arriving on a VC (3745 or 3746 port), and switch all the traffic on that VC to an outbound VC. All forms of encapsulation are supported as the frame contents (apart from the frame header) are not examined by the FRFH function.

The FRFH function can switch VCs between ports on the 3745 and ports on the 3746 (both machines must be connected by an internal or external link).

The definitions for frame switching are loaded from NCP into the 3745 and 3746, or from CCM for the 3746 adapters controlled by the NNP.

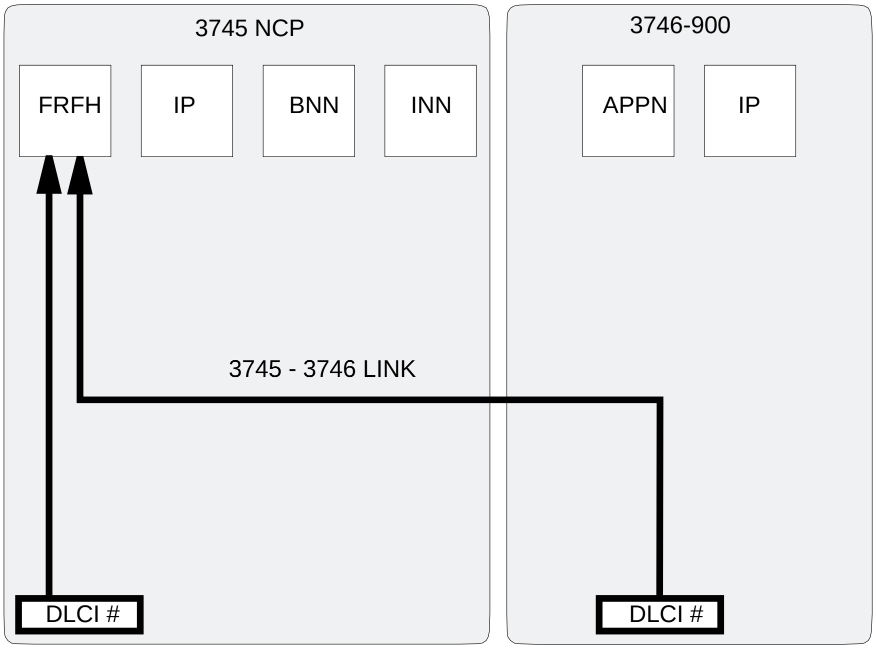

Figure 7-5. 3745 Frame Handler Function

IP over Frame Relay

3745 and 3746 IP traffic is sent using NLPID encapsulation with NLPID=X‘CC’, and uses the RFC1490 routed frame format. For IP the bridged frame format is not supported. The IP SAP X‘AA’ is used for LAN traffic but is not used for IP over frame relay. IP traffic from 3746 adapters is passed to the 3746 IP component, while traffic from the 3745 adapters is passed to the 3745 IP component.

IP traffic on 3746 adapters can be multiplexed with APPN and BNN traffic on the same DLCI. SNA traffic is encapsulated with NLPID=X‘08’ (routed frame format) or NLPID=X‘80’ (bridged frame format). Traffic for the 3746 APPN CP and NCP BNN function is distinguished by the SAP values in incoming SNA frames.

Figure 7-6 on page 7-14 shows the IP support.

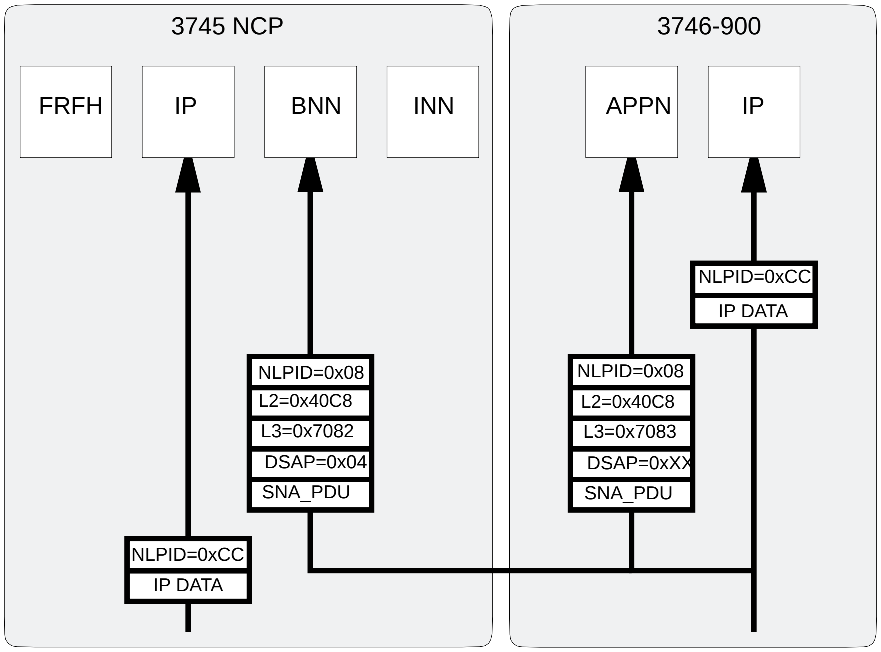

Figure 7-6. 3745 and 3746 IP over Frame Relay

INN over Frame Relay

Frame relay INN traffic can be either in the RFC1490 routed or bridged frame format. INN traffic on 3745 ports can share a DLCI with IP traffic for the NCP. This traffic can be distinguished by the NLPIDs used (see Figure 7-7). INN traffic over frame relay (routed or bridged frame format) always uses DSAP=X‘04’ and SSAP=X‘04’. Traffic from 3746 ports can also be passed to the NCP INN function.

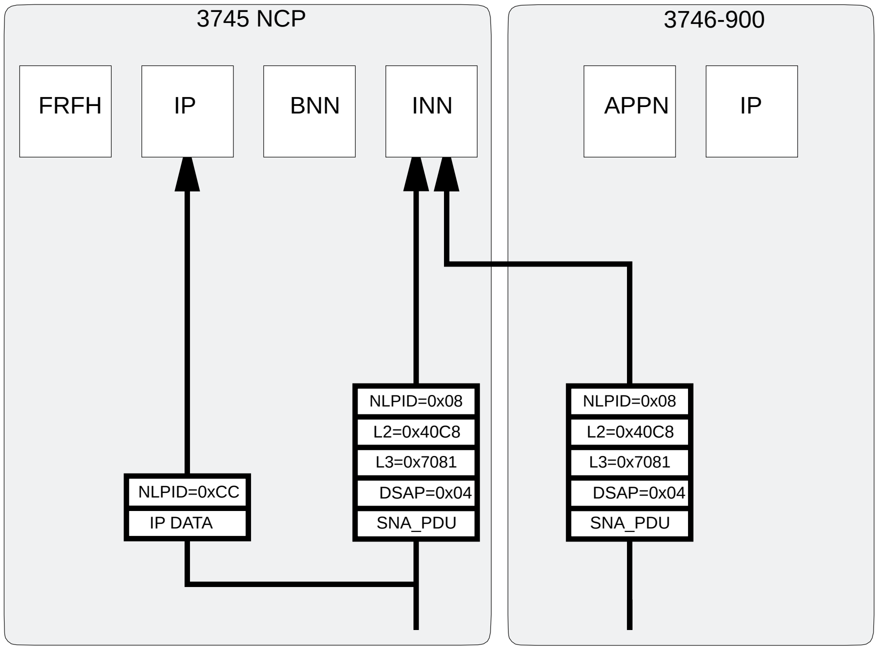

Figure 7-7. 3745 INN Traffic - Routed Frame Format

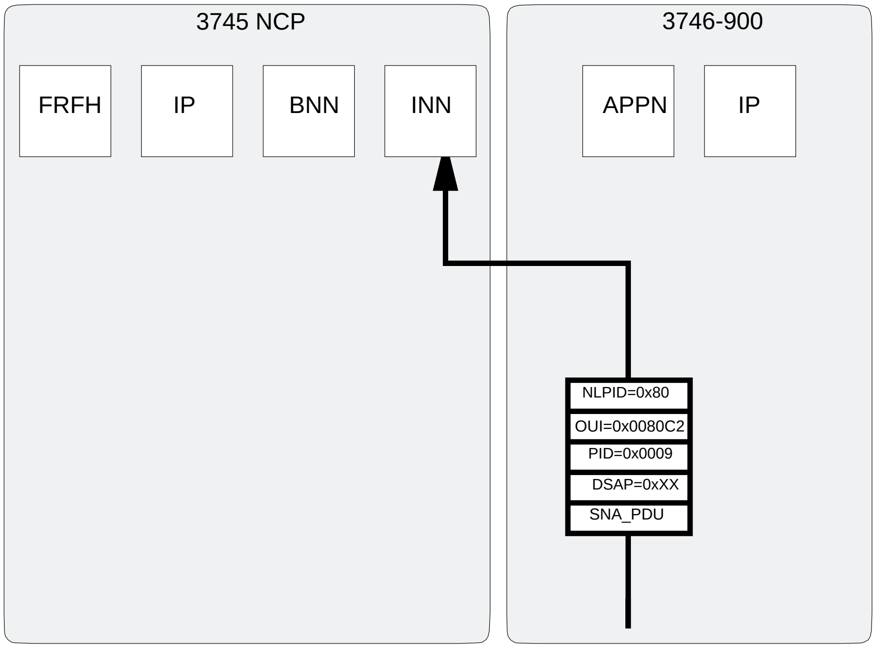

Figure 7-8. 3745 INN Traffic - Bridged Frame Format

BNN and APPN over Frame Relay

BNN and APPN traffic from 3745 adapters uses a separate DLCI and cannot be multiplexed with other traffic. BNN and APPN traffic from 3746 adapters can be multiplexed with IP traffic on the same DLCI (see Figure 7-9 and Figure 7-10 on page 7-16).

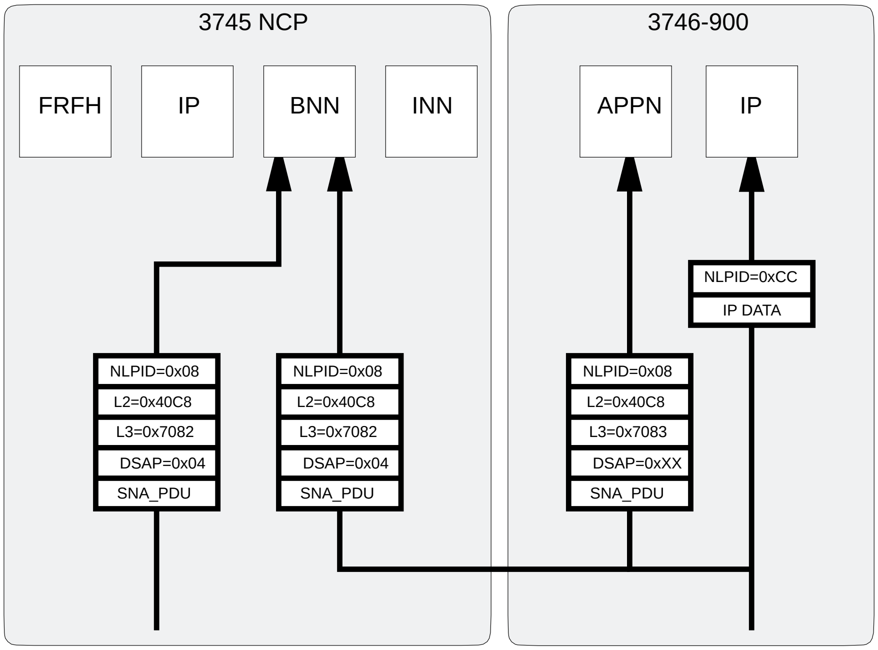

Figure 7-9. BNN/APPN Traffic - Routed Frame Format

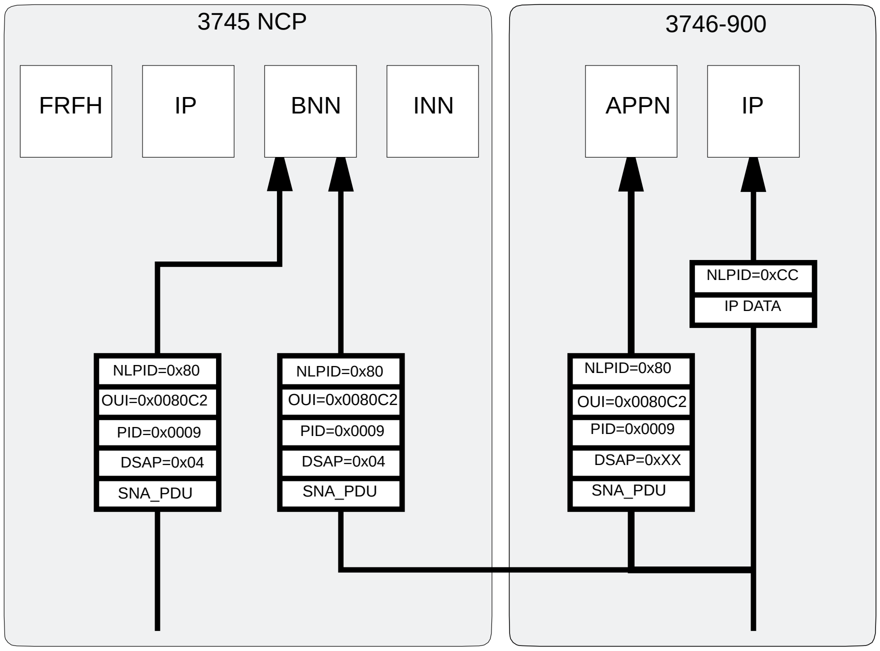

Figure 7-10. BNN/APPN Traffic - Bridged Frame Format

FRTE Subarea Connections (FRTE SA)

A subarea node is either a PU type 4 (NCP) or a PU type 5 (VTAM). FRTE SA connections apply to connections between PU4s (NCPs).

FRTE subarea connections are supported by NCP and the 3746-900 in the same way as the token-ring subarea connections. A physical line running frame relay is used to multiplex subarea logical connections. The subarea logical lines are defined as leased lines, each with a single T4 node. The 3746-900 provides an IEEE 802.2 connection-oriented attachment to the adjacent subarea node. NCP provides the transmission group, XID, path control, physical unit management and services, and the DLCI to be used for the connection.

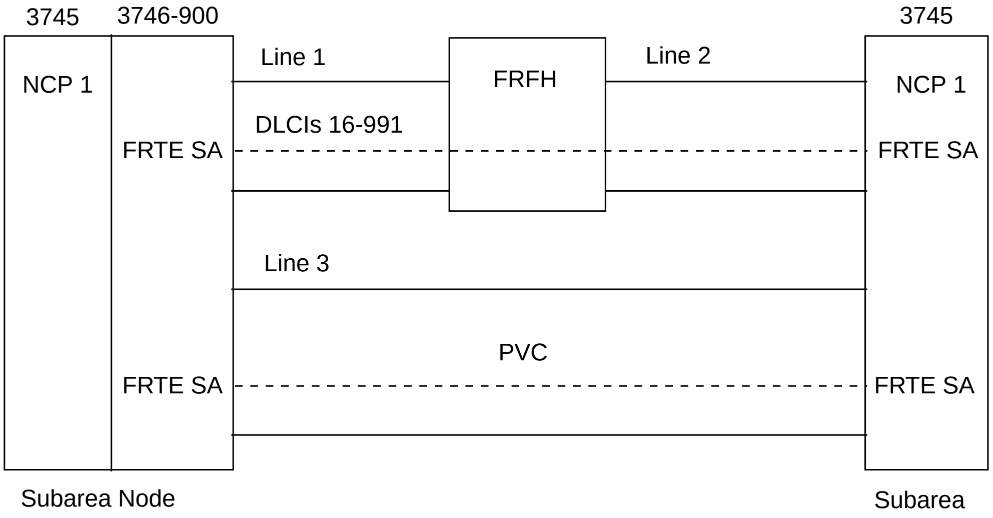

Figure 7-11. Frame-Relay Terminating Equipment: Subarea Connections

Figure 7-11 shows an FRTE SA between two NCPs that is supported on a direct connection between two NCPs as on line 3 or through an intermediate network as on line 1 and line 2.

Mixed media multilink transmission groups (MLTGs) are supported over an FRTE SA subport of the 3746-900. In Figure 7-11, the two connections between NCP 1 and NCP 2 may belong to the same TG.

Frame relay congestion management is provided by the 3746-900 when:

- A frame with the backward explicit congestion notification (BECN) bit set is received, and the sending rate is decreased.
- A frame with the forward explicit congestion notification (FECN) bit set is received, and the BECN bit is set in the next frame sent.

NCP provides frame discard eligibility (DE) support if selected at NCP configuration time. When DE is supported by the transport network, the 3746-900 uses the DE bit in the frame relay packet header. DE is set for packets containing data frames, but it is never turned on (set) for packets containing SNA control frames.

FRTE Peripheral Node Connections (FRTE PN)

As with FRTE subarea connections, FRTE peripheral connections (BAN, BNN, APPN) are supported by NCP/3746-900 in the same way as token-ring peripheral connections. The peripheral connections are defined as switched logical lines.

The frame relay *boundary access node* (BAN) function enables the 3745 and 3746-900 to communicate with IBM 2210 and 6611 routers and their SNA downstream physical units (PUs) using:

Dynamic route selection

> The 3745/3746-900, in conjunction with ACF/NCP Version 7 Release 3, dynamically routes the SNA flow from the downstream PUs to the appropriate destination, thus eliminating the need for an additional router adjacent to the 3745/3746-900.

Multiple stations over the same DLCI

> The frame relay BAN function uses the RFC 1490 bridged-frame format. This BAN support of the 3745, 3746-900, 2210, and 6610 uses MAC address multiplexing, which allows a practically unlimited number of stations to use the same DLCI number.
>
> The number of stations using the same data link connection identifier (DLCI) is only limited by the bandwidth of the frame-relay link between the router and the 3745 or 3746-900.

Multiple DLCIs over the same frame-relay link

> Though only one DLCI is usually required between the 3745 or 3746-900 and the router, the frame-relay BAN supports multiple DLCIs over the same frame-relay link.

The frame-relay *boundary network node* function (BNN) allows the 3745 and the 3746-900 to route SNA traffic for frame-relay attached equipment, such as the IBM 2217 Nways Multiprotocol Concentrator, the IBM 3174 Establishment Controller, or the IBM PS/2 with Route Expander. Multiple SNA physical units, connected to an IBM 3174 Establishment Controller, can access the 3745 or the 3746-900 over a single DLCI. This function is called service access point (SAP) multiplexing. The frame-relay BNN function uses RFC 1490 routed format.

The *frame relay APPN function* allows the 3745/3746-900, in conjunction with VTAM and NCP, to be a composite network node (APPN CNN).

As with the frame relay subarea connections, the 3746-900 provides the IEEE 802.2 connection-oriented services and NCP provides the remainder of the necessary functions. The DLCI may or may not be given to the 3746-900 by NCP depending on where the connection request originates. For a host-initiated connection (call-out), the DLCI is provided by NCP. For an end-user device initiated connection (call-in), the DLCI is provided by the end-user device.

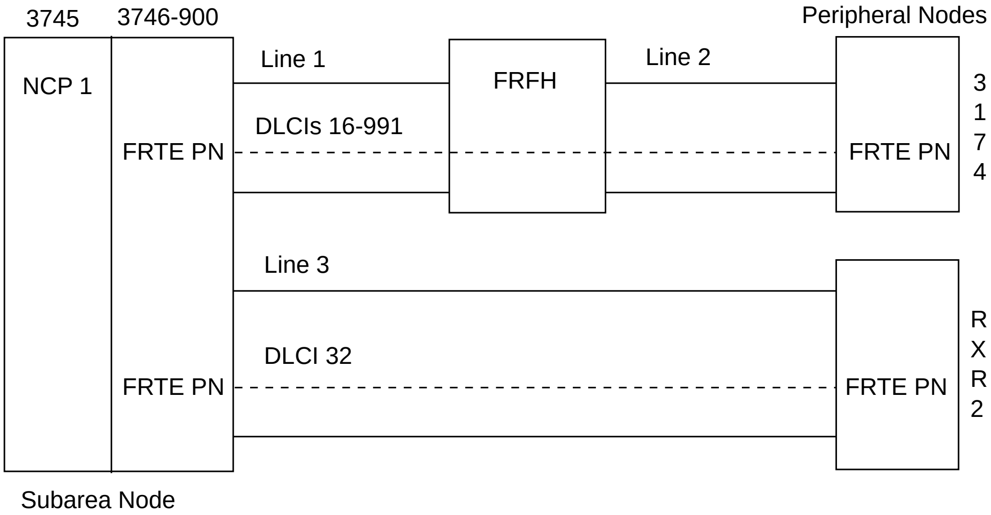

Figure 7-12. Frame Relay Terminating Equipment: Peripheral Node

FRTE PN is supported on a direct connection as on Line 3 as shown in Figure 7-12, or through an intermediate network.

Frame relay congestion management is provided by the 3746-900 when:

- A frame with BECN bit set is received. The sending rate is decreased.
- A frame with FECN bit set is received. The BECN bit is set in the next frame sent.

NCP provides frame-discard eligibility (DE) support if selected at NCP configuration time. When DE is supported by the transport network, the 3746-900 uses the DE bit in the frame relay packet header. The DE bit is set for packets containing data frames, but it is never turned on (set) for packets containing SNA control frames.

Both PU 2 and PU 2.1 devices are supported. The NCP/3746-900 provides support for both subarea and peripheral connections over the same 3746-900 frame relay physical line.

Frame Relay Switching Equipment (FRSE) Functions

Frame Relay switching equipment support includes frame relay frame handler (FRFH) support, dynamic reconfiguration of FRFH subports and routes, alternate physical links (substitute subports), frame relay congestion, and performance management.

Previously, NCP was responsible for the definition, activation, deactivation, and dynamic reconfiguration of FRFH subports. These subports could be on 3745 or 3746 serial adapters. This support is now also in NN.

NCP V7.5 permits FRSE between NCP base frame lines and the NCP lines in the 900 frame. Note that an NN node can only use lines in the 950 frame.

Note: FRFH functions could also be run over connections between token-ring adapters on the 3745. This is not possible on the ODLC TIC3 adapters of the 3746.

Support has now been added allowing the NNP of the 3746 to configure and control FRFH functions. This allows the 3746-950 to support FRFH functions.

For FRFH functions, the 3746-900 provides the actual routing functions (see Figure 7-20 on page 7-34). Routing frame relay data between 3746-900 ports is supported. The 3746-900 ports may be attached to the same or to different CLPs.

The same line can support both FRFH and FRTE subports.

The FRFH function can be used to switch any type of traffic. Incoming frames are switched to their outbound port without any regard for the data encapsulated in the frames.

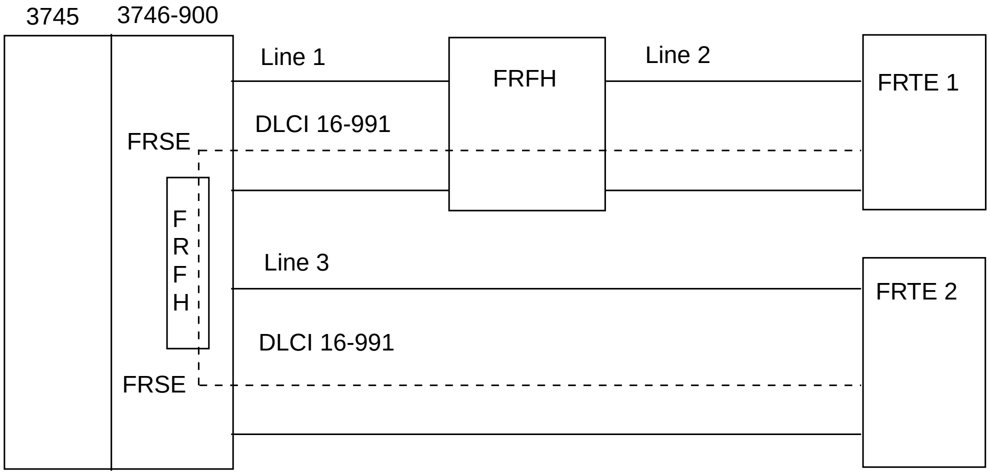

Figure 7-13. Frame Relay Switching Equipment (FRSE or FRFH)

An FRTE can be directly attached as in Line 3 or through an intermediate network as on Line 1 in Figure 7-13. Frame relay congestion management is provided by the 3746-900:

- When the first level of congestion occurs, FECN is set. The BECN bit is never set.

  Note: This is different from the 3745. If there are frames available flowing in the reverse direction, the 3745 sets the BECN bit.

- In case of severe congestion, frames are discarded. The frames with the DE bit set are discarded first.

Frame Relay Substitute Subport

An *FRSE SET* is a set of two to four subports (A,B,C,D) where A and B are defined as the *primary partners*. We may substitute C for A or D for B. The path from C to D is NOT allowed. A diagram is provided below in Figure 7-14 on page 7-21.

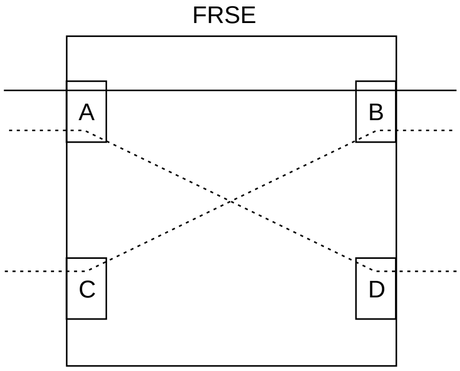

Figure 7-14. Frame Relay FRSE Set

A substitute subport provides non disruptive route switching (NDRS) to the end user.

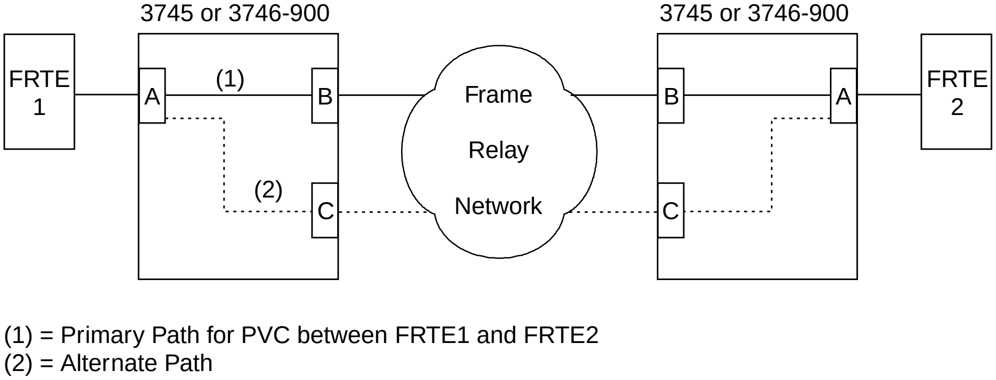

Figure 7-15. Frame Relay Substitute Subport

A substitute subport can be defined for an FRFH subport. When a substitute subport is defined at both ends of a PVC, this provides an alternate PVC between paired 3745 or 3746-900 switching nodes as shown in the Figure 7-15. In both boxes, the primary path is defined from subport A to subport B. When the path from subport B of the left box to the subport B of the right box fails, an alternate path is used between subport C of both boxes. Local management interface (LMI) is required at subports B and C. This is also supported by the 3746-950.

The substitute subport provides a nondisruptive route switching for the PVC established between FRTE1 and FRTE2 in the previous figure. When the primary path is reestablished, the traffic automatically switches back to it.

Frame Relay Local Management Interface (LMI) Support

The 3746-900 LMI support conforms to ANSI T1.617 Annex D and ITU-T Q.933 Annex A.

Frame relay LMI support is provided entirely by the 3746-900, except for the definition of the LMI timers and thresholds. The timer and threshold values are defined in the NCP and passed to the 3746-900 at activation time. Events that are detected via LMI (for example, LMI status change or reaching an LMI error threshold) may result in failure of the associated SNA resources that are handled by the NCP. The link integrity verification (LIV) tests and full status polling are handled entirely by the 3746-900.

NCP determines the mode of LMI support at activation of the LMI subport only when LMI support is defined for that physical line. NCP will determine the mode based on the LMI frames received dynamically at the time of connection establishment.

The different LMI modes are supported in the same way as for 3745 attached resources. The behavior is dynamically discovered at connection establishment time.

The modes that are possible are listed below:

- Network to Network Interface (NNI): The LIV inquiries and the PVC status inquiries and responses are sent by both partners on the frame relay interface. This is also called LMI bidirectional protocol. Both partners must support this mode.
- Network to User Interface: The partner is a DTE that does not support NNI. The partner polls the 3746-900 and does not answer the LIV enquiries sent by the 3746-900.
- User to Network Interface (UNI): The partner is a DCE that does not support NNI. The partner answers the status inquiries sent by the 3746-900.
- No LMI: The partner does not answer the status request and does not send status inquiries to the 3746-900.

Since NCP negotiates dynamically with the remote end, it is possible to find a misconfiguration: ITU vs ANSI, one at one end of the line and the other at the other end. The result will be "no LMI".

NN lines can either negotiate as above or be forced into the NUI mode. If they are misconfigured we will wait in an LMI PENDED STATE and the line will not be available for use. Redefinition will be necessary.

System Definitions

The configuration of the NCP 3746-900 frame relay resources is similar to that of the 3745. It is created via the network definition facility (NDF) during NCP generation.

- Physical lines represent frame relay ports. The maximum frame size supported is 8250 bytes (as it was for the 3745).
- The LMI is represented as a physical station PU type 1.
- Additional physical stations are used to represent frame relay subports used for FRSE. These PUs can be dynamically added or deleted from VTAM without requiring NCP regents or loads.
- A logical line and the associated logical station (PU type 4) represent a FRTE SA.
- A pool of logical lines and associated logical stations are available for dynamic use at connection setup time for FRTE PN (BAN,BBN or APPN).
- FRSESET is used to define a switching path inside the 376-900 between two of its subports. The FRSESET paths can be modified from VTAM without NCP regents or reloads.
- Substitute subports are defined via the FRSESET definition statement. NDF is used to define the communication rate for each DLCI.

Configuration of NNP frame relay resources by CCM is straightforward. Screens are presented for port, LMI, DLCI and APPN definitions.

Congestion Control

Frame Relay Standards

Congestion control can be defined as a set of mechanisms incorporated to attain certain network performance objectives, particularly in the peak periods, while optimizing or improving the network resource requirements. It aims to minimize the number of occurrences of user-perceived congestion. Frame relaying networks should not allow users to monopolize network resource usage at the expense of other users. Congestion control includes both congestion avoidance and congestion recovery mechanisms.

The service offered by a frame relaying service is the transparent and unacknowledged transfer of frames. The user data received will be like the data sent except the address and FCS field, which can be modified by the network. The network does not guarantee message delivery. Therefore, frames may be dropped. A frame relaying network experiencing congestion will either inform its users about the congestion, assuming the users will take appropriate action (not detailed in the frame relay standards!) to relieve the congestion, or it simply discards frames.

Frame relay networks using *out of band* congestion signaling report congestion by sending explicit congestion control messages on a dedicated DLCI. In addition to frames originating from remote end stations and LMI message sent by the network, end stations may also receive Consolidated Link Layer Management (CLLM) messages generated by the network reporting congestion. The use of CLLM is not widespread.

Frame relay networks using *in-band* congestion signaling report congestion by using bits in the frame address field. End station will receive no other frames (the exception being LMI messages) than those frames sent by another end station.

The network is able to inform end stations about congestion by using two fields in the frame address field. For this purpose the *forward explicit congestion notification (FECN)* bit and the *backward explicit congestion notification (BECN)* bit have been reserved. The FECN bit will be set in frames flowing in the direction in which the network is experiencing congestion. The BECN field will be set in frames flowing in the opposite direction in which the network is experiencing congestion.

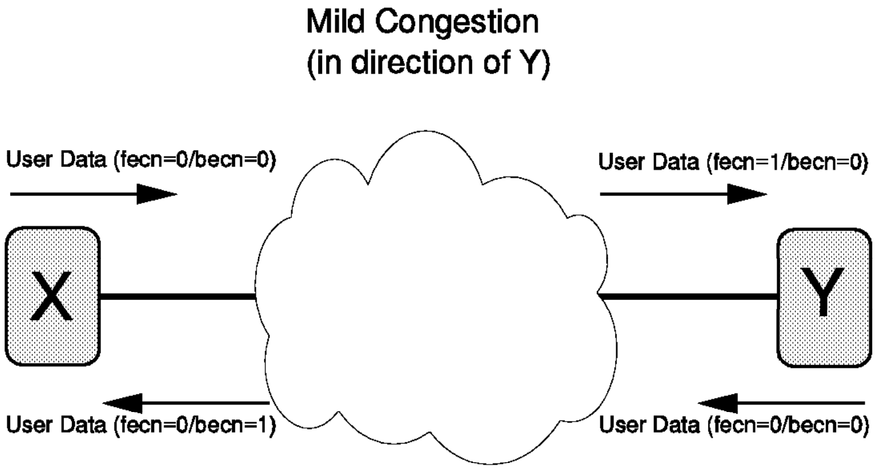
Mild Congestion
(in direction of Y)
User Data (fecn=0/becn=0)
User Data (fecn=1/becn=0)
X
Y
User Data (fecn=0/becn=1)
User Data (fecn=0/becn=0)

Figure 7-16. In-Band Congestion Signaling

The consequence is that if traffic on a specific virtual circuit is uni-directional, only the receiving station will be informed about the congestion and not the transmitting station, which may be the cause of the congestion.

FECN and BECN congestion indicators are usually set by the network only. However, in particular cases they may be set by end stations as well. As an example, ACF/NCP Version 7 Release 2 sets BECN in the first frame to be transported after a frame with FECN has been received. This informs the other end of the PVC about the congestion, allowing it to decrease its transfer rate helping the network to relieve the congestion.

FECN/BECN will be set during mild congestion, while the network is still able to transfer frames. A frame relay network will usually start discarding frames during severe congestion. End stations are able to "prioritize" their traffic by using the *discard eligibility (DE)* in the address field of the frame header. The network will start to discard frames with the DE field set first; however, frame delivery is not guaranteed and there is nothing in the frame relay standards that restrains networks from discarding frames without the DE bit set.

The frame relay standards do not specify the conditions under which the FECN/BECN bits will be set and when frames with or without DE will be discarded. It is assumed, but not enforced, that end stations reduce their information transfer rate upon detection of network congestion. Congestion control based on the discarding of frames or the use of FECN/BECN bits, and relying on the "good behavior" of end stations, has therefore been considered inadequate to networks providing a frame relaying service and additional provisions have been defined.

Committed Information Rate (CIR) and Burst Sizes

The 3746-950 provides CIR support which is *not* available under NCP. The maximum number of bits per seconds that an end station can transmit into the network is bounded by the *access rate* of the user network interface. The access rate is limited by the line speed of the user network connection and established by subscription.

The maximum committed amount of data that a user may offer to the network is defined as *committed burst size (B_c)*. B_c is a measure for the volume of data for which the network will guarantee message delivery under normal conditions. It is measured during the *committed rate measurement interval (T_c)*.

End stations are allowed to transmit data in excess of the committed burst rate. The *excess burst size (B_e)* has been defined as the allowed amount of data by which a user can exceed B_c during the committed measurement rate interval T_c. If spare capacity exists, the network will forward the data to its destination. The network, however, is free to mark the data as discard eligible (DE).

The *committed information rate (CIR)* has been defined as the allowed amount of data that the network is committed to transfer under normal conditions. The rate is averaged over an increment of time T_c. The CIR is also referred to as *minimum acceptable throughput*. Remember that CIR is implemented only for NN and IP 950 lines.

B_c and B_e are expressed in bits, T_c in seconds. The access rate and CIR in bits per second. CCM allows to define B_c, B_e, T_c, DLCI while the CIR is computed. The access rate is valid for each user network interface. For B_c, B_e and CIR incoming and outgoing values can be distinguished. If the connection is symmetrical the values in both directions are the same. For permanent virtual circuits B_c (incoming and outgoing), B_e(incoming and outgoing) and CIR (incoming and outgoing) are defined at subscription time. They are negotiated for SVCs at call establishment time. T_c is calculated as depicted in Table 7-1.

Table 7-1. Measurement Interval Calculation

CIR	**B_c**	**B_e**	**Measurement Interval (T_c)**
> 0	> 0	> 0	$T_c = B_c/CIR$
> 0	> 0	0	$T_c = B_c/CIR$
0	0	> 0	T_c = (B_e/Access Rate)[2]

Note:

Table depicts the valid parameter configurations. Other configurations are for further study.

When the two communicating end stations have different access rates the network may define a smaller T_c value.

Individual CIRs on a physical connection are always lower than the access rate; however, the sum of CIRs defined can be larger than the access rate. An example could be a network connection with an access rate of 256 Kbps on which three virtual circuit have been defined: two having a CIR of 128 Kbps each, one having a CIR of 64 Kbps.

Optimal values for the above parameters depend on network implementation, availability of spare network capacity, charging methods, type of user device and performance required. These are only a few considerations and careful study of the total network is required (as well as the more immediate future changes that can be anticipated).

Networks may mark frames above B_c with discard eligible (DE) but have plenty of spare capacity to transport the frame, or may instead have limited capacity and discard excessive frames immediately. Networks may mark frames above B_c+B_e with discard eligible (DE), and possibly transport it, or just drop the frames as suggested by ITU-T I.370.

The Network manager always tries to balance costs and performance, and examines the frame relay service provider charging schemes. Networks may implement fixed charging depending on access rate, a scheme dependent on CIR, B_c and B_e or more sophisticated schemes, for example charging on number of bits transported and charging progressively for data above B_c or B_c+B_e. Depending on the charging scheme employed, subscribing to high values of CIR, B_c and B_e, may lead to high networking costs. It should be examined if the performance gain, if any, counterbalances the additional networking expenses.

Many devices have limited control over the volumes of data they send into the network. Assuming that flow control mechanisms implemented on top of the layer 2 core function are not inhibiting data transfer, data will be transmitted with a speed up to the network access rate. If the device has only one DLCI active, or has (temporarily) data to send for one DLCI only, the data rate on a single DLCI may be equal to the network access rate. If the sum of committed and excess burst size (B_c+B_e) is lower than the access rate times T_c, the network may decide to discard frames. In this situation it may be advisable to give all DLCIs the following values:

B_c+B_e = Access rate*T_c

Depending on functions implemented on top of the layer 2 core functions, lost frames may be quickly detected and recovered from. This may be a time-consuming activity that severely impacts performance. In the latter case, subscribing to high values of CIR, B_c and B_e is important.

Communication Rate (CR) and Committed Information Rate (CIR)

The 3746 provides communication rate (CR) support. A part of the access rate (physical line speed) is assigned by the user to each station. IP traffic *per DLCI* is represented by *one* station. The total bandwidth available is split between the stations. This capability differs from the committed information rate (CIR), which is defined as the information rate which the network is committed to transfer under normal conditions over a DLCI.

If all the stations, at any point in time, require more bandwidth than is available, then each station is limited to their user-predefined bandwidth. In case of overflow, the data for those stations that create the overflow are kept in a software queue. They will be transmitted at the next opportunity. If the overflow on the DLCI lasts too long, the data in excess is discarded. The stations that create the overflow are paced and slowed down to their communication rate, while the other ones continue to get their communication rate.

When CR is implemented, if the total physical bandwidth is not fully used the unused bandwidth is available for stations that may then exceed their CR in making use of the unused bandwidth.

For private networks where one wishes to fully utilize all the bandwidth, implementing CR is preferable. For public networks, where CIR is a cost factor and constraint, it is preferable to implement CIR.

The assignment of the communication rate to the stations is via the CCM at 3746 configuration time. Note that CR and CIR can *not* be shared on the same physical line. Also SNA line sharing is *not* possible when the CP defines the link as having a given CIR. This information is passed by the NNP to the 3746 CLPs at activation time.

When CR is enabled for a given line:

- SNA and APPN FRTEs echo incoming FECNs as BECNs; incoming BECNs reduce XMIT windows as per DYNWIND definitions.
- IP FRTEs ignore incoming FECNs and BECNs.
- FRFHs transport FECNs and BECNs transparently. They do not set BECNs.
- FRTEs and FRFHs set FECNs whenever there is congestion on the transmit physical line.

When CIR is enabled for a given line:

- FRTEs echo incoming FECNs as BECNs.
- FRTEs and FRFHs set FECNs whenever there is congestion on the transmit physical line. They use Bc as a CR to manage the contention.
- FRFHs transport FECNs and BECNs transparently.
- FRFHs will also set FECNs as soon as the traffic received from a partner subport exceeds that DLCIs CIR during a Tc period. This will also happen when the delay introduced by the FRFH gets too large.

Throughput is optimized when CIR is enabled by tuning to just under the level at which the network sets FECNs and BECNs. All FRTEs will adjust their output rate between *minimum information rate* (MIR) and *excess information rate* (EIR). This is called *adaptive-CIR* (A-CIR) and is based on a unique tuning algorithm. This is opposed to CR, which handles the first physical hop and does *not* look for the logical bottleneck within the network. This minimizes queues and delays throughout the network and saves bandwidth on the first hop for other DLCIs that still have end-to-end bandwidth available.

The following provides further detail:

EIR = 0 means the same as $B_c = 0$.

MIR is set at either:

$25 * B_c/T_c$

$MIR = 9.6$ Kbps if
$B_c = 0$.

The EIR can be calculated using the following formula:

$EIR = (B_c + B_e)/T_c$

Also, when $B_c = 0$, the EIR becomes B_e/T_c. This is different from the 2216 or 2210, which do not define T_c and would then commit the whole physical bandwidth as EIR.

FRFHs do not implement A-CIR. They set:

$CIR = (_c + B_e)/T_c$

or, if the first hop is congested they set8colon

$CIR = B_c/T_c$

Having FRFHs implement CR and FRTEs implement CIR will move delays and congestion to the endpoints and optimize the utilization of the network backbone. But it is best that both endpoints implement CIR simultaneously because of the setting of FECNs and BECNs. Note that CIR values need not be the same for these endpoints if there is unbalanced traffic.

Figure 7-17 illustrates how A-CIR is used to maximize the use of available bandwidth between two network endpoints. Every 3.2. seconds a new CIR is computed based on current network congestion (BECNs). (The time interval, 3.2 seconds, was determined pragmatically and took into consideration the anticipated system turnaround time (the value was chosen to exceed it). This permits learning about current network conditions that resulted from the previous network settings. You will note that this permits the A-CIR curve to stabilize around the available network bandwidth.

You can also see in the graph how the A-CIR curve quickly approaches the network's available bandwidth even after that bandwidth changes. The broken variable CIR curve oscillates and makes measurably less efficient use of the bandwidth.

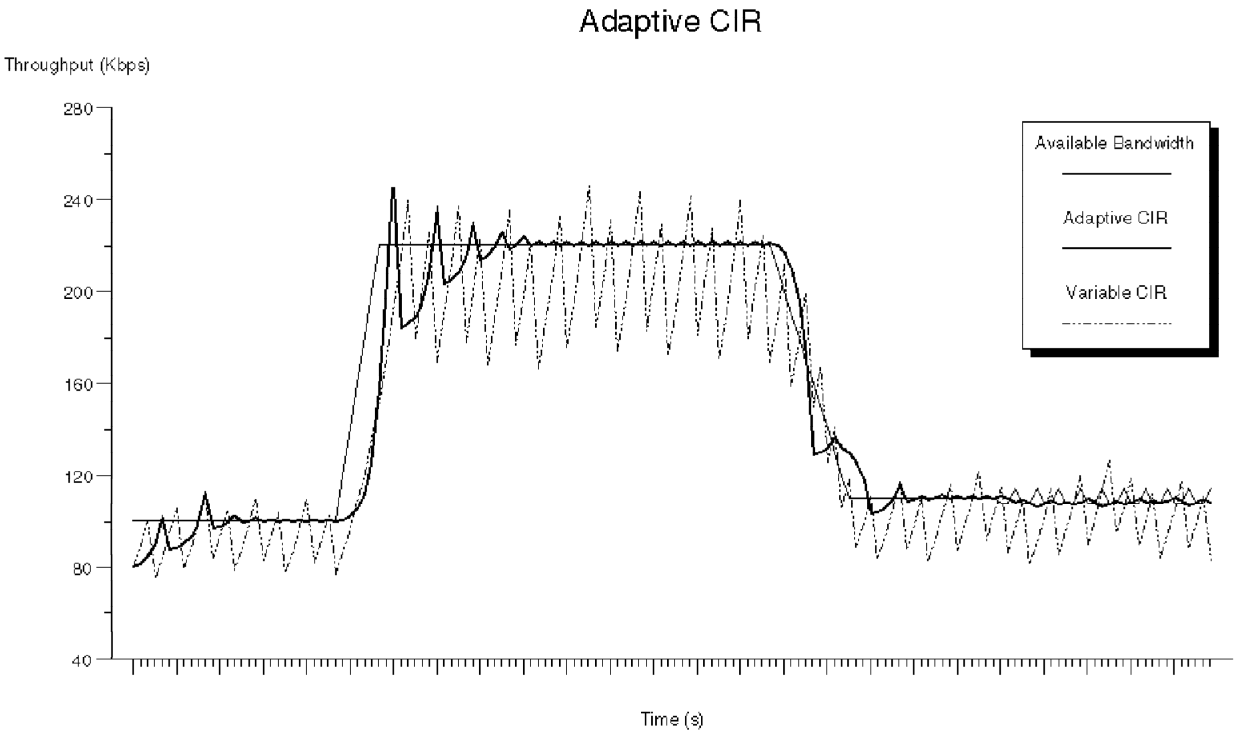

Figure 7-17. Adaptive-CIR vs Available Bandwidth

Data points were calculated for the above graph using the respective algorithms for A-CIR and CIR. Available bandwidth was varied to simulate possible network conditions. The interval is 3.2 seconds as mentioned above.

The formula is below:

$$\text{New CIR/Old CIR} = 1 +/- \ 2^{N-P}$$

where P is *PRECISION* which is configured and preset to 8. Its possible range is 6 - 10 and N is a runtime variable that is dynamically incremented with a range of 0 - 5. The sign +/- is determined by BECNs. If BECNs are received during the last 1.6

second interval, the sign is - and the CIR is effectively decreased, allowing network congestion to dissipate. If no BECNs are received the sign is set to +.

When CIR is updated with the same sign as the previous update, N is incremented by 1 until it reaches its maximum value of 5. When the sign changes, N is reset to 0. This is done after computing the new CIR when the CIR is less than the previous CIR and before new CIR computation in the other case.

For the 3746 adapters, either COMRATE or CIR can be specified at the port level. The default value is CIR disabled. This means COMRATE is enabled.

3746 Bandwidth Reservation System (BRS)

The Bandwidth Reservation System (BRS) allows the customer to define how bandwidth should be assigned between protocols within a DLCI.

As per Figure 7-18 on page 7-30 and other descriptions, BRS works on top of CIR. It does not have priority capability. The BRS implementation in the 3746 by the NNP supports a mechanism whereby each of three traffic types, SNA(APPN, DLUR, HPR-ERP), HPR¬ERP (non-ERP), and five different IP sockets can be assigned a portion of a DLCI's CIR (bandwidth). The balance, from 100%, should be assigned to the remaining IP traffic. These five sockets do not have to be well known and are assigned at configuration time.

Figure 7-18 on page 7-30 shows how 3746 BRS works. When CIR is enabled, T_c, B_c, and B_e can be defined for each DLCI. B_c and B_e may be equal to zero, but B_c plus B_e must be greater than zero. B_c and B_e are expressed in multiples of DATABLK. For FRFH and FRTE DLCI's:

T_c The measurement interval (T_c) is defined at the DLCI level. The default value is 0.1s. The value is specified in tenths of a second (1-255).

B_c Committed burst size in units of DATABLK (0-64).

B_e Excess burst size in units of DATABLK (0-64).

Then for FRTE DLCIs, each protocol can be assigned a percentage of those CIR values:

```
CIR parameters per DLCI: Bc, Be, Tc

BRS DLUR/APPN/HPR (ERP): T%
    HPR (non-ERP):      U%
    IP:                 V%
    Socket1:            W%
    .
    .
    .
    .
    Socketn:            X%
```

The percentages can be defined between 0% and 100%. A zero value means no bandwidth is reserved for that protocol, which means all that protocol traffic will be discarded when congestion occurs. If a protocol has no BRS defined, then that protocol's traffic does not participate in BRS and all that traffic is transferred whatever the level of congestion.

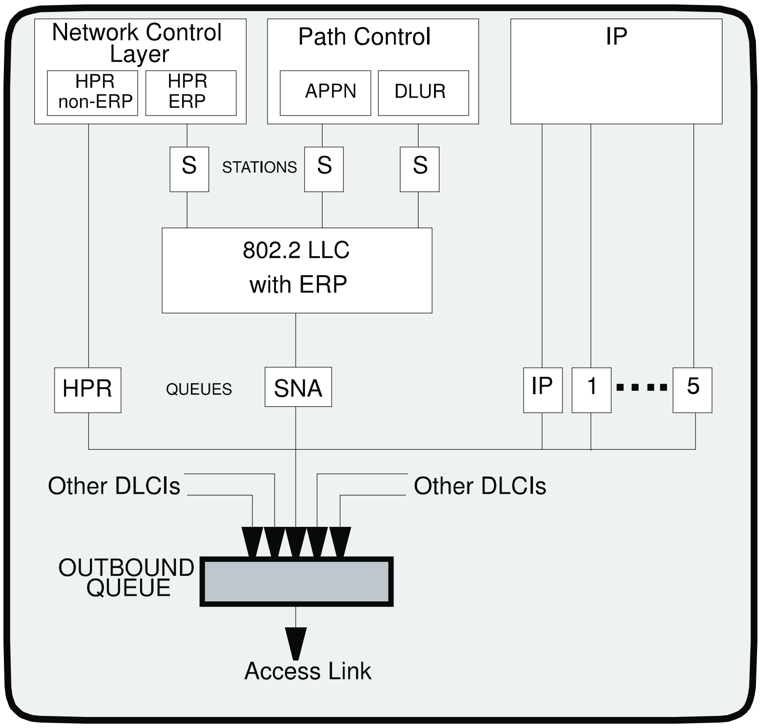

Figure 7-18. 3746 BRS

BRS is only used when there is more traffic to send than the available bandwidth. 3746 BRS is not used for FRFH DLCIs.

MAE Bandwidth Reservation System (BRS)

On serial connections, frame relay (FR) and PPP, the MAE implements a *Bandwidth Reservation (BRS)* mechanism. BRS enables the network administrator to reserve portions of the bandwidth of a circuit for specific types of data, differentiate between urgent, high, normal, and low-priority traffic within that bandwidth, and therefore favor the transmission of the highest-priority data.

BRS allows you to decide which packets to drop when demand (traffic) exceeds supply (throughput) on a network connection. Bandwidth reservation is not used until more than 100% of the available bandwidth is requested.

Bandwidth reservation *reserves* transmission bandwidth for a network connection. This reservation feature allocates minimum percentages of total connection bandwidth for specified classes of traffic.

BRS Components: BRS uses the following mechanisms to differentiate between traffic types and then to queue that traffic.

Circuit Classes

Frame relay interfaces can be grouped into circuit classes and each circuit class is assigned a percentage of the frame relay interface's bandwidth. The sum of bandwidths reserved per link must be less than 100%. A *default* class is defined per frame relay interface and cannot be deleted. The bandwidth assigned to the DEFAULT class can be changed.

Traffic Classes

Bandwidth reservation guarantees bandwidth for specific types of encapsulated traffic (classes) identified by either the protocol type or a filter. Traffic classes are defined for each PPP interface and each frame relay circuit.

BRS supports the following protocols:

- IP
- ARP
- IPX
- Bridging
- SNA/APPN-ISR (BAN and BNN)
- APPN-HPR (BAN and BNN)
- Appletalk
- DECnet IV
- Banyan Vines
- OSI/DECnet V

Note: By default, all protocols/applications are assigned to the default class with priority normal.

BRS also supports the following filters:

- IP tunneling
- SDLC tunneling over IP (SDLC Relay)
- Rlogin
- Telnet
- SNA/APPN-ISR
- APPN-HPR
- SNMP
- IP Multicast
- DLSw
- MAC Address (through MAC filtering tags)
- MAC Filters
- NetBIOS
- Network-HPR
- High-HPR
- Medium-HPR
- Low-HPR
- X.25 Transport Protocol (XTP)

Using either the type of protocol or a filter to differentiate between traffic types, traffic can be assigned to one of the *traffic classes*.

The reserved percentages for each class are a minimum slice of bandwidth for the network connection. When the network is operating to capacity, messages in any one class can be transmitted only until they use the configured bandwidth allocated for the class. In this case, additional transmissions are held until other bandwidth transmissions have been satisfied.

Priority Levels

Within each traffic class, the traffic can also be assigned a *priority level*. When BRS transmits packets for a traffic class, all packets with urgent priority are sent first, then all high priority, then all normal priority, and then all low priority. The following priority levels are defined:

- Urgent (U)
- High (H)
- Normal (N)
- Low (L)

Figure 7-19 shows three traffic classes. Each traffic class has its own set of data, which has been given a priority (shown by the four queues: urgent(U), high (H), normal (N), and low (L).

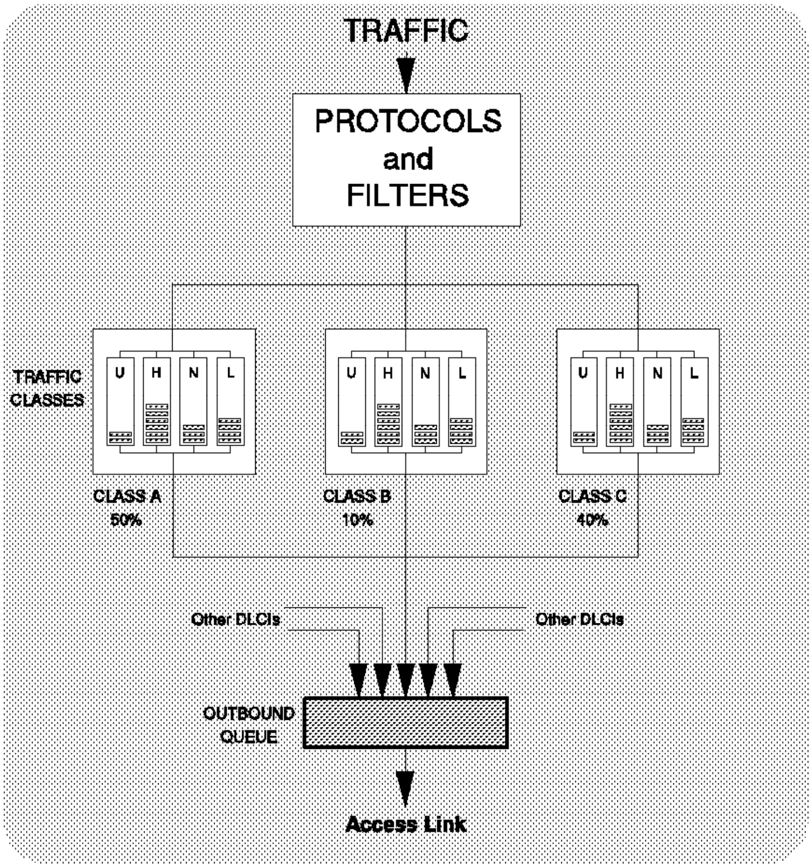

Figure 7-19. 3746 Multiaccess Enclosure BRS

Both *orphan circuits* (that is circuits that are not configured but are learned via LMI) and configured circuits with BRS explicitly disabled, use a default queueing mechanism where all frames are assigned to a default traffic class at the circuit level, and the circuits are assigned to the default circuit class.

Figure 7-19 on page 7-32 shows three traffic classes. Class A is assigned 50% of the bandwidth available to that DLCI, class B is assigned 10%, and class C is assigned 40%. Traffic bound for the DLCI shown is differentiated by the previously discussed protocol types or filters, and is assigned one of the four priorities. In the 2210 and 2216, each traffic class has a queue for each priority level.

BRS Support of APPN Traffic: When SNA/APPN-ISR is assigned to a traffic class, either APPN-ISR traffic that is being routed by the router's APPN code or SNA or APPN-ISR traffic that is being bridged will be assigned to this class. This is why SNA/APPN-ISR shows up as a protocol (the *routed* case) and as a filter (the *bridged* case). To identify SNA/APPN-ISR traffic that is being bridged, the BRS code looks for any bridging frames that use a DSAP or SSAP of 0x04, 0x08, 0x0C and a LLC (802.2) control field value that is *not* the un-numbered information (UI) type (that is *not* 0x03).

If SAPs other than 0x04, 0x08, or 0x0C are used for SNA/APPN-ISR bridge traffic, a sliding window MAC filter can be created to identify and tag SNA/APPN traffic. Using the BRS MAC filtering support, MAC filter tags can be assigned to a traffic class and priority.

When APPN-HPR is assigned to a traffic class, the BRS code looks for any bridging frames that use a DSAP or SSAP of 0x04, 0x08, 0x0C, and 0xC8 and a LLC (802.2) control field value that is equal to the un-numbered information (UI) type (that is 0x03).

If the user wants to differentiate between HPR-HPR traffic depending on its transmission priority then the user can use the following HPR filters:

Network-HPR
Used for HPR traffic that is using the network transmission priority.

High-HPR
Used for high transmission priority.

Medium-HPR
Used for medium transmission priority.

Low-HPR
Used for low transmission priority.

This means that one of the above HPR transmission filters can be assigned to a different traffic class and/or priority than the other APPN HPR traffic.

Summary of 3745/3746-900 SNA NCP Routing and Frame-Relay Switching

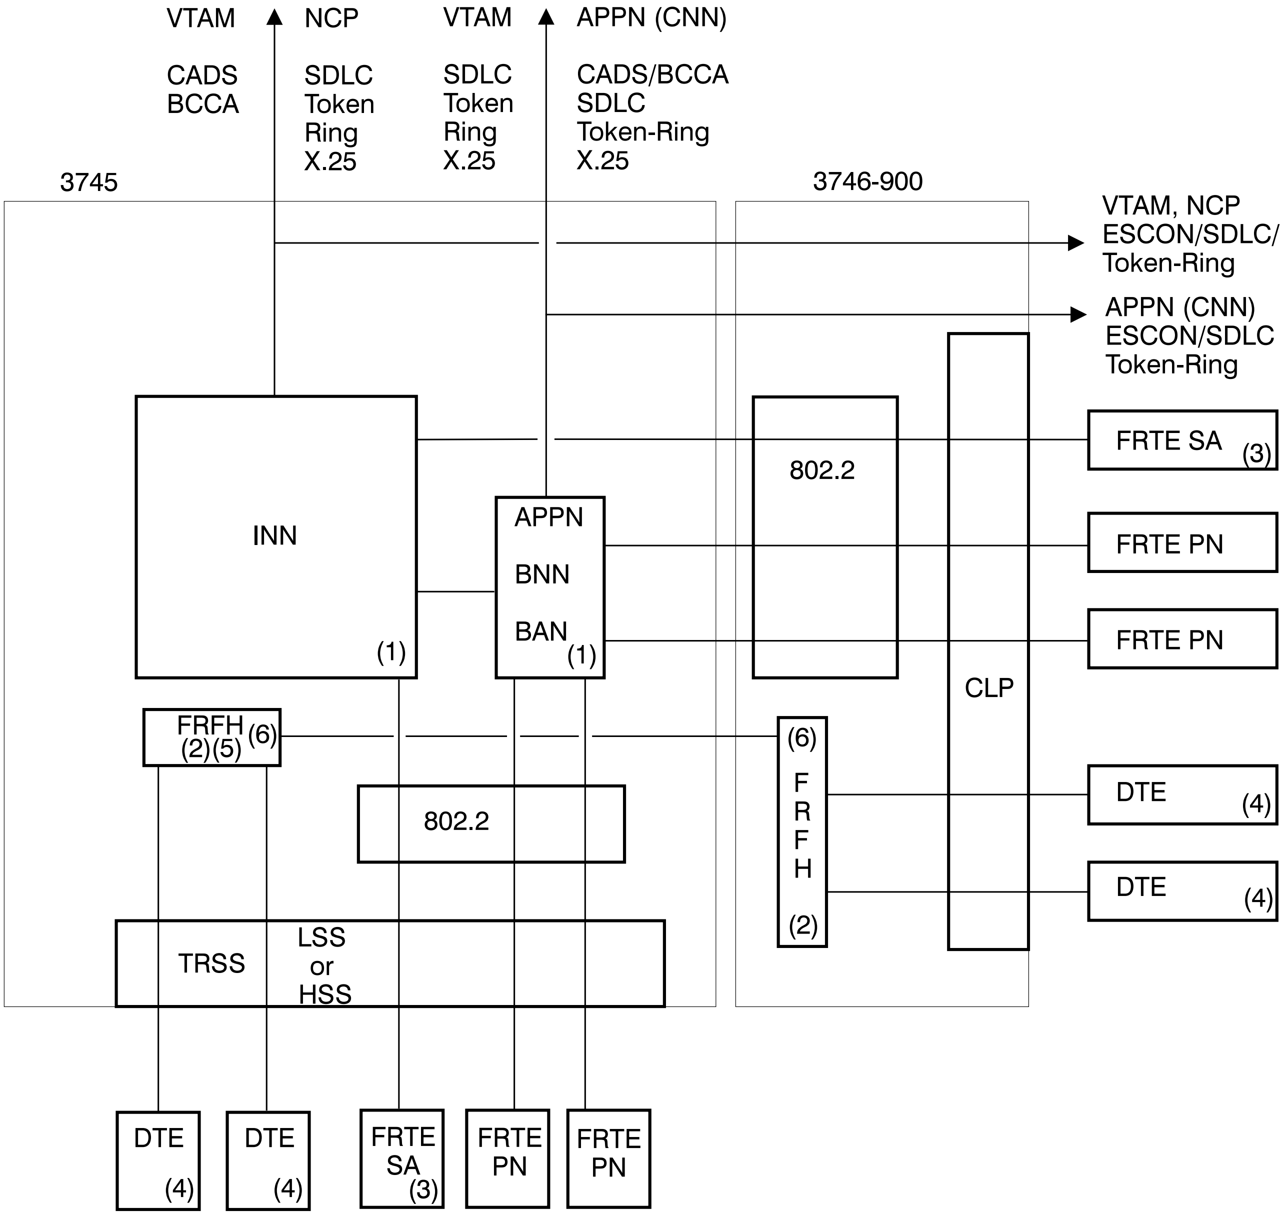

Figure 7-20. Summary of 3745/3746-900 SNA SNA Routing and Frame-Relay Switching

Notes:

1. SNA Routing layer.
2. Switching layer.
3. 3745 or 3745/3746-900.
4. These DTEs can be FRTE SA, FRTE PN, or any FRTE or FRFH.
5. FRFH is also supported over the 3745 TRSS (token-ring adapter) in addition to the LSS and HSS.
6. FRFH between a 3745 subport and a 3746-900 subport (NCP V7 R4).

3745 and 3746-900 FRFH Details

A DLCI used for FRFH has to be either predefined in NCP or dynamically added from VTAM.

For NCP V7 R4

In a NCP V7 R4 frame relay node, the frames can be switched between any pair of the following:

- 3745 LSS or HSS port
- 3745 TIC2 port
- 3746-900 CLP port

See Figure 7-20 on page 7-34.

When using substitute subports, the primary and secondary subports can be any type of the above ports.

For NCP V7 R2 and NCP V7 R3

When using substitute subports, both the primary and the substitute subport must be part of the 3746-900 or of the 3745 (for 3745/NCP frame relay). There is no frame switching between a 3745 line or a TIC2 port and a 3746-900 line. However, for SNA frame relay traffic, the routing functions of ACF/NCP are used to route the SNA traffic (subarea, peripheral) between a 3745 line and a 3746-900 line.

Table 7-2 shows the possible frame switching from either of the following:

- 3745 LSS/HSS/TIC2 subport to a 3745 LSS/HSS/TIC2 subport
- 3746-900 CLP subport to a 3746-900 CLP subport

Also refer to FRFH in Figure 7-20 on page 7-34.

Table 7-2. Subport FRFH for FR Switched Traffic with NCP V7 R2 and V7 R3

From Subport	**To Subport**			
	3745 Primary (LSS/HSS) (TIC2)	**3746-900 Primary (CLP)**	**3745 Substitute (LSS/HSS) (TIC2)**	**3746-900 Substitute (CLP)**
3745 Primary (LSS/HSS/TIC2)	Yes	No	Yes	No
3746-900 Primary (CLP)	No	Yes	No	Yes
3745 Substitute (LSS/HSS/TIC2)	Yes	No	N/A	N/A
3746-900 Substitute (CLP)	No	Yes	N/A	N/A

Note: N/A = Substitute subport to substitute subport is not applicable

There are no such limitations with NCP V7 R4 or V7 R5.

Functions not Supported by NCP

There is no NCP support for frame relay over TIC3s in the 3746-900 and no NCP support for IP over 3746-900 frame relay lines.

Functions Supported by the 3746 Nways Multiprotocol Controllers - models 900 or 950

Orphanage Support

The 3746 NNP supports *orphanage*. An *orphan circuit* is a PVC that is not defined. When an undefined DLCI is discovered thru LMI, the 3746 can learn the IP address of this DLCI by sending an Inverse ARP to acquire the partner IP addresses over active DLCIs. LMI is required to support orphan circuits. If a partner FR DTE does not support Inverse ARP, the DLCI used for IP must be statically defined in CCM. The 3746 answers all Inverse ARP messages received.

Note: When the 3746 is directly connected with no intermediate network to equipment that only supports the LMI user side, LMI should be set to NUI the 3746 and the attached equipment, such as a 2210 or 2216, must be configured with *orphanage OFF*. This is illustrated in Figure 7-21 and Figure 7-22.

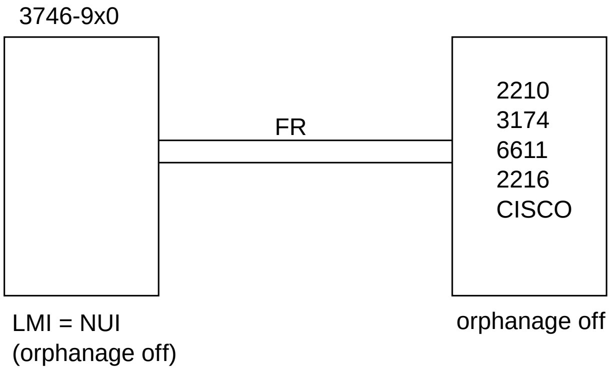

Figure 7-21. Direct Connection with a Router

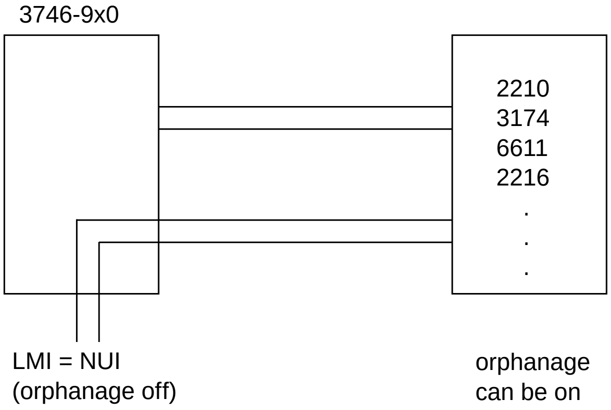

Figure 7-22. FRSE and Direct Connection

Note: *Orphanage OFF* will require explicit definition of the DLCIs but will save cycles at session setup time. Using *orphanage* will save on definitions but will generate more network traffic.

The 3746 NNP never sends ARP messages. The 3746 answers ARP messages received;the ARP answer is sent by the 3746 on the same DLCI as the ARP was received. This could be different from the DLCI number found inside the ARP message itself. The format of the ARP message is defined in RCF 1490. When the 950 is connected to a frame relay cloud LMI should be set to adaptive and orphanage can be set on as shown in Figure 7-23

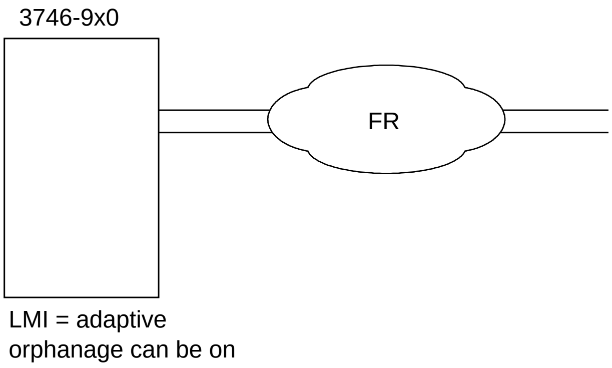

Figure 7-23. Not Point to Point

FRTE Peripheral Node Connections (FRTE PN)

FRTE PN is supported on a direct connection as shown on Line 3, or through an intermediate network as shown between Lines 1 and 3 in Figure 7-24.

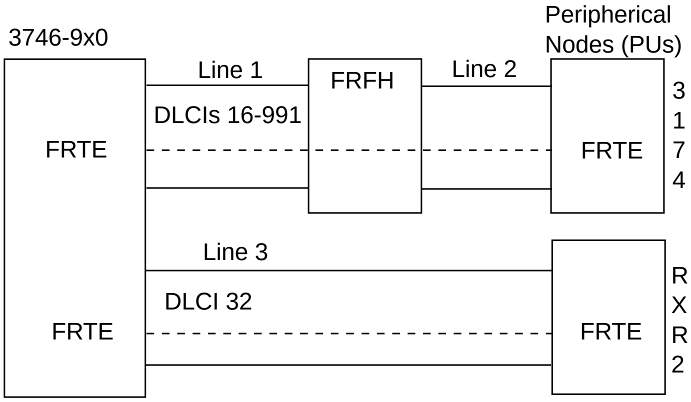

Figure 7-24. Frame Relay Terminating Equipment Peripheral Node

Frame relay congestion management is provided by the 3746 when:

- A frame with BECN bit set is received; the sending rate is decreased.
- A frame with FECN bit set is received; the BECN bit is set in the next sent frame.

The 3746 provides frame-discard eligibility (DE) support if selected at CCM configuration time. When DE is supported by the transport network, the 3746 uses the DE bit in the frame-relay packet header. The DE can be set for packets containing data frames, but it is never turned on (set) for packets containing SNA or APPN/HPR control frames.

Both PU 2 and PU 2.1 devices are supported with either format:

- RFC 1490 routed
- RFC 1490 bridged

Frame Relay Local Management Interface (LMI) Support

The 3746 LMI support conforms to ANSI T1.617 Annex D and ITU-T Q.933 Annex A.

Frame relay LMI support is provided entirely by the 3746 for the 950 NNP or in the NCP generation for the 3746-900. The LMI timer and threshold values are defined using CCM. They are passed to the 3746-9x0 CLPs at activation time. Events that are detected via the LMI (for example, LMI status change or reaching an LMI error threshold) may result in failure of the associated resources that are handled by the NNP.

The different LMI modes are dynamically discovered at connection establishment time as follows:

- **Network to Network Interface (NNI)**: The link integrity verification (LIV) inquiries and the PVC status inquiries and responses are sent by both partners on the frame relay interface. This is also called LMI bidirectional protocol.
- **Network to User Interface (NUI)**: The partner is a DTE that does not support NNI. The partner polls the 3746-9x0 and does not answer the LIV enquiries sent by the 3746-9x0.
- **User to Network Interface (UNI)**: The partner is a DCE that does not support NNI. The partner answers the status inquiries sent by the 3746.
- **No LMI**: The partner does not answer the status request and does not send status inquiries to the 3746-9x0.

System Definitions

The configuration of 3746-9x0 frame relay resources is done using CCM or in the NCP generation:

- The frame relay port of the 3746-9x0 contains the LMI parameters (ANSI versus CCITT, timers, thresholds).
- Frame relay DLCI parameters.
- Stations (SNA, APPN/HPR) used for FRTE. Logical SNA stations attached via DLUR may be defined in the VTAM/DLUS Switched Major Node.
- IP over frame relay.

Maximum Number of Resources Supported Per CLP

For information about the resources supported by a CLP please refer to “Communication Line Processor Connectivity” on page 18-2.

Frame-Relay Tuning Recommendation

A Note of Warning

> Proper operation of the 3746 requires correct setting of the configuration parameters. When making frame-relay line definitions with the CCM, this section should be used for guidance.

NCP Frame-Relay Tuning Recommendations

For additional information that you cannot find in this guide, refer to *IBM 3746-900 and NCP Version 7, Release 2*, GG23-4464 (a “redbook”).

- **Link Transmission Group (TGCONF in the PU Statement)**

  `TGCONF` specifies whether this subarea link station is in a multilink or single-link transmission group. The default is:

  ```
  TGCONF=MULTI
  ```

  For single-link transmission groups, code:

  ```
  TGCONF=SINGLE
  ```

  Coding `TGCONF=SINGLE` for a transmission group that contains a single SDLC, token-ring, or frame relay line can improve NCP performance for the transmission group. Therefore, it is recommended to code `TGCONF=SINGLE` for any transmission group that will contain a single frame relay connection.

- **FRTE Resources (LOCALTO or T1TIMER in LINE Statement)**

  `LOCALTO` or `T1TIMER` specifies the T1 reply timer (for example, the time a 3745/3746-900 FRTE will wait to receive the acknowledgment from the frame sent to its FRTE partner). When coding the T1 timer, the time taken by a frame to reach the remote partner and the time taken for the reply to come back must be taken into account. This time is the round trip delay and is mainly dependent on:

  – The speed of the lines on each hop

  – The number of hops to reach the partner

  – The level of congestion that can create queuing in the transiting nodes

  Therefore, the minimum value to code is twice the time taken by a frame of an average size to be transmitted from the FRTE (that you are defining) to the FRTE partner.

  For example, for an average frame size of 256 bytes, a speed of 19.2 Kps, and a path with two hops (FRTE to intermediate FRFH and FRFH to FRTE) the minimum transmit time is:

  ```
  (256 × 8) ÷ 19200) × 2 = approximately
  0.2 seconds
  ```

  .

  In addition, you may assume an average queue of four frames in the sending FRTE and the intermediate FRFH (transit node). This means an overall transmit time of:

  ```
  0.2 + (4 × 0.2) = 1 second.
  ```

Therefore the time taken to send a frame and receive the reply is 2 seconds. There will be additional queueing during peak periods and additional processing time. To avoid unnecessary retries it is recommended that the LOCALTO value be set to three times the minimal value calculated in the above fashion. In the above example use a `LOCALTO` of approximately 6 seconds.

- **Blocking Factor (BLOCK in the PU Statement)**

  `BLOCK`, in the PU statement, can be used to improve the throughput of the frame relay INN links and reduce the processing load of the CLP. NCP can support multiple path information units (PIUs) in frame relay frames when they are routed over transmission group links. `BLOCK` specifies the maximum frame size and the maximum number of PIUs per frame.

  Blocking occurs in NCP when a high traffic load leads to queuing in NCP before transfer to the 3746-900. For NCP queueing to occur, one of the following must take place:

  – Queueing in the 3746-900 because the line speed is not large enough to transmit all the traffic for all the DLCIs

  – The partner NCP node does not acknowledge fast enough

  – The intermediate network is congested

  If you have blocking, defining `BLOCK` can improve the performance in the 3746-900 (as there are fewer frames to transmit) and can result in better line utilization (as there is less overhead due to frame relay headers).

  Note: Blocking was originally created for SDLC lines running on LSS and HSS lines, which have an inter-frame gap of flags between frames. This does not apply to 3746-900 lines.

  If `BLOCK` is coded, NCP verifies that the sender of the XID2 supports PIU blocking. If not, NCP uses one frame per PIU. If the sender of the XID2 supports a smaller block size, NCP adjusts its block size to the value in the XID2.

  Code the blocking factors at each side of the transmission group links connecting two subarea nodes. The values recommended are:

  ```
  BLOCK=(4096,16)
  ```

  This means that each SDLC frame is up to 4096 bytes long and will contain up to 16 PIUs.

  Blocking is also dependent on the line speed. With a speed less than or equal to 64 Kbps, the values recommended are:

  ```
  BLOCK=(2048,8)
  ```

- **Interframe gap (ADDIFG in the LINE statement)**

  When the line adapter of the partner controller is experiencing a high rate of overrun errors, code:

  ```
  ADDIFG=YES
  ```

  This increases the gap between frames and decreases the frame rate.

  The ADDIFG support was enhanced in NCP V7 R5. There are additional definitions possible for a frame relay physical link:

  NO Transmit a minimum of one flag between frames.

YES Transmit a minimum number of flags between frames as determined by NDF using the line speed.

NNN Transmit at least NNN number of flags between frames, where NNN is a decimal value from 1 to 255.

For more details about these values, refer to the *NCP Resource Definition Reference* manual.

- **Communication Rate (COMRATE in the PU Statement**)
- **Communication Rate (DATABLK in the LINE Statement**)

  In the 3745 and 3746-900 the input is not policed, that is, the DTE can enter data up to the lined speed with any distribution on the DLCIs. For instance, one DLCI can enter data at line speed and take the whole bandwidth for it.

  The 3745 and 3746-900 applies an *outbound* congestion control mechanism called *communication rate* (CR) during transmission.

  The communication rate function in the 3745 and 3746-900 allows you to prioritize virtual circuits (VCs) on a physical link. When the available transmission bandwidth is fully used, the CR allows each virtual circuit (DLCI) to have a guaranteed portion of the physical link bandwidth.

  A numeric value (n), specified in COMRATE, is assigned to each virtual circuit. This value represents the number of transmission units sent for this virtual circuit at each of its transmit opportunities when the outbound link is congested. The DATABLK keyword defines the size of the transmission unit. By assigning different (n) values to different virtual circuits, the virtual circuits with higher (n) values will be able to use more of the available bandwidth than those with a lower (n) value.

Communications Rate Definition

The CR is defined by a numeric value (n) on the COMRATE keyword at the station level (FRSE or FRTE). This value is multiplied by the value of the DATABLK keyword value, defined at the physical link level, to determine the maximum number of bytes sent by this station at each transmit opportunity:

- COMRATE = (FULL|NONE,n):
  - FULL|NONE applies only to FRTE.
    - FULL = DE bit is off in all frames
    - NONE = DE bit is on in data frames.
  - n applies to both FRTE and FRSE.
    n = the relative priority of this DLCI. n ranges from 1, low, to 64, high, with a default value of 1.
- DATABLK = transmission unit size in bytes
  DATABLK defines the minimum transmission unit allocated to one DLCI of this physical link. It is coded in the Physical LINE definition statement.

  DATABLK ranges from 265 to 16372 with a default value of 2048.

Note: When a frame larger than DATABLK is sent and n= 1, the frame is not segmented before being sent. For example, if DATABLK = 500 and a one kilobyte frame is sent, the DLCI is serviced as if n = 2.

Communications Rate Value Recommendations

There is a relation between the:

- Line speed of a physical port
- Number of DLCIs defined on this port
- Communication rate (CR) defined for each DLCI

In case of congestion, the 3746-900 will transmit, for each DLCI, the number of bytes defined by CR with a round-robin service mechanism. This means that the maximum time elapsed between each opportunity to send a CR for a DLCI is the total time required to send a CR for each DLCI:

```
(Sum (n*DATABLK)) * 8 / line speed
```

This value represents the maximum duration taken to send a CR for a given DLCI when all the virtual circuits are exceeding their share of the bandwidth. It is the time the CR burst may wait in a node on a busy transmitting line. It is recommended that this not exceed 3 seconds for any low speed line. The recommended maximum duration calculated for higher speed lines is less. Recommendations are below:

- 0.4 s for E1 speed
- 0.5 s for T1 speed
- 1 s for 256Kbps
- 2 s for 64Kbps
- 3 s for 19.2Kbps

Communications Rate and the Transmit Window

In case of congestion, the data to be transmitted may be held up in the outbound queue before it is transmitted over the line. The time elapsed in the queue should be taken into consideration when calculating the overall round trip delay that is used to define the T1 timer. Refer to FRTE Resources (LOCALTO or T1TIMER in LINE Statement) on page 7-39.

For FRTE, the burst of data at each transmission opportunity (`n * DATABLK`) should be at least as large as the average amount of data transmitted in a window (`MAXOUT * average frame size`).

Note: If XID2 or XID3 is used, `MAXOUT` is taken from the XID field defined as the maximum number of I-frames that can be received by the XID sender before an acknowledgment is sent.

For type 1 and type 2.0 nodes that do not send XID3s, `MAXOUT` is taken from the VTAM PU statement.

Communications Rate Coherence across the Network

The bandwidth allocated to a permanent virtual circuit (PVC) should be the same all along the PVC, that is, for a given PVC, the portion of the bandwidth allocated to the PVC all along the path, `(n / Sum (n))*SPEED`, must be the same on each transmitting leg of each hop.

For example, take a congested line with a speed of 19.2Kbps and two DLCIs defined, the first DLCI with n = 2 and the second with n = 3. If `DATABLK` = 1000, the first DLCI will be allowed 2000 bytes (16000 bits) and the second DLCI will be allowed 3000 bytes (24000 bits) out of the next 5000 bytes to be sent.

3746 Frame-Relay Tuning Recommendations

Frame-Relay Port: DLC Parameters

When the line adapter of the partner controller is experiencing a high rate of overrun errors, code YES for the Interframe gap (ADDIFG). This increases the gap between frames (that is, the number of flags sent between frames), and decreases the frame transmit rate.

Frame-Relay LMI Parameters

LMI Mode, ANSI versus ITU-T: Be sure to define the same LMI mode in the 3746 and in the partner attached equipment. The attached equipment may not have the same default mode as the 3746, which is ITU-T (ITU-T was previously CCITT).

LMI Echo:

This is a unique feature of the 3745 and the 3746. If you are connected to any other equipment other than a 3745 or a 3746, you must keep the default value, which is Neither.

LMI Timers: When you don't use the default values you must code t392 on one side higher than t391 on the other side. A recommended ratio is 1.5 between them, such that:

t392 = 1.5 × t391.

Maximum Number of DLCIs: It is important to code the actual maximum number of DLCIs that are used on this interface when attaching devices that only support the LMI User to Network Interface (UNI) such as the 2210, 2216, or 3174. When attaching devices that only support UNI LMI, the 3746 sends a PVC full status message each time the DTE polls the 3746 with a full status inquiry. If the maximum number of DLCIs is specified too high, the PVC full status message is too long and wastes the bandwidth.

Direct Connection with a Router: When there is no intermediate network for the LMI on the 3746-9x0 side, we recommend coding the LMI as *NUI*, which in turn forces *orphanage off* for both the 3746-9x0 and the router. Remember that orphanage off will require explicit DLCI definitions. The directly connected router should have orphanage off. If the router can not turn orphanage off we should take care to code a minimum number of DLCIs for this connection. This can be done using the MAXDLCI.

Frame-relay Frame Handler Interface: When the 3746-9x0 functions as a frame handler, we recommend coding the LMI as *adoptive* and specifying *orphanage on*.

Frame Handler and Direct Connection: When the 3746-9x0 functions both as a frame handler as well as direct connection, we recommend coding the LMI as *NUI*, which forces *orphanage off*. The remote can have orphanage on; this is a designer choice.

Figure 7-25. Frame Handler and Direct Connection

Frame-Relay DLCI/COMRATE Parameters

Communication Rate: The 3746 9x0 applies a congestion control mechanism called *communication rate* (CR) for both NCP and NNP resources during transmission.

The communication rate function in the 3746 allows you to prioritize stations on a physical link. When the available transmission bandwidth is fully used, the CR allows each station to have a guaranteed portion of the physical link bandwidth.

A numeric value (n) is assigned to each station. This value represents the number of transmission units, defined by `DATABLK`, sent for this station at each of its transmit opportunities when the link is congested. By assigning different (n) values to different stations, if the bandwidth is totally used the station with higher (n) values will use more of the available bandwidth than the other stations.

`DATABLK` is defined in the Frame-Relay Port-DLC Parameters screen. The numeric (n) value per station is defined in the Frame-Relay DLCI Parameters screen in CCM. The value defined for APPN/HPR is applicable to each APPN/HPR or SNA (DLUR) station of the DLCI. The value defined for IP is applicable to all IP traffic over the DLCI.

Notes:

1. When a frame larger than `DATABLK` is sent and n = 1, the frame is not segmented before being sent. For example, if `DATABLK = 500` and a one kilobyte frame is sent, the DLCI is serviced as if n = 2.
2. When there is less than `n * DATABLK` in a queue, the 3746 will send what is queued for this station if the window permits. The 3746 does not wait for the amount of data queued for this station to reach `n * DATABLK` before sending data.
3. For NCP/VTAM, the range for COMRATE is 1 to 64, which means:

   ```
   COMRATE (BYTES) = (1 to 64) * DATABLK
   ```

   In the CCM, the value for COMRATE is shown and set in bits, this means:

   ```
   COMRATE (bits) = (1 to 64) * DATABLK * 8
   ```

   This explains why the default value in the CCM screen is set to 16384 bits (default DATABLK=2048 bytes * 8).

Communication Rate Value Recommendations: The following factors influence the Communication Rate value used:

Line Speed

There is a relation between the:

- Line speed of a physical port (SPEED)
- Number of DLCIs and stations defined on this port
- Communication rate (CR) defined for each station

In case of congestion, the 3746 will transmit, for each station. the number of bytes defined by CR with a round-robin service mechanism. This means that the maximum time elapsed between each opportunity to send a CR for a station is the total time required to send one CR for every station:

```
(Sum (n*DATABLK)) * 8 / SPEED)
```

This value represents the maximum duration taken to send a CR for a given station when all the stations are exceeding their share of the bandwidth. It is the time the CR burst may wait in a node on a busy transmitting line and must not exceed t seconds. This value must be lower than 0.4 s for high-speed lines, 1.5 s for medium-speed lines, and must never exceed 3 s for low-speed lines.

Transmit Window

In case of congestion, data to be transmitted may be held up in the outbound queue before it is transmitted over the line. The time elapsed in the queue should be taken into consideration when calculating the overall round trip delay that is used to define the T1 timer. Refer to "Frame-Relay DLC Parameters for FRTE Stations" on page 7-46.

For FRTE, the burst of data in each transmission (n * DATABLK) should be at least as large as the average amount of data transmitted in a window (MAXOUT * average frame size):

```
(n * DATABLK) ≥ (MAXOUT * average frame size)
```

MAXOUT is defined as the maximum number of I-frames that can be received by the XID sender before an acknowledgment is sent.

To give priority to the DLCIs that transport interactive traffic over DLCIs that transport batch traffic, it is recommended that you assign:

- High windows (10 to 20) and high COMRATE to interactive traffic
- Small windows (3 to 6) and small COMRATE to batch traffic

Notes:

1. If XID3 is used, the 3746 uses whichever is the lowest MAXOUT, either the value received in the XID3 or defined in the configuration file by the CCM.

   The 3746 will set the field of the transmitted XID3 to the value of MAXIN defined by the CCM.

2. For type 1 and type 2.0 nodes that do not send XID3s, MAXOUT is taken from the CCM FRTE Station-DLC Parameters 2/2 dialog window.

MAXFRAME and DLCIs

- It is recommended that MAXFRAME be set at no more than 2K bps for reasons of congestion control and CRC effectiveness. This is set at the line level. When LMI is enabled, LMI exchanges require 5 bytes per DLCI. Thus MAXFRAME/5 provides another limit on the number of DLCIs per physical line. (Since there is other data in the LMI exchange, this number should be further reduced by 5.)

- CLP level constraints suggest that 500 DLCIs should be an upper limit for the physical frame relay lines they support.

- In NUI LMI configurations, MAX_DLCI can be used to limit the number of DLCIs, since each DLCI contributes to full status exchanges and I_ARP traffic.

- The burst of data in each transmission (n * DATABLK) is recommended to be at least four times as large as MAXFRAME in order for the communication rate mechanism to work when there is congestion.

Communication Rate Coherence across a 3746 Network: The bandwidth allocated to a permanent virtual circuit (PVC) should be the same all along the PVC, that is, for a given PVC the portion of the bandwidth allocated to the PVC all along the path `((n / Sum (n)) * Speed)` must be the same on each transmitting leg of each hop.

For example, take a congested line with a speed of 19.2Kbps and two DLCIs defined, the first DLCI with n = 2 and the second with n = 3. If `DATABLK = 1000`, the first DLCI will be allowed 2000 bytes (16000 bits) and the second DLCI will be allowed 3000 (24000 bits), so that each will get its proportional use of the physical line.

Frame-Relay DLC Parameters for FRTE Stations

T1TIMER: `T1TIMER` specifies the T1 reply timer, for example, the time a 3746-950 FRTE will wait to receive the acknowledgment for the frame sent to its FRTE partner. When coding the T1 timer, the time taken by a frame to reach the remote partner and the time taken for the reply to come back must be taken into account. This time is the round trip delay and is mainly dependent on:

- The speed of the lines on each hop
- The number of hops to reach the partner
- The level of congestion that can create queuing in the transiting nodes

Therefore, the minimum value to code is twice the time taken by a frame of an average size to be transmitted from the 3746 that you are defining to the FRTE partner.

For example, for an average frame size of 256 bytes, a speed of 19.2 Kbps, and a path with two hops (FRTE to intermediate FRFH and FRFH to FRTE), the minimum transmit time is:

```
(256 × 8) ÷ 19200) × 2 = approximately 0.2 seconds
```

In addition, you may assume an average queue of four frames in the sending FRTE and the intermediate FRFH (transit node). This means an overall transmit time of:

```
0.2 + (4 × 0.2) = 1 second.
```

Therefore the time taken to send a frame and receive the reply is 2 seconds. To take into account additional queueing during peak periods and prevent useless retries, use three times this value (for example use a `T1TIMER` of approximately 6 seconds).

RETRIES: The total number of times that a frame can be sent is `(m + 1) × (n + 1)`.

The corresponding total time that can elapse before using all of the retries is `(m × (n + 1) × T1TIMER) + (n × t)` where:

- m = the number of retries per retry sequence
- n = the number of retry sequences
- t = the pause between retry sequences.

T2TIMER and N3: These parameters permit the delay of sending isolated RR (Receiver Ready). An RR is sent either when the T2TIMER elapses or when n3 frames have been received. In the meantime, when the traffic is not uni-directional there is a high chance that the acknowledgment is sent within a data frame (802.2 “piggybacking”). When you decrease the number of isolated RRs, less CLP processing power is required to handle the 802.2 LLC and the line utilization is better.

Notes:

1. T2 and n3 are recommended for file transfer.
2. Caution should apply when using T2 and n3 when there is interactive traffic because the response time may be degraded.

IP over Frame Relay Parameters

Orphanage is *disabled* when LMI=NO. When the attached router is directly connected and supports only the LMI user side, then LMI should be set to NUI, which in turn disables orphanage.

VTAM DLCADDR for Frame Relay DTE

When the frame relay DTE connection with a dependent PU is initiated from VTAM, the following connection parameters must be defined in the VTAM Path definition statement by the DLCADDR keyword. Refer to the *VTAM Resource Definition Reference* for the syntax.

A frame relay connection for an SNA peripheral resource has these elements:

1. DLC type: FRPVC
2. Port name
3. DSAP (destination service access point)
4. DLCI (data link connection identifier of the FR PVC)
5. SSAP (source service access point)
6. Destination MAC address (if BAN attached)

Code a DLCADDR keyword for each element as following the syntax:

```
DLCADDR=(subfield_id,type,string)
```

The procedure to set the SNA peripheral resources follows:

Step 1. To specify a frame relay DLC type, code:

```
DLCADDR=(1,C,FRPVC)
```

Step 2. Identify the port name of the frame relay physical line:

```
DLCADDR=(2,I,portname)
```

where `portname` is the port name defined in CCM. The 3746 NN expects the port name in ASCII. The ASCII type is indicated by I and requires:

- For VTAM V4 R2, PTF UW28497
- For VTAM V4 R3, PTF UW28498.

Step 3. To specify the DSAP, code:

```
DLCADDR=(3,X,hh)
```

where hh is the SAP of the remote frame relay device in hexadecimal. The value must be even, in the range X'02' to X'FE'.

Step 4. To specify the DLCI of the frame relay PVC, code:

```
DLCADDR=(4,D,nnn)
```

or

```
DLCADDR=(4,X,hhh)
```

where nnn and hhh are the local DLCI number of the remote frame relay device in decimal and hexadecimal respectively. The value must be in the range of 16 to 991 in decimal, that is, X'10' to X'3DF' in hexadecimal.

Step 5. To specify the SSAP, code:

```
DLCADDR=(5,X,hh)
```

where hh is the SAP of the source NNP in hexadecimal. If you choose to specify the source SAP, the value must be the same as the local SAP defined for the frame relay port - DLC parameters.

Step 6. To specify the destination MAC address for the BAN, code:

```
DLCADDR=(6,X,hhhhhhhhhhhh)
```

where hhhhhhhhhhhh is the destination MAC address of the BAN attached peripheral device on the LAN which connects it to the remote BAN router. The range is from X'1' to X'7FFFFFFFFFFF'.

For example, suppose you need to make outgoing connections from DLUS to a device using a frame relay PVC with a DLCI of X'20' and a DSAP of 8 on the port name FRP001. If the connection is with BAN, code:

```
DLCADDR=(1,C,FRPVC)              Frame relay PVC
DLCADDR=(2,I,FRP001)             Port name
DLCADDR=(3,X,08),                DSAP=8
DLCADDR=(4,D,32)                 DLCI=32
DLCADDR=(6,X,400000071088)       Remote MAC address.
```

Note that the SSAP need not be coded, as the frame relay port local SAP will be used.

LMI T391 and T392 Timers

The T391 timer on the user side must be lower than the T392 timer on the network side. A good ratio is T392 = 1.5xT391.

Note: NCP decides on the user and/or network side support dynamically. The NCP defaults are T391 = 10s and T392 = 15. These are usually adequate; however, they should be checked when connecting remote NCPs via frame relay.

Frame Relay 802.2 Values

The following describes the 3746 support for setting 802.2 timers for frame relay connections.

802.2 parameters can be set for stations individually. These values are used when the 3746 established dial-out connections. For dial-in connections, the 3746 uses hard-coded values. These values are listed below (timer values are in tenths of a second):

```
rsap                  = 0x04;
retrans_t1_timer      = 0x20;
ack_wait_t2_timer     = 0x00;
inact_err_ti_timer    = 0x60;
xmit_window_cnt       = 0x08;
receive_window_cnt    = 0x01;
max_retry_count       = 0x10;
second_level_retry    = 0x00;
pause_between_retry   = 0x00;
rnr_limit             = 0x708;
infinit_retry         = 0x01;
hpr_sap_value         = 0x04;
hpr_timer_t1          = 0x06;
hpr_timer_ti          = 0x30;
hpr_max_retry_count   = 0x03;

window_increment      = 0x04;
window_dec_for_cgst   = 0x01;
window_dec_for_lost   = 0xFF;
discard_eligibility   = 0x01;
echo_defeat           = 0x00;
pu_type               = PUTYPE_2_0;
```

File Transfer with IP over Frame Relay

For the best possible throughput for file transfer over IP connections, use the CCM to specify a maximum BTU, or maximum IP transmission unit greater than the default value of 2052. This should take into account the transmission quality of the lines that will be traversed.

NetDA/2 Frame Relay Support

IBM's “Network Design and Analysis/2 (NetDA/2)” is an OS/2-based tool for SNA subarea and APPN topology design and evaluation. It also generates PATH and COS tables and provides assistance to those implementing networks with frame relay, High-Performance Routing (HPR), or asynchronous transfer mode (ATM) connections.

The *NetDA/2 V1R4 Design Tool Guide and Tutorial* redbook SG24-4225 describes how to use NetDA/2. A chapter of this redbook describes using NetDA/2 to create NCP statements for frame relay networks.

NetDA/2 offers the following frame relay capabilities:

- Allows definition of existing PVCs in standard NetDA/2 tables
- Generates PVCs according to user-defined parameters
- Assigns DLCIs automatically but also allows you to control how DLCIs are assigned
- Displays frame relay nodes on the Network view
- Lists PVCs and highlights them on the Network view
- Produces reports on network PVCs, which can be arranged by node or PVC

- Adds PVCs to the TRUNKS table (registers PVCs as TGs) so that they can be used in route control
- Combines PVCs into MLTGs
- Integrates PVCs into existing subarea routes, thereby “shortening” the routes. When NetDA/2 integrates PVCs into routes, it registers the PVCs as TGs. When you select **Integrate**, NetDA/2 attempts to integrate only unregistered PVCs into routes.

NetDA/2 will not only help in the design of a future frame relay network, but also in the evaluation of an already installed one. Using utilities available in SSP V4R3, the user can process the NCP definitions and create files that are readable by NetDA/2. After NetDA/2 has processed these files, the model will contain the labels of all the resources involved in the frame relay definitions, physical and logical, including the DLCIs and MLTGPRIs. Once the model includes all this data, NetDA/2 can produce reports detailing the path followed by each PVC. The user can quickly verify the accuracy of the model by requesting NetDA/2 to produce a skeleton of the NCPs; the definitions created should match the NCP sources.

The frame relay definitions are quite complex and the user can easily make mistakes. A common one is to create PVCs that will loop; the data may go more than once through the same NCP. NetDA/2 will detect these loops and produce a report that will highlight those PVCs.

NetDA/2 can also show in a map or diagram the path followed by the different PVCs.

Chapter 8. Point-to-Point Protocol (PPP) Overview

The Point-to-Point Protocol (PPP) provides communication between stations connected on a simple point-to-point link and has three main components:

1. A method for encapsulating IP datagrams over serial links.
2. A Link Control Protocol (LCP) for establishing, configuring, and testing the data-link connection.
3. A family of PPP Network Control Protocols (PPPNCP) for establishing and configuring different network-layer protocols.

The 3746 Nways Multiprotocol Controller provides the following support for PPP:

- Base PPP and Line Control Protocols (LCP) (RFC 1331).
- PPP uses LIC11 and LIC12 and is supported for *leased* lines only, not for *switched* lines.
- Compressed TCP/IP headers (RFC 1144).
- IP Control Protocol (IPCP) (RFC 1332).
- PPP Link Group MIB (RFC 1471).

PPP in Operation

To establish communication over a point-to-point link:

1. Each end of the link must first send LCP packets to configure and test the data link.
2. After the link has been established and optional facilities have been negotiated as required by the LCP, PPP must send PPPNCP packets to choose and configure one or more network layers.
3. Once each of the chosen network layers has been configured, datagrams from each network layer protocol can be sent over the link.

Once the link has been established, the two stations can start communicating.

The PPP Frame

The unit of communication is the PPP frame which consists of the following items:

- A flag, indicating the start of the frame
- An address field
- A control byte
- A two-byte protocol field
- The information being communicated between the two stations
- Padding
- FCS (Frame Check Sequence)
- A flag, indicating the end of the frame

The Protocol field is defined by PPP and is *not* a field defined by HDLC. All protocol field values are odd, and assigned so that the least significant bit of the most significant octet is 0. Protocol field values are shown in Table 8-1 on page 8-2.

Table 8-1. PPP Protocol Field Values

Protocol Field value	Protocol Name
0x0021	IP data frame
0x002B	IPX data frame
0x0031	Bridging data frame
0x004B	APPN ISPR data frame
0x004D	APPN HPR data frame
0x8021	IP control protocol
0x802B	IPX control protocol
0x8031	Bridging control protocol
0x804B	APPN ISPR control protocol
0x804D	APPN HPR control protocol

Note: For IP over PPP, the IP Control Protocol is the *only one* that is supported.

LCP Options Supported by the 3746 Nways Multiprotocol Controller

The following LCP options are supported by the 3746 PPP:

- Maximum Receive Unit (MRU). The minimum value for the MRU is such that 576 data bytes plus the PPP header will fit into one packet. The default value for the MRU is 1500. The maximum value for an MRU is 4096 bytes.
- Async Control Character Map (ACCM). This option is used for control character transparency on async links. The PPP handler will always accept this option, but will ignore it. It will never generate this option.
- Magic number. The PPP handler will always request this option, but will accept a rejection from the other end. It will always accept this option.

LCP Options Not Supported by the 3746 Nways Multiprotocol Controller

The 3746 PPP network handler will always reject these options and never generate them:

- Authentication protocol
- Link Quality Monitoring (LQM)
- Protocol field compression
- Address and control field compression
- 32-bit FCS

IP over PPP Protocol Support

The 3746 Nways Multiprotocol Controller supports IP over PPP:

IP Compression Protocol
Van Jacobson (VJ) Compressed TCP/IP is supported for headers only, not for data. Up to 16 slots will be supported by the PPP handler for VJ compression.

IP address
The PPP handler will always report its IP address to the other side in a configuration request message. If the other side reports its IP address, the PPP handler simply stores it for future use. It does not check for equality between the IP addresses assigned to the two ends of the link.

If the other end of the link requests its own address to be reported to it (by sending a zero value as its own address), the PPP handler will reject the option.

Configuring a PPP Port

This section describes the parameters that must be set with Controller Configuration and Management (CCM) to configure a PPP port. Refer to Controller Configuration and Management help panels for default values and valid ranges of values.

Port Type

This parameter defines the PPP line type as leased, that is, a dedicated, non-switched line. The line is activated as soon as the port is active.

Interface

This parameter defines the 3746 interface type for the PPP port as:

V.24
This defines the interchange circuits between 3745 data terminal equipment (DTE) and data circuit terminating equipment (DCE).

V.35
This defines the list of definitions for interchange circuits between 3746 data terminal equipment (DTE) and data circuit-terminating equipment (DCE) at various data rates.

X.21
This defines the interface between 3746 data terminal equipment (DTE) and public data networks, for digital leased (switched lines not supported by the 3746) synchronous services.

Clocking

Clocking is a process that uses clock pulses to synchronize the transmission of data and control characters during synchronous communication. Clocking is also used for controlling the number of data bits transmitted across a telecommunications link during a specified period. The clocking types are:

Internal clocking
Provided by a scanner in the communications controller. This type of clocking is also known as *business machine clocking*.

Direct clocking
A variation of the internal clocking provided by communication controllers. Clocking is provided by the CLP, which propagates it to the data terminating equipment (DTE).

External clocking
Provided by a modem (DCE) or other attached device.

Data Rate

This parameter defines the rate at which the link connected to the modem operates as either *high* or *low*.

Speed

This parameter defines the speed of the link in kilobits per second. Valid values are 9.6, 19.2, 38.4, 55.855, 64, 256, 512, 1024, 1536, 2048.

Transmit NRZI

This parameter defines operation of the data terminating equipment at each end of the link as NRZI (Non Return to Zero Inverted). NRZI is a data encoding method in which ones are represented by a change in the signal level or polarity and zeros are represented by no change.

It is sometimes called *marking* because only the ones (the mark signals) are explicitly encoded. NRZI is also sometimes known as *non-return-to-zero change-on-ones* (NRZ-1) or *non-return-to-reference* encoding.

Maximum Transmission Unit

This defines the number of bytes per frame of data that can be transmitted or received over the interface (the size of the largest packet that can be received).

Note: The value specified for MTU for end stations must be the same as the value specified for this parameter.

Retry Timer

This is the amount of time that elapses before transmission of the request to open the link and the request to close the link is timed out. Expiration of this timer causes a timeout and the halting of configure-request and terminate-request packet transmission. The timer is expressed in milliseconds.

Config Tries

This defines the number of times that the 3746 LCP or PPPNCP sends configure-request packets to a peer station to establish the opening of a PPP link. The sending of configure-request packets indicates the need to open an LCP connection with a specified set of configuration options.

The retry timer starts after a configure-request packet is transmitted. This is done to guard against packet loss.

NAK Tries

This defines the number of times that the 3746 LCP or PPPNCP sends configure-nak packets (“nak” means not acknowledged) to a peer station while attempting to open a PPP link.

Upon receiving configure-request packets with some unacceptable configuration options, LCP and PPPNCP send configure-nak packets. These packets are sent to refuse the offered configuration options and to suggest modified, acceptable values.

Terminate Tries

This parameter specifies the number of times that the 3746 LCP or PPPNCP sends terminate-request packets to a peer station to close a PPP link.

The retry timer is started after a terminate-request packet is transmitted. This is done to guard against packet loss.

Down Time

Down time specifies the time allowed to elapse before the 3746 IP router declares an interface to be down (inactive).

The normal maintenance packet interval is three seconds, and four maintenance failures can occur before the interface is declared inactive.

Down time is used primarily when tunneling LLC traffic over an IP network using OSPF. If an interface goes down, OSPF cannot detect it fast enough because of the time taken for an interface to be declared down. In this situation, LLC sessions begin to time out.

If you specify a lower value for down time, OSPF can quickly detect that an interface is down. This allows an alternative route to be quickly established and prevents the LLC sessions from timing out.

The same value for down time must be specified at each end of a serial link, otherwise the link may not come up.

High down time values cause the router to ignore transient connection problems, *low* down time values cause the router to react more quickly.

Enable Timer (ENABLTO)

This parameter defines the timeout used for detecting the failure of the data-set-ready signal line of a modem (or equivalent) that is turned on when the telecommunications link connected to the modem is enabled (for nonswitched links), or when a dialing operation is complete (for switched links).

Magic Number

This function provides a way of detecting looped-back links in serial line configurations.

With this parameter enabled, the link uses the system clock as a random number generator. The numbers generated are known as *magic numbers*.

If this parameter is enabled, when LCP (Link Control Protocol) receives a *configure request*, the received magic number is compared with the magic number in the last configure request that was sent to the peer.

If the magic numbers are *different*, the link is not considered to be looped back. If the magic numbers are *identical*, the PPP handler attempts to bring the link down and up again, in order to renegotiate the magic numbers.

IP Compression

The 3746 PPP supports Van Jacobson Compressed TCP/IP for header compression. You should enable this if the PPP link is running with a low baud rate.

If you are configuring PPP on a high-speed link, leave this parameter disabled. This is because the time needed for compression and decompression is greater than the time saved *with* this parameter enabled.

Number of Slots

This parameter defines the number of IP headers that are saved for reference, when determining the type of IP compression that is enabled on the PPP link.

This Internet Protocol Control Protocol (IPCP) option is active across the entire link. IPCP configures, enables, and disables the IP protocol at both ends of the PPP link.

Send IP Address

The send IP address function determines whether or not the local IP address is sent to the remote end of the link. Enable this parameter if the other end of the link requires the IP address. PPP then sends the local IP address if the other end of the link requests it.

Request IP Address

The request IP address function determines whether or not the IP address is requested from the remote end of the link.

Performance Tuning

File Transfer with IP over PPP

Using CCM, specify a maximum IP transmission unit (MTU) greater than the default value of 2052, according to the transmission quality of the lines.

Chapter 9. X.25 Overview

X.25 is a packet switching technology allowing data transmission at low, medium or high speed (up to 2 Mbps). X.25 guarantees data transmission without error or loss, and packets arrive at their destination in the order that they were sent.

The X.25 ITU-T recommendation specifies the three following OSI layers:

- Layer 1 : Physical layer
- Layer 2 : Data link layer, called Link Access Procedure Balanced (LAP-B)
- Layer 3 : Network layer, called Packet Layer Protocol (PLP).

For X.25, the physical layer is the same as for any other WAN protocol implemented in the CLP (frame relay, SDLC or PPP). The X.25 support implements the X.25 layer 2 (LAP-B) and 3 (PLP). It also implements another layer above PLP for adapting SNA, APPN/HPR, and IP protocols to the X.25 layers.

Figure 9-1 shows how the X.25 functions are implemented in the 3746-9x0. The CLP can be under NCP or NNP control.

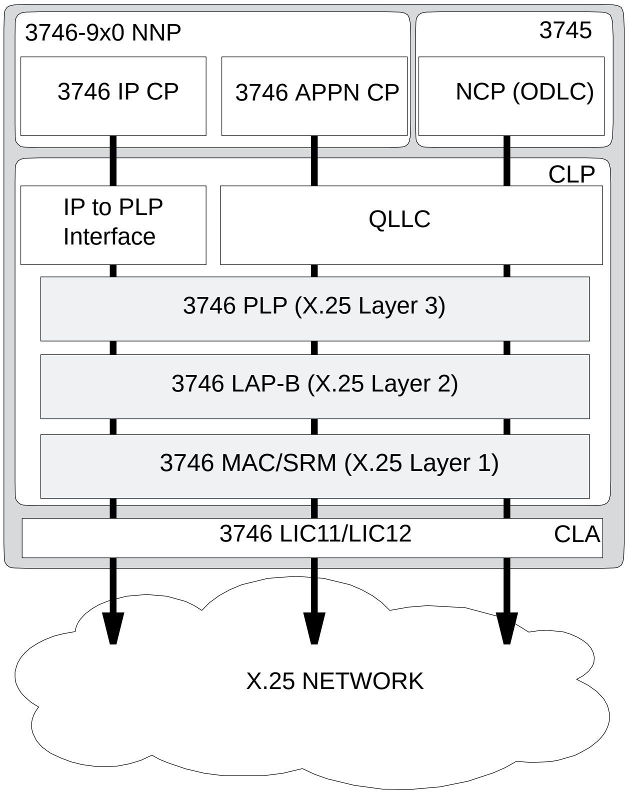

Figure 9-1. 3746-9x0 Native X.25 Implementation

X.25 Supports

NPSI Support (NPSI ODLC)

The Communication Line Processors (CLP) support ITU-T X.25 protocol in conjunction with the X.25 NCP Packet Switching Interface (NPSI) program running with ACF/NCP in the 3745. This allows the 3746-900 to carry all traffic flows supported by NPSI, SNA, and non-SNA, over connections to an X.25 private or public network. The CLP performs only the physical layer (layer 1), the other layers (X.25 and above) are performed by NPSI.

The rest of this chapter, except when otherwise indicated, does not apply to the NPSI support in the 3746-900 (NPSI ODLC). For further information refer to the NPSI documentation, in particular *X.25 Planning and Installation*, SC30-3470.

Native X.25 Support (FC 5030)

The X.25 support feature (feature code 5030) allows the CLP to perform the X.25 layers, the QLLC layer and the IP-to-PLP interface in addition to the physical layer (refer to Figure 9-1 on page 9-1). The following traffic is supported:

- **NCP traffic (X.25 ODLC)**

  Starting from Version 7 Release 4, NCP is able to flow SNA traffic over X.25 without requiring NPSI:

  – Subarea node traffic

  – Peripheral node traffic

  – Low entry network (LEN) traffic

  – APPN traffic, when NCP acts together with VTAM as a composite network node (CNN)

  In summary, any SNA traffic that NCP is capable of handling can flow over an X.25 line attached to the 3746-9x0 Communication Controller.

  The interface layer between X.25 and NCP implements the QLLC protocol

  Note: The short hold mode is not supported over a subarea SVC.

- **APPN/HPR traffic**

  The APPN/HPR control point, residing in the network node processor (NNP), is able to flow APPN/HPR traffic over an X.25 line without requiring NCP:

  – SNA subarea node traffic (DLUR traffic)

  – HPR traffic

  – APPN traffic

  In summary, any traffic that the APPN/HPR control point is capable of handling can flow over an X.25 line attached to the 3746-9x0 Communication Controller.

  The interface layer between X.25 and the APPN/HPR control point implements the QLLC protocol.

- **TCP/IP traffic**

  The IP control point, residing in the network node processor (NNP), is able to flow any TCP/IP traffic over an X.25 line.

The following TCP/IP routing protocols are supported over X.25:

- RIP
- OSPF

  The neighbors around the X.25 network have to be configured manually through CCM.
- BGP

The SNMP MIBs for X.25 are also implemented. Refer to “SNMP” on page 9-11.

The interface layer between X.25 and the IP control point conforms to the RFC 1356.

Functions Supported

The 3746-900 and 950 can be either a DTE:

- Attached to a packet switched data network (PSDN) that conforms to the ITU-T Recommendation X.25 (1993).
- A DTE directly attached to another DTE, that conforms to the ISO 7776 (layer 2) and ISO 8208 (layer 3) standards.

The functions supported are:

Layer 1 The physical layer can be either leased X.21, V.24 or V.35.

The maximum speed for an X.25 line is 2 Mbps.

Layer 2 (LAP-B) The (data) link layer conforms to X.25 LAP-B single link procedure (SLP).

- The modulo can be either 8 or 128.
- Piggy-backing technique is used at the frame level (that is, use of an information frame to acknowledge another information frame flowing in the reverse direction) with use of the timer T2. This technique decreases the number of RR frames flowing over the line and, therefore, improves the performance.

Layer 3 (PLP) The packet layer features the following characteristics:

- Switched Virtual Circuit (SVC) and Permanent Virtual Circuit (PVC) are supported.
- The X.25 support of the 3746-9x0 complies with the X.25 versions 1980, 1984, 1988 and 1993.
- X.25 addresses with the TOA/NPI format are supported.
- The DTE/DCE role is negotiated at restart time in case of a direct DTE to DTE attachment, in conformance with the ISO 8208 standard.
- The modulo at packet level is either 8 or 128.
- Piggy-backing technique is used at the packet level (that is, use of a data packet to acknowledge another data packet flowing in the reverse direction) whenever possible. This technique decreases the number of RR

packets flowing over the line and, therefore, improves the performance.

- Data packet segmentation and re-assembly is performed when the PIU size (SNA) or MTU size (IP) is greater than the packet size.
- Optional user facilities can be included in a call request packet (specified by X.25 or not).
- Up to 16 bytes long user data can be included in the Call User Data (CUD) field of a call request packet.
- A data packet with the delivery confirmation bit (D-bit) is never sent. When such a packet is received, the D-bit is ignored.
- Interrupt packet is not supported.

Above Layer 3 (PLP)

- For SNA, the QLLC layer implements the loading and activation of a remote NCP over an X.25 subarea PVC or SVC controlled by NCP.
- Performance monitoring (for LAPB and PLP) and accounting through the Network Performance Monitor (NPM).

Performance Enhancements

- For X.25 SNA, the data throughput of the 3746-900 is multiplied by a factor of up to 10 compared to NPSI [1], allowing a 3746 to support up to 1000 packets per second (128 bytes/packet) on each CLP.
- For X.25 SNA traffic, the processor (CCU) load of the 3745 attached to the 3746-900 is reduced by up to 90% compared to NPSI.
- For IP traffic a 3746 is able to support up to 1000 packets per second on each CLP.
- X.25 lines can be used very efficiently, close to 100% utilization, at every speed, up to 2.048 Mbps.

To get precise X.25 performance figures, use the 3745/6 configurator (CF3745).

CLP Lines

Any port of a LIC11 or LIC12 can be defined as an X.25 line. An X.25 line can coexist with any line supporting other protocols on the same CLP (frame relay, SDLC, PPP, ISDN). It can also coexist with NPSI ODLC lines (3746-900 only).

Refer to “Communication Line Processor Connectivity” on page 18-2 for information about the total of resources supported per CLP.

[1] The improvement factor varies, depending on the network environment and traffic characteristics (message size, packet size, etc)

X.25 Line Sharing

Figure 9-2 shows an example of how X.25 is implemented in the 3746-900, and it also shows how lines can be shared.

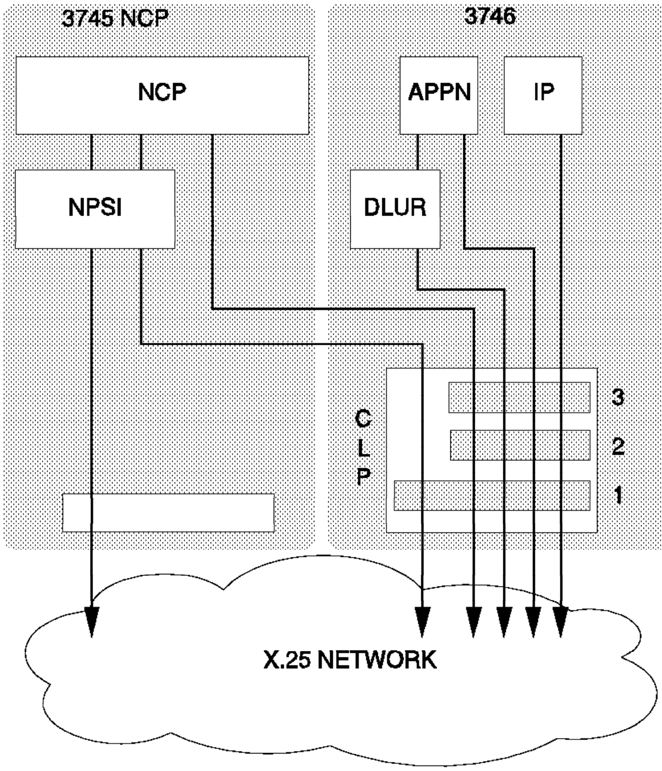

```
Legend:
1. Physical Layer
2. DLC Layer (LAP-B + PLP)
3. QLLC (LLC3), or "IP to PLP" Interface

Note:
NPSI traffic cannot share the same physical
line with NCP (X.25 ODLC), APPN and IP traffic.
```

Figure 9-2. 3746-900 X.25 Line Sharing

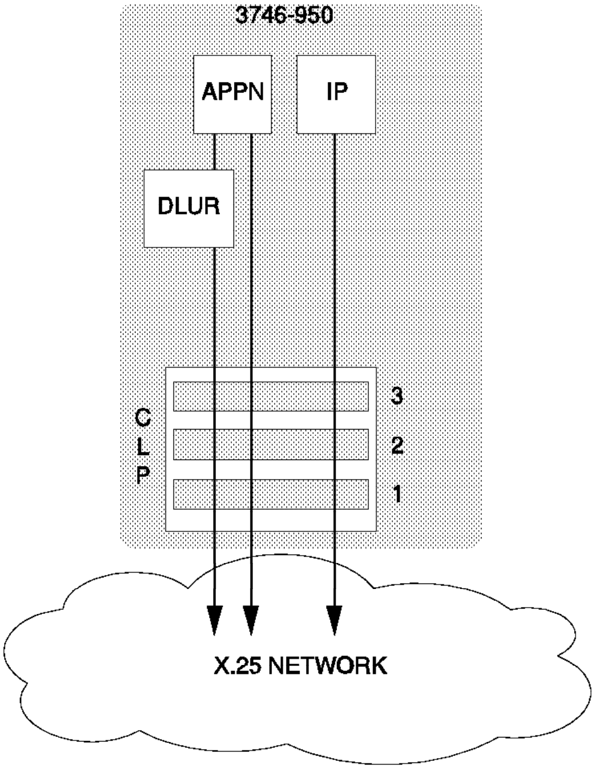

Figure 9-3. 3746-950 X.25 Line Sharing

The following traffic types can flow simultaneously over an X.25 line:

- Subarea SNA
- APPN/HPR
  - SNA (DLUR)
  - HPR
  - APPN.
- TCP/IP.

This means that either or both 3746 control points (APPN/HPR and IP), and *one* NCP (3746-900), can activate or de-activate an X.25 line.

Notes:

1. In the case of 3745-41A or -61A (dual-CCU), only one NCP can control an X.25 line at a time.
2. In a 3746-900, the X.25 line can additionally carry NCP-controlled traffic (X.25 ODLC), but not NPSI traffic (NPSI ODLC).
3. When the line is shared by NCP in addition to APPN/HPR and/or IP, the line must be configured in CCM, for both APPN/HPR and IP, and in NDF for NCP. The configuration parameters in CCM and NDF must be identical (see also “Resource Activation” on page 9-7).

How to Deliver Incoming Calls

Since an X.25 line can be shared by several control points (network node, IP and NCP), an incoming call should be delivered to the right control point. This is performed by means of two pieces of information received in the incoming call packet. These are:

- The first byte of the Call User Data (CUD) field, called the *protocol identifier*
- The last two digits of the called DTE address, called the *subaddress.*

  Note: If the called DTE address contains only one digit, the subaddress is considered as 0 followed by the received digit.

The following rules apply to deliver an incoming call:

- If the protocol identifier is X'CC', the call is delivered to the IP control point, in accordance with the RFC 1356.
- If there is no protocol identifier (no CUD field was received in the incoming call) the call is delivered to the IP control point, to support the old IP routers (before RFC 1356).
- If the protocol identifier is X'E3' or X'EB' (subarea node traffic), the call is delivered to the active NCP, whatever the subaddress value is.
- If the subaddress received in the incoming call exactly matches with one coded at port level in a control point (APPN/HPR or NCP), the call is delivered to that control point.
- The call is delivered to the active NCP, when there is no called DTE address in the incoming call or if the received subaddress does not match with the APPN one, on condition that no subaddress is defined in NCP (NPADTEAD not coded). This allows backward compatibility.

When the call cannot be routed to a control point, it is rejected and an alert is sent.

Resource Activation

This section applies only to a 3746-900.

When the line is shared by NCP in addition to APPN and/or IP, two sets of configuration parameters are needed. At activation time, each is checked. All the parameters must be the same in the two sets. If not, the values provided in the older activation are used. Those provided in the later activation are ignored, but the activation itself is not rejected. There are two exceptions:

NPA eligible Can be the same or different. Each control point can ask for NPM independently.

Local DTE address Must be different. The last two digits are used as a subaddress to route the incoming calls and therefore must be different. If the subaddresses are equal, the later activation is rejected.

Subaddress Allocation

Since subaddress can be used to deliver an incoming call, a subaddress must be allocated to each control point when the line is being shared by NCP and APPN/HPR (this does not apply to IP).

For NCP, the subaddress is coded as the last two digits of the NPADTEAD operand of the X25.MCH statement. If this operand is coded with only one digit, it is considered as equal to 0 followed by that digit.

For APPN/HPR, the subaddress is coded as the last two digits of the DTE address at the port level. It is recommended to allocate the APPN/HPR subaddress as follows:

- The subaddresses, allocated to NCP and to APPN/HPR, must be different. If not, the second activation of the X.25 line is rejected.

  For a 3746-900, allocate a value different from the last two digits of the NPADTEAD operand. This ensures backward compatibility with NCP, when the NPADTEAD operand is coded. This assumes that NPADTEAD contains the actual DTE address as known by the network.

- The local DTE address must be at least two digits long.

The IP control point cannot be allocated a subaddress. Subaddress is never used to route a call to the IP control point (the protocol identifier is used).

The subaddresses, allocated to NCP and to APPN, must be different. If not, the later activation of the X.25 line is rejected.

Note: The X.25 networks generally provide subaddress in two different ways:

- The subaddress is included within the address, for example the last two digits.
- The subaddress consists of all the digits in addition to the address. It is then referred to as the complementary address.

Ask your network provider about the subaddress.

Implementation Details

Protocol Identifier

The protocol identifier is the first byte of the Call User Data (CUD) field. The 3746 Network Node includes the following values in the call request packets that it sends:

- X'C3' for an APPN/HPR or SNA peripheral node traffic, if the X.25 version is 1980.
- X'CB' for an APPN/HPR or SNA peripheral node traffic, if the X.25 version is not 1980.
- X'E3' for a SNA subarea node traffic (NCP), if the X.25 version is 1980.
- X'EB' for a SNA subarea node traffic (NCP), if the X.25 version is not 1980.
- X'CC' for an IP traffic whatever the X.25 version is.

These values can be overridden by CCM, when defining the switched X.25 station, or by the VTAM switched major node (DLCADDR operand) in case of DLUR and X.25 ODLC.

The 3746 can receive without differentiation:

- X'C3' or X'CB' for an APPN/HPR or SNA peripheral node traffic. The 3746 behavior is the same, whether it receives either a X'C3' or X'CB'.
- X'E3' or X'EB' for a SNA subarea node (NCP) or DLUR traffic. The 3746 behavior is the same, whether it receives either a X'E3' or X'EB'.

TOA/NPI Address Format

The 1993 version of the X.25 recommendation has introduced a new format for the X.25 addresses. It is referred to as the TOA/NPI format. This format is needed particularly in case of interworking with an ISDN network. With this format, the X.25 addresses can contain more than 15 digits, unlike previous versions. The address digits are preceded by two fields:

- Type Of Address (TOA) field

  It is a one-decimal digit field. Possible values are:

  0 Network dependent number
  1 International number
  2 National number
  5 Alternative address (can be used only to place an outgoing call)

- Numbering Plan Identification (NPI) field

  It is a one-decimal digit field. Possible values are:

  1 Address as defined in the ITU-T E.164 recommendation (ISDN and telephony numbering plan)
  3 Address as defined in the ITU-T X.121 recommendation (for public X.25 networks)

Example: TOA/NPI fields = 23 means an X.121 national number.

X.25 Security Considerations

APPN/HPR and IP Control Point

Applies for each virtual circuit controlled by the APPN/HPR or IP control points.

You can define the X.25 port to accept any incoming call. In this case there is no verification of the incoming calls. In addition, the corresponding station is created dynamically at the call time, if there is no station configured by CCM and having the same remote DTE address.

In the other case, each incoming call is checked. The calling DTE address must correspond exactly to one of the remote DTE addresses coded at the station level in CCM (including the TOA and NPI fields, if any). If it does not, the call is rejected. In addition each station is created when the X.25 port is activated to ensure that the incoming call is not rejected due to a lack of CLP storage needed to create the corresponding station.

In case of DLUR traffic, the VTAM security mechanism (the VERIFY and VERID operands; refer to “NCP”) is not used, but the APPN/HPR mechanism is used instead. The calling DTE address is checked against the remote DTE address of the station is pre-defined in the 3746 via CCM if the X.25 port is configured to not accept incoming calls.

NCP

Applies for each virtual circuit controlled by NCP.

If you do not need to check incoming calls, code VERIFY=NONE or do not code the VERIFY operand in the VTAM switched major node. All incoming calls will then be accepted.

If the VERIFY operand is coded with IN (or INOUT), VTAM checks that the calling DTE address received in an incoming call corresponds to one coded in the VERID operand. The VERID operand must be coded in decimal exactly as the expected calling DTE address. It must then include the TOA (Type Of Address) and NPI (Numbering Plan Identification) fields, if any.

Checking through VERIFY/VERID operands applies to peripheral node traffic, as well as subarea node traffic.

Network Management

Fault Management

Network Management Vector Transport (NMVT) alerts are sent to NetView/390 or NetView for AIX®, depending on which control point (NCP, 3746 APPN/HPR, or 3746 IP) activated the X.25 line. If activated by:

- NCP, the alerts are sent to NetView/390 via the NCP.
- 3746 APPN/HPR control point, the alerts are sent to NetView/390 via the APPN control point.
- 3746 IP control point, the alerts are sent to NetView for AIX via the IP control point. In that case the NMVT alerts are enveloped within an SNMP trap.

If the line is shared between multiple control points, each activating control point sends its own alerts to NetView or NetView for AIX. Therefore, the same problem may be reported by up to three alerts.

Accounting and Performance Monitoring through NPM

A 3746 Nways Multiprotocol Controller can transfer accounting and performance monitoring data to the Network Performance Monitor (NPM) for:

- Data Link level (LAP-B) performance monitoring
- Packet level (PLP) performance monitoring
- Virtual circuit level accounting

When and only when the X.25 line is activated by the 3746 APPN/HPR control point or NCP, the performance monitoring counters count the overall/traffic, including the IP traffic. When the X.25 line is activated by the 3746 IP control point only, there is no performance monitoring through the NPM.

Performance counters for IP are also reported to NetView for AIX. Refer to "SNMP."

There is no virtual circuit accounting for IP traffic, even if the X.25 line has been activated by the APPN/NCP control point or NCP.

SNMP

Configuration parameters and performance counters can be displayed through SNMP, using the MIB Version II (described in RFC 1213). The 3746 IP router implements the SNMP MIBs described in the following RFC:

- RFC 1381 (MIB for LAP-B)
- RFC 1382 (MIB for PLP)

Each MIB element can be only read.

X.25 Port Subscription to an X.25 Network

Each X.25 port must be subscribed to in the public or private X.25 network.

Table 9-1 gives the X.25 parameters that a 3746-9x0 X.25 port can subscribe to. Although the 3746-9x0 itself does not need any X.25 network services, each of its X.25 ports must have a 'contract' with an X.25 private or public network that guarantees the value of certain X.25 parameters. This section can help you in subscribing an X.25 port of the 3746-9x0 communication controller.

Table 9-1. X.25 LAPB subscription Parameters

Subscription Parameter	Subscription	Note	Section in X.25 (see note 1)
Extended sequence numbering (modulo 128)	Yes	2	2.1.4
Retransmission timer (X.25 parameter T1)	Yes	3	2.4.8.1
Acknowledgement timer (X.25 parameter T2)	Yes	3	2.4.8.2
Max number of retransmission (parameter N2)	Yes	3	2.4.8.4
Maximum frame size (X.25 parameter N1)	Yes	4	2.4.8.5
Frame window size (X.25 parameter K)	Yes	2	2.4.8.6
MultiLink Procedure (MLP)	No		2.5

Legend

No Do not subscribe to this parameter.

Yes You can subscribe to this parameter. This impacts the 3746-9X0 operation and therefore requires generally a parameter in the NCP or CCM configuration.

Notes :

1 Version 1993 of the X.25 recommendation.

2 For this parameter the DTE and DCE values must be exactly the same. You must therefore configure the X.25 port with the subscribed value.

3 The parameters T1, T2 and N2 need not to be the same for the DTE and DCE. However the network values must be known in order to ensure they are compatible with the DTE values or to calculate some DTE values. Especially the network T1 value is needed to calculate the DTE T2 value.

4 The parameter N1 need not to be the same for the DTE and DCE. However it is recommended to configure and subscribe the same value. Besides don't subscribe a value less than 135 bytes (including the address, control and FCS fields) or 131 bytes (excluding these fields).

Table 9-2. X.25 PLP subscription Parameters

Subscription Parameter	Subscription	Note	Section in X.25 (see note 1)
Logical channel ranges for SVCs	Yes	2	Annex A
Permanent virtual circuits (PVC)	Yes	3	Annex A
On-line registration facility	No		6.1
extended packet sequence numbering (modulo 128 at packet level)	Yes		6.2
D-bit modification	No		6.3
Packet retransmission (reject packet)	No		6.4
Incoming calls barred	Any	4	6.5
Outgoing calls barred	Any	4	6.6
One-way logical channel outgoing	Any		6.7
One-way logical channel incoming	Any		6.8
Non standard default packet size	Yes	5	6.9
Non standard default window size	Yes	6	6.10
Default throughput class assignment	Any		6.11
Flow control parameter negotiation	Any	7	6.12
Throughput class nego subscription	Any		6.13
CUG (Closed User Group) subscription	Any	8	6.14.1
CUG with IA (Incoming Access) subscription	Any	9	6.14.3
CUG with OA (Outgoing Access) subscription	Any	10	6.14.2
CUG with IA and OA subscription	Any	11	6.14.2-3
CUG with IA subscription w/o preferential	Any	12	6.14.3
CUG with OA subscription w/o preferential	Any	12	6.14.2
CUG IA and OA subscription w/o preferential	Any	12	6.14.2-3
Incoming call barred within CUG	Any	13	6.14.4
Outgoing call barred within CUG	Any	14	6.14.5
Bilateral CUG subscription	No		6.15.1
Bilateral CUG OA subscription	No		6.15.2
Fast select acceptance	No		6.17
Reverse charging acceptance	Any	15	6.19
Local charging prevention	Any	16	6.20
NUI subscription	Any	17	6.21.1
NUI override	Any	17	6.21.2
Charging information related facilities	No		6.22
ROA subscription	Any	18	6.23.1
Hunt group	Any	19	6.24
Call redirection (CRD)	Any	20	6.25.1
ICRD prevention subscription	Any	21	6.25.4
Call deflection subscription	No		6.25.2.1
TOA/NPI address subscription	Yes	22	6.28
Alternative address usage subscription	Yes	23	6.28

Legend

No Do not subscribe to this parameter.

Yes You can subscribe to this parameter. This impacts the 3746-9X0 operation and therefore requires generally a parameter in the NCP or CCM configuration.

Any You can subscribe to this parameter which do not impact the 3746-9X0 operation. There is generally no impact in the NCP or CCM configuration.

Notes :

1 Version 1993 of the X.25 recommendation.

2 If you intend to use SVCs, you must define with the network administrator the logical channel ranges that will be used for SVCs, that is the lowest and highest values for the one-way incoming, two-way and one-way outgoing logical channels. You have to configure exactly the subscribed values.

3 If you intend to use PVCs, you must define the parameters of each PVC with the network administrator, especially the identification of the remote end, logical channel number at each end, packet size and window size which will apply over the PVC, etc.. You have to configure exactly most subscribed values.

4 Rather than to subscribe to this facility, it is preferable to assign with the network administrator all the logical channel numbers for SVCs to the one-way incoming or one-way outgoing channel range.

5 It is recommended to subscribe to this parameter and choose the largest packet size possible. This improves performances.

6 It is recommended to subscribe to this parameter and choose the largest window size possible. This improves performances, because the 3746-9x0 does not systematically acknowledge each packet. It does so when the window minus one is reached.

7 Flow control parameter negotiation allows for negotiation of the packet and/or window sizes. It is recommended to subscribe to this parameter.

8 Closed User Group (CUG) facility allows to form groups of DTEs with restricted access. A DTE can subscribe to one or more CUGs.

9 If the X.25 port subscribes to a CUG with Incoming Access (IA), it can communicate with any DTE within the same CUG and receive calls from any DTE which does not belong to any CUG.

10 If the X.25 port subscribes to a CUG with Outgoing Access (OA), it can communicate with any DTE within the same CUG and place calls to any DTE which does not belong to any CUG.

11 If the X.25 port subscribes to a CUG with Incoming and Outgoing Access (IA and OA), it can communicate with any DTE within the same CUG and receive calls from or place calls to any DTE which does not belong to any CUG.

12 A preferential CUG is a kind of default CUG. For instance, communicating within the preferential CUG does not require the CUG selection facility in the call request packet. CUG subscription without preferential CUG is mainly meant to gateways between two X.25 networks using the X.25 protocol. Therefore if you need the CUG facility, choose to subscribe to this function with a preferential CUG.

13 If the X.25 port subscribes this facility, it will not receive calls from any DTE that belongs to the given CUG.

14 If the X.25 port subscribes this facility, it cannot initiate calls to any DTE that belongs to the given CUG.

15 Subscribe to this parameter if you accept that the X.25 port be charged when requested by the calling DTE through the reverse charging facility in the call request. If this parameter is not subscribed to, the DTE will never receive incoming calls with reverse charging.

16 When the X.25 port subscribes to the Local Charging Prevention, all the calls it initiates will be charged to the called DTE and all calls it receives will be charged to the calling DTE.

17 Subscribe to the NUI (Network User Identification) facility for the following purposes:

- Detailed billing (the network user identifier is included in the bill)
- Security (access through a switched network)
- Network management (override of the subscription parameters by a set of parameters depending on the network user identifier provided).

Associated to the NUI facility subscription, one or more network user identifiers are also agreed to.

If the X.25 port has subscribed to the NUI facility in the network, each call request must include the NUI selection facility, which provides the network user identifier that applies to the call to be established. This can be performed either through CCM or the VTAM switched major node (DLCADDR operand).

The 3746-9x0:

- Does not support this facility in the call accepted packet.
- Overrides all subscribed parameters, except packet size and window size.

18 Subscribe the X.25 port to the ROA (Recognized Operating Agency) facility, if you want that each internetwork call be routed, whenever possible, through a predefined sequence of ROA transit networks.

This sequence can be overriden by the ROA selection facility included in the call request.

Remark that the X.25 version 1993 has changed the name from RPOA (Recognized Private Operating Agency).

19 Hunt group consists in distributing calls on several X.25 lines. It is not supported by the 3756-9x0, but it can be subscribed on the network side. In that case, only the incoming calls are distributed on the X.25 lines. If there is a need to initiate calls over these lines, assign the X.25 stations to the X.25 lines (static assignment), in order to balance as far as possible the outgoing call traffic.

20 Subscribe an X.25 port of the 3746-9x0 to the call redirection, if you want incoming calls be redirected to another port of the same 3746-9x0 or another DTE:

- when the port is out of order (e.g. failed or not activated)
- when the port is busy (no logical channel available)
- systematically (for instance, during an interim period, when migrating an application)

21 Subscribe the X.25 port to the Internetwork Call Redirection or Deflection (ICRD) prevention, if you want to prevent an outgoing call being redirected or deflected, except when the alternate DTE belongs to the same network as either the X.25 port (calling DTE) or the originally called DTE.

This subscription parameter can be overriden by the ICRD status selection facility included in the call request.

22 This facility allows the X.25 port to receive a calling DTE address with the TOA/NPI format in incoming calls. If this facility is not subscribed to, the calling DTE address is generally not included in incoming calls, when it exceeds 15 digits. This facility is mainly meant for ISDN interworking.

Since the 3746-9x0 supports the TOA/NPI format, it is recommanded to subscribe to this facility, especially:

- if the X.25 port needs to communicate with another DTE attached to an ISDN (directly or through a terminal adapter) and working in packet mode
- if the X.25 port itself is attached to an ISDN through a terminal adapter.

23 This facility allows the X.25 port to include a non X.25 address, for instance a mnemonic, OSI (NSAP), MAC or IP address, in the call request packet.

| The 3746-9X0 requires that the alternative address be limited to 8 bytes.

How to Configure IP over an X.25 Network

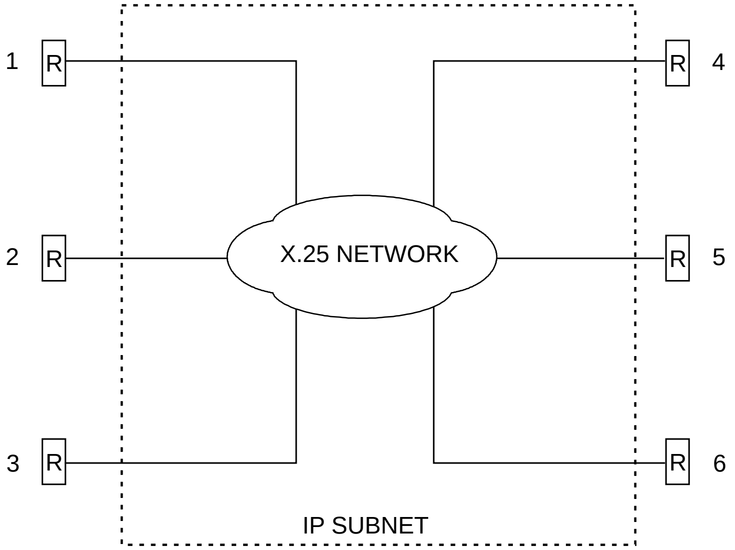

Figure 9-4. IP over X.25 Configuration Sample

The simplest way to configure IP over an X.25 network is to create one IP subnet that includes all the X.25 ports of routers and only the X.25 ports. Refer to Figure 9-4 on page 9-16. So each X.25 port is allocated an IP address within this subnet. Once this is performed, IP stations that can reach all the remote routers should be created for each X.25 port in the network. Therefore if there are N routers around the X.25 network, you have to create N-1 IP stations in each router, hence N(N-1)/2 stations for the whole network.

Each IP station binds in fact an X.25 address with an IP address, the X.25 address being that of the remote port in the X.25 network and the IP address that of the remote port in the IP subnet. The IP stations are configured through CCM. They allow therefore for router interconnection. For instance from router R2, you have to create a station towards R5, which provides router R2 with the X.25 and IP addresses of router R5. If this station does not exist, it is not possible to reach R5 from R2.

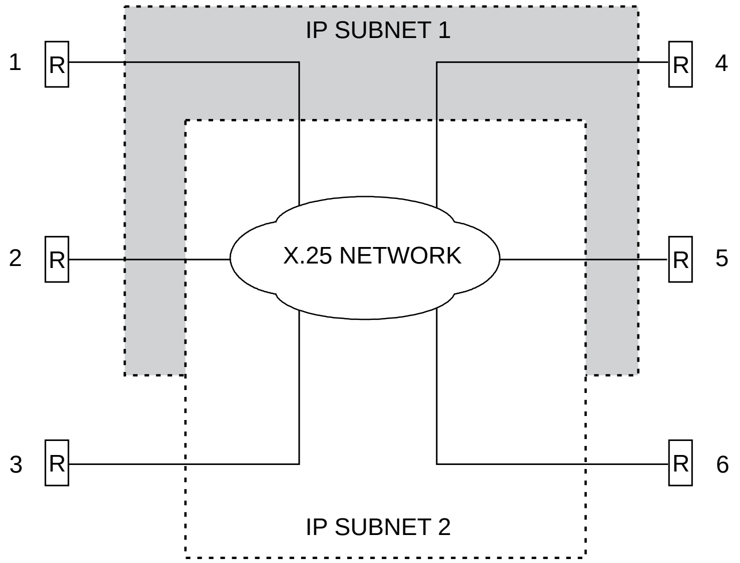

Figure 9-5. Another Configuration Sample

Another way to configure IP over an X.25 network is to create several IP subnets (for instance two, as in the sample of Figure 9-5 on page 9-17) including together all the X.25 ports, each subnet including only X.25 ports. At least one port must belong to the two subnets. Otherwise there would be no connectivity between the routers of each subnet. In the above sample two routers belong to the two subnets (R2 and R5). These two routers must be given two IP addresses, one in each subnet. Each other router must be given one IP address in the subnet it belongs to.

For an X.25 port that belongs to more than one IP subnets, an IP station binds one X.25 address with one or more IP address, the X.25 address being that of the remote port in the X.25 network and the IP address(es) that(those) of the remote port in the IP subnet(s). Example: When you define for the port R2 the IP station towards R5, you have to provide the X.25 address of the router R5 and two IP addresses, those of R5 in each subnet.

When there are more than one IP subnets for one X.25 network, be careful of the following:

- No SVC will be established between two routers that do not belong to the same IP subnet. So the IP traffic from R1 to R6, for instance, will flow from R1 to R2 (or to R5), then forwarded from R2 (or R5) to R6. Each packet flows therefore twice in the X.25 network and is paid twice. In addition the router R2 (or R5) could be overloaded unnecessarily.
- If a PVC already exists in the X.25 network, its two ends must belong to the same IP subnet. Otherwise no IP traffic could flow through this PVC.
- If OSPF is used over the X.25 network, allocate different weights to the two links R2-R5 (one link in each subnet) in order to avoid that two virtual circuits be set up between R2 and R5 with an unnecessary load balancing between them.

Choose therefore to define only one IP subnet for an X.25 network, except for specific reasons, for instance in order to decrease the number of IP stations to configure or when the IP address space prevents all the X.25 routers be included in the same IP subnet. Anyway the traffic between IP subnets should be low enough in order not to overcharge the X.25 traffic.

In summary:

- An X.25 port is allocated one or more IP addresses.
- Each IP address of a given X.25 port must belong to a different IP subnet.
- For each X.25 port, one or more IP stations must be configured.
- Each IP station is configured with the X.25 address of the remote router.
- Each IP station is configured with one or more IP addresses, those of the remote router.
- Each IP address configured for a given station must be unique for all the stations related to the same port. This means that an IP address of an X.25 port can be reach only through one X.25 address.
- Two IP addresses configured for a given station cannot belong to the same IP subnet.
- Each IP address configured for a given station must belong to the IP subnet of one IP address of the port.

Configuration

X.25 APPN and X.25 IP Configuration

Configuration of 3746-9x0 X.25 APPN and X.25 IP resources is done via CCM. For more information refer to *Controller Configuration and Management User's Guide*, SH11-3081.

X.25 ODLC Configuration

Configuration of the 3746-900 X.25 resources is done via the Network Definition Facility, that is, during NPSI and NCP generation.

- A physical line is defined by the X25.MCH definition statement.
- A PVC is defined by the X25.LINE and X25.PU definition statements. A peripheral node is defined by the X25.PU definition statement.
- A SVC is defined by either:
  - X25.LINE and X25.PU definition statements
  - X25.SVC definition statement

  For more information, refer to the *NCP X.25 Planning and Installation* manual, SC30-3470.

For DTE-to-DTE SVCs

For X.25 ODLC SVCs that can be set up by either of the DTEs, the SVC must be defined with `CALL=INOUT`.

X25.OUFT Statement

For X.25 ODLC, each X25.OUFT definition statement defines an entry in the CCU Optional User Facilities Table (OUFT), one per network. The OUFT table of a network defined by the X.25.NET definition statement is shared by NPSI and X.25 ODLC when both are running in the CCU.

Specific X.25 Parameters

NCP Parameters for X.25 ODLC

- **STATION=(a,b)** in the **X25.MCH Definition Statement**

  a is the LAPB role: DTE or DCE
  b is the PLP role: DTE, DCE or NEG (Negotiable)

- **MAXPKTL=(a,b)** in the **X25.VCCPT Definition Statement**
  - For PVCs:
    - a is the packet size out
    - b is the packet size in
  - For SVCs:
    - a is the packet size from the calling to the called DTE
    - b is the packet size from the called to the calling DTE
- **VWINDOW=(a,b)** in the **X25.VCCPT Definition Statement**
  - For PVCs:
    - a is the window size out
    - b is the window size in
  - For SVCs:
    - a is the window size from the calling to the called DTE
    - b is the window size from the called to the calling DTE

NPSI Parameters Not Used in X.25 ODLC Environment: If you use your NPSI definition statements for the NCP generation, the following parameters, if used, are ignored:

- BRKCON
- CAUSE
- CCX DELAY
- DCI
- DM
- INSLOW
- ITRACE (NPSI LAPB trace)
- SPNQLLC

If you use your NPSI definition statements for the NCP generation, the following parameters, if used, are automatically changed as follows:

- MAXPUI, replaced by the standard NCP parameter TRANSFR
- ACTIVTO is forced to 0
- LCN0 is forced to NOT USED
- RESTPVC is forced to YES
- SHM is forced to NO

Note: The inactivity timer can be coded as the operand T3 or T4 whether the LAPB role (operand STATION) is defined as DTE or DCE.

How to Define an Outgoing Call

X.25 ODLC and NPSI ODLC

All the parameters needed to place an outgoing X.25 call, in particular the remote DTE address, are provided in the VTAM DLCADDR operand of the PATH definition statement in the VTAM switched major node. Refer to the *NCP X.25 Planning and Installation Guide* SC30-3470 for the X.25 specifics, to the *VTAM Resource Definition Reference* for a complete syntax of the DLCADDR operand.

X.25 APPN (without DLUR), and X.25 IP

All the parameters needed to place an outgoing X.25 call, in particular the remote DTE address, are provided in the CCM when defining an APPN or IP switched station. refer to the CCM User's Guide for more information.

X.25 DLUR

As for X.25 ODLC, all the parameters needed to place an outgoing X.25 call are provided in the VTAM DLCADDR operand of the PATH definition statement in the VTAM switched major node. But in this case the DLCADDR is coded slightly differently from X.25 ODLC. A X.25 SVC connection for a DLUR resource has these elements:

1. DLC type: X25SVC (required)
2. Port name (required)
3. Protocol Identifier (Optional)
4. Called DTE address (required)
5. Calling DTE address (optional)
6. Optional User Facilities
7. Optional Call User Data (CUD) field

Code a DLCADDR keyword for each element using the following syntax:

```
DLCADDR=(subfield_id,data_type,data_string)
```

where:

`subfield_id`
: Specifies the subfield identification number:

1	DLC type identifier
2	Port name
4	Protocol Identifier
21	Called DTE address
22	Calling DTE address
30	User facilities
61	Call User Data Field, bytes 1 to 15.

`data_type`
: Specifies one of the following data types:

C Specifies EBCDIC characters. The value is sent in EBCDIC to the DLUR.

I Specifies ASCII characters. The value is converted into ASCII before being sent to the DLUR.

X Specifies hexadecimal, two digits per byte. The value coded is sent as is to the DLUR

D Specifies decimal. The value coded is converted in binary before being sent to the DLUR.

BCD Specifies Binary Coded Decimal, two digits per byte. The value coded is sent as is to the DLUR.

`data_string`
The actual data sent to the DLUR.

Use the following procedure to code `DLCADDR`:

Step 1. To specify an X.25 SVC DLC type, code:

```
DLCADDR=(1,C,X25SVC)
```

Step 2. Identify the port name of the X.25 port:

```
DLCADDR=(2,I,portname)
```

where `portname` is the port name defined in CCM. The 3746 network node expects the port name in ASCII. The ASCII type is indicated by I and requires for:

- VTAM V4 R2 that PTF UW28497 is installed
- VTAM V4 R3 that PFT UW28498 is installed

Step 3. If you need to specify a protocol identifier different from `X'CB'`, code:

```
DLCADDR=(4,X,hh)
```

where `hh` is the protocol identifier inserted as the first byte of the call user data field in the sent call request packet. This field is optional and should only be used when the called DTE does not support receiving the value `X'CB'` in the first byte of the call user data field. The only values you can code are `X'CB'` or `X'C3'`. The value `X'CB'` is the default value used when you omit to code this field.

Step 4. Specify the X.25 Called DTE address. For a:

Non-TOA/NPI address, code `DLCADDR=(21,X,000Ldddd)`, where:

- `L` is, in binary, the number of decimal digits of the Called DTE address, the maximum value is F.
- `dddd` is the called DTE address coded in BCD (Binary Coded Decimal: decimal digits coded onto 4 bits and 2 digits in each byte), the maximum number of digits is 15. If the address contains an odd number of digits, the last byte is padded with a zero on the right.

TOA/NPI address, code `DLCADDR=(21,X,80LLtndddd)`, where:

- `LL` is, in binary, the number of decimal digits of the Called DTE address. This number includes the TOA and NPI subfields `tn` described below but does not include this length field itself, the maximum value is `X'12'`.
- `t` is a nibble that contains the type of address (TOA). The four bit value is:

```
B'0000' = Network dependent number
B'0001' = International number
B'0010' = National number
B'0101' = Alternative address, limited to 8 bytes
```

- n contains in four bits the Numbering Plan Identifier (NPI)

```
B'0001' = Address as defined in ITU-T Recommendation E.164
B'0011' = Address as defined in ITU-T Recommendation X.121
```

- dddd is the called DTE address coded in Binary Coded Decimal (BCD: decimal digits coded onto 4 bits and 2 digits in each byte), the maximum number of digits is 17. If the address contains an odd number of digits, the last byte is padded with a zero on the right.

Step 5. Specify the X.25 Calling DTE address. For a:

Non-TOA/NPI address, code DLCADDR=(22,X,000Ldddd) with Ldddd as in Step 4 on page 9-21.

The calling DTE address is optional. It is included by the X.25 network if you do not define it.

TOA/NPI address, code DLCADDR=(22,X,80LLtndddd) with LLtndddd as in Step 4 on page 9-21.

Step 6. Specify the Optional User Facilities (if needed,) code:

```
DLCADDR=(30,X,hhhh)
```

where hhhh are the X.25 user facilities as coded in an X.25 call packet, but the facilities length field must not be coded.

Step 7. Specify the data to insert in the Call User Data field of a sent Call Request packet, code:

```
DLCADDR=(61,X,hh)
```

where hh is the data in hexadecimal that is inserted in the Call User Data field at offset 1, just after the Protocol Identifier specified in the subfield 4. A maximum of 30 hexadecimal digits can be defined; the number of hexadecimal digits must be even. This field can be specified using any of the data_type defined above: C, I, BCD, D, X. When specified in characters (data_type C or I), the length cannot exceed 15 characters.

For example, suppose you need to make an outgoing connection from DLUS to a DTE with a non-TOA/NPI X.121 address = 0492114583, using a X.25 SVC on the port name X25P01 and asks for reverse charging. You have to code in the PATH definition statement:

```
DLCADDR=(1,C,X25SVC),                  X.25 SVC
DLCADDR=(2,I,X25P01),                  Port name
DLCADDR=(21,X,000892114583),           Called DTE address
DLCADDR=(30,X,0101),                   Facility, reverse charging
DLCADDR=(61,C,CUDATA)                  CUD data
```

Note: The protocol identifier need not be coded when the called DTE support QLLC with SNA diagnostic codes indicated by a protocol identifier of X'CB', which is the default value for an outgoing connection.

How to Define an X.25 PVC for DLUR

Each PVC is defined in the CCM as a leased station, which contains the PVC characteristics (LC number, packet and window sizes, QLLC retry count and timer). The advantage of this method is to define all the X.25 parameters in a single place.

In the VTAM switched major node, you have to define a PU statement to provide VTAM with mainly the ID block and ID number (operands IDBLK and IDNUM). No VTAM PATH statement is required.

Each PVC must be activated using CCM (this is not possible using VTAM).

Performance and Tuning

Note: This section applies only to X.25 ODLC and NPSI ODLC.

Operating Modulo 8 and Modulo 128 Lines

Make sure the frame level Window Size parameter is set to 7 for modulo 8 lines, or more for the modulo 128 lines, if possible. Otherwise, you may limit the number of non-acknowledged transmitted frames to less than these lines are capable of handling.

Line Utilization

The X.25 lines attached to a 3746-900 can be used at media speed. This means that:

- The CLP can send frames separated by only a single flag when the data rate coming from the CCU is high.
- The CLP can receive frames at the line speed and is not subject to overrunning the input buffer, as happens for the 3745 LSS and HSS adapters. There are enough input buffers to receive one complete window of frames. When the CCU or CLP cannot handle the rate, normal X.25 flow control at the PLP level can lower the pace of the partner as necessary.

CCU Utilization for X.25 ODLC

The CCU utilization is independent of the speed and the active parameter definitions of the X.25 lines. The CCU is only able to process a certain number of PIUs. The same amount of traffic, whether over four X.25 ODLC lines running at 64K bps or over a single X.25 ODLC line running at 256 Kbps, loads the CCU almost the same amount.

CCU Utilization for X.25 Lines (ODLC and NPSI)

An application builds messages that are called request/response units (RUs). These RUs are then packaged in a physical information unit (PIU). When the PIU is larger than the partner maximum received basic transmission unit (BTU) size, the PIU is segmented by NCP into PIU segments. These PIU segments are then packaged in X.25 packets and finally an X.25 frame that contains one packet is sent. When the PIU segment is larger than the packet size, the PLP layer splits the segment into smaller packets.

To decrease CCU and CLP utilization you may:

1. Use the exception response mode that creates the least number of responses when it is supported by the destination application.
2. Increase the pacing window sizes, or work with no pacing when possible.
3. Use the largest segment size that the partner devices support, for example, 521 or 1033. This is either defined in MAXDATA or received in XID3.
4. Fit an RU into a PIU.
5. Fit a PIU into a packet.
6. Use the largest packet size possible, for example 512 or 1024.
7. Increase packet window size above 2 and code PLPIGGYB=YES in the X25.MCH definition statement.

Table 9-3 shows which items in the above list are the most important for tuning the CCU and CLP hardware running the two different implementations of X.25 lines.

Table 9-3. CCU and CLP vs X.25 ODLC and NPSI Lines

Type of X.25 Line	CCU	CLP
ODLC	1,2,4,5	3,4,6,7
NPSI	1,2,3,4,5,6,7	-

For example, peripheral devices using an RU size of 480 gives an FID2 PIU of 489 bytes. If you use a packet size of 512, the initial data from the application will not be segmented into SNA PIU segments and will not be segmented into X.25 packets. This lack of segmenting reduces the CCU overhead.

Segmentation Example

The following example of SNA segmentation is shown in Figure 9-6 on page 9-25:

- The application RU is 600 bytes.
- BFRS = 240 (NCP buffer size)
- MAXDATA = 521
- The PIU of 609 bytes therefore occupies three NCP buffers.

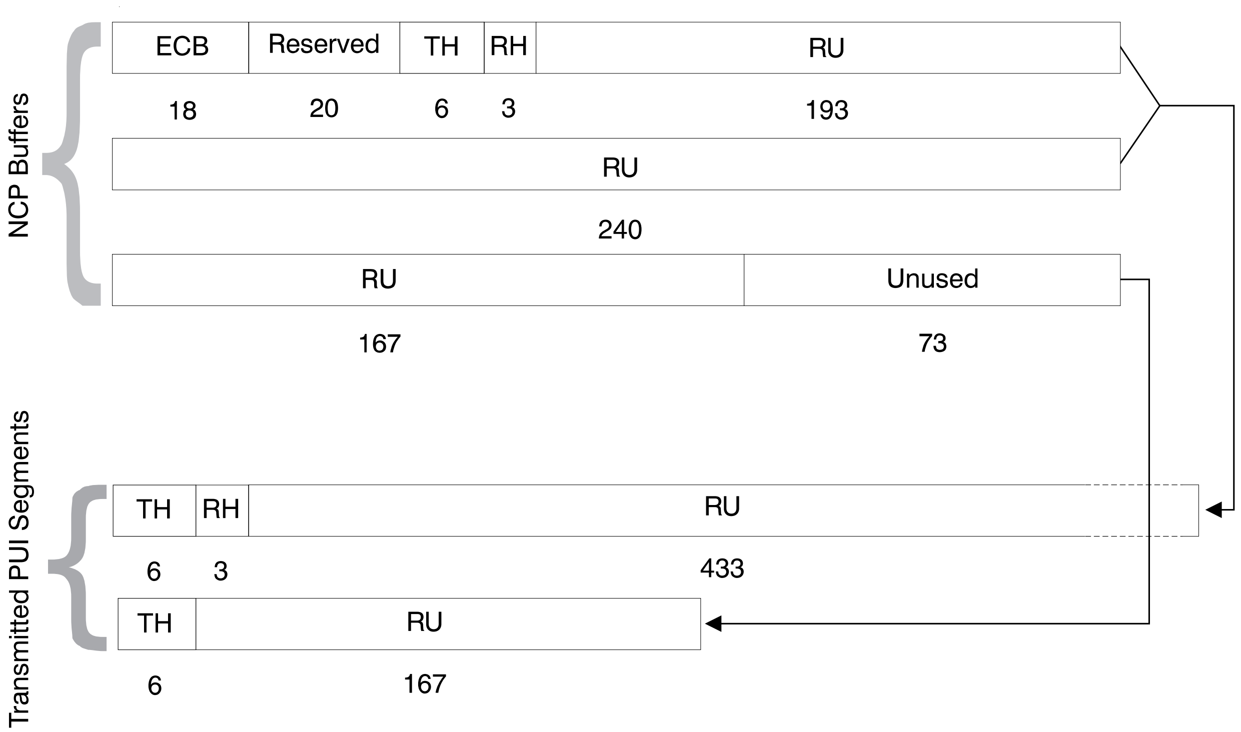

Figure 9-6. NCP Segmentation Example

Legend:

ECB	Event control block
RH	Request/response header
RU	Request/response unit
TH	Transmission header

Before transmission, NCP divides the 609 byte PIU (600 bytes for the RU and 9 bytes for the TH/RH) into two PIU segments. Each segment contains a whole (integer) number of buffers. The RU data is not moved, that is, the NCP buffers are not broken up so that the PIU segments can be filled up to their maximum of `MAXDATA`.

This is done in the following manner:

- The first segment PIU is built with as many buffers as possible while leaving the segment less than or equal to `MAXDATA` bytes long. In this example, the first segment length is TH + RH + RU, which equals 442 bytes:
  RU = 193 + 240 = 433, TH = 6 bytes, and RH = 3 bytes.
- The second segment is built with the remaining RU data + TH, which equals 173 bytes:
  RU = 167, TH = 6 bytes.

  If necessary, the PIU segments are adjusted according to the X.25 packet size. The number of actual packets sent depends on the X.25 packet size:

  – If the X.25 packet size is 128:

    - The first segment is split into four packets: 3 of 128 bytes and 1 of 58 bytes.
    - The second segment is split into two packets: 1 of 128 bytes and 1 of 45 bytes.

    Thus six data packets are sent.

- If the X.25 packet size is 256:
  - The first segment is split into two packets: 1 of 256 bytes and 1 of 186 bytes.
  - The second segment is not split: 1 packet of 173 bytes.

  Thus three data packets are sent.
- If the X.25 packet size is 512 or above, each PIU segment will be sent in a single packet.

  Thus only two data packets are sent.

For the same sized application RU, as the number of X.25 packets used decreases, the amount of CCU and CLP resources (time and memory) used to transmit data decreases.

3746 IP/APPN over X.25 Configuration Example

This section gives an example of how to configure X.25 ports, and how to configure IP to use these ports.

Refer to “3746 Machine A” on page 9-28 and “3746 Machine B” on page 9-35 for information on how to configure the most important parameters for the configuration shown in Figure 9-7.

Assumptions:

- Direct DTE to DTE attachment (no X.25 network).
- The ports are shared between APPN and IP, on both machines.
- Two virtual circuits are defined: one PVC dedicated to APPN and one SVC dedicated to IP.

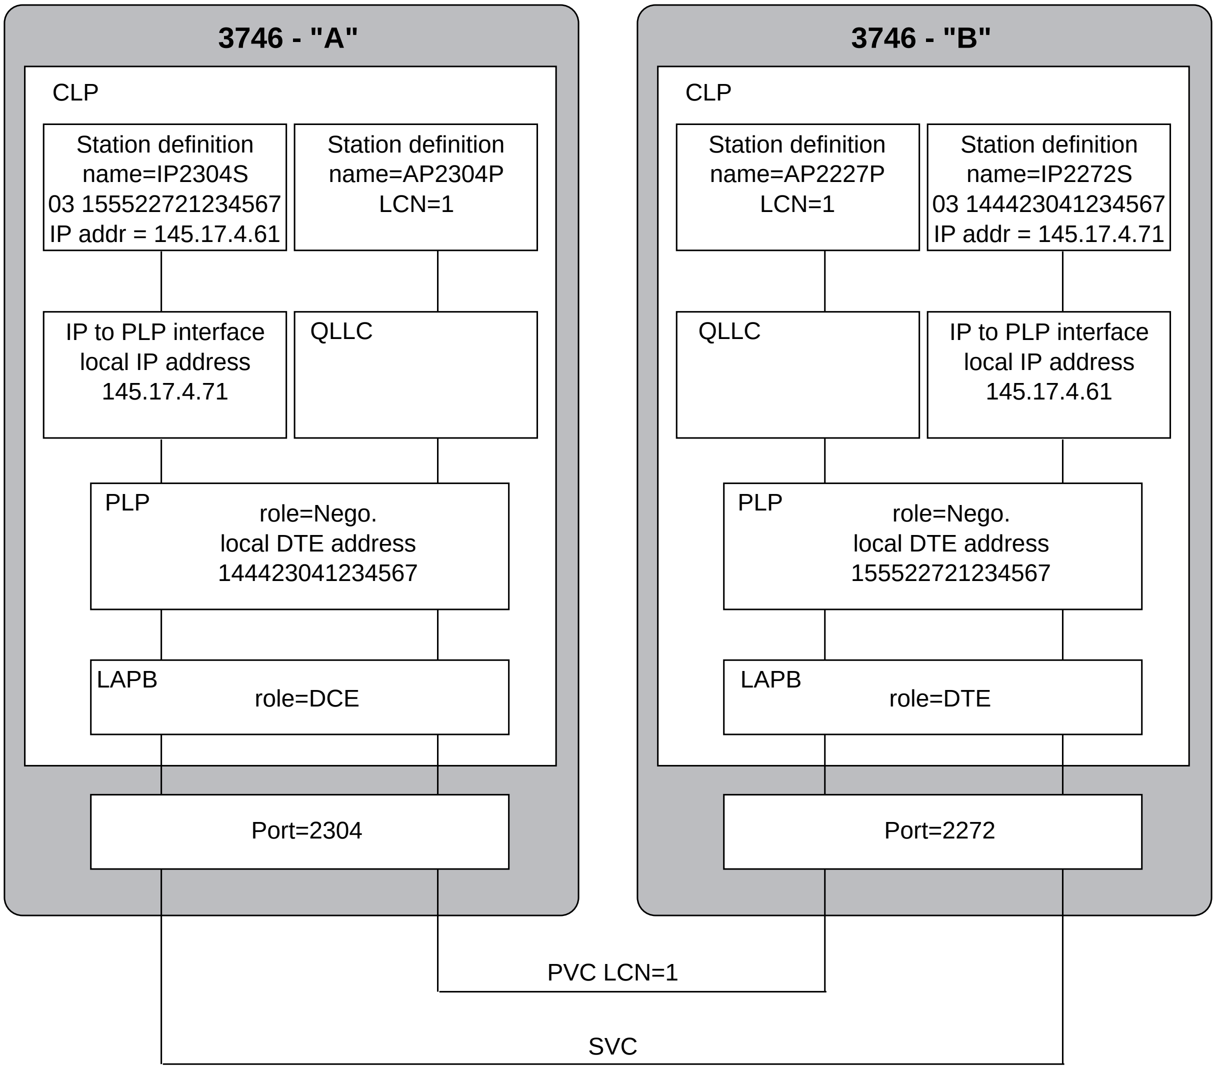

3746 - "A"
CLP
Station definition name=IP2304S 03 155522721234567 IP addr = 145.17.4.61
Station definition name=AP2304P LCN=1
IP to PLP interface local IP address 145.17.4.71
QLLC
PLP role=Nego. local DTE address 144423041234567
LAPB role=DCE
Port=2304
3746 - "B"
CLP
Station definition name=AP2227P LCN=1
Station definition name=IP2272S 03 144423041234567 IP addr = 145.17.4.71
QLLC
IP to PLP interface local IP address 145.17.4.61
PLP role=Nego. local DTE address 155522721234567
LAPB role=DTE
Port=2272
PVC LCN=1
SVC

Figure 9-7. IP/APPN over X.25 Example Network

3746 Machine A

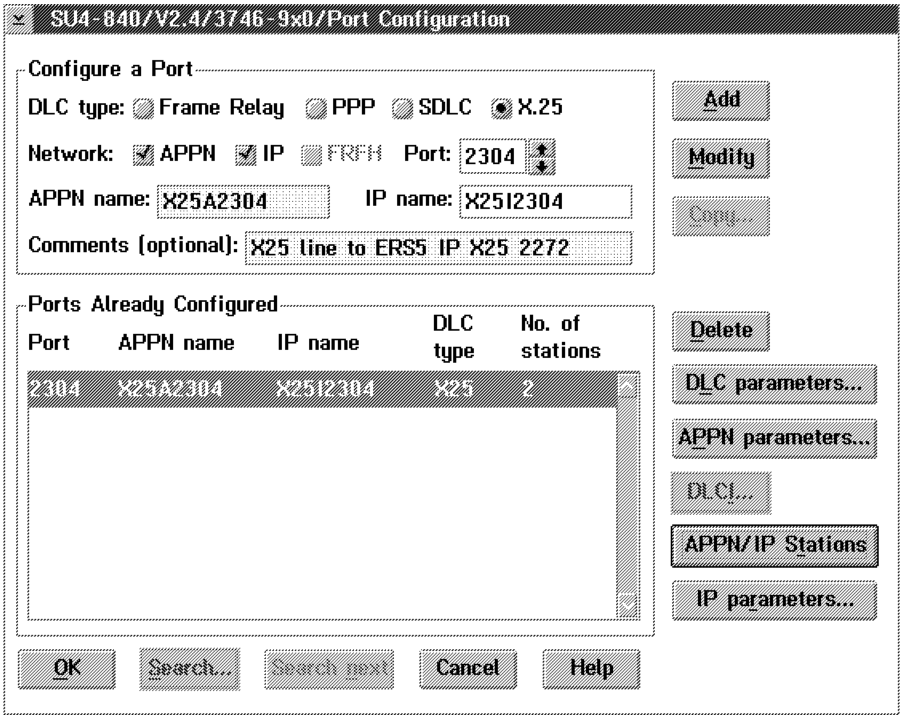

SU4-840/V2.4/3746-9x0/Port Configuration

Configure a Port

DLC type: ○ Frame Relay ○ PPP ○ SDLC ● X.25

Network: ☑ APPN ☑ IP □ FRFH Port: 2304

APPN name: X25A2304 IP name: X25I2304

Comments (optional): X25 line to ERS5 IP X25 2272

Ports Already Configured

Port	APPN name	IP name	DLC type	No. of stations
2304	X25A2304	X25I2304	X25	2

Figure 9-8. Port 2304 Configuration

Step 1. Select the port to be configured, in this case **2304**.

Step 2. Select **X.25** and the other parameters.

Step 3. Select **Add** to register the port.

Step 4. Select **DLC Parameters**.

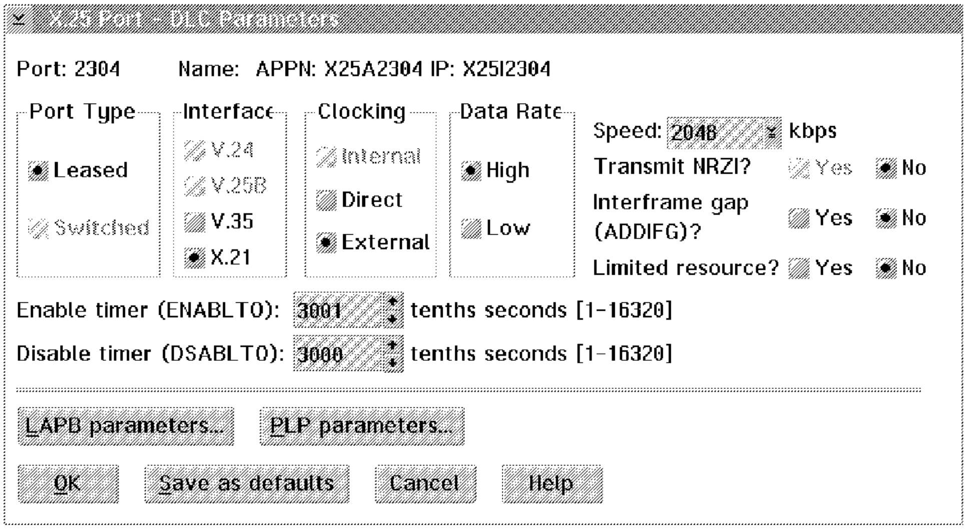

X.25 Port - DLC Parameters

Port: 2304 Name: APPN: X25A2304 IP: X25I2304

Port Type: ● Leased ○ Switched

Interface: ○ V.24 ○ V.25B ○ V.35 ● X.21

Clocking: ○ Internal ○ Direct ● External

Data Rate: ● High ○ Low

Speed: 2048 kbps

Transmit NRZI? ○ Yes ● No

Interframe gap (ADDIFG)? ○ Yes ● No

Limited resource? ○ Yes ● No

Enable timer (ENABLTO): 3001 tenths seconds [1-16320]

Disable timer (DSABLTO): 3000 tenths seconds [1-16320]

Figure 9-9. DLC Parameters for Port 2304

Step 5. Select the right parameters (interface, speed, etc..)

Step 6. Select **LAPB Parameters**.

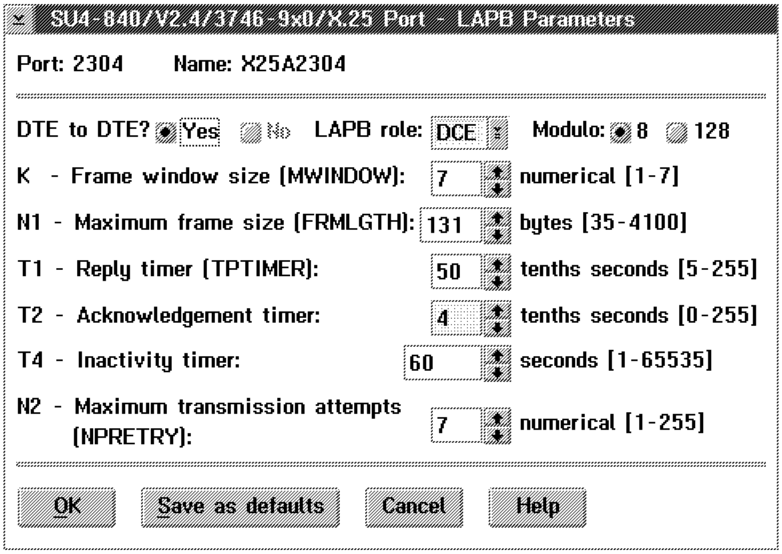

SU4-840/V2.4/3746-9x0/X.25 Port - LAPB Parameters

Port: 2304 Name: X25A2304

DTE to DTE? ● Yes ○ No LAPB role: DCE Modulo: ● 8 ○ 128

K - Frame window size (MWINDOW): 7 numerical [1-7]

N1 - Maximum frame size (FRMLGTH): 131 bytes [35-4100]

T1 - Reply timer (TPTIMER): 50 tenths seconds [5-255]

T2 - Acknowledgement timer: 4 tenths seconds [0-255]

T4 - Inactivity timer: 60 seconds [1-65535]

N2 - Maximum transmission attempts (NPRETRY): 7 numerical [1-255]

Figure 9-10. LAPB Parameters

Step 7. Set DTE to DTE? to **YES**, LAPB Role. to **DCE** and LAPB Modulo to **8**.

Step 8. Configure the other LAPB parameters if needed.

Step 9. Select **OK** to return to the DLC Parameters dialog and save the parameters.

Step 10. From this window, select **PLP Parameters**.

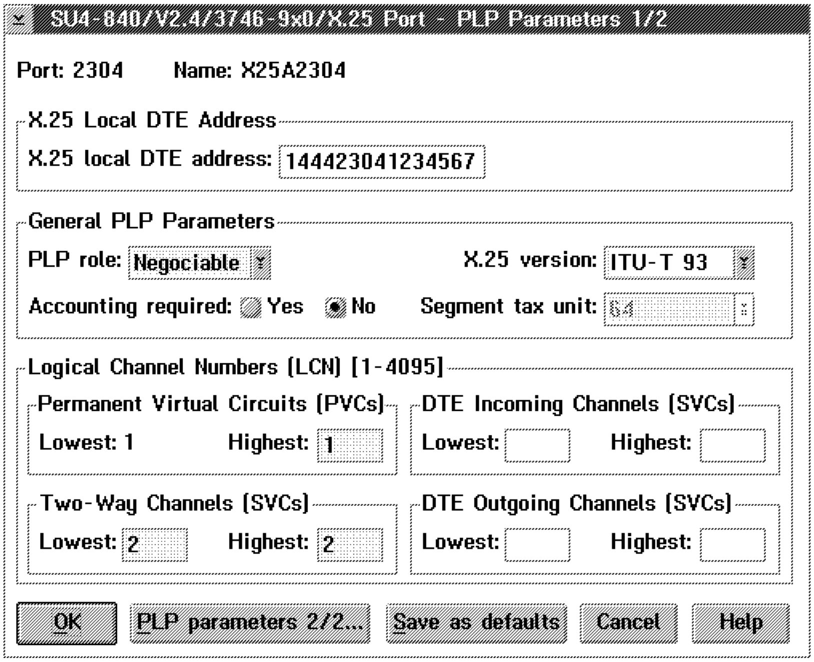

SU4-840/V2.4/3746-9x0/X.25 Port - PLP Parameters 1/2

Port: 2304 Name: X25A2304

X.25 Local DTE Address

X.25 local DTE address: 144423041234567

General PLP Parameters

PLP role: Negociable X.25 version: ITU-T 93

Accounting required: ○ Yes ● No Segment tax unit: 64

Logical Channel Numbers (LCN) [1-4095]

Permanent Virtual Circuits (PVCs)

Lowest: 1 Highest: 1

DTE Incoming Channels (SVCs)

Lowest: Highest:

Two-Way Channels (SVCs)

Lowest: 2 Highest: 2

DTE Outgoing Channels (SVCs)

Lowest: Highest:

Figure 9-11. PLP Parameters 1/2

Step 11. Set Local DTE Address to **144423041234567**.

Step 12. Configure the Logical Channel Numbers.

In the DTE-to-DTE environment, the logical channel numbers must be the same on both ends of the same physical line. When connected to an X.25 network, the logical channel numbers must be those subscribed to in the network.

Step 13. Select **PLP Parameters 2/2**.

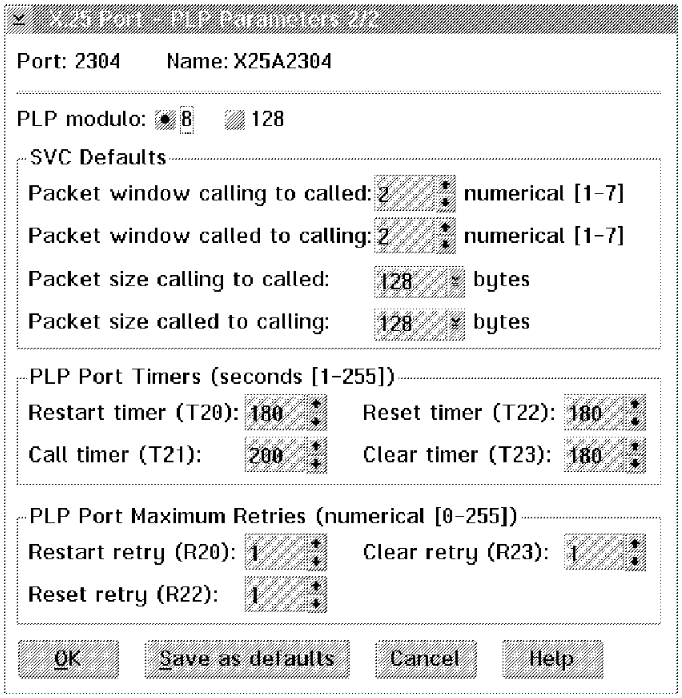

X.25 Port - PLP Parameters 2/2

Port: 2304 Name: X25A2304

PLP modulo: ● 8 ○ 128

SVC Defaults

Packet window calling to called: 2 numerical [1-7]

Packet window called to calling: 2 numerical [1-7]

Packet size calling to called: 128 bytes

Packet size called to calling: 128 bytes

PLP Port Timers (seconds [1-255])

Restart timer (T20): 180 Reset timer (T22): 180

Call timer (T21): 200 Clear timer (T23): 180

PLP Port Maximum Retries (numerical [0-255])

Restart retry (R20): 1 Clear retry (R23): 1

Reset retry (R22): 1

Figure 9-12. PLP Parameters 2/2

Step 14. Set PLP Modulo to **8**.

Step 15. Change the packet sizes and/or window sizes, if needed.

Step 16. Select **OK** three times to return to the Port Configuration dialog and save the parameters in the corresponding windows.

Step 17. From the Port Configuration window, select **IP Parameters**.

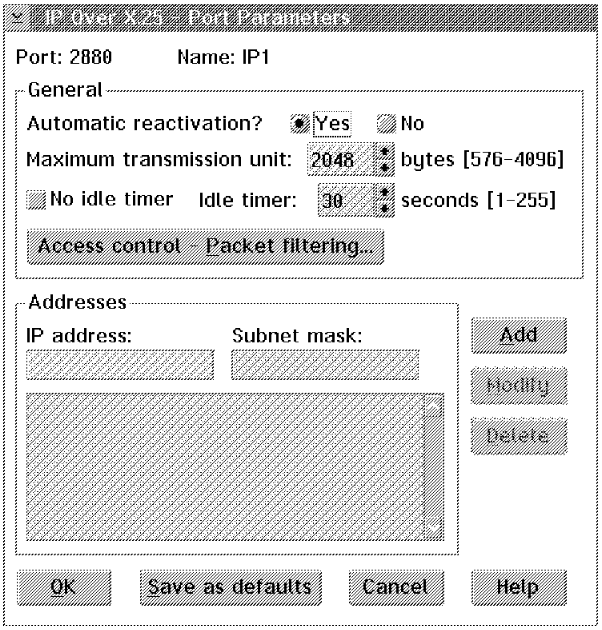

IP Over X.25 - Port Parameters

Port: 2880 Name: IP1

General

Automatic reactivation? ● Yes ○ No

Maximum transmission unit: 2048 bytes [576-4096]

□ No idle timer Idle timer: 30 seconds [1-255]

Addresses

IP address:

Subnet mask:

Figure 9-13. IP Port Parameters

Step 18. Set the MTU size for this port.

Step 19. Configure the IP address of this port to 145.17.4.71

Step 20. Configure the Subnet mask to 255.255.255.0. Port 2272 of machine B must have the same value, because both ports belong to the same IP subnet.

Step 21. Select **Add** to register this address

Step 22. Select **OK** to return to the Port Configuration dialog and save the IP parameters.

Step 23. From this window, select **APPN/IP Stations** to define the IP and APPN link stations in the adjacent node.

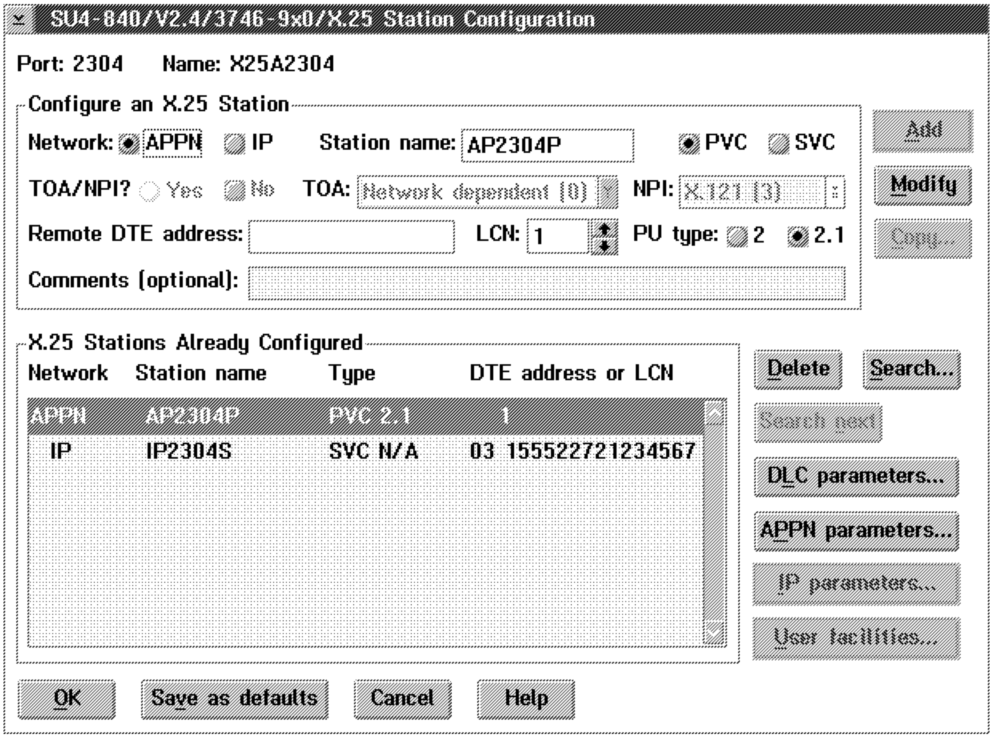

SU4-840/V2.4/3746-9x0/X.25 Station Configuration

Port: 2304 Name: X25A2304

Configure an X.25 Station

Network: ● APPN ○ IP Station name: AP2304P ● PVC ○ SVC

TOA/NPI? ○ Yes ● No TOA: Network dependent (0) NPI: X.121 (3)

Remote DTE address: LCN: 1 PU type: ○ 2 ● 2.1

Comments (optional):

X.25 Stations Already Configured

Network	Station name	Type	DTE address or LCN
APPN	AP2304P	PVC 2.1	1
IP	IP2304S	SVC N/A	03 155522721234567

Figure 9-14. X.25 Station Configuration

Step 24. Configure the APPN station:

- Select **APPN**.
- Set Station name to **AP2304P**.
- Select PVC.
- Set LCN to **1**.
- Set PU to **2.1**.

Step 25. Select **Add** to register this station.

Step 26. Select **APPN parameters** to define additional parameters for this station, especially the HPR support, adjacent node identification or MLTG parameters, DLUR parameters.

Step 27. From the X.25 Station Configuration window, configure now the IP station:

- Select **IP**.
- Set Station name to **IP2304S**.
- Select **SVC**.
- Select **TOA/NPI?**.
- Set TOA to **Network dependent (0)**.
- Set NPI to **X.121 (3)**.
- Set Remote DTE address to **155522721234567**.

Step 28. Select **Add** to register this station.

Step 29. Select **IP parameters** to define the destination IP address. Note that the IP parameters window must be selected for each IP station, if more than one IP station is to be configured.

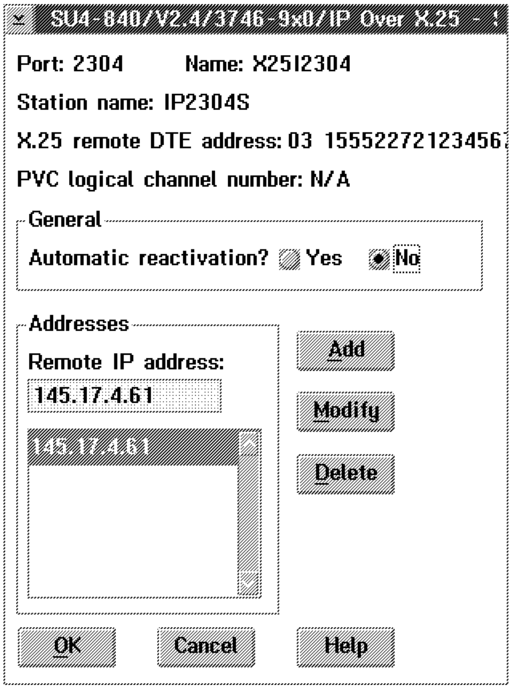

SU4-840/V2.4/3746-9x0/IP Over X.25 - !

Port: 2304 Name: X25I2304

Station name: IP2304S

X.25 remote DTE address: 03 155522721234567

PVC logical channel number: N/A

General

Automatic reactivation? ○ Yes ● No

Addresses

Remote IP address:

145.17.4.61

145.17.4.61

Figure 9-15. Remote DTE IP Address

Step 30. Enter `145.17.4.61` in the Remote IP address field.

Step 31. Select **Add** to register this address.

Step 32. Select **OK** to return to X.25 Station Configuration dialog and save parameters.

Step 33. Select **User Facilities**.

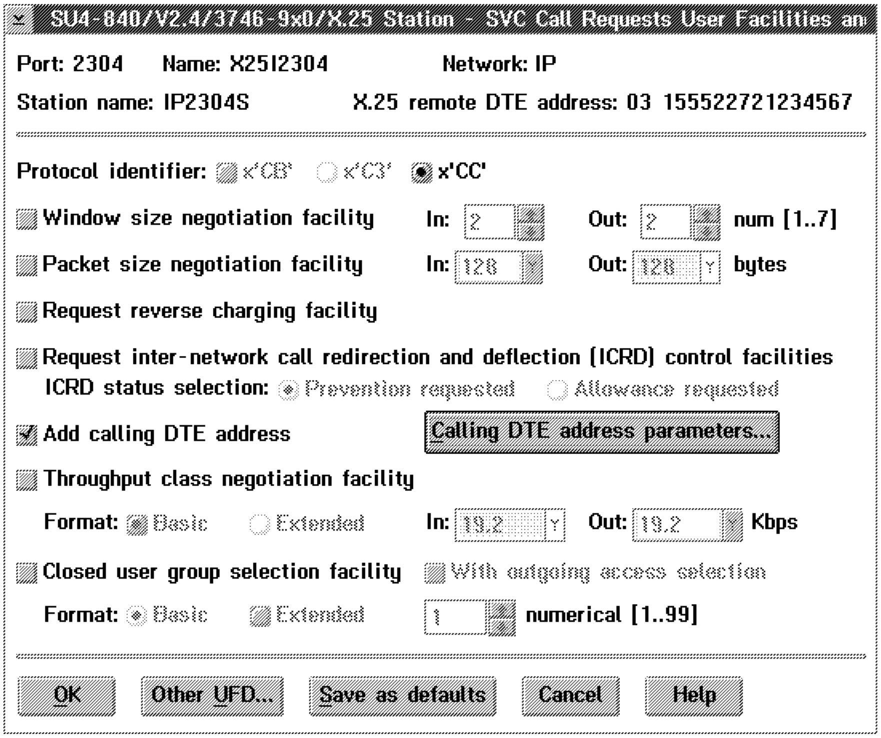

SU4-840/V2.4/3746-9x0/X.25 Station - SVC Call Requests User Facilities an

Port: 2304 Name: X25I2304 Network: IP

Station name: IP2304S X.25 remote DTE address: 03 155522721234567

Protocol identifier: ○ x'CB' ○ x'C3' ● x'CC'

☐ Window size negotiation facility In: 2 Out: 2 num [1..7]

☐ Packet size negotiation facility In: 128 Out: 128 bytes

☐ Request reverse charging facility

☐ Request inter-network call redirection and deflection (ICRD) control facilities

ICRD status selection: ● Prevention requested ○ Allowance requested

☑ Add calling DTE address

☐ Throughput class negotiation facility

Format: ● Basic ○ Extended In: 19.2 Out: 19.2 Kbps

☐ Closed user group selection facility ☐ With outgoing access selection

Format: ● Basic ○ Extended 1 numerical [1..99]

Figure 9-16. SVC CAll Requests, User Facilities and Data

Step 34. Select **Add Calling DTE Address**.

Remark that the calling DTE address must be included in the call request in the DTE to DTE environment.

Step 35. Select **Calling DTE Address Parameters**.

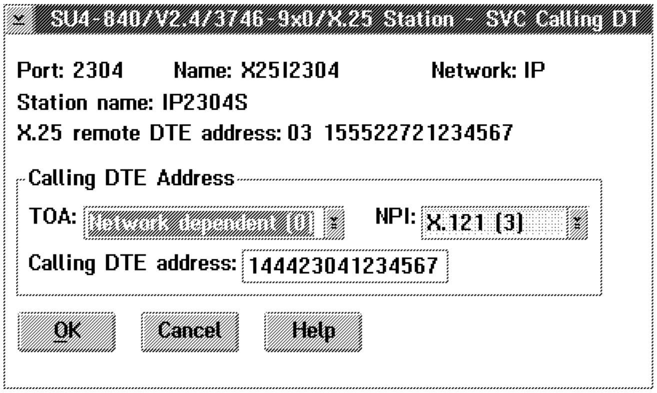

SU4-840/V2.4/3746-9x0/X.25 Station - SVC Calling DT

Port: 2304 Name: X25I2304 Network: IP

Station name: IP2304S

X.25 remote DTE address: 03 155522721234567

Calling DTE Address

TOA: Network dependent (0) NPI: X.121 (3)

Calling DTE address: 144423041234567

Figure 9-17. SVC Calling DTE Address

Step 36. Check that the proposed value for the calling DTE address is right. This one is the local DTE address configured in the PLP Parameters 1/2 dialog, in addition to default values for TOA and NPI, because the remote DTE address has been configured with the TOA/NPI format in the X.25 Station Configuration window.

This concludes the parameter definitions for port 2304.

3746 Machine B

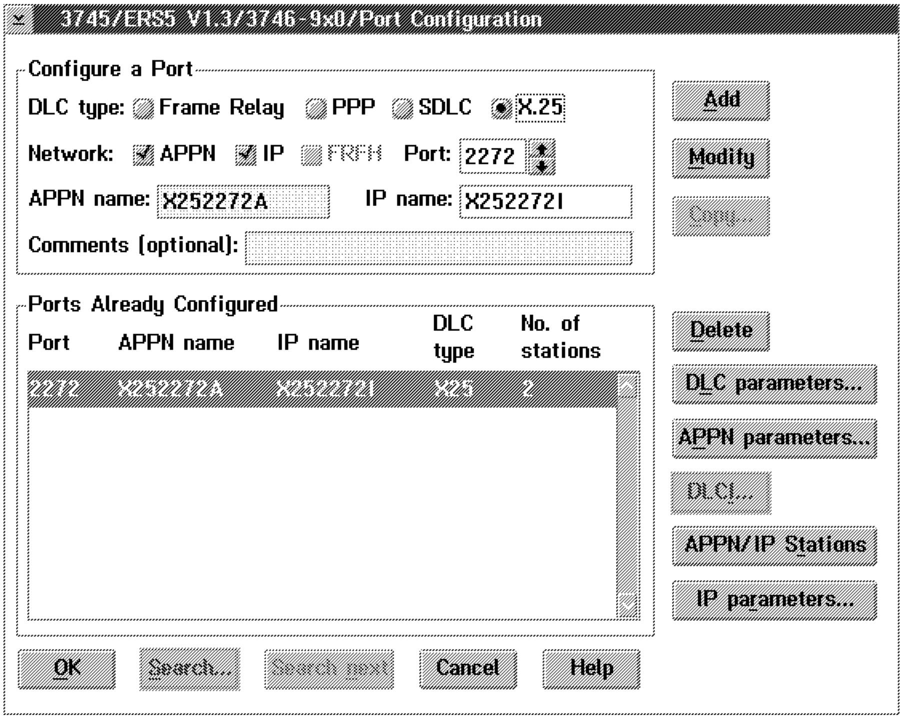

Configure a Port

DLC type: ○ Frame Relay ○ PPP ○ SDLC ● X.25

Network: ☑ APPN ☑ IP □ FRFH Port: 2272

APPN name: X252272A IP name: X252272I

Comments (optional):

Ports Already Configured

Port	APPN name	IP name	DLC type	No. of stations
2272	X252272A	X252272I	X25	2

Figure 9-18. Port 2272 Configuration

1. Select the port to be configured, in this case **2272**.
2. Select **X.25** and the other parameters.
3. Select **Add** to register the port.
4. Select **DLC Parameters**.

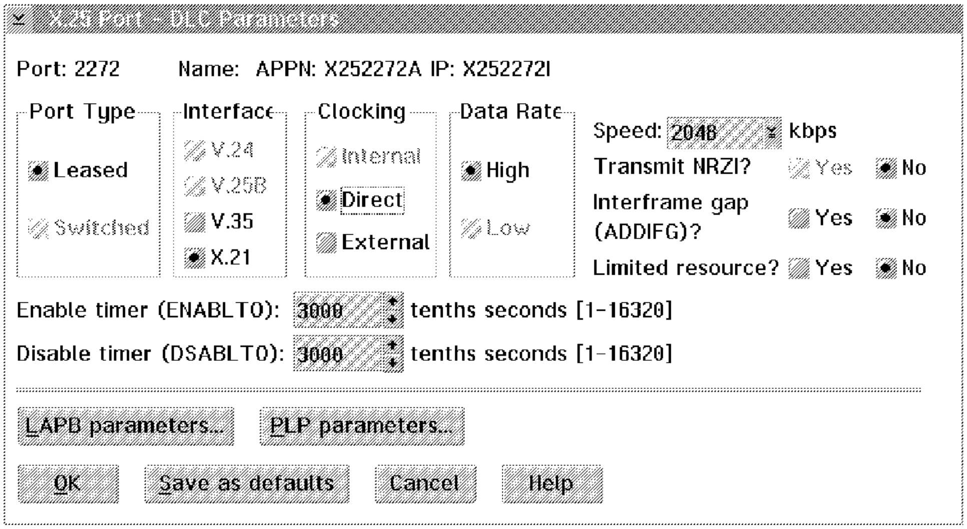

Port: 2272 Name: APPN: X252272A IP: X252272I

Port Type: ● Leased ○ Switched

Interface: ○ V.24 ○ V.25B ○ V.35 ● X.21

Clocking: ○ Internal ● Direct ○ External

Data Rate: ● High ○ Low

Speed: 2048 kbps

Transmit NRZI? ○ Yes ● No

Interframe gap (ADDIFG)? ○ Yes ● No

Limited resource? ○ Yes ● No

Enable timer (ENABLTO): 3000 tenths seconds [1-16320]

Disable timer (DSABLTO): 3000 tenths seconds [1-16320]

Figure 9-19. DLC Parameters for Port 2272

5. Select the right parameters (interface, speed, etc..)

6. Select **LAPB Parameters**.

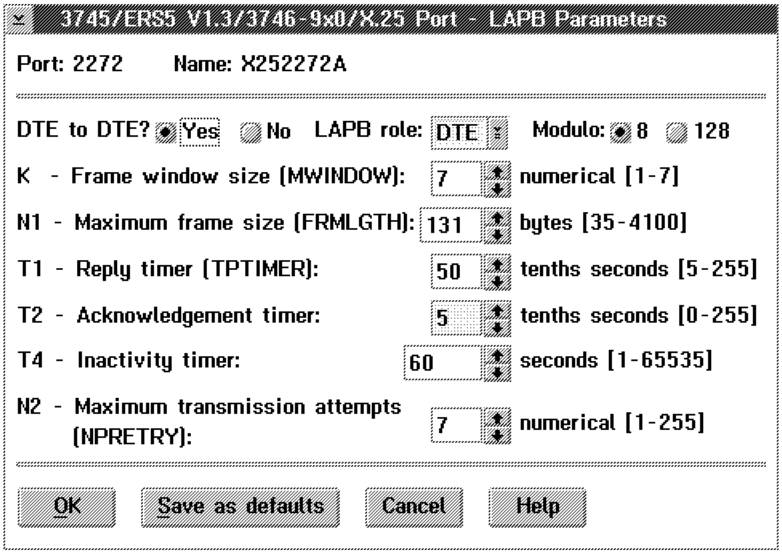

Port: 2272 Name: X252272A

DTE to DTE? ● Yes ○ No LAPB role: DTE Modulo: ● 8 ○ 128

K - Frame window size (MWINDOW): 7 numerical [1-7]

N1 - Maximum frame size (FRMLGTH): 131 bytes [35-4100]

T1 - Reply timer (TPTIMER): 50 tenths seconds [5-255]

T2 - Acknowledgement timer: 5 tenths seconds [0-255]

T4 - Inactivity timer: 60 seconds [1-65535]

N2 - Maximum transmission attempts (NPRETRY): 7 numerical [1-255]

Figure 9-20. LAPB Parameters

7. Set *DTE to DTE?* to **YES**, *LAPB Role* to **DTE** and *LAPB Modulo* to **8**.
8. Configure the other LAPB parameters if needed.
9. Select **OK** to return to the DLC Parameters dialog and save the parameters.
10. From this window, select *PLP Parameters*.

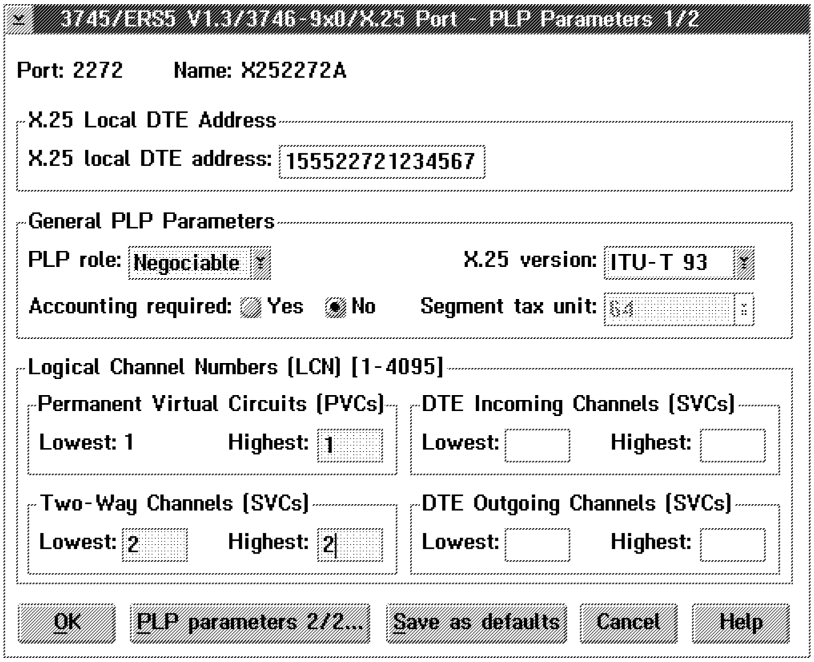

Port: 2272 Name: X252272A

X.25 Local DTE Address

X.25 local DTE address: 155522721234567

General PLP Parameters

PLP role: Negociable X.25 version: ITU-T 93

Accounting required: ○ Yes ● No Segment tax unit: 64

Logical Channel Numbers (LCN) [1-4095]

Permanent Virtual Circuits (PVCs)

Lowest: 1 Highest: 1

DTE Incoming Channels (SVCs)

Lowest: Highest:

Two-Way Channels (SVCs)

Lowest: 2 Highest: 2

DTE Outgoing Channels (SVCs)

Lowest: Highest:

Figure 9-21. PLP Parameters 1/2

11. Set *Local DTE Address* to **155522721234567**.

12. Configure the *Logical Channel Numbers*.

    In the DTE-to-DTE environment, the logical channel numbers must be the same on both ends of the same physical line. When connected to an X.25 network, the logical channel numbers must be those subscribed to in the network.

13. Select **PLP Parameters 2/2**.

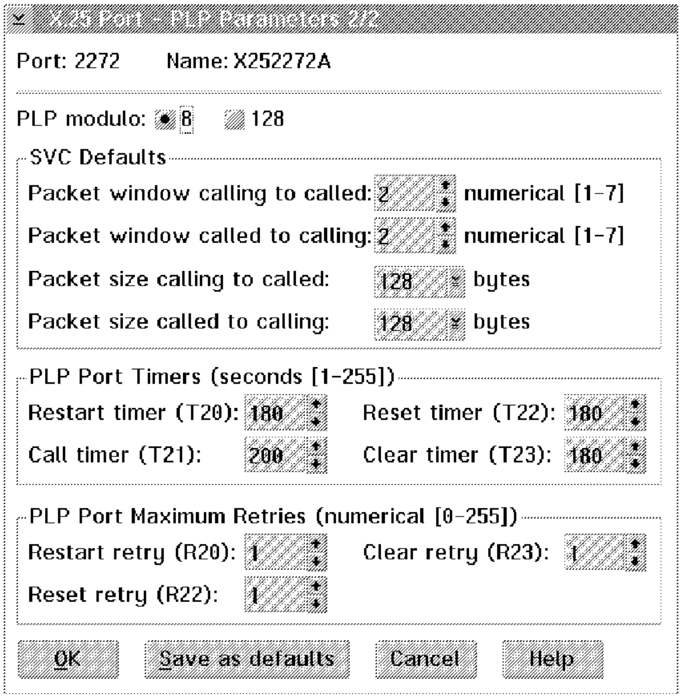

X.25 Port - PLP Parameters 2/2

Port: 2272 Name: X252272A

PLP modulo: ● 8 ○ 128

SVC Defaults

Packet window calling to called: 2 numerical [1-7]

Packet window called to calling: 2 numerical [1-7]

Packet size calling to called: 128 bytes

Packet size called to calling: 128 bytes

PLP Port Timers (seconds [1-255])

Restart timer (T20): 180 Reset timer (T22): 180

Call timer (T21): 200 Clear timer (T23): 180

PLP Port Maximum Retries (numerical [0-255])

Restart retry (R20): 1 Clear retry (R23): 1

Reset retry (R22): 1

Figure 9-22. PLP Parameters 2/2

14. Set PLP Modulo to **8**.

15. Change the packet sizes and/or window sizes, if needed.

16. Select **OK** three times to return to the Port Configuration dialog and save the parameters in the corresponding windows.

17. From the Port Configuration window, select **IP Parameters**.

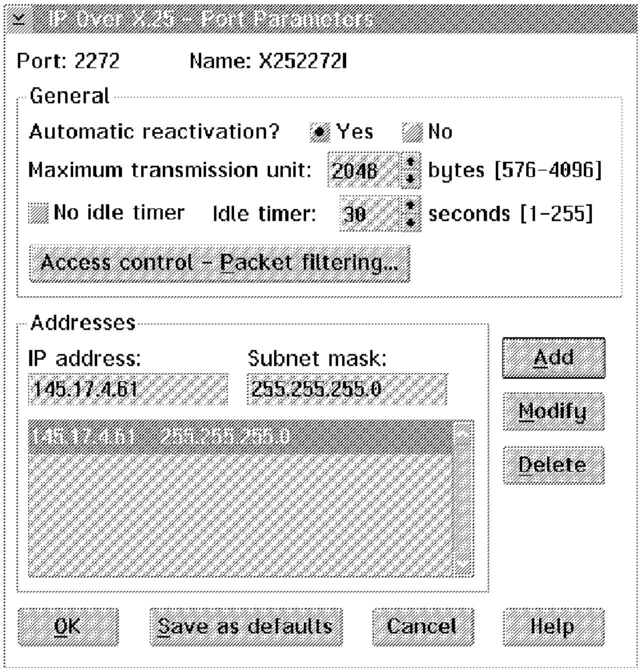

IP Over X.25 - Port Parameters

Port: 2272 Name: X252272I

General

Automatic reactivation? ● Yes ○ No

Maximum transmission unit: 2048 bytes [576-4096]

□ No idle timer Idle timer: 30 seconds [1-255]

Addresses

IP address: 145.17.4.61

Subnet mask: 255.255.255.0

145.17.4.61 255.255.255.0

Figure 9-23. IP Port Parameters

18. Set the MTU size for this port.
19. Configure the IP address of this port to 145.17.4.61.
20. Configure the Subnet mask to 255.255.255.0. Port 2304 of machine A must have the same value, because both ports belong to the same IP subnet.
21. Select **Add** to register this address
22. Select **OK** to return to the Port Configuration dialog and save the IP parameters.

23. From this window, select **APPN/IP Stations** to define the IP and APPN link stations in the adjacent node.

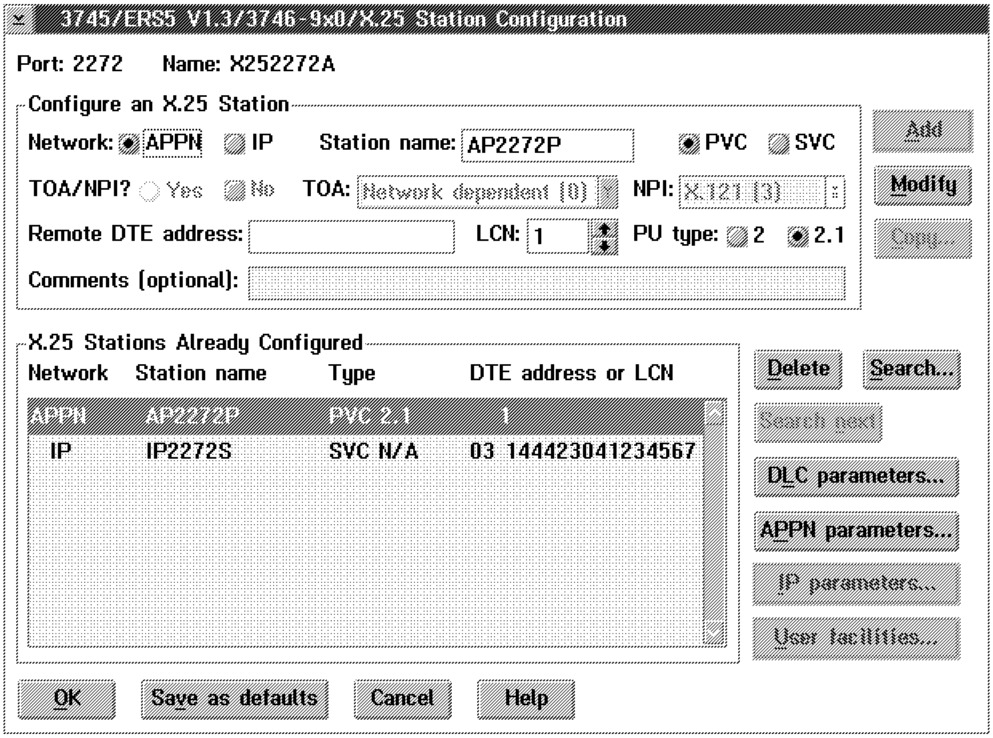

3745/ERS5 V1.3/3746-9x0/X.25 Station Configuration

Port: 2272 Name: X252272A

Configure an X.25 Station

Network: ● APPN ○ IP Station name: AP2272P ● PVC ○ SVC

TOA/NPI? ○ Yes ○ No TOA: Network dependent (0) NPI: X.121 (3)

Remote DTE address: LCN: 1 PU type: ○ 2 ● 2.1

Comments (optional):

X.25 Stations Already Configured

Network	Station name	Type	DTE address or LCN
APPN	AP2272P	PVC 2.1	1
IP	IP2272S	SVC N/A	03 144423041234567

Figure 9-24. X.25 Station Configuration

24. Configure the APPN station:
    - Select **APPN**.
    - Set Station name to **AP2272P**.
    - Select **PVC**.
    - Set LCN to **1**.
    - Set PU to **2.1**.

25. Select **Add** to register this station.

26. Select **APPN parameters** to define additional parameters for this station, especially the HPR support, adjacent node identification or MLTG parameters, DLUR parameters.

27. From the X.25 Station Configuration window, configure now the IP station:
    - Select **IP**.
    - Set Station name to **IP2272S**.
    - Select **SVC**.
    - Select **TOA/NPI?**.
    - Set *TOA* to **Network dependent (0)**.
    - Set NPI to **X.121 (3)**.
    - Set Remote DTE address to **155522721234567**.

28. Select **Add** to register this station.

29. Select **IP parameters** to define the destination IP address. Note that the IP parameters window must be selected for each IP station, if more than one IP station is to be configured.

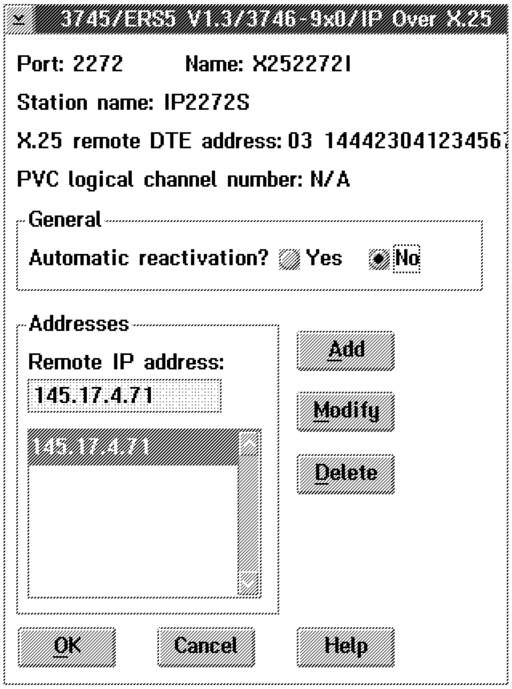

3745/ERS5 V1.3/3746-9x0/IP Over X.25

Port: 2272   Name: X252272I

Station name: IP2272S

X.25 remote DTE address: 03 14442304123456

PVC logical channel number: N/A

General

Automatic reactivation? ○ Yes ● No

Addresses

Remote IP address: 145.17.4.71

145.17.4.71

Figure 9-25. Remote DTE IP Address

30. Enter `145.17.4.71` in the Remote IP address field.
31. Select **Add** to register this address.
32. Select **OK** to return to 25 Station Configuration dialog and save parameters.

33. Select **User Facilities**.

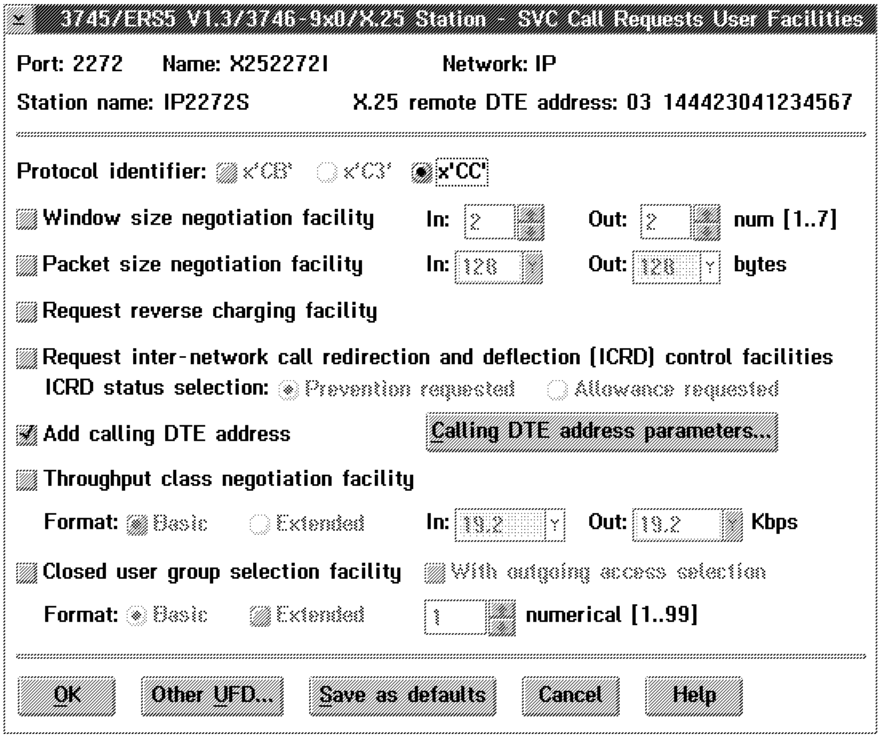

3745/ERS5 V1.3/3746-9x0/X.25 Station - SVC Call Requests User Facilities

Port: 2272 Name: X2522721 Network: IP

Station name: IP2272S X.25 remote DTE address: 03 144423041234567

Protocol identifier: □ x'CB' ○ x'C3' ● x'CC'

□ Window size negotiation facility In: 2 Out: 2 num [1..7]

□ Packet size negotiation facility In: 128 Out: 128 bytes

□ Request reverse charging facility

□ Request inter-network call redirection and deflection (ICRD) control facilities

ICRD status selection: ● Prevention requested ○ Allowance requested

☑ Add calling DTE address

□ Throughput class negotiation facility

Format: ● Basic ○ Extended In: 19.2 Out: 19.2 Kbps

□ Closed user group selection facility □ With outgoing access selection

Format: ● Basic □ Extended 1 numerical [1..99]

Figure 9-26. SVC Call Requests User Facilities and Data

34. Select **Add Calling DTE Address**.

    Remark that the calling DTE address must be included in the call request in the DTE to DTE environment.

35. Select **Calling DTE Address Parameters**.

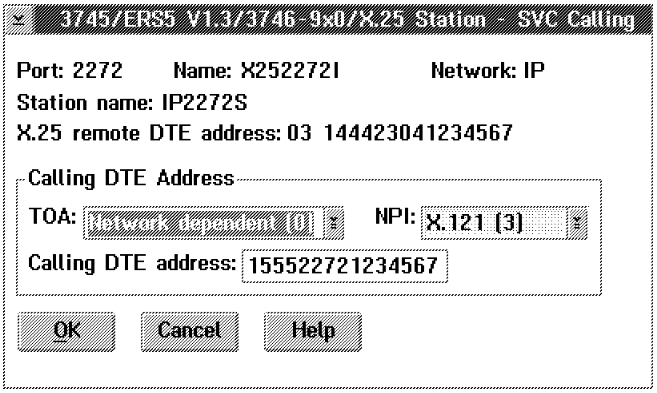

3745/ERS5 V1.3/3746-9x0/X.25 Station - SVC Calling

Port: 2272 Name: X2522721 Network: IP

Station name: IP2272S

X.25 remote DTE address: 03 144423041234567

Calling DTE Address

TOA: Network dependent (0) NPI: X.121 (3)

Calling DTE address: 155522721234567

Figure 9-27. SVC Calling DTE Address

36. Check that the proposed value for the calling DTE address is right. This one is the local DTE address configured in the PLP Parameters 1/2 dialog, in addition to default values for TOA and NPI, because the remote DTE address has been configured with the TOA/NPI format in the X.25 Station Configuration window.

This concludes the parameter definitions for port 2272.

Chapter 10. ISDN Adapters

Integrated services digital network (ISDN) is an Open Systems Interconnection (OSI) protocol for digital telecommunications network, based on the 64 Kbps circuit switched technology. ISDN supports multiple services including voice and data transmission. Some examples of devices that might have an ISDN interface are a digital telephone, an integrated digital voice or data terminal, and digital facsimile equipment. ISDN is advantageous for data transmission, when you do not need a constant, dedicated connection because you only pay for the duration of a call.

There are two ways to connect a 3746 Model 900 to an ISDN:

- Through the LIC16, which is a native ISDN adapter; see Chapter 18, “Serial Line Adapters” on page 18-1
- Through a terminal adapter which is an external box; see “ISDN Terminal Adapter” on page 10-14.

The 3746-9x0 offers also the ISDN connectivity for non-NCP controlled traffic through the 3746 MultiAccess Enclosure (MAE). Refer to Chapter 23, “Multiaccess Enclosure ISDN Support” on page 23-1.

Figure 10-1 on page 10-2 shows how NCP controlled traffic can access ISDN networks.

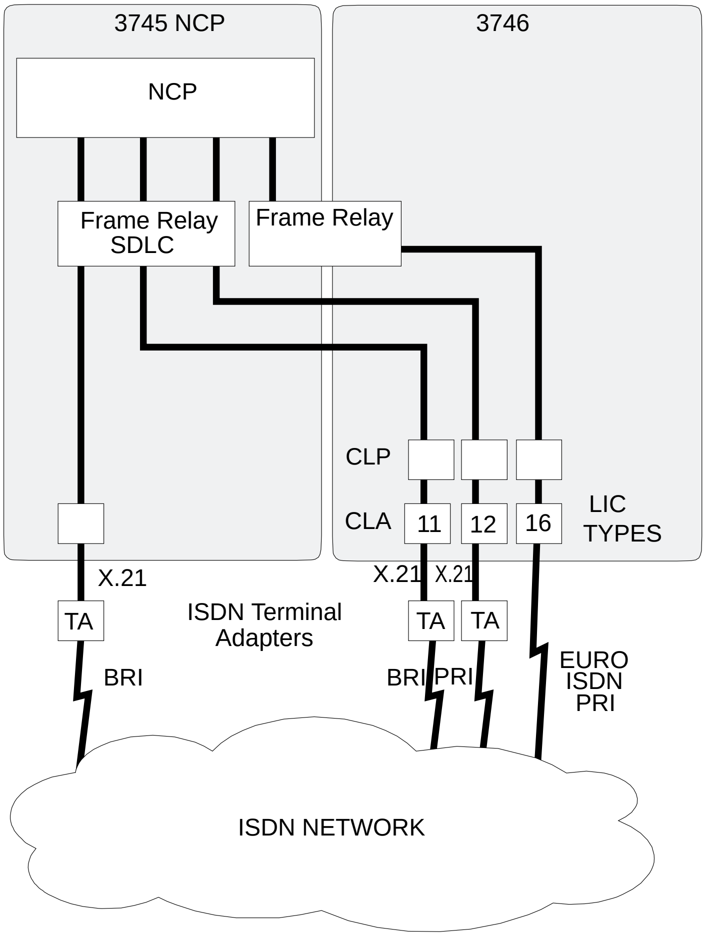

Figure 10-1. ISDN on the 3746-900 for NCP Controlled Traffic

Native Adapter Through LIC16

Primary Rate Interface Support

The 3746 Model 900 supports frame relay over ISDN Primary Rate Interface (PRI) conforming to Euro-ISDN standard.

An ISDN PRI port of the 3746 Model 900 (LIC16) provides one ISDN D channel, which is reserved for ISDN signalling, and 30 ISDN B channels for simultaneous full-duplex transport of user data at 64 Kbps. The ISDN B channels of a PRI port are separately used to establish connections with one or multiple remote equipment. The 3746 Model 900 does not support ISDN connections that include multiple B channels, such as H0, H11, or H12 channel.

A B channel connection, when established, constitutes a 64 Kbps end-to-end pipe. The 3746 Model 900 uses the frame relay protocol over this pipe to transport user data.

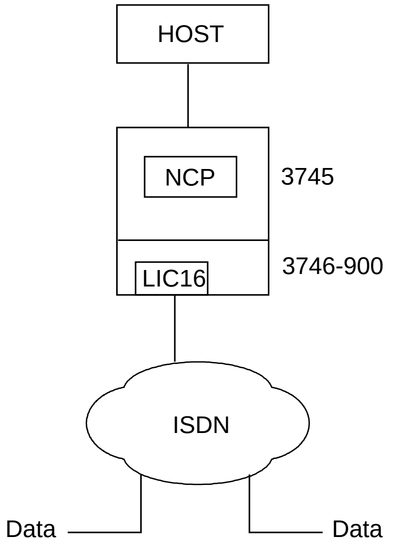

Figure 10-2. ISDN on the 3746-900

SNA Connectivity

Starting with NCP V7R5, the ISDN PRI ports of the 3746 Model 900 allow ACF/NCP to route SNA traffic to remote equipment supporting fuame relay over ISDN, such as the IBM 2210 Nways Multiprotocol Router, the IBM 2216 Nways Multiaccess Connector, or another IBM 3746 Model 900.

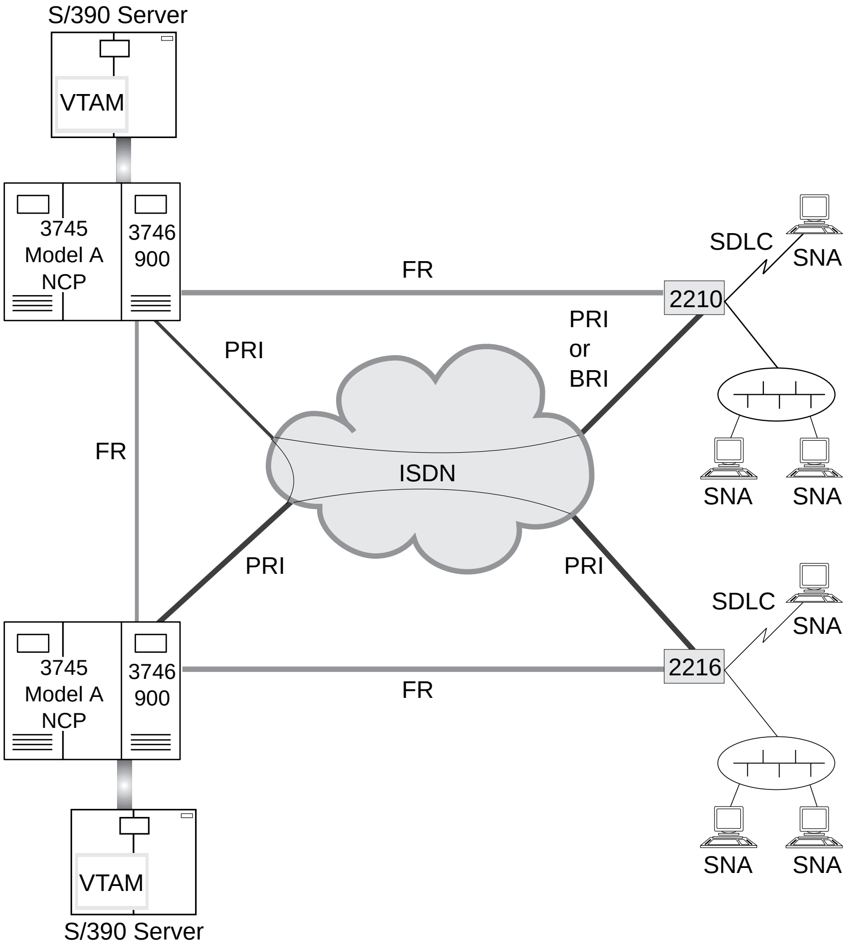

Legend:
PRI Primary Rate Interface (30B + D channels)
BRI Basic Rate Interface (2B + D channels)
FR Frame Relay (via leased connections or frame-relay network)

Figure 10-3. Connectivity to ISDN Network

Although the 3746 Model 900 supports connections including only one B channel, multiple ISDN B channels connecting two 3746 Model 900s can be used as a single logical connection (MultiLink Transmission Group) to provide high bandwidth for communication between two ACF/NCPs.

The 3746 Model 900 can automatically call remote equipment over ISDN. For incoming calls over ISDN, the ISDN number and subaddress of the calling party is passed to VTAM for possible verification via user exit routine. Figure 10-3 represents a sample network that uses frame relay and ISDN connections, both primary (PRI) and basic (BRI). The 3746 Model 900 PRI ports enable the following SNA connections over ISDN:

ISDN subarea links

Two NCPs can exchange data over a frame relay subarea link which passes through the ISDN network.

Multiple subarea links over multiple B-channels between two NCPs can be aggregated using the Mixed Media Multilink TG (MMMLTG) feature.

A 3746 Model 900 with LIC16 can connect with a remote 3745 or 3746 Model 900 attached to the network via a terminal adapter, for example a 7820, or via a BAN device, such as a 2210 and a token-ring. In that case the APAR IR34013 is required on the NCP V7R5.

ISDN peripheral links

A NCP can exchange data with peripheral devices over frame relay peripheral links that pass through the ISDN network.

Two NCPs can also connect together over a frame relay APPN peripheral link, when each belongs to a Composite Network Node (CNN).

Functions Supported

Attachment to an Integrated Services Digital Network (ISDN) through a LIC16 provides the following functions:

Call on demand

When equipment does not need to be permanently connected to the 3746 Model 900, ISDN can be used to established connections only for the duration of data transmissions. The 3746 Model 900 can initiate or receive ISDN calls.

Refer to “How an Incoming Call is Processed” on page 10-6.

In case of a subarea node traffic, you can define a MultiLink Transmission Group (MLTG) to the subarea node and as many ISDN connections within the MLTG as you need. The bandwidth of the subarea link can reach up to 30 x 64 = 1920 Kbps through ISDN.

Backup through ISDN

Refer to “Backup through ISDN” on page 10-7.

Bandwidth on demand

Refer to “Bandwidth on demand” on page 10-9.

Calling party verification

Refer to “Security Considerations” on page 10-10.

Remote loading and activation

A NCP can load and activate a remote NCP through a subarea node connection established over an ISDN.

Performance monitoring

The Network Performance Monitor (NPM) program can monitor the traffic on each B-channel, exactly as if the B-channel was a frame relay physical line.

There is no performance monitoring for the D-channel, because it flows a very low traffic (less than 0.5 Kbps compared to its 64 Kbps speed, for an average call duration of one minute and the 30 B-channels permanently used).

Accounting

The data of each ISDN call, such as the call duration, called and calling party numbers, called and calling subaddresses, and cause for call clearing, are reported to the Network Performance Monitor (NPM) program for accounting purposes. This allows you to check and dispatch the carrier's charges, as well as to have detailed statistics at the call level.

How an Incoming Call is Processed

An ISDN incoming call can be destined to the NCP (call on demand) or initiate an automatic backup (refer to “Automatic Backup” on page 10-7). So an incoming call is checked for both call on demand and backup as follows:

1. If a called party subaddress is present in the ISDN incoming call setup and if it is identical to the local party subaddress configured for the ISDN line, then the call is accepted and transmitted to the NCP (call on demand).

2. Each frame relay line eligible for the automatic backup over this ISDN line is examined.

   If a called party subaddress is present in the ISDN incoming call setup and if it is identical to the local party subaddress configured for the frame relay line being checked, then the call is accepted and the backup for that frame relay line starts. If not identical, the next frame relay line is checked.

   If the called party subaddress that is received in the incoming call does not correspond to any frame relay line that is eligible for automatic backup, the call is rejected.

   Note that the called party subaddress identifies the frame relay line for which the backup is being performed. Therefore each frame relay line, eligible for automatic backup and selectable through the called subaddress, must be allocated a *unique* subaddress. One subaddress must be additionally allocated for the calls on demand. The latter identifies the NCP.

3. If there is no called party subaddress in the incoming call setup, the called party number is checked. If the called party number, received in the incoming call setup, is identical to the local party number configured for the ISDN line, then the call is accepted and transmitted to the NCP (call on demand).

4. If there is no called party subaddress in the incoming call setup, and if the called party number, received in the incoming call, is identical to the local party number configured for the frame relay line being checked, then the call is accepted and the backup for that frame relay line starts. Otherwise the next frame relay line is checked.

   If the called party number that is received in the incoming call, does not correspond to any frame relay line that is eligible for automatic backup, the call is rejected.

   Note that the called party number identifies the frame relay line for which the automatic backup is being performed, when the called party subaddress is not provided by the remote end.

   Therefore each frame relay line, eligible for backup and selectable through the called party number, must be allocated a *unique* number. This requires to subscribe either the Direct Dialling Inward (DDI) or Multiple Subscriber Number (MSN) supplementary service to the network, in order to have as many ISDN numbers as frame relay lines eligible to automatic backup (plus one for the calls on demand).

In summary, the 3746 Model 900 uses two methods to detect whether an incoming ISDN call is for call on demand or for automatic backup and, in the latter case, to identify the line to be backed up:

- Subaddress method
- Number method

Checking subaddresses is the preferred method, whenever possible, because ISDN networks always transparently carry the subaddress. On the other hand, ISDN networks can change the format of the ISDN number, for example, from a national to an international number. Also some networks (for example, the French Numeris network) can truncate the called party number. If you intend to use the number based-method, ask your service provider how the network delivers the called party number.

The subaddress method can also help avoid unexpected backup attempts if used as a kind of password. Refer to “Security Considerations” on page 10-10.

If you plan to use the subaddress method, refer to “Allocating Subaddresses” on page 10-12.

Backup through ISDN

Back p of a subarea or peripheral line is possible through ISDN either automatically or by means of C-list or operator commands.

Automatic Backup

A frame relay line attached to a 3746 Model 900 can be automatically backed up through an ISDN line attached to the same 3746 Model 900. When the frame relay line fails, the 3746 Model 900 automatically reroutes the frame relay traffic over a B-channel of the ISDN line.

Backup of a frame relay line over ISDN provides the following benefits:

- Backup is automatic, requiring neither a CLIST nor manual intervention to operate.
- Backup is non-disruptive. The frame relay terminating equipment recovers any lost frames. Frame relay resources are not impacted. A resource can be activated or de-activated on the failed frame relay line.

Since the remote end of the failed frame relay line must be aware of the backup, because it participates in establishing the ISDN connection, it must be a 3746 Model 900 or have equivalent backup capability.

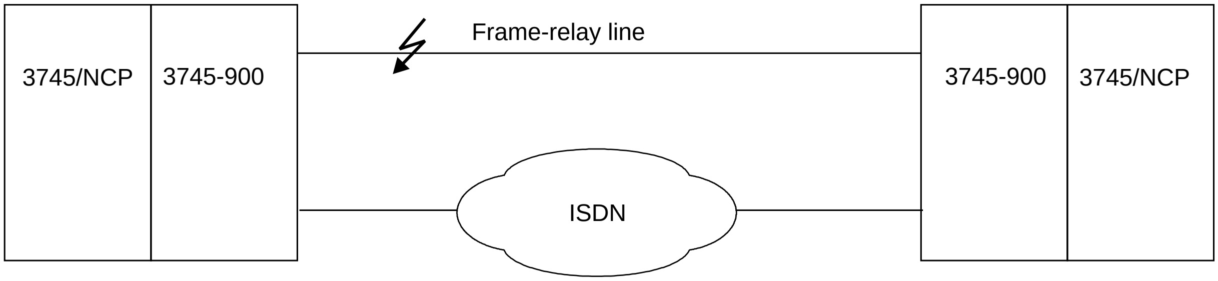

Figure 10-4. Automatic Frame Relay Backup over an ISDN

How automatic backup operates: One end of the frame relay line is configured as the backup initiator and the other as the backup recipient. When the backup initiator detects a failure, it triggers the backup function by placing an ISDN call over one of the two ISDN lines predefined to back up this frame relay line.

The failure can be:

- An external hardware failure (such as line, or modem)
- An internal hardware failure (such as LIC, or ARC)
- A LMI failure

When an ISDN call arrives, it is first checked to determine whether it is a call on demand or a backup call, according to the algorithm described in “How an Incoming Call is Processed” on page 10-6.

When the backup recipient detects a failure, it waits for the backup call. If it does not arrive within about one minute, it enters the normal failure processing by reporting the error to NCP.

For further information on the automatic backup function, refer to the NCP Resource Definition Guide (SC31-6223-05).

RDG V7R6 (which is not published yet), Otherwise APAR doc on NCP V7R5.

Manual or Automated Backup of a Subarea Link

To backup through ISDN a subarea link without traffic disruption, define the subarea link to be backed up as part of a Mixed Media MultiLink Transmission Group (MMMLTG) together with another link defined as a “hot standby” link (TGCONF=STANDBY) and as many ISDN links as necessary to carry all of the subarea traffic, see Figure 10-5.

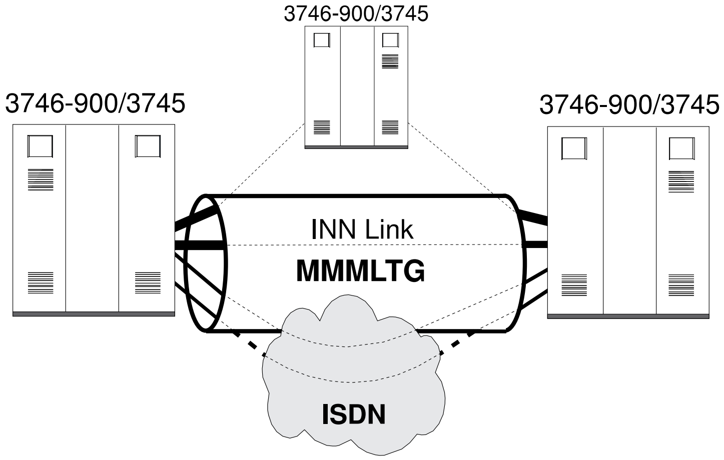

Figure 10-5. MMMLTG with ISDN Backup

This “hot standby” link must be activated, but does not carry any MMMLTG traffic, except for a short period during the ISDN backup establishment. So it can be normally used for other purposes, as shown on Figure 10-5.

The subarea link to be backed up can carry SDLC or frame relay traffic.

If the subarea link fails, the MMMLTG does not fail because its traffic immediately begins to flow over the standby link.

Failure of the subarea link causes an alert to be sent to NetView. This alert can trigger a C-list, which establishes one or more ISDN connections between the two 3746-900s. As soon as established, the whole subarea traffic flows through ISDN, no longer through the standby link. So the backup of a subarea link can be automated through a C-list without any traffic disruption. The standby link prevents disruption of the traffic during the setup of ISDN connections.

Up to 30 B-channels (up to 1.92Mbps) can be used for backup per ISDN port.

Note: If no standby line is specified in the MMMLTG, backup is still possible (see Figure 10-6 on page 10-10). But in this case, the backup is disruptive, because there is no link available in the MMMLTG during the ISDN link setup.

When the MMLTG leased line goes up again, NCP sends a CONTACTED message to VTAM, after the XID exchange. This message can trigger another CLIST that releases all the ISDN connections (B-channels) being used for backup. Thus, the switch back can also be automated through a CLIST without any traffic disruption.

Manual or Automated Backup of a Peripheral Line

Three configurations are possible:

1. Frame relay leased line to an IBM 2210 or 2216 router.

   In this case use the WAN reroute capability of the 2210 or 2216.

2. Access line to a frame relay network.

   The frame relay network must provide a backup function through ISDN.

3. Access to a IBM 2210 or 2216 router via a frame relay network.

   In this case only one DLCI must be subscribed to the network (or several DLCIs, if all terminate to the same 2210 or 2216). In addition the DLCI number must be the same on both sides of the network.

   The WAN reroute capability of the 2210 or 2216 is also used for backup.

In cases 1 and 3, the ISDN call can be initiated either by the 2210 or the 3746-900.

Backup can be automated through a CLIST triggered by the NetView alert caused by the line failure. Only one B-channel (64 Kbps) can be used for backup. Backup is disruptive.

Bandwidth on demand

When the traffic rate exceeds the capacity of permanent connections, temporary connections over ISDN can be used to provide additional bandwidth. Therefore, the maximum bandwidth that might be required at peak traffic time does not need to be permanently available over leased connections. For NCP-to-NCP traffic between two 3745/3746-900s (see Figure 10-3 on page 10-4), the frame relay connection can be complemented at peak hours by one or multiple ISDN B channel connections. The Mixed Media MultiLink Transmission Group (MMMLTG) support of NCP allows the frame relay and all the ISDN connections to be aggregated as a single logical connection between the two controllers. See Figure 10-6 on page 10-10.

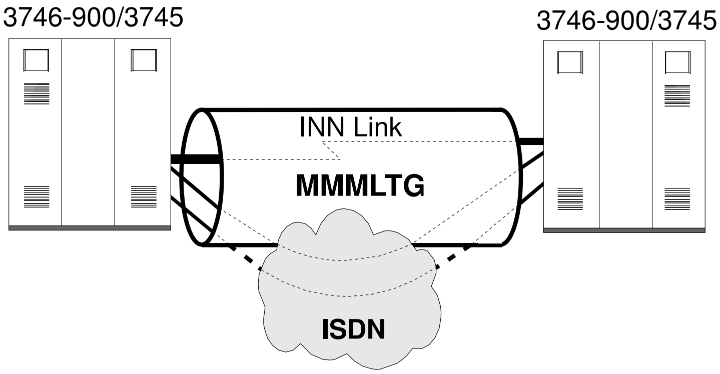

Figure 10-6. ISDN Bandwidth on Demand

When additional bandwidth is needed, establish the ISDN links between the two subarea nodes. See Figure 10-6 on page 10-10. The B-channel connections are in parallel with the normal subarea links to give additional bandwidth to the MMMLTG. When the additional bandwidth is no longer needed, the ISDN connections are to be released.

Increasing or decreasing bandwidth using ISDN channels does not disrupt the MMMLTG traffic.

Bandwidth on demand over ISDN can be automated through CLISTS either:

- Scheduled at a given time of day, or
- Started at the beginning and end of jobs requiring bulk data transfer

The subarea link that bandwidth is to be added to can carry SDLC or frame relay traffic.

Up to 1.92 Mbps (30 B-channels) can be added to a subarea link per ISDN port.

Security Considerations

An unexpected ISDN call could break down connections over a frame relay line eligible for backup, if the called party subaddress or called party number, that is included in the ISDN call setup, is identical to the local party subaddress or local party number configured for that frame relay line. It's the reason why it is important to protect against unexpected or malicious calls.

For that purpose, it is recommended that you use the subaddress method described in section “How an Incoming Call is Processed” on page 10-6. The subaddresses can be chosen long enough to act as a kind of password.

However if the subaddress method is not sufficient or if you don't plan to use it, you can improve security by writing a Configuration Services XID exit routine (ISTEXCCS) at VTAM level.

When VTAM receives a request contact for a B-channel station, it invokes this exit routine, which can use the DLC Connection Data Control Vector (CV57) coming from the 3746-900 through NCP to check the incoming call parameters. The CV57 contains mainly the calling party number and the calling party subaddress, which were received from the network in the incoming call. The exit routine can therefore check these call parameters and abort the call if necessary.

For ISDN, the DLC type at offset 2 is X'05'. The specific subfields for ISDN are:

X'05' which contains the remote party number (calling party number).

X'09' which contains the remote subaddress (calling party subaddress).

X'0A' which contains the allocated bandwidth (in bits/sec). It therefore contains 64000 for one B-channel.

X'0F' which imbeds the DLCI information. This subfield contains other subfields. Refer to the NCP documentation.

The format of the subfields X'05' and X'09' is shown in Table 10-1 and Table 10-2 on page 10-12.

Table 10-1. Remote Party Number Subfield (X'05')

Offset (dec)	Length	Format or value	Contents Description
0	1	hex	Subfield length (from 6 to 20) (see note 1)
1	1	X'05'	Subfield key
2	1	X'00'	Subfield flags
3	1	hex	Party number length (see note 2)
4	1	bit 0... .... .xxx xxxx	TON/NPI fields Reserved Type Of Number (TON): 000 : Unknown 001 : International number 010 : National number 011 : Network specific number 100 : Subscriber number Numbering Plan Identification (NPI): 0000 : Unknown 0001 : ISDN/Telephony numbering plan (E.164) 1000 : National standard numbering plan 1001 : Private numbering plan
5	max 15	IA5	Number digits (see note 3)

Note:
Each subfield is optional. For instance, if the incoming setup message does not include the calling party number, the subfield X'05' is not present in the CV57. If neither the calling party number nor the calling party subaddress is present in the incoming call, none of the subfields X'05' and X'09' is present in the CV57.

1. The subfield length includes the length field itself, that is it equals the sum of the party number length plus 4. The minimum length of the subfield X'05' assumes that the calling party number consists of only one digit.
2. The party number length does not include the length field itself.
3. The digits are decimal and expressed in IA5 characters (that is, like ASCII).

Table 10-2. Remote Party Number Subfield (X'09')

Offset (dec)	Length	Format or value	Contents Description
0	1	hex	Subfield length (from 6 to 25) (see note 1)
1	1	X'09'	Subfield key
2	1	X'00'	Subfield flags
3	1	hex	Subaddress length (see note 2)
4	1	bit	
		`1... ....`	Reserved
		`.xxx ....`	Type of subaddress: 000 : NSAP (specified in X.213 or ISO 8348 addendum 2) 010 : User-specified.
		`.... x...`	Odd/even indicator (used only if subaddress is user-specified and coded in BCD): 0 : even number 1 : odd number.
		`.... .000`	Reserved
5	max 20		Subaddress information

Note:
Each subfield is optional. For instance if the incoming setup message does not include the calling party number, the subfield X'05' is not present in the CV57. If neither the calling party number nor the calling party subaddress is present in the incoming call, none of the subfields X'05' and X'09' is present in the CV57.

1. The subfield length includes the length field itself, that is, it equals the sum of the subaddress length plus 4. The minimum length of the subfield X'09' assumes that the subaddress information field is only one byte long.
2. The subaddress length does not include the length field itself.

Allocating Subaddresses

If you intend to use the subaddress method, described in “How an Incoming Call is Processed” on page 10-6, you have to allocate a subaddress to:

- Each NCP of your network that might receive ISDN calls
- Each frame relay line that is eligible for automatic backup

This subaddress is coded in the ISDNLSA operand.

The subaddress format is either user specified or a NSAP (Network Service Access Point) specified by ISO (standard ISO8348 Addendum 2) or ITU-T (recommendation X.213).

Each subaddress must be *unique*, at least within the customer network. One subaddress identifies one NCP or one frame relay line eligible for automatic backup.

Note: Usage of a NSAP ensures uniqueness of subaddresses. There is, however, one exception for the AFI=50, because the IDI field of the NSAP is not present for this AFI. Refer to Figure 10-7 on page 10-13. The DSP field of the NSAP must be unique within the customer network, whatever the AFI value is.

Format of the NSAPs

There are several formats for a NSAP, depending on the Authority and Format Identifier (AFI) that is included in the first octet of the NSAP. The AFI specifies both the format of the rest of NSAP and which authority allocates the IDI. Depending on the AFI, the DSP and IDI fields are coded with binary, BCD or IA5 (similar to ASCII) syntax. The length of NSAP cannot exceeds 20 octets.

There are two possible formats of NSAP:

1. AFI=50

   For this AFI, there is no IDI. The DSP field must be decimal digits (therefore from 0 to 9), coded with IA5 syntax (similar to ASCII and specified in ITU-T recommendation T.50).

   For example : `50 30313233343536373839`

   Note that, for clarity, the AFI and DSP fields are separated by a blank character.

2. AFI=45

   For this AFI, the IDI contains the ISDN number allocated to the port. The international format, as defined in recommendation E.164, is used. The length of the IDI field is 8 octets, while the E.164 number can be up to 15 digits long. The digits of this number are coded in Binary Coded Decimal syntax (BCD: 2 decimal digits per octet). The E.164 number is padded with leading semi octet 0000 to obtain the maximum length (15 digits). Semi-octet value 1111 is used as a pad after the final semi octet to obtain an integral number of octets. The decimal digits of the DSP is coded with BCD syntax too.

   For example : `45 000033492115220F 0123456789`

   Note that, for clarity, the AFI, IDI and DSP fields are separated by blank characters.

Other formats of NSAP can also be used.

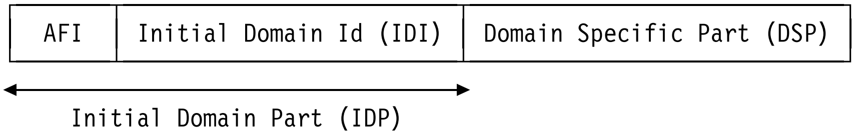

`AFI : Authority and Format Identifier (one octet long)`

Figure 10-7. General Format of NSAPs

Configuration

To configure ISDN resources, refer to the *NCP V7R5 Resources Definition Guide*, and *NCP V7R5 Resources Definition Reference*. The following terms define NCP ISDN resources:

ISDN D Channel
A full-duplex digital channel that carries signaling information to the ISDN network. An NCP D-channel definition represents the physical line.

ISDN B Channel
A full-duplex digital channel that carries user data through the ISDN network. An NCP B-channel definition represents a time slot on an ISDN interface. Each B-channel has an access rate of 64Kbps.

ISDN Terminal Adapter

In addition to the native ISDN adapter (LIC16), ISDN terminal adapters can connect the 3746-900 to remote SNA PUs over an ISDN network for the call-on-demand and backup functions. The X.21 call setup procedure can be used in the 3746-900 to automatically call out (autocall) or answer (autoanswer) via an ISDN network. The framing is SDLC after a call is set up.

The dial digits for a call out are defined in the VTAM switched major node. A call out is initiated by either:

- An application for call on demand
- A CLIST for dial backup

When the call is initiated by the remote DTE, the 3746-900 has to answer the call.

In both modes (autodial or autoanswer), the 3746-900 acts as an X.21 serial interface. The terminal adapter translates the X.21 switched protocol into ISDN call setup messages and vice versa.

The ISDN terminal adapter attaches to a X.21 cable (for the LIC11, an X.21 ARC), which is connected to a 3746-900 LIC11 or LIC12.

This line is defined as a switched line in NCP by coding `DIAL=YES` in the GROUP definition statement. In this statement the `X21NKWT` and `X21SW` are not required, as they are only defined for an X.21 network connection.

Basic Rate Interface

The LIC11 is compatible with Basic Rate Interface (BRI) terminal adapters such as the IBM 7820. Two 64 Kbps LIC11 lines can attach the two 64 Kbps ports of the IBM 7820. See Figure 10-8.

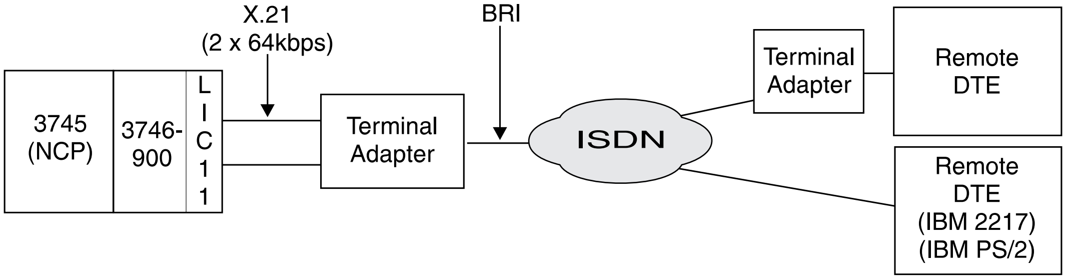

Figure 10-8. BRI-ISDN Terminal Adapter

Primary Rate Interface

The LIC11 and LIC12 are compatible with Primary Rate Interface (PRI) terminal adapters such as the Hitachi HN-5110-24. A LIC11 line or LIC12 line is required for each connected DTE destination:

- Using a LIC11 with 64 Kbps lines, the number of ISDN B-channels that can be used by the terminal adapter (see Figure 10-9) depends on how the terminal adapter is connected:
  - For a T1 or J1 terminal adapter-ISDN network connection, up to 24 ARC can attach up to 24 ports at 64 Kbps.
  - For an E1 terminal adapter-ISDN network connection, up to 30 ARCs can attach up to 30 ports at 64 Kbps.
  - The LIC11 can be connected to terminal adapter ports with different speeds: 64, 128, 192, or 256 Kbps.

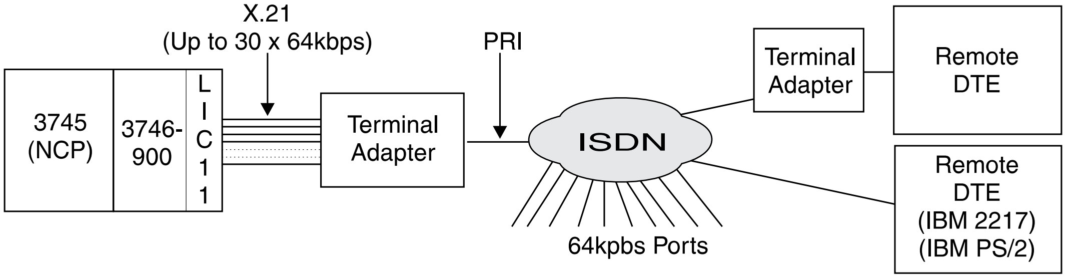

Figure 10-9. PRI-ISDN Terminal Adapter with a LIC11

  The number of lines between the 3746-900 and the terminal adapter corresponds to the maximum number of remote DTEs simultaneously connected. The speed of the line attached to the 3746-900 corresponds to the speed supported by the remote DTE attachment. For example, if the line speed is 128 Kbps, the remote DTE is either attached:
  - To a terminal adapter by a 128 Kbps line
  - Directly to the ISDN network over two B channels (2 x 64 Kbps).
- Using a LIC12, a single line attaches the 3746-900 to the PRI terminal adapter at any speed up to E1 (including the various H-channel speeds). See Figure 10-10.

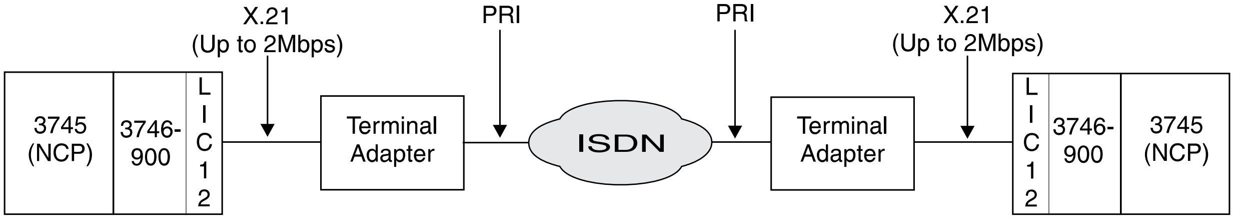

Figure 10-10. PRI-ISDN Terminal Adapter with a LIC12

The line between the LIC12 and the PRI-ISDN terminal adapter supports only clear channel traffic. The PRI-ISDN terminal adapter converts the LIC12 interface (clear channel) into n x B-channels on the PRI ISDN port. For example, if the X.21 interface runs at 1536 Kbps, n = 24.

- The terminal adapters at both ends of the network are responsible for data synchronization to keep the data arriving at one of the DTE interfaces in the same order as its partner DTE sent it. This data synchronization is accomplished through “bonding” protocols when multiple B channels are used. The two terminal adapters must, therefore, have compatible bonding protocols.

Direct Call

The terminal adapter can also be attached to the 3746-900 over a nonswitched V.35 or X.21 interface. In this case, the terminal adapter can:

- Set up a call to a remote DTE. It can do this automatically when the line is activated. This *direct call* function uses a called number that is predefined in the terminal adapter.
- Automatically answer a call from a remote DTE without involving the 3746-900.

The line is defined as a nonswitched line in NCP by coding `DIAL=NO` (the default value) in the GROUP definition statement.

Chapter 11. ESCON Overview

Enterprise System Connectivity (ESCON) fiber optic channels can be installed in either the 3746-9x0 base or expansion frames, or as an adapter in the 3746 Multiaccess Enclosure (MAE). These two options use different non-interchangeable hardware adapters. This chapter gives an overview of these options.

Further details can be found in Chapter 15, “ESCON Adapters” on page 15-1, Chapter 22, “ESCON Channel Adapter” on page 22-1, and Chapter 28, “Configuring the MAE ESCON Channel Adapter” on page 28-1.

This chapter assumes that you understand ESCON architecture and the generation processes for IOCP (and for NCP, if you plan to share base 3746-900 ESCON adapters between the 3746 Nways Multiprotocol Controller control points and NCP running in the 3745).

Note: MAE ESCON adapters are not supported by NCP.

You should also know TCP/IP if you plan to use ESCON adapters for the 3746 IP support or for NCP IP support running in the 3745 CCU.

A 3746-950 or 3746-900 can house up to 16 ESCON Channel Adapters (ESCAs), which allow a 3746 Nways Multiprotocol Controller to communicate with PU type 2.1 nodes, such as VTAMs and TPFs (Transaction Processing Facility), and TCP/IP MVS hosts (PU type 1). The MAE supports up to four ESCON adapters.

The 3746 Network Node supports parallel APPN/HPR Transmission Groups (TGs) over ESCAs in the same way as NCP does for the 3745 parallel channel adapters type 6 (CADS) and type 7 (BCCA), and for the 3746-900 ESCAs.

In a 3746-900, the ESCA also allows NCP to communicate with TPFs (PU type 2.1), VTAMs (PU types 5 and 2.1), and TCP/IP MVS hosts (PU type 1).

The MAE's 1-port ESCA (FC3287) supports three channel access protocols:

- Multi-Path Channel+ (MPC+)
- LAN Channel Station (LCS)
- Link Services Architecture (LSA)

The ESCA of the MAE provides access to SNA-based host applications (such as CICS and DB2) or TCP/IP-based host applications (such as Web Server and FTP). Each ESCON Channel Adapter provides access to Multiple Image Facility (EMIF)-capable hosts or Logical Partitions (LPAR), up to a maximum of either 16 for IP or HPR traffic, or 32 for APPN traffic. It will also support the traffic on the new ATM, TR and Ethernet adapters found in the MAE.

The base frame 9x0 ESCON adapters are more effective for SNA traffic, while the MAE adapters provide a better solution for IP and APPN traffic. In a similar fashion the path length is longer when traffic from MAE adapters is defined to flow over hardware in the base 9x0 frame.

Additional information on planning for ESCON cable installations can be found in the publication, *Fiber Optic Link Planning*, GA23-0367.

Summary of 3745 and 3746 Channel Options

The following table gives an overview of the channel options supported by the 3745 and 3746.

Table 11-1. Summary of 3745 and 3746 Channel Options

	3745/NCP	3746-900	3746-950	3746-MAE 1
Servers	Multiple S/390	Multiple S/390, RS/6000	Multiple S/390, RS/6000	Multiple S/390, RS/6000
Host Operating Systems	MVS, VM, VSE	MVS, VM, VSE, AIX	MVS, VM, VSE, AIX	MVS, VM, VSE, AIX
Channel Adapters	16 Parallel	16 ESCON	16 ESCON	4 ESCON
Protocols	CDLC	CDLC	CDLC	MPC+, LSA, LCS
LPARs	16	256	256	128
SNA Routing	ISR/HPR	ISR/HPR	ISR/HPR	ISR/HPR
IP Routing	RIP	RIP	RIP	RIP

Notes:

1. 1 These are in addition to the 3745 and 3746 limits

ESCON Fiber and Host Link Sharing

Figure 11-1 and Figure 11-2 on page 11-4 illustrate fiber and host link sharing.

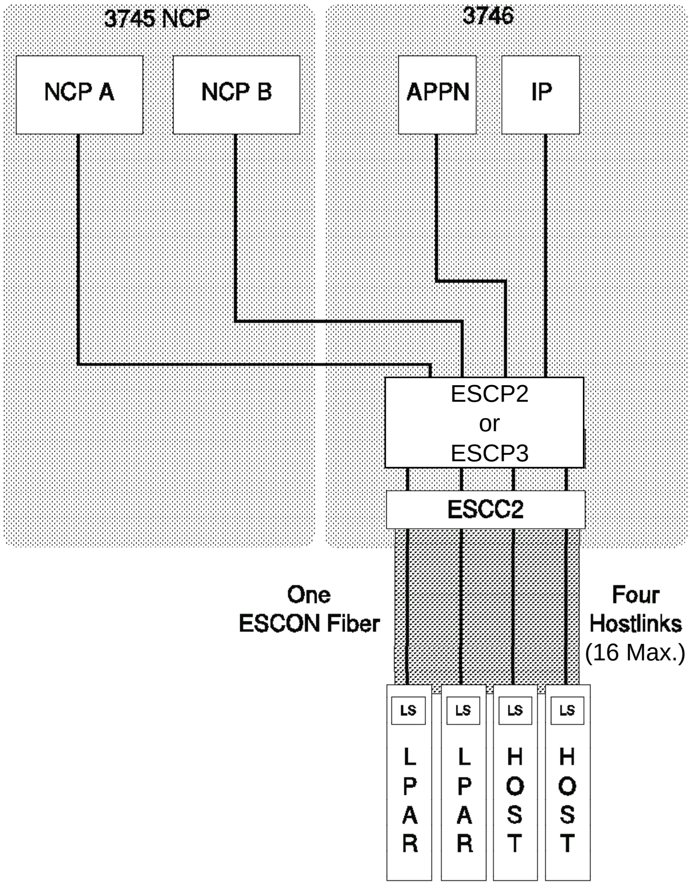

| *Figure 11-1. Sharing at the Fiber Level*

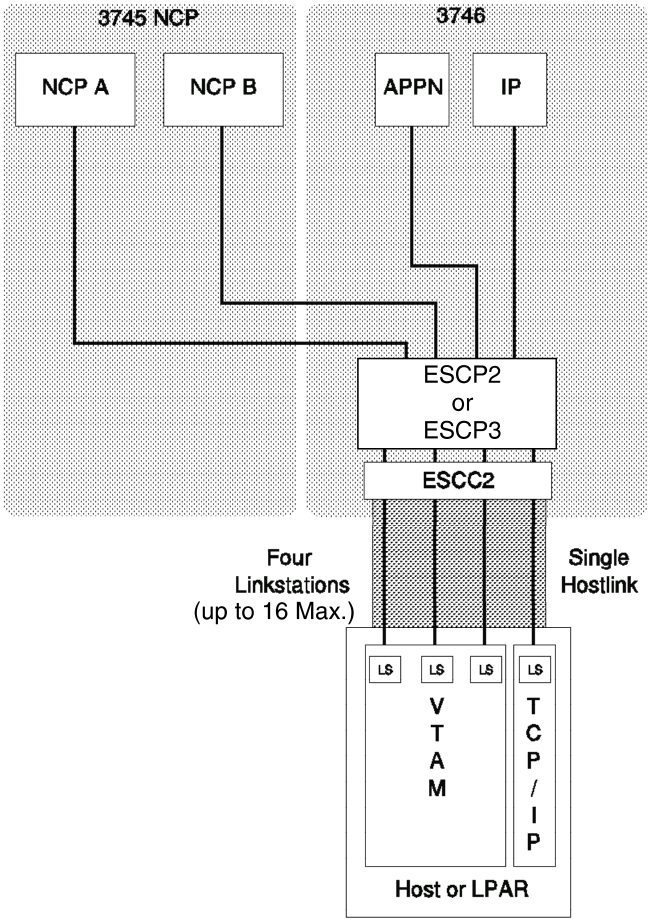

Figure 11-2. Sharing at the Host Link Level

3746 Base Frame ESCON Hardware

The ESCA consists of two features (see Figure 14-3 on page 14-6 and Figure 14-4 on page 14-7):

1. **ESCON channel processor type 3** (ESCP3 - feature number 5523), which provides the channel data link control. The ESCP3 is an enhancement of ESCP2 and provides higher processing power, more storage. ESCP3 supports 3746 IP router and APPN/HPR network node functions.

2. **ESCON channel coupler type 1** (ESCC - feature number 5501), or the **ESCON channel coupler type 2** (ESCC2 - feature number 5502), which contains the interface to the ESCON multimode, duplex fiber optic channel cable. The ESCC2 is recommended, as it provides the same functions as the ESCC but with increased performance in throughput and channel utilization. Any ESCC is field upgradable to an ESCC2. The ESCC is not available with the 3746-950.

   There is one ESCON channel coupler per ESCON channel processor.

ESCA: Distances

The ESCON channel couplers support the standard ESCON fiber distance (3 km). This is equally true for the ESCON adapter in the MAE. They do not themselves support the ESCON eXtended Distance Feature (XDF). However, longer distances can be reached via an ESCD using the ESCON XDF. An S/390 can be reached up to 23 km away or, via two cascaded ESCDs each with the XDF, up to 43 km away.

Table 11-2 gives the maximum 3746-9x0 to S/390 distance for various ESCON configurations. It is valid for ESCON adapters in the MAE. Also see Figure 11-3.

Table 11-2. Maximum 3746-9x0 to S/390 Distances

Extended Distance Links	**Direct Host Connection** km (miles)	**One ESCON Director** km (miles)	**Two Cascaded ESCON Directors** km (miles)
0	3 (1.8)	6 (3.7)	9 (5.5)
1	-	23 (14.3)	26 (16.1)
2	-	-	43 (26.7)

Figure 11-3 illustrates the maximum 3746-9x0-to-S/390 distances.

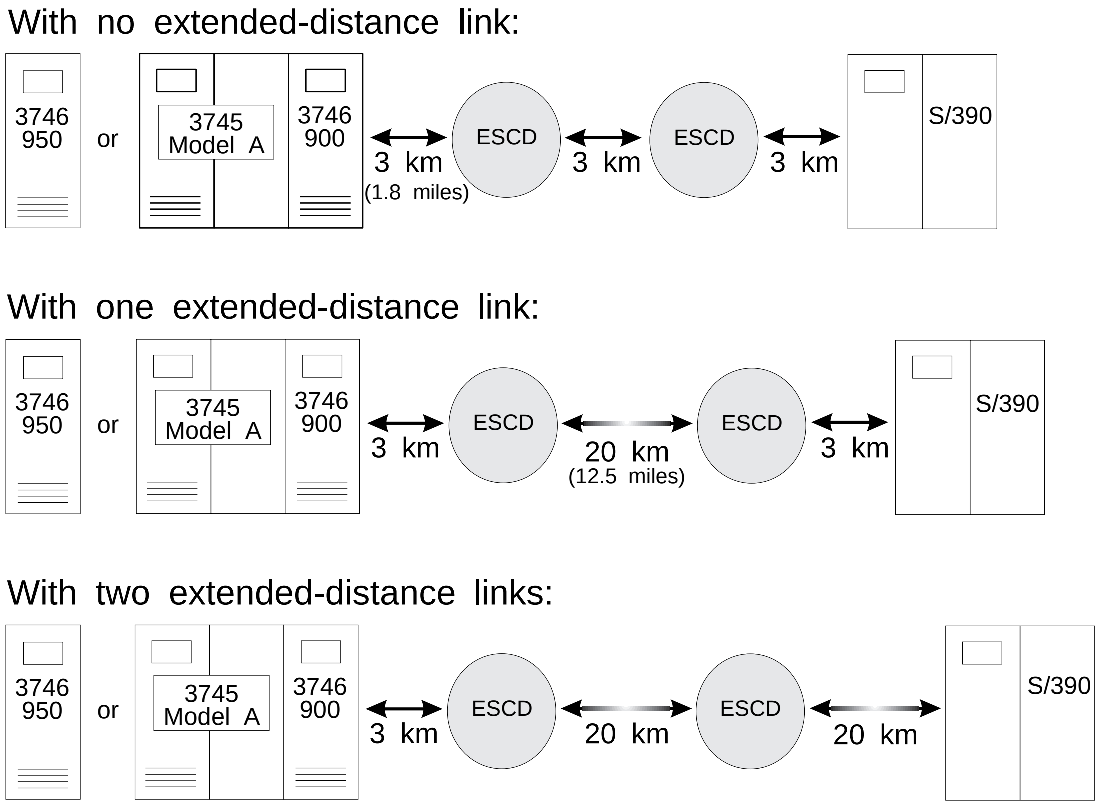

Figure 11-3. Extended ESCON Support. This figure shows the different possible distances between the 3746-9x0 and the S/390 through the ESCON support.

ESCA Sharing

The following section shows how the 3746, 3745 with NCP, and the MAE can share ESCAs, either in the base 3746 frame or in the MAE.

Note: The diagrams are meant to show how the routing occurs, but do not necessarily show all components involved. Depending upon whether ISR or HPR is used, and whether inter-adapter or intra-adapter is used, the session connectors or NCL layers will be in different locations. The same applies for IP routing.

3746 ESCA Sharing

Figure 11-4 shows which traffic can use the 3746 native ESCON adapters. The APPN and IP components are shown as being connected together. In the 06/30/97 release of the MAE, this connection will involve using an external token-ring connections between the 3746 frame and the MAE. In a later release, this will be replaced by an internal hardware connection.

The numbers in the following list correspond to the numbers in the diagram:.

1 Subarea traffic, or APPN traffic (when the 3745/NCP is part of a CNN), can directly access the 3746 ESCON adapters. This traffic can enter the 3745 via its adapters, or the 3746 adapters.

2 APPN, DLUR, or IP traffic, entering the 3746 through its adapters, can be routed to the 3746 ESCON adapters.

3 APPN, DLUR, or IP traffic, entering the 3746 through MAE adapters, is routed by the APPN or IP routing functions to the APPN or IP function of the 3746 base, where it is then passed to the ESCON adapters.

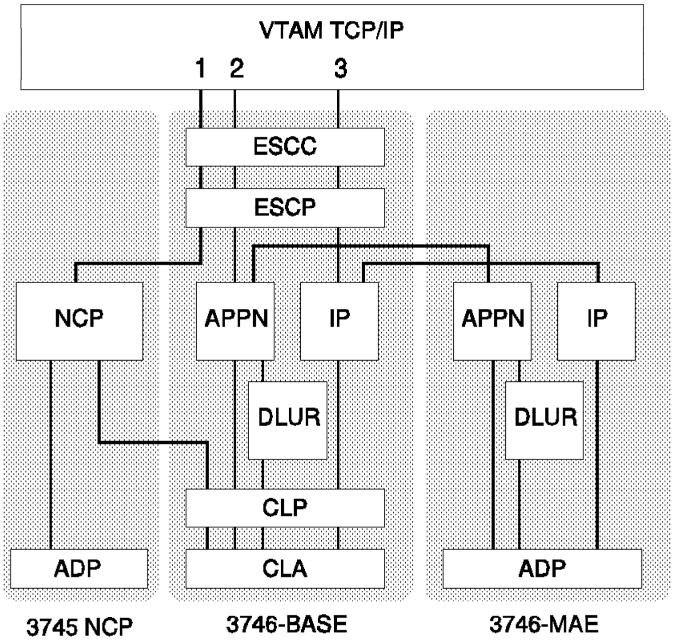

Figure 11-4. 3746 Base ESCON Connectivity

Subarea traffic entering the MAE through LAN or ATM adapters, can be routed to the 3746 base ESCON adapters. The following methods are available:

1. Figure 11-5 shows traffic entering the MAE (3), which is then bridged to the 3745 or 3746. This traffic could either be bridged over the dedicated token-ring connections between the MAE and 3746 frame, or over a different token-ring, either to a 3746 (4), or to a 3745 (5) token-ring adapter. The traffic is then routed by the APPN (3746 NNP or NCP/CNN) or IP functions as in Figure 11-4 on page 11-6
2. DLSw traffic entering the MAE (3) can be used to transport subarea BNN traffic, this traffic can then be directed to token-ring connections on the MAE. As with the previous example, this can be the dedicated ring, or a different ring connecting a 3745 or the 3746 to its MAE.

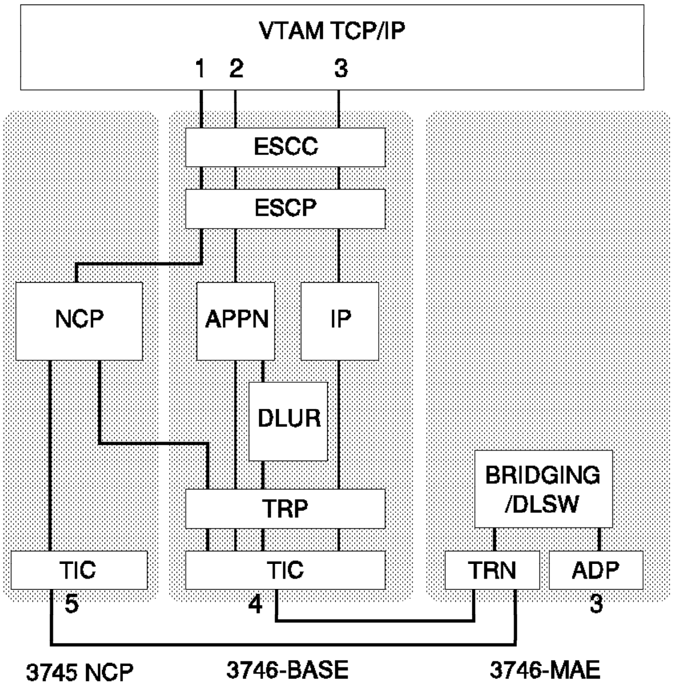

Figure 11-5. MAE to 3746 Bridging and DLSw for Subarea Traffic

MAE ESCA Sharing

Figure 11-6 shows which traffic can use the 3746 MAE ESCON adapters. The numbers in the following list correspond to the numbers in the diagram:

1 3745 APPN traffic (when the 3745/NCP is part of a CNN), entering either through 3745 or 3746 adapters, 3746 APPN and DLUR traffic, and MAE APPN traffic can be routed by the APPN functions to the MAE ESCON adapters.

APPN traffic is routed through either an internal or an external APPN link between the 3745 and the 3746 frame.

2 3746 IP traffic, and MAE IP traffic is routed by the IP routing functions to the MAE ESCON adapters.

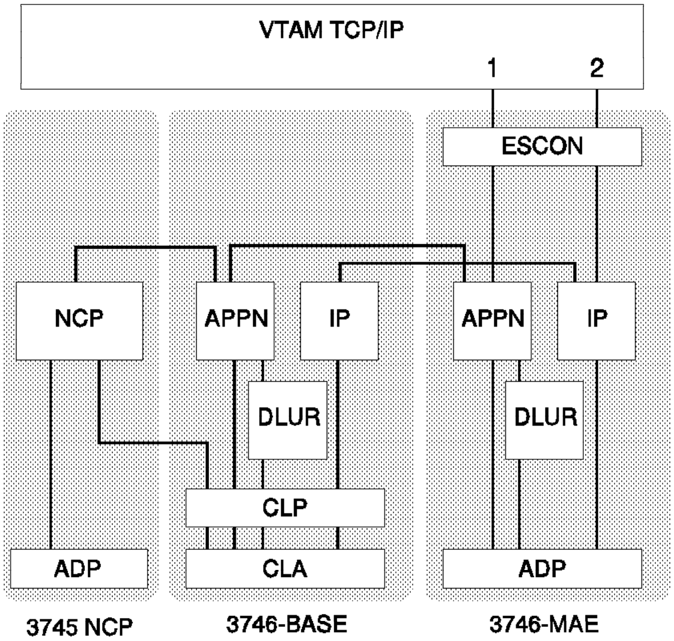

Figure 11-6. 3746 MAE ESCON Connectivity

3745 Parallel Channel Adapter Sharing

To complete the section on sharing of ESCAs, Figure 11-7 shows which traffic can use the 3745 parallel channel adapters. The numbers in the following list correspond to the numbers in the diagram:.

1 3745 APPN traffic (when the 3745/NCP is part of a CNN), entering either through 3745 or 3746 adapters, and 3746 and MAE APPN and DLUR traffic can be routed to the 3745 parallel channel adapters.

APPN traffic is routed through either an internal or an external APPN link between the 3745 and the 3746 frame.

2 3746 IP traffic, and MAE IP traffic is routed by the IP routing functions via an external connection to a 3745 adapter where the NCP IP function routes it to the 3745 parallel channel adapters.

As in Figure 11-5 on page 11-7, traffic can also be bridged or send from DLSw to the 3745 via token-ring, this traffic can also then be transported over the 3745 parallel channels.

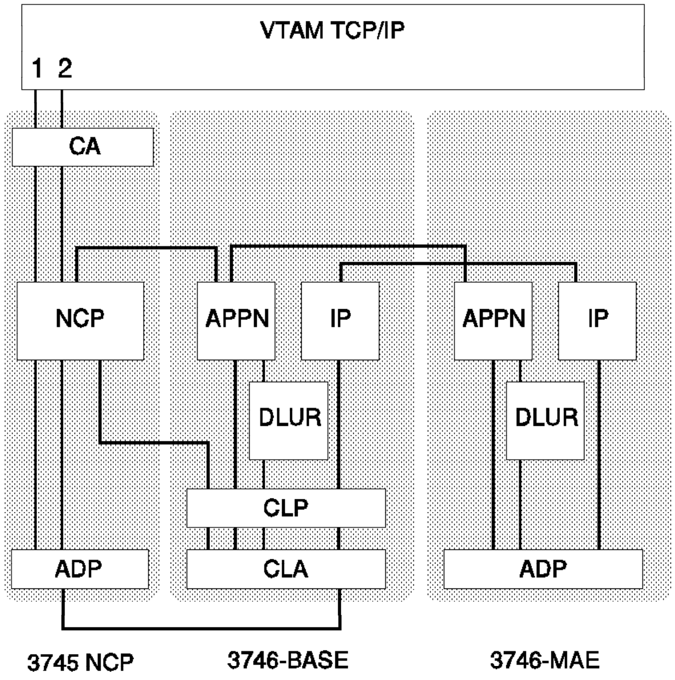

Figure 11-7. 3746 Parallel Channel Connectivity

Chapter 12. 3745 and 3746 Installation and Upgrade Scenarios

The possible network migration paths for the 3746 Nways Multiprotocol Controller Models 900 and 950 are given in the *3745 Communication Controller Models A, 3746 Expansion Unit Model 900, 3746 Nways Multiprotocol Controller Model 950: Overview*, GA33-0180.

This chapter describes possible hardware and microcode installation scenarios that may be used to implement a 3746 Nways Controller. The scenario best suited to your needs depends on your existing hardware and on the configuration to be installed. The scenarios are grouped as follows:

1. Installation of a new 3746-950:
   - APPN/HPR
   - APPN/HPR and IP
2. Upgrade of an existing 3746-900 (subarea or composite network node) to a 3746 Nways Multiprotocol Controller Model 900 operating as a network node and/or IP router.
3. Installation of a new 3746 Nways Multiprotocol Controller Model 900 shared between NCP and the 3746 network node and the IP router.
4. Upgrade of a 3746-900 network node/IP router to a 3746-950.

Note: Scenarios 1, 2, and 3 require the installation of the Controller Expansion (Feature 5023). One or two controller expansions can be installed.

Other possible scenarios (for example, going from an NCP-controlled 3746-900 to a 3746-950) can be constructed using the above examples.

The interruption times given in the scenarios:

- Assume that two IBM Customer Engineers (CEs) are doing the installation and perform some tasks in parallel to optimize the installation time.
- Do not include the time needed to restart your network. Network restart times vary considerably depending on the size of the installation (number of ports, physical units, and SNA/APPN sessions) and should be taken into consideration when deciding which scenario is best for you.
- Assume that your current service processor is used for the new 3746 Nways Multiprotocol Controller. If an additional service processor is installed, you need to add about four hours to the time for the scenario. The 3745 traffic is not interrupted during this additional time.

The scenario that best fits your needs depends on:

- The type of installation you want (a new machine or upgrading an existing one).
- Which machine (the 3745 or 3746 Nways Multiprotocol Controller) you need to have back in operation first.

IBM does not recommend modifying existing network connections during these operations. Changes not related to these operations may increase the duration of the 3745 and/or 3746 interruption and testing. Configuration changes not dependent on the 3746 Nways Multiprotocol Controller features should be scheduled at a different time.

Type 1 Scenario: Installation of a New 3746-950

Installation of a new 3746-950 takes the IBM CEs about ten hours.

Notes:

1. If a service processor is already available at the installation site, it takes about seven hours.
2. If a second network node processor (B) is installed, this will add about half an hour.
3. If Ethernet bridges are installed (3746 only) this will add about one hour for installation and customization *for each* Ethernet bridge.

Step	Performed By	Description
1	User	Refer to “Your Task Responsibilities as a Customer” on page xlviii: 1. Physical planning. 2. Software definitions and tuning. 3. Filling out plugging sheets, if necessary. 4. NetView definitions for VTAM and the MOSS-E. 5. Controller, service processor, and network node processor definitions (including the second network node processor, if any). 6. Remote console definitions, if necessary.
2	IBM CE	1. Install the 3746-950. 2. Install the controller expansion. 3. Install the second controller expansion if two have been ordered.
3	IBM CE	1. Install the service processor and microcode (licensed internal code). 2. Customize the service processor.
4	IBM CE	Upgrade the microcode (see note).
5	IBM CE	1. Install the network node processor A. 2. Customize the network node processor A. If applicable: 3. Install the network node processor B. 4. Customize the network node processor B.
6	IBM CE	Install Ethernet bridges, if any.
7	IBM CE	Test the 3746-950.
8	User	Network startup (3746-950).

Note: If you are an RSF user, an analog telephone line must be installed before the controller installation to ensure that the RSF link can be tested at the time of installation and that the latest microcode changes are loaded via this connection during installation.

Type 2 Scenarios: Installation of the Network Node Processor on your 3746-900

The time required for this procedure depends on the type of service processor system unit installed on the 3745:

1. Stand-alone (tower) or rack-mountable service processor
2. Desktop service processor. The desktop unit requires additional time because it has to be upgraded.

Scenario 2.1: Network Node Processor Installation Using the Stand-Alone (Tower) or Rack-Mountable Service Processor

The steps performed by the IBM CEs take about 5.5 hours.

Notes:

1. If a second network node processor (B) is installed, this will add about half an hour.
2. If Ethernet bridges are installed (3746 only), this will add about one hour for installation and customization *for each* Ethernet bridge.

The 3746-900 is not operational during this time.

Step	Performed By	Description	3745 Interruption	3746-900 Interruption
1	User	Refer to "Your Task Responsibilities as a Customer" on page xlviii: 1. Physical planning. 2. Software definitions and tuning. 3. Filling out plugging sheets, if necessary. 4. NetView definitions for VTAM and the MOSS-E. 5. Controller, service processor, and network node processor definitions (including the second network node processor, if any). 6. Remote console definitions, if necessary.	No	No
2	IBM CE	1. Install the controller expansion 2. Install the network node processor.	No[1 page 12-4]	Yes
3	IBM CE	1. Exchange type 1 processors for type 2 processors (see note 2 on page 12-4). 2. Customize the service processor.	No	Yes
4	IBM CE	Upgrade the microcode (licensed internal code).	No	Yes
5	IBM CE	1. Copy control point microcode onto the network node processor A. 2. Customize the network node processor A. If applicable: 3. Copy control point microcode onto the network node processor B. 4. Customize the network node processor B.	No	Yes
6	IBM CE	Install Ethernet bridges, if any.	—	Yes
7	IBM CE	Test the 3746-900.	No	Yes

Step	Performed By	Description	3745 Interruption	3746-900 Interruption
8	User	Network startup (3746-900).	No	Yes

Notes:

1. Traffic that does not depend on the 3746-900 is not interrupted during the rest of the scenario. Of course, any 3746-900 traffic is halted until the 3746-900 is put back into operation.
2. The time for this scenario assumes that five processors are upgraded to type 2. A processor takes about 15 minutes to upgrade, so changing more or less than five processors may change the total time needed for the installation.

Scenario 2.2: Network Node Processor Installation Using the Desktop Service Processor

The steps performed by the IBM CEs take about 7.5 hours. The 3746-900 is not operational during this time.

Notes:

1. If a second network node processor (B) is installed, this will add about half an hour.
2. If Ethernet bridges are installed (3746 only), this will add about one hour for installation and customization *for each* Ethernet bridge.

Step	**Performed By**	**Description**	**3745 Interruption**	**3746-900 Interruption**
1	User	Refer to "Your Task Responsibilities as a Customer" on page xlviii: 1. Physical planning. 2. Software definitions and tuning. 3. Filling out plugging sheets, if necessary. 4. NetView definitions for VTAM and the MOSS-E 5. Controller, service processor, and network node processor definitions (including the second network node processor, if any). 6. Remote console definitions, if necessary.	No	No
2	IBM CE	1. Backing up the hard drive. 2. Installing a larger hard drive. 3. Restoring the hard drive.	No[1]	Yes
3	IBM CE	1. Install the controller expansion. 2. Install the network node processor.	No	Yes
4	IBM CE	1. Exchange type 1 processors for type 2 processors (see note). 2. Customize the service processor.	No	Yes
5	IBM CE	Upgrade the microcode (licensed internal code).	No	Yes
6	IBM CE	1. Copy control point microcode onto the network node processor A. 2. Customize the network node processor A. If applicable: 3. Copy control point microcode onto the network node processor B. 4. Customize the network node processor B.	No	Yes
7	IBM CE	Install Ethernet bridges, if any.	—	Yes
8	IBM CE	Test the 3746-900.	No	Yes
9	User	Network startup (3746-900).	No	Yes

Notes:

1. Traffic that does not depend on the 3746-900 is not interrupted during the rest of the scenario. Of course, any 3746-900 traffic is halted until the 3746-900 is put back into operation.
2. The time for this scenario assumes that five processors are upgraded to type 2. A processor takes about 15 minutes to upgrade, so changing more or less than five processors may change the total time needed for the installation.

Type 3 Scenario: 3745 X10 Migration to a Model X1A plus a 3746 Nways Multiprotocol Controller Model 900

In this scenario:

1. A 3745 Model 210, 310, 410, or 610 is migrated to a Model 21A, 31A, 41A, or 61A and a 3746 Nways Multiprotocol Controller Model 900 is added to the 3745 Model A.
2. There is only one 3745 interruption for hardware installation.
3. The average 3745 down-time is about four hours. However, this figure may vary depending on the controller configuration.

Notes:

1. If a second network node processor (B) is installed, this will add about half an hour.
2. If Ethernet bridges are installed (3746 only), this will add about one hour for installation and customization ***for each*** Ethernet bridge.

Step	Performed By	Description	3745 Interruption
1	User	For NCP resources: 1. Generate a load module, called “NCP1”, with NCP V6 R2[1] (with the 3746-900 Feature). **Do not include** the new 3746-900 resources. 2. Generate a second load module, called “NCP2“, with NCP V6 R2[1] (with the 3746-900 Feature) **including** the new 3746-900 resources. 3. Load “NCP1” on the 3745. 4. Test current network configuration[1].	Yes
2	User	For 3746 Nways Multiprotocol Controller resources (APPN/HPR, IP): 1. Use the Controller Configuration and Management to generate the configuration for the planned 3746 Nways Multiprotocol Controller resources. 2. Prepare the 3746 Nways Multiprotocol Controller configuration diskette for the IBM CE.	No
3	IBM CE	1. Install service processor. 2. Install 3746-900 hardware. 3. Perform the pre-installation setup procedures.	No
4	IBM CE	1. Install the controller expansion. 2. Install network node processor A. 3. Customize network node processor A. If applicable: 4. Install network node processor B. 5. Customize network node processor B.	No
5	IBM CE	Install Ethernet bridges, if any.	Yes
6	IBM CE	1. Install 3745 Model A hardware. 2. Connect 3746-900 to 3745.	Yes (2 hours)

Step	Performed By	Description	3745 Interruption
7	IBM CE	1. Connect service processor to 3745 and 3746-900. 2. Test RSF link to IBM support center[3]. 3. Run diagnostics.	Yes (2 hours)
8	User	1. Load “NCP2” on 3745 Model A. 2. Activate and test the 3746 Nways Multiprotocol Controller and NCP resources[1,2].	Yes

Notes:

1. Minimum NCP level. The lengths of the testing and stabilization phases are user-defined, but they need to be defined before the installation starts.
2. During this period of testing and stabilization you can perform operator training.
3. If you are an IBM RSF user, an analog telephone line must be installed before the controller installation to ensure that the RSF link can be tested at the time of installation and that the latest microcode changes can be loaded via this connection during installation.

Type 4 Scenarios: Upgrade of a 3746 Nways Multiprotocol Controller Model 900 to a 3746-950

There are two different ways of doing this migration, depending on your network priority. You can choose one of the following ways to migrate:

1. With the 3745 back into operation as soon as possible.
2. With the 3746-950 operational as soon as possible.

The following procedures assume that there are no changes to the communication processors and couplers.

Scenario 4.1: 3745 Back in Operation as soon as Possible

The scenario migrates a 3746-900 with the network node processor already installed to a 3746-950.

The 3745 is operational after about three hours.

The 3746-950 is operational after about six hours.

Step	Performed By	Description	3745 Interruption	3746-900 Interruption
1	User	Refer to “Your Task Responsibilities as a Customer” on page xlviii and use the CCM as necessary: 1. Physical planning. 2. Software definitions and tuning. 3. Filling out plugging sheets, if necessary. 4. NetView definitions for VTAM and the MOSS-E 5. Controller, service processor, and network node processor definitions. 6. Remote console definitions, if necessary.	No	No
2	IBM CE	Physically detach the 3746-900 from the 3745.	Yes	Yes
3	IBM CE	Return the 3745 to operation: 1. Verify the CDF 2. Run diagnostics.	Yes	Yes
4	User	Network startup (3745).	Yes	Yes
5	IBM CE	Convert the 3746-900 frame into a 3746-950.	No	Yes
6	IBM CE	Upgrade the microcode (see note).	No	Yes
7	IBM CE	Customize the network node processor.	No	Yes
8	IBM CE	Test the 3746-950.	No	Yes
9	User	Network startup (3746-950).	No	Yes

Note: If you are an RSF user, an analog telephone line must be available for the controller before the installation to ensure that the RSF link can be tested at the time of installation and that the latest microcode changes can be loaded via this connection during installation.

Scenario 4.2: 3746-950 Operational as soon as Possible

The scenario migrates a 3746-900 with the network node processor already installed to a 3746-950.

The 3746-950 is operational after about 3.5 hours.

The 3745 is operational after about six hours.

Step	Performed By	Description	3745 Interruption	3746-900 Interruption
1	User	Refer to "Your Task Responsibilities as a Customer" on page xlviii: 1. Physical planning. 2. Software definitions and tuning. 3. Filling out plugging sheets, if necessary. 4. NetView definitions for VTAM and the MOSS-E 5. Controller, service processor, and network node processor definitions. 6. Remote console definitions, if necessary.	No	No
2	IBM CE	Physically detach the 3746-900 from the 3745.	Yes	Yes
3	IBM CE	Convert the 3746-900 frame into a 3746-950.	Yes	Yes
4	IBM CE	Upgrade the microcode (see note).	Yes	Yes
5	IBM CE	Customize the network node processor.	Yes	Yes
6	IBM CE	Test the 3746-950.	Yes	Yes
7	User	Network startup (3746-950).	Yes	Yes
8	IBM CE	Return the 3745 to operation: 1. Verify the CDF. 2. Run diagnostics.	Yes	No
9	User	Network startup (3745).	Yes	No

Note: If you are an RSF user, an analog telephone line must be available for the controller before the installation to ensure that the RSF link can be tested at the time of installation and that the latest microcode changes can be loaded via this connection during installation.

Type 5 Scenarios: Migration to 3745 A Models

Three model upgrade scenarios are proposed by IBM to allow migration from the Models 130, 150, 160, 170, 210, 310, 410, or 610 to the Models A with, optionally, the addition of a 3746-900.

Details of possible migration paths between the different models are given in the *3745 Models A Overview*, GA33-0180.

The six upgrade scenarios are:

1. Migration to a Model 21A, 31A, 41A, or 61A without the addition a 3746-900.
2. Migration to a Model 21A, 31A, 41A, or 61A with the addition of a 3746-900.

3. Migration to a Model 21A, 31A, 41A, or 61A with addition of a 3746-900, as in scenario 2, but all tests are performed in a single step at the end of the migration.
4. Migration to a Model 17A without the addition a 3746-900.
5. Migration to a Model 17A with 3746-900 that has two interruptions.
6. Migration to a Model 17A with the addition of a 3746-900 as in scenario 5, but all tests are performed in a single step at the end of the migration.

Note: The times given in the scenarios do not include the time:

- Needed to restart your network. Network restart times vary considerably depending on the size of the installation and should be taken into consideration when deciding which scenario is best for you.
- That the IBM service representative is at your site performing the installation and migration. Many tasks can be performed while your network is operational.

In scenarios 1, 3, 4, and 6 there is only a single interruption of your network operation. Scenarios 2 and 5 include several interruptions of your network operation but allow testing software and hardware between each migration step. Because of the additional intermediate testing, scenario 2 or 5 is recommended over scenario 3 or 4 by IBM. You must select the scenario that best meets the availability needs of your communication network.

IBM recommends that Step 1 in the scenarios be performed very early in the installation phase to stabilize the software before the hardware arrives. This will minimize changing both the hardware and software at the same time.

Also, IBM does not recommend modifying existing network connections during this upgrade. Changes not related to this upgrade will increase the duration of the 3745 interruption and testing. Configuration changes not dependent on the new 3746-900 features should be scheduled at a different time.

Scenario 5.1: Migration to a Model 21A, 31A, 41A, or 61A without 3746-900

The following table summarizes the migration steps and indicates when a 3745 interruption is required.

In scenario 1, the average 3745 down-time for hardware installation is at least three hours. This figure may vary depending on the communication controller configuration.

Step	Performed by	Description	3745 Interruption
1	User	1. Generate a load module using NCP V6 R2[1]. 2. Load it on the 3745. 3. Test current network configuration[2].	Yes
2	IBM CE	1. Install service processor.	No
3	IBM CE	1. Install Model A hardware. 2. Connect service processor to 3745. 3. Test RSF link to IBM support center[4]. 4. Run diagnostics.	Yes (3 hours)
4	User	1. Load Step 1 load module on Model A. 2. Test network configuration[2,3].	Yes

Notes:

1. V6 R2 is the minimum NCP level that can be used. Higher releases or versions may be required, depending on the network functions implemented.
2. The lengths of the testing and stabilization phases are user-defined, but they need to be defined before the installation starts.
3. During this period of testing and stabilization you can use the MOSS-E tutorials and demonstrations for operator training.
4. An analog telephone line must be installed before the communication controller installation to ensure that the RSF link can be tested at the time of installation. Updated microcode changes are loaded via this connection during installation.

Scenario 5.2: Migration to a Model 21A, 31A, 41A, or 61A with 3746-900 (Two Interruptions)

The following table summarizes the migration steps and indicates when 3745 interruptions are required. In scenario 2, the average 3745 down-time for hardware installation is about six hours. This figure may vary depending on the communication controller configuration.

Step	Performed by	Description	3745 Interruption
1	User	1. Generate a load module, called “NCP1”, with NCP V6 R2[1] (with the 3746-900 Feature). **Do not include** the new 3746-900 resources. 2. Generate a second load module, called “NCP2”, with NCP V6 R2[1] (with the 3746-900 Feature) **including** the new 3746-900 resources. 3. Load “NCP1” on the 3745. 4. Test current network configuration[2].	Yes
2	IBM CE	1. Install service processor.	No
3	IBM CE	1. Install Model A hardware. 2. Connect service processor to 3745. 3. Test RSF link to IBM support center[4]. 4. Run diagnostics.	Yes (3 hours)
4	User	1. Load the “NCP2” on Model A. 2. Test network configuration[2].	Yes
5	IBM CE	1. Install 3746-900 hardware. 2. Connect service processor to 3746-900 and test.	No
6	IBM CE	1. Connect 3746-900 to 3745. 2. Run diagnostics.	Yes (3 hours)
7	User	1. Load “NCP2” on Model A. 2. Activate and test 3746-900 resources[2,3].	Yes

Notes:

1. V6 R2 is the minimum NCP level that can be used. Higher releases or versions may be required, depending on the network functions implemented and the hardware contained in the 3746-900. For information about programming requirements refer to the *3745 Models A: Overview*, GA33-0180.
2. The lengths of the testing and stabilization phases are user-defined, but they need to be defined before the installation starts.
3. During this period of testing and stabilization you can use the MOSS-E tutorials and demonstrations for operator training.
4. An analog telephone line must be installed before the communication controller installation to ensure that the RSF link can be tested at the time of installation. Updated microcode changes are loaded via this connection during installation.

Scenario 5.3: Migration to a Model 21A, 31A, 41A, or 61A with 3746-900 (One Interruption)

The main scenario is as follows:

1. Only one 3745 interruption for hardware installation.
2. The average 3745 down-time is about four hours However, this figure may vary depending on the communication controller configuration.

Step	Performed by	Description	3745 Interruption
1	User	1. Generate a load module, called “NCP1”, with NCP V6 R2[1] (with the 3746-900 Feature). **Do not include** the new 3746-900 resources. 2. Generate a second load module, called “NCP2“, with NCP V6 R2[1] (with the 3746-900 Feature) **including** the new 3746-900 resources. 3. Load “NCP1” on the 3745. 4. Test current network configuration[2].	Yes
2	IBM CE	1. Install service processor. 2. Install and configure the 3746-900 hardware. 3. Perform pre-installation setup procedures. 4. Run diagnostics on the 3746-900.	No
3	IBM CE	1. Install Model A hardware. 2. Connect 3746-900 to 3745.	Yes (2 hours)
4	IBM CE	1. Connect service processor to 3745 and 3746-900. 2. Test RSF link to IBM support center[4]. 3. Run diagnostics on the 3745 - 3746-900 link.	Yes (2 hours)
5	User	1. Load “NCP2” on the Model A. 2. Activate and test 3746-900 resources[2,3].	Yes

Notes:

1. V6 R2 is the minimum NCP level that can be used. Higher releases or versions may be required, depending on the network functions implemented and the hardware contained in the 3746-900. For information about programming requirements refer to the *3745 Models A: Overview*, GA33-0180.
2. The lengths of the testing and stabilization phases are user-defined, but they need to be defined before the installation starts.
3. During this period of testing and stabilization you can use the MOSS-E tutorials and demonstrations for operator training.
4. An analog telephone line must be installed before the communication controller installation to ensure that the RSF link can be tested at the time of installation. Updated microcode changes are loaded via this connection during installation.

Scenario 5.4: Migration to a Model 17A without 3746-900

The following table summarizes the migration steps and indicates when a 3745 interruption is required.

In scenario 4, the average 3745 down-time for hardware installation is six and a half hours. This figure varies depending on the communication controller configuration.

Step	Performed by	Description	3745 Interruption
1	User	1. Generate a load module using NCP V6 R3[1]. 2. Load it on the 3745. 3. Test current network configuration[2].	Yes
2	IBM CE	1. Install service processor.	No
3	IBM CE	1. Install Model A hardware. 2. Connect service processor to 3745. 3. Test RSF link to IBM support center[4]. 4. Run diagnostics.	Yes (6.5 hours)
4	User	1. Load Step 1 load module on Model A. 2. Test network configuration[2,3].	Yes

Notes:

1. V6 R3 is the minimum NCP level that can be used. Higher releases or versions may be required, depending on the network functions implemented.
2. The lengths of the testing and stabilization phases are user-defined, but they need to be defined before the installation starts.
3. During this period of testing and stabilization you can use the MOSS-E tutorials and demonstrations for operator training.
4. An analog telephone line must be installed before the communication controller installation to ensure that the RSF link can be tested at the time of installation. Updated microcode changes are loaded via this connection during installation.

Scenario 5.5: Migration to a Model 17A with 3746-900 (Two Interruptions)

In scenario 5, the average 3745 down-time for hardware installation is about eight hours. This figure may vary depending on the communication controller configuration.

Step	Performed by	Description	3745 Interruption
1	User	1. Generate a load module, called “NCP1”, with NCP V6 R3[1] (with the 3746-900 Feature). **Do not include** the new 3746-900 resources. 2. Generate a second load module, called “NCP2”, with NCP V6 R3[1] (with the 3746-900 Feature) **including** the new 3746-900 resources. 3. Load “NCP1” on the 3745. 4. Test current network configuration[2].	Yes
2	IBM CE	1. Install service processor.	No
3	IBM CE	1. Install Model A hardware. 2. Connect service processor to 3745. 3. Test RSF link to IBM support center[4]. 4. Run diagnostics.	Yes (6.5 hours)
4	User	1. Load the “NCP2” on Model A. 2. Test network configuration[2].	Yes
5	IBM CE	1. Install 3746-900 hardware. 2. Connect service processor to 3746-900 and test.	No
6	IBM CE	1. Connect 3746-900 to 3745. 2. Run diagnostics.	Yes (1.5 hours)
7	User	1. Load the “NCP2” on the Model A. 2. Activate and test 3746-900 resources[2,3].	Yes

Notes:

1. V6 R3 is the minimum NCP level that can be used. Higher releases or versions may be required, depending on the network functions implemented and the hardware contained in the 3746-900. For information about programming requirements refer to the *3745 Models A: Overview*, GA33-0180.
2. The lengths of the testing and stabilization phases are user-defined, but they need to be defined before the installation starts.
3. During this period of testing and stabilization you can use the MOSS-E tutorials and demonstrations for operator training.
4. An analog telephone line must be installed before the communication controller installation to ensure that the RSF link can be tested at the time of installation. Updated microcode changes are loaded via this connection during installation.

Scenario 5.6: Migration to a Model 17A with 3746-900 (One Interruption

The main changes in scenario 6 from scenario 5 are:

1. Only one 3745 interruption for hardware installation instead of two.
2. The average 3745 down-time is about seven hours instead of eight. However, this figure may vary depending on the communication controller configuration.
3. All installation tests are performed at the end of the migration scenario.

Step	Performed by	Description	3745 Interruption
1	User	1. Generate a load module, called "NCP1", with NCP V6 R3[1] (with the 3746-900 Feature). **Do not include** the new 3746-900 resources. 2. Generate a second load module, called "NCP2", with NCP V6 R3[1] (with the 3746-900 Feature) **including** the new 3746-900 resources. 3. Load "NCP1" on the 3745. 4. Test current network configuration[2].	Yes
2	IBM CE	1. Install service processor. 2. Install 3746-900 hardware. 3. Perform pre-installation setup procedures.	No
3	IBM CE	1. Install Model A hardware. 2. Connect 3746-900 to 3745.	Yes (6 hours)
4	IBM CE	1. Connect service processor to 3745 and 3746-900. 2. Test RSF link to IBM support center[4]. 3. Run diagnostics.	Yes (1 hour)
5	User	1. Load "NCP2" on the Model A. 2. Activate and test 3746-900 resources[2,3].	Yes

Notes:

1. V6 R3 is the minimum NCP level that can be used. Higher releases or versions may be required, depending on the network functions implemented and the hardware contained in the 3746-900. For information about programming requirements refer to the *3745 Models A: Overview*, GA33-0180.
2. The lengths of the testing and stabilization phases are user-defined, but they need to be defined before the installation starts.
3. During this period of testing and stabilization you can use the MOSS-E tutorials and demonstrations for operator training.
4. An analog telephone line must be installed before the communication controller installation to ensure that the RSF link can be tested at the time of installation. Updated microcode changes are loaded via this connection during installation.

Chapter 13. Configuration Scenarios

3746 Multiaccess Enclosure: Availability and Backup

In an environment where high availability is required, redundant 3746 Multiaccess Enclosures can be used for backup. Various backup scenarios will be examined on the following pages. They are not meant to be exhaustive but should provide for discussion of typical recovery configurations. They may require redundant 3746 Multiaccess Enclosure and ESCON adapters.

LCS and LSA Examples

To recover from host failures, a secondary host and a path to that host are required. In Figure 13-1, if Host A fails, users have access to Host B through 3746-MAE2. When Host A recovers, user access to Host A can be restored. This is just one example of recovering from a host failure. For host failures, other factors should also be considered, such as availability of applications and data on the secondary host.

To recover from 3746/MAE failures, a second 3746/MAE must be in the path to your hosts. In Figure 13-1, when 3746-MAE1 fails, users of Host A could be rerouted through 3746-MAE2, because there are two paths to Host A. When the 3746-MAE1 recovers, users of Host A could be rerouted back through 3746-MAE1. There is only one path to Host B; and if the 3736-MAE2 fails, users will not be able to access Host B. This could be resolved by connecting and configuring Host B to the 3746-MAE1. Then when 3746-MAE2 recovers, user access to Host B could be restored. Again, other variations of this example could be employed to recover from 3746 Multiaccess Enclosure failure.

The backup configuration in Figure 13-1 on page 13-2 is designed as follows:

- Host A has access to LAN Segment A through 3746-MAE1. It has access to LAN Segment B through either the 3746-MAE2 or 3746-MAE1.
- Host B has access to LAN Segment A and to LAN Segment B through 3746-MAE2.
- Both of the 3746/MAEs have the same node address. The node addresses can be different for TCP/IP, but not for SNA.
- The IP address for both the 3746/MAEs is the same.
- LAN Segment A and LAN Segment B are connected by a bridge or router.

  For TCP/IP, a router cannot be used. Using a router would, by design, require the IP addresses for 3746-MAE1 and 3746-MAE2 to be different. The IP addresses for the 3746/MAEs must be the same so that the LAN devices can rediscover the route to the secondary 3746/MAE.

Figure 13-1 on page 13-2 and Figure 13-2 on page 13-3 show sample token-ring LAN segments, but the same principles apply to all the LAN types supported by the 3746/MAE.

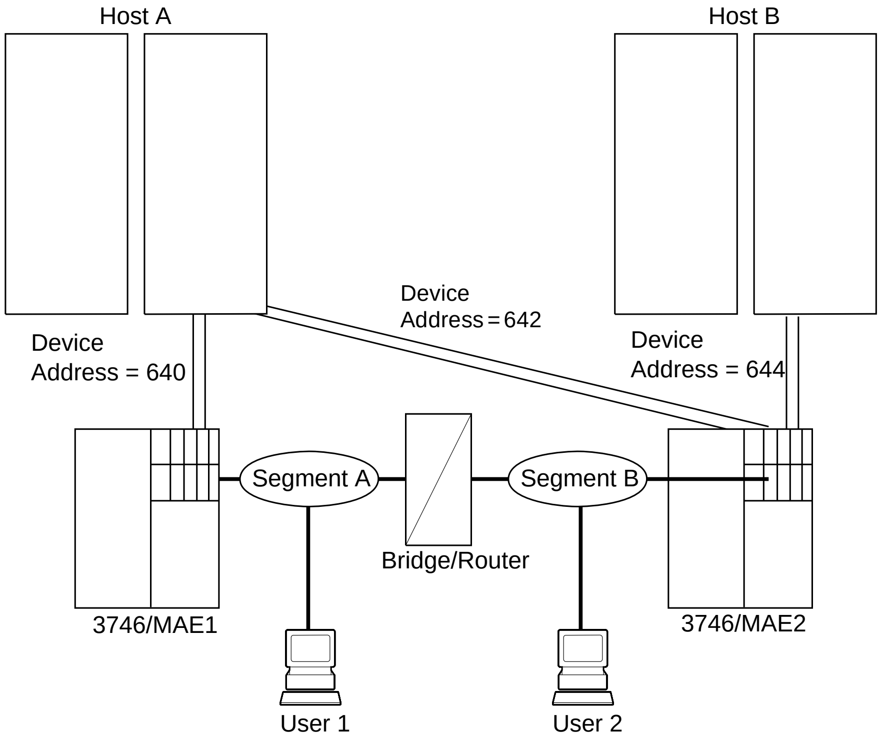

Figure 13-1. 3746/MAE and Host Backup Using Two LAN Segments

Figure 13-2 provides the same host and 3746/MAE backup as Figure 13-1 on page 13-2, except:

- 3746-MAE1 and 3746-MAE2 are connected to the same LAN segment.
- The primary 3746/MAE must finish IPL before the secondary 3746/MAE is restarted. Then the secondary LAN adapter will not become active on the LAN, because the node addresses are the same.

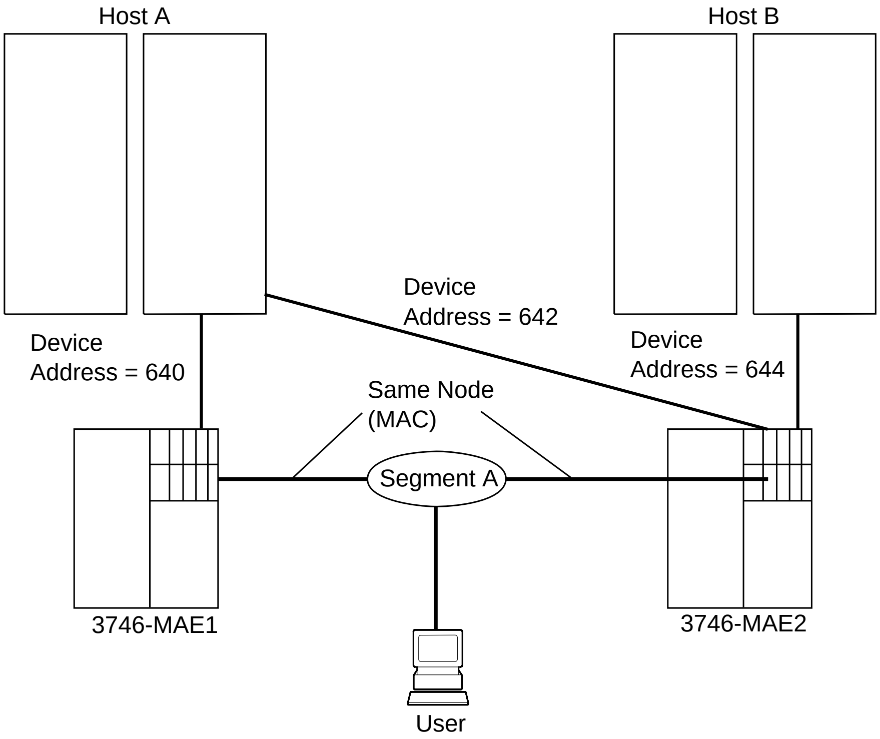

Figure 13-2. 3746/MAE and Host Backup Using One LAN Segment

Increasing Availability in VTAM Environments

To demonstrate backup in the VTAM environment using the system in Figure 13-1 on page 13-2, assume that both 3746/MAEs and hosts are operational and that 3746-MAE1 and Host A are being used as the primary system. Only the XCA major node for 3746-MAE1 is active. Refer to “VTAM Control Blocks Used to Configure LSA at the Host” on page 22-27 for more information.

When 3746-MAE1 fails:

1. From Host A, deactivate the XCA major node for 3746-MAE1.
2. From Host A, activate the XCA major node for 3746-MAE2.
3. Host A users experience a temporary session loss, but connect through 3746-MAE2 using the same node address.

When Host A fails:

1. 3746-MAE1 detects the inactive host.
2. From Host B, activate the XCA major node for 3746-MAE2.
3. Host A users experience a temporary session loss, but connect to Host B through the 3746-MAE2 using the same node address.

In both cases, after the failure is corrected, sessions can be moved from 3746-MAE2 to 3746-MAE1 in a similar manner.

To demonstrate backup in the VTAM environment using the system in Figure 13-2 on page 13-3, assume that both hosts and both 3746/Multiaccess Enclosures are operational, and that Host A and 3746-MAE1 are being used as the primary system.

When 3746-MAE1 fails:

1. From 3746-MAE1, issue a 2216 command to remove 3746-MAE1 connection from the LAN (stop the LAN-A adapter). This can be done from the service processor.
2. From 3746-MAE2, issue a 2216 command to attach the 3746-MAE2 adapter to the LAN (start the LAN-B adapter) This would be from its service processor.
3. From Host A, activate the XCA major node for 3746-MAE2.
4. Host A users experience a temporary session loss, but reconnect through 3746-MAE2 using the same node address.

When Host A fails:

1. From the 3746-MAE1, issue a 2216 command to remove 3746-MAE2 from the LAN (stop the LAN-A adapter). This would be from its SP.
2. From 3746-MAE2, issue a 2216 command to attach 3746-MAE2 to the LAN (start the LAN-B adapter). This would be from its SP.
3. From Host B, activate the XCA major node for the 3746-MAE2.
4. Host A users experience a temporary session loss, but connect to Host B through 3746-MAE2 using the same node address.

In both cases, after the failure is corrected, sessions can be moved from 3746-MAE2 to 3746-MAE1 in a similar manner.

Increasing Availability in a TCP/IP Environment

To demonstrate backup in a TCP/IP environment using the system in Figure 13-1 on page 13-2, assume that both 3746/MAEs and hosts are operational and that 3746-MAE1 and Host A are being used as the primary system.

The TCP/IP Profile data set in the host includes the following statements:

Note: The default name for the TCP/IP profile data set is TCPIP.PROFILE.TCPIP for MVS and PROFILE TCPIP for VM.

- DEVICE statements for both 3746/MAEs attached to Host A (subchannels 640 and 642)
- LINK statements for the LAN adapters in both 3746/MAEs
- A HOME statement only for the LAN-A adapter
- A GATEWAY statement for your routing table that includes the LAN-A adapter
- A START statement for the device using the LAN-A adapter.

When 3746-MAE1 fails, the OBEYFILE command can be executed from Host A to give users access to Host A through 3746-MAE2. You will need to issue a PING command to one device *after* the obey file has been used. You will need an obey file that contains:

- A STOP statement for the LAN-A adapter

  The STOP statement deactivates *all lan* adapters associated with that device. Therefore, if another LAN adapter was defined in 3746-MAE1 for device 640, a STOP for device 640 would stop both the LAN-A adapter and the other adapter. Depending on the type of 3746/Multiaccess Enclosure failure, you may have to stop the LAN-A adapter using a 2216 command and execute the OBEYFILE command again.

- A HOME statement for the LAN-B adapter with the same IP address that is used for the LAN-A adapter
- A GATEWAY statement for your routing table that includes the LAN-B adapter
- A START statement for the device using the LAN-B adapter

A similar obey file can be used to restore user access to Host A when it recovers.

These steps can also be followed using the system in Figure 13-2 on page 13-3.

TCP/IP host backup is not illustrated in Figure 13-1 on page 13-2 or Figure 13-2 on page 13-3. Because each TCP/IP host requires its own LAN adapter, you will not be able to recover from TCP/IP host failures as you can with VTAM hosts. You can, however, add a LAN adapter to either 3746/MAE (or both if you want a backup for both hosts) and define a connection between the host and LAN adapters.

- These new adapters will not have the same node (or MAC) address as the adapters for your primary hosts (shown in Figure 13-1 on page 13-2 and Figure 13-2 on page 13-3).
- The adapters for the backup host will need to use the same IP addresses that you used for the primary hosts. You can switch the IP/LAN adapter association when one host fails in a similar manner as you do for 3746/MAE failures described above.

For more information about TCP/IP configuration, refer to “Host Definition Planning” on page 22-19.

MPC+ Examples

Figure 13-3 shows an example of backup for Multi-Path Channel+ (MPC+).

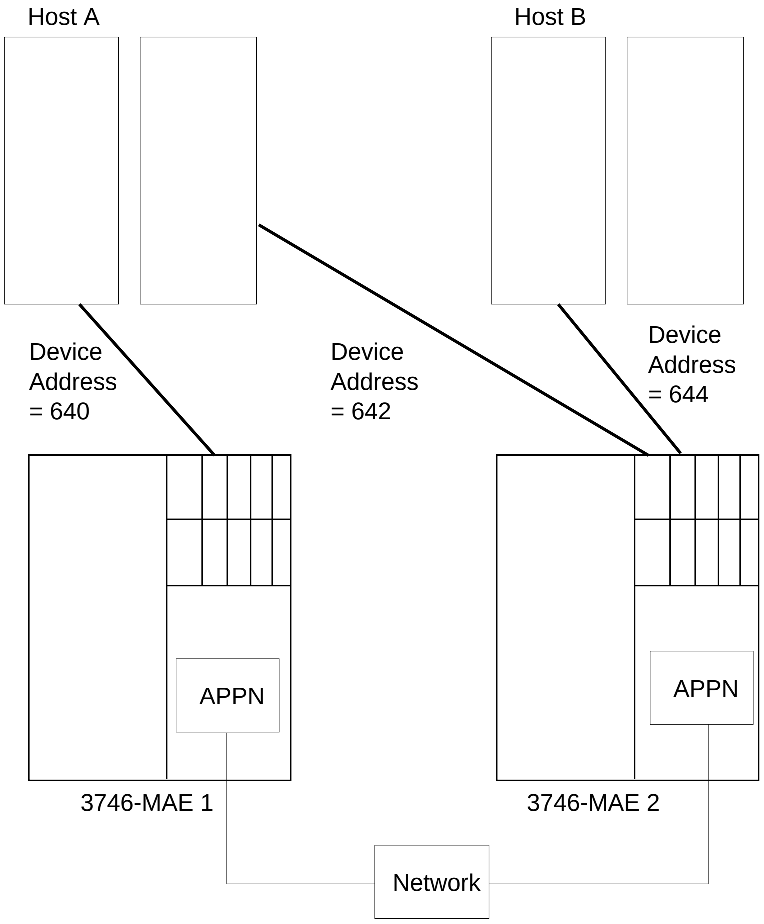

Figure 13-3. 3746/MAE and Host Backup - Multi-Path Channel+

If one of the 3746/MAEs goes down, any HPR traffic should route around that down box. It is a property of APPN HPR to route around the down box. For example, if 3746-MAE1 goes down, the APPN HPR will path switch around the down box and the session that may have been going between Host A and the network will now all go though 3746-MAE2.

To do the above you just need to make sure that 3746-MAE2 has an APPN PORT/LINK set up to both Host A and Host B.

List of Abbreviations

AB	area border
ACF	Advanced Communications Function
ACF/VTAM	Advanced Communications Function for the Virtual Telecommunications Access Method
ANR	automatic network routing
APPN	advanced peer-to-peer networking
ARB	adaptive rate-based flow/congestion control
ARC	active remote connector
ARP	address resolution protocol
AS	autonomous system
ASB	autonomous system border
ASE	autonomous system external
ASCII	American National Standard Code for Information Interchange
AUTO	automatic
BECN	backward explicit congestion notification
BER	box event record
BGP	border gateway protocol
BOOTP	bootstrap protocol
bps	bits per second
BRS	bandwidth reservation system
BSC	binary synchronous communication
C&SM	communications and system management
CBSP	control bus and service processor
CCITT	Comité Consultative International Télégraphique et Téléphonique The International Telegraph and Telephone Consultative Committee
CCU	central control unit
CD	carrier detector
CDF-E	configuration data file - extended
CE	customer engineer
CF3745	3745 and 3746 Configurator and Performance Model
CHPID	channel path id
CIDR	classless inter-domain routing
CIR	committed information rate
CLIST	command list
CLA	communication line adapter
CLP	communication line processor
CM	Communications Manager
CNN	composite network node
CNM	communication network management
COS	cost of service
CP	control point
CR	communications rate
CSU	customer service unit
DCAF	Distributed Console Access Facility
DCE	data circuit-terminating equipment
DDS	digital data service
DE	discard eligibility
DLC	data link control
DLCI	data link connection identifier
DLUR	dependent LU requester
DLUS	dependent LU server
DMUX	double multiplex circuit
DSU	data service unit
DTE	data terminal equipment
DX	duplex
EBCDIC	extended binary-coded decimal interchange code
EBN	extended border node
EC	engineering change
EMIF	ESCON Multiple image Facility
EN	end node
EP	emulation program
EPO	emergency power OFF
ESCA	ESCON channel adapter
ESCC	ESCON channel coupler
ESCD	ESCON Director
ESCON	Enterprise Systems Connection
ESCP	ESCON processor
FDX	full duplex
FECN	forward explicit congestion notification
FRFH	frame relay frame handler
FRSE	frame relay switching equipment
FRTE	frame relay terminating equipment

HCD	Hardware Configuration Definition
HDX	half duplex
HI	high
HLA	host link address
HONE	Hands-On Network Environment
HPR	high performance routing
ICMP	internet control message protocol
IML	initial microcode load
INN	intermediate network node or IBM Information Network
IOCP	Input/Output Configuration Program
IP	internet, or internetwork, protocol
IPL	initial program load
IPR	Installation Planning Representative
ITU-T	International Telecommunications Union - Telecommunications (ex-CCITT)
KB	kilobyte (processor storage) 1KB = 2^{10} bytes (1 024 bytes)
Kbps	kilobits bits per second (speed or communication volume per second) 1Kbps = 1 000 (one thousand) bits per second
LAA	locally administered address
LAN	local area network
LCB	line connection box
LCBB	line connection box base
LCBE	line connection box expansion
LCP	link control protocol
LDM	limited distance modem
LED	light emitting diode
LIB n	line interface board type n
LIC n	line interface coupler type n
LSA	link state advertisement
LIU n	line interface coupler unit type n
LIV	link integrity verification
LMI	local management interface
LPAR	logical partition
LPDA	Link Problem Determination Aid
LQ	line quality
LU	logical unit
MAC	medium access control
MAU	medium attachment unit
MB	megabyte (processor storage) 1MB = 2^{20} bytes (1 048 576 bytes)
Mbps	megabits per second (speed or communication volume per second) 1Mbps = 1 000 000 (one million) bits per second
MCL	microcode change level
MES	miscellaneous equipment specification
MIB	management information base
MIH	missing interrupt handler
MLC	machine level control
MLTG	multi-link transmission group
MOSS-E	maintenance and operator subsystem - extended
MTP	multipoint
MUX	multiplex circuit
MVS	multiple virtual storage
NAU	network addressable unit
NMBA	nonbroadcast multiaccess
NCP	Network Control Program
NDRS	non-disruptive route switching
NGMF	NetView Graphic Monitor Facility
NN	network node
NNP	network node processor
NPM	NetView Performance Monitor
NRZI	non-return-to-zero inverted
NVT	network virtual terminal
ODLC	outboard data link control
OSPF	open shortest path first
PBN	peripheral border node
PCI	Peripheral component interconnect
PEP	partitioned emulation program
PING	packet internet groper
PN	peripheral node
PPP	point-to-point protocol
PPPNCP	point-to-point network control protocol
PTP	point-to-point
PTT	post, telegraph, and telephone
PU	physical unit
PVC	permanent virtual circuit
QUAL	quality
RCV	receive clock

RETAIN	Remote Technical Assistance Information Network
RFS	ready for sending
RIP	routing information protocol
ROS	read-only storage
RSF	remote support facility
RTP	rapid transport protocol
RTS	request to send
SDLC	Synchronous Data Link Control
SMUX	single multiplex circuit
SNBU	switched network backup
SNI	SNA network interconnection
SNMP	simple network management protocol
SPAU	service processor access unit
SRC	service reference code
S/S	start-stop
SVC	switched virtual circuit
TC	test control
TCM	Trellis code modulation
TCP	transmission control protocol
TG	transmission group
THRES	threshold
TICn	token-ring interface coupler type n
TIM	Time Services
TOS	type of service
TPF	Transaction Processing Facility
TRA	token-ring adapter
TRP	token-ring processor
UDP	user datagram protocol
UTP	unshielded twisted pair
VTAM	Virtual Telecommunications Access Method
XID	exchange station identification
XMIT	transmit

Glossary

This glossary defines new terms used in this manual. It also includes terms and definitions from the *IBM Dictionary of Computing,* SC20-1699.

adaptive rate-based flow and congestion control (ARB). A function of High Performance Routing (HPR) that regulates the flow of data over an RTP connection by adaptively changing the sender's rate based on feedback on the receiver's rate. It allows high link utilization and prevents congestion before it occurs, rather than recovering after congestion has occurred.

Advanced Communication Function (ACF). A group of IBM licensed programs. principally VTAM programs. TCAM*, NCP, and SSP, that use the concepts of Systems Network Architecture (SNA), including distribution of function and resource sharing.

Advanced Communications Function for the Virtual Telecommunications Access Method (ACF/VTAM). An IBM licensed program that controls communication and the flow of data in an SNA network. It provides single-domain, multiple-domain, and interconnected network capability.

advanced peer-to-peer networking (APPN). Data communications support that routes data in a network between two or more advanced program-to-program communications (APPC) systems that do not need to be adjacent.

automatic network routing. A function of High Performance Routing (HPR) that is provides a low-level routing mechanism that requires no intermediate storage.

channel adapter (CA). A communication controller hardware unit used to attach the controller to a host processor.

communication controller. A device that directs the transmission of data over the data links of a network; its operation may be controlled by a program executed in a processor to which the controller is connected or it may be controlled by a program executed within the device. For example, the IBM 3745 and 3746 Nways Multiprotocol Controllers.

communications manager. A function of the OS/2 Extended Edition program that lets a workstation connect to a host computer and use the host resources as well as the resources of the other personal computers to which the workstation is attached, either directly or through a host system. The communications manager provides application programming interfaces (APIs) so that users and develop their own applications.

configuration data file - extended (CDF-E). A 3746 Nways Multiprotocol Controller MOSS-E file that contains a description of all the hardware features (presence, type, address, and characteristics).

communications management configuration host node. The type 5 host processor in a communications management configuration that does all network-control functions in the network except for the control of devices channel-attached to a data host nodes. Synonymous with communications management host. See also data host node.

control panel. A panel that contains switches and indicators for the customer's operator and service personnel.

control program. A computer program designed to schedule and to supervise the execution of programs of the controller.

control subsystem. The part of the controller that stores and executes the control program, and monitors the data transfers over the channel and transmission interfaces.

Customer Engineer. See IBM service representative

data circuit-terminating 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 between the data terminal equipment (DTE) and the line. For example, a modem is a DCE.

Note: The DCE may be a stand-alone equipment or integrated in the 3745.

data terminal equipment (DTE). That part of a data station that serves as a data source, data link , or both, and provides for the data communication control function according to protocols. For example, the 3174 and PS/2s are DTEs.

data host node. In a communication management configuration, a type 5 host node that is dedicated to processing applications and does not control network resources, except for its channel adapter-attached or communication adapter-attached devices. Synonymous with data host. See also communications management configuration host node.

direct attachment. The attachment of a DTE to another DTE without a DCE.

ESCON channel. A channel having an Enterprise System Connection* channel-to-control-unit I/O interface that uses optical cables as a transmission medium.

ESCON channel adapter (ESCA). A communication controller hardware unit used to attach the controller to a host via ESCON fiber optics. An ESCA consists of an ESCON channel processor (ESCP) and an ESCON channel coupler (ESCC).

ESCON channel coupler (ESCC). A communication controller hardware unit which is the interface between the ESCON channel processor and the ESCON fiber optic cable.

ESCON channel processor (ESCP). A communication controller hardware unit which provides the channel data link control for the ESCON channel adapter.

Distributed Console Access Facility. (1) This program product provides a remote console function that allows a user at one programmable workstation (PS/2) to remotely control the keyboard input and monitor the display of output of another programmable workstation. The DCAF program does not affect the application programs that are running on the workstation that is being controlled. (2) An icon that represents the Distributed Console Access Facility.

Enterprise Systems Connection (ESCON). A set of IBM products and services that provides a dynamically connected environment within an enterprise.

Host. See host processor

host processor. (1) A processor that controls all or part of a user application network. (2) In a network, the processing unit where the access method for the network resides. (3) In an SNA network, the processing unit that contains a system services control point (SSCP). (4) A processing unit that executes the access method for attached communication controllers.

High Performance Routing (HPR). An extension of APPN that provides faster traffic throughput, lower delays, and lower storage overheads.

IBM service representative. An individual in IBM who does maintenance services for IBM products or systems. Also called the IBM *Customer Engineer*.

initial microcode load (IML). The process of loading the microcode into an adapter, the MOSS, or the service processor.

Internet. (1) A wide area network connecting disparate networks using the internetwork protocol (IP) (2) A public domain wide area network connecting thousands of disparate networks in industry, education, government and research. The Internet uses TCP/IP as the standard for transmitting information.

Internet address. The numbering system used in IP internetwork communications to specify a particular network, or a particular host on that network with which to communicate.

Internet Control Message Protocol (ICMP). A protocol used by a gateway to communicate with a source host, for example, to report an error in a datagram. It is an integral part of the Internetwork Protocol (IP).

Internetwork Protocol. A protocol that routes data from its source to its destination in an internet environment. It is also called the *Internet Protocol*.

Internetwork. Any wide area network connecting more than one network.

initial program load (IPL). The initialization procedure that causes the 3745 control program (NCP) to begin operation.

LAN-attached console. A PS/2 attached to the token-ring LAN that has the service processor attached. It is used to operate remotely the MOSS and MOSS-E functions.

IP router. A device that enables an Internetwork Protocol (IP) host to act as a gateway for routing data between separate networks.

line interface coupler (LIC). A circuit that attaches up to four transmission cables to the controller (from DTEs, DCEs or telecommunication lines).

locally administered address. In a local area network, an adapter address that the user can assign to override the universally administered address.

maintenance and operator subsystem - extended (MOSS-E). The licensed internal code loaded on the service processor hard disk to provide maintenance and operator facilities to the user and IBM service representative.

microcode. A program that is loaded in a processor (for example, the MOSS processor) to replace a hardware function. The microcode is not accessible to the customer.

modem (modulator-demodulator). See DCE.

Multiple Virtual Storage (MVS). Multiple Virtual Storage, consisting of MVS/System Product Version 1 and the MVS/370 Data Facility Product operating on a System/370* processor.

NetView. An IBM licensed program used to monitor a network, manage it, and diagnose its problems.

nonswitched line. A connection between systems or devices that does not have to be made by dialing. The connection can be point-to-point or multipoint. The line can be leased or private. Contrast with *switched line.*.

ping. A simple IP application that sends one or more messages to a specified destination host requesting a reply. Usually used to verify that the target host exists, or that its IP address is a valid address.

remote console. A PS/2 attached to the 3746 Nways Multiprotocol Controller either by a switched line (with modems) or by one of the communication lines of the user network.

Remote Technical Assistance Information Network (RETAIN).

service processor. The processor attached to a 3745, 3746-900, and 3746-950 via a token-ring LAN.

remote support facility (RSF). RSF provides IBM maintenance assistance when requested via the public switched network. It is connected to the IBM RETAIN database system.

service representative. See IBM service representative

services. A set of functions designed to simplify the maintenance of a device or system.

switched line. A transmission line with which the connections are established by dialing, only when data transmission is needed. The connection is point-to-point and uses a different transmission line each time it is established. Contrast with *nonswitched line*.

Synchronous Data Link Control (SDLC). A discipline for managing synchronous, code-transparent, serial-by-bit information transfer over a link connection. Transmission exchanges may be duplex or half-duplex over switched or nonswitched links. The configuration of the link connection may be point-to-point, multipoint, or loop. SDLC conforms to subsets of the Advanced Data Communication Control Procedures of the American National Standards Institute and High-Level Data Link Control (HDLC) of the International Standards Organization.

synchronous transmission. Data transmission in which the sending and receiving instruments are operating continuously at substantially the same frequency and are maintained, through correction, in a desired phase relationship.

token-ring adapter (TRA) type 3. 3746-900 and 3746-950 line adapter for IBM Token-Ring Network, composed of one token-ring processor card (TRP2), and two token-ring interface couplers type 3 (TIC 3s).

token-ring interface coupler type 2 (TIC2). A circuit that attaches an IBM Token-Ring network to the 3745.

token-ring interface coupler type 3 (TIC3). A circuit that attaches an IBM Token-Ring network to the 3746-900 or 3746-950.

user access area. A specific area in the controller where the customer can install, remove, change, or swap couplers and cables without IBM assistance.

universally administered address. In a local area network, the address permanently encoded in an adapter at the time of manufacture. All universally administered addresses are unique.

user application network. A configuration of data processing products, such as processors, controllers, and terminals, for data processing and information exchange. This configuration may use circuit-switched, packet-switched, and leased-circuit services provided by carriers or PTT. Also called a *user network.*

V.24, V.35, and X.21. ITU-T (ex-CCITT) recommendations on transmission interfaces.

Bibliography

Customer Documentation for the 3745 (Models 210, 310, 410, 610, 21A, 31A, 41A, and 61A), and 3746 (Model 900)

Table X-1 (Page 1 of 4). Customer Documentation for the 3745 Models X10 and X1A, and 3746 Model 900

This customer documentation has the following formats:		
Books	Online	Books and Diskettes
Finding Information		
		3745 Models A and 3746 Books Starting with engineering change (EC) F12380, all of the books in the 3745 Models A and 3746 library are available on the CD-ROM that contains the Licensed Internal Code (LIC) for this EC.
📖	SA33-0172	***IBM 3745 Communication Controller Models 210 to 61A IBM 3746 Expansion Unit Model 900*** ***Customer Master Index***[1] Provides references for finding information in the customer documentation library.
Evaluating and Configuring		
📖	GA33-0092	***IBM 3745 Communication Controller Models 210, 310, 410, and 610*** ***Introduction*** Gives an introduction of the IBM Models 210 to 610 capabilities. For Models A refer to the *Overview*, GA33-0180.
📖	GA33-0180	***IBM 3745 Communication Controller Models A***[2] ***IBM 3746 Nways Multiprotocol Controller Models 900 and 950*** ***Overview*** Gives an overview of connectivity capabilities within SNA, APPN, and IP networking.
📖	GA33-0457	***IBM 3745 Communication Controller Models A***[2] ***IBM 3746 Expansion Unit Model 900 Models 900 and 950*** ***Planning Guide*** Planning for: • Field upgrades • Service processor and alert management configuration • Network integration (NCP, APPN, and IP control) • Physical installation.

Table X-1 (Page 2 of 4). Customer Documentation for the 3745 Models X10 and X1A, and 3746 Model 900

Preparing Your Site		
📖	GC22-7064	***IBM System/360, System/370, 4300 Processor*** ***Input/Output Equipment Installation Manual-Physical Planning*** (Including Technical News Letter GN22-5490) Provides information for physical installation for the 3745 Models 130 to 610. For 3745 Models A and 3746 Model 900, refer to the *Planning Guide*, GA33-0457.
📖	GA33-0127	***IBM 3745 Communication Controller*** ***Models 210, 310, 410, and 610*** ***Preparing for Connection*** Helps for preparing the 3745 Models 210 to 610 cable installation. For 3745 Models A refer to the *Connection and Integration Guide*, SA33-0129.
Preparing for Operation		
📖	GA33-0400	***IBM 3745 Communication Controller All Models***[3] ***IBM 3746 Nways Multiprotocol Controller*** ***Models 900 and 950*** ***Safety Information***[1] Provides general safety guidelines.
📖	SA33-0129	***IBM 3745 Communication Controller All Models***[3] ***IBM 3746 Nways Multiprotocol Controller Model 900*** ***Connection and Integration Guide***[1] Contains information for connecting hardware and integrating network of the 3745 and 3746-900 after installation.
📖	SA33-0416	***Line Interface Coupler Type 5 and Type 6*** ***Portable Keypad Display*** ***Migration and Integration Guide*** Contains information for moving and testing LIC types 5 and 6.
📖 💾	SA33-0158	***IBM 3745 Communication Controller All Models***[3] ***IBM 3746 Nways Multiprotocol Controller Model 900*** ***Console Setup Guide***[1] Provides information for: • Installing local, alternate, or remote consoles for 3745 Models 130 to 610 • Configuring user workstations to remotely control the service processor for 3745 Models A and 3746 Model 900 using: – DCAF program – Telnet Client program.
Customizing Your Control Program		
📖	SA33-0178	***Guide to Timed IPL and Rename Load Module*** Provides VTAM procedures for: • Scheduling an automatic reload of the 3745 • Getting 3745 load module changes transparent to the operations staff.
Operating and Testing		

Table X-1 (Page 3 of 4). Customer Documentation for the 3745 Models X10 and X1A, and 3746 Model 900

📖	SA33-0098	***IBM 3745 Communication Controller All Models[4]*** ***Basic Operations Guide[1]*** Provides instructions for daily routine operations on the 3745 Models 130 to 610.
📖	SA33-0177	***IBM 3745 Communication Controller Models A[2] IBM 3746 Nways Multiprotocol Controller Model 900*** ***Basic Operations Guide[1]*** Provides instructions for daily routine operations on the 3745 Models 17A to 61A, and 3746 Model 900 operating as an SNA node (using NCP), APPN/HPR Network Node, and IP Router.
📖	SA33-0097	***IBM 3745 Communication Controller All Models[3]*** ***Advanced Operations Guide[1]*** Provides instructions for advanced operations and testing, using the 3745 MOSS console.
🖥	On-line Information	***Controller Configuration and Management Application*** Provides a graphical user interface for configuring and managing a 3746 APPN/HPR Network Node and IP Router, and its resources. Is also available as a stand-alone application, using an OS/2 workstation. Defines and explains all the 3746 Network Node and IP Router configuration parameters through its online help.
📖	SH11-3081	***IBM 3746 Nways Multiprotocol Controller Models 900 and 950*** ***Controller Configuration and Management: User's Guide[5]*** Explains how to use CCM and gives examples of the configuration process.
Managing Problems		
📖	SA33-0096	***IBM 3745 Communication Controller All Models[3]*** ***Problem Determination Guide[1]*** A guide to perform problem determination on the 3745 Models 130 to 61A.
🖥	On-line Information	***Problem Analysis Guide*** An online guide to analyze alarms, events, and control panel codes on: • IBM 3745 Communication Controller Models A[2] • IBM 3746 Nways Multiprotocol Controller Models 900 and 950.

Table X-1 (Page 4 of 4). Customer Documentation for the 3745 Models X10 and X1A, and 3746 Model 900

📖	SA33-0175	***IBM 3745 Communication Controller Models A[2]*** ***IBM 3746 Expansion Unit Model 900*** ***IBM 3746 Nways Multiprotocol Controller Model 950*** ***Alert Reference Guide*** Provides information about events or errors reported by alerts for: • IBM 3745 Communication Controller Models A[2] • IBM 3746 Nways Multiprotocol Controller Models 900 and 950.

[1] Documentation shipped with the 3745.
[2] 3745 Models 17A to 61A.
[3] 3745 Models 130 to 61A.
[4] Except 3745 Models A.
[5] Documentation shipped with the 3746-900.

Additional Customer Documentation for the 3745 Models 130, 150, 160, 170, and 17A

Table X-2. Additional Customer Documentation for the 3745 Models 130 to 17A

This customer documentation has the following format:		
Books		
Finding Information		
📖	SA33-0142	***IBM 3745 Communication Controller Models 130, 150, 160, 170, and 17A IBM 3746 Nways Multiprotocol Controller Model 900*** ***Customer Master Index***[1] Provides references for finding information in the customer documentation library.
Evaluating and Configuring		
📖	GA33-0138	***IBM 3745 Communication Controller Models 130, 150, and 170*** ***Introduction*** Gives an introduction about the IBM Models 130 to 170 capabilities, including Model 160. For Model 17A refer to the *Overview*, GA33-0180.
Preparing Your Site		
📖	GA33-0140	***IBM 3745 Communication Controller Models 130, 150, 160, and 170*** ***Preparing for Connection*** Helps for preparing the 3745 Models 130 to 170 cable installation. For 3745 Model 17A refer to the *Connection and Integration Guide*, SA33-0129.

[1] Documentation shipped with the 3745.

Customer Documentation for the 3746 Model 950

Table X-3 (Page 1 of 2). Customer Documentation for the 3746 Model 950

This customer documentation has the following formats:		
Books	Online	Books and Diskettes
Finding Information		
		3745 Models A and 3746 Books Starting with engineering change (EC) F12380, all of the books in the 3745 Models A and 3746 library are available on the CD-ROM that contains the Licensed Internal Code (LIC) for this EC.
Preparing for Operation		
📖	GA33-0400	***IBM 3745 Communication Controller All Models*[1]** ***IBM 3746 Expansion Unit Model 900*** ***IBM 3746 Nways Multiprotocol Controller Model 950*** ***Safety Information*[2]** Provides general safety guidelines
Evaluating and Configuring		
📖	GA33-0180	***IBM 3745 Communication Controller Models A*[3]** ***IBM 3746 Nways Multiprotocol Controller*** ***Models 900 and 950*** ***Overview*** Gives an overview of connectivity capabilities within SNA, APPN, and IP networking.
📖	GA33-0457	***IBM 3745 Communication Controller Models A*[2]** ***IBM 3746 Expansion Unit Model 900*** ***Models 900 and 950*** ***Planning Guide*** Planning for: • Field upgrades • Service processor and alert management configuration • Network integration (NCP, APPN, and IP control) • Physical installation.

Table X-3 (Page 2 of 2). Customer Documentation for the 3746 Model 950

Operating and Testing		
📖💾	SA33-0356	***IBM 3746 Nways Multiprotocol Controller Model 950*** ***User's Guide***[2] Explains how to: • Carry out daily routine operations on Nways controller • Install, test, and customize the Nways controller after installation • Configure user's workstations to remotely control the service processor using: – DCAF program – Telnet client program.
🖥	On-line information	***Controller Configuration and Management Application*** Provides a graphical user interface for configuring and managing a 3746 APPN/HPR network node and IP Router, and its resources. Is also available as a stand-alone application, using an OS/2 workstation. Defines and explains all the 3746 Network Node and IP Router configuration parameters through its on-line help.
📖	SH11-3081	***IBM 3746 Nways Multiprotocol Controller Models 900 and 950*** ***Controller Configuration and Management: User's Guide***[2] Explains how to use CCM and gives examples of the configuration process.
Managing Problems		
🖥	On-line information	***Problem Analysis Guide*** An on-line guide to analyze alarms, events, and control panel codes on: • IBM 3745 Communication Controller Models A[3] • IBM 3746 Nways Multiprotocol Controller Models 900 and 950.
📖	SA33-0175	***IBM 3745 Communication Controller Models A***[3] ***IBM 3746 Expansion Unit Model 900*** ***IBM 3746 Nways Multiprotocol Controller Model 950*** ***Alert Reference Guide*** Provides information about events or errors reported by alerts for: • IBM 3745 Communication Controller Models A[3] • IBM 3746 Nways Multiprotocol Controller Models 900 and 950.

[1] Models 130 to 61A.
[2] Documentation shipped with the 3746-950
[3] 3745 Models 17A to 61A.

Related Manuals

Related Manuals for 3745

The following documents are indispensable for planning for your 3745 Communication Controllers Models A:

- *IBM 3745 Communication Controller: Console Setup Guide*, GA33-0158
- *IBM 3745 Communication Controller Models A: Overview*, GA33-0180.

Be sure to use the latest editions of the above documents.

Also helpful are:

- *Planning for Integrated Networks*, SC31-8062
- *Planning and Reference for NetView, NCP, and VTAM*, SC31-7122.
- *Virtual Telecommunications Access Method V3 R4: Resource Definition Reference*, SC31-6438

The following Enterprise Systems Connection (ESCON) documents may be helpful:

- *Introducing the Enterprise Systems Connection*, GA23-0383
- *Enterprise Systems Connection Migration*, GA23-0383
- *Planning for Enterprise Systems Connection Links*, GA23-0367
- *Introducing Enterprise Systems Connection Directors*, GA23-0363.

The following *IBM International Technical Support Centers* “redbooks” are generally very helpful:

- *Frame Relay Guide*, GG24-4463
- *3746-900 and NCP Version 7 Release 2*, GG24-4464.

The following Network Control Program (NCP) documents may be helpful:

- For NCP V6 R2:
  - *Network Control Program V6 R2: Migration Guide*, SC31-6216
  - *Network Control Program V6 R2, ACF/SSP V3 R8, EP R11: Resource Definition Guide*, SC31-6209-01
  - *Network Control Program V6 R2, ACF/SSP V3 R8, EP R11: Resource Definition Reference*, SC31-6210-01
  - *Network Control Program V6 R2: Planning and Implementation Guide*, GG24-4012
  - *Network Control Program V6 R2, ACF/SSP V3 R8, EP R11: Library Directory*, SC31-6215.
- For NCP V6 R3:
  - *Network Control Program V6 R3: Migration Guide*, SC31-6217
  - *Network Control Program V6 R3, ACF/SSP V3 R9, EP R11: Resource Definition Guide*, SC31-6209-02
  - *Network Control Program V6 R3, ACF/SSP V3 R9, EP R11: Resource Definition Reference*, SC31-6210-02 Guide,
  - *Network Control Program V6 R3, ACF/SSP V3 R9, EP R11: Library Directory*, SC31-6218.
- For NCP V7 R1:
  - *Network Control Program V7 R1: Migration Guide*, SC31-6219
  - *Network Control Program V7 R1, ACF/SSP V4 R1, EP R12: Resource Definition Guide*, SC31-6223-00
  - *Network Control Program V7 R1, ACF/SSP V4 R1, EP R12: Resource Definition Reference*, SC31-6224-00
  - *Network Control Program V7 R1, ACF/SSP V4 R1, EP R12: Library Directory*, SC31-6220.
- For NCP V7 R2:
  - *Network Control Program V7 R2, ACF/SSP V4 R2, EP R12: Generation and Loading Guide*, SC31-6221.
  - *Network Control Program V7 R2: Migration Guide*, SC31-6258-00
  - *Network Control Program V7 R2, ACF/SSP V4 R2, EP R12: Resource Definition Guide*, SC31-6223-01

- *Network Control Program V7 R2, ACF/SSP V4 R2, EP R12: Resource Definition Reference*, SC31-6224-01
- *Network Control Program V7 R2, ACF/SSP V4 R2, EP R12: Library Directory*, SC31-6259.

- For NCP V7 R3:
  - *Network Control Program V7 R3: Migration Guide*, SC31-6258-01
  - *Network Control Program V7 R3, ACF/SSP V4 R3, EP R12: Resource Definition Guide*, SC31-6223-02
  - *Network Control Program V7 R3, ACF/SSP V4 R3, EP R12: Resource Definition Reference*, SC31-6224-02
  - *Network Control Program V7 R3, ACF/SSP V4 R3, EP R12: Library Directory*, SC31-6262.
- For NCP V7 R4:
  - *Network Control Program V7 R4: Migration Guide*, SC30-3786
  - *Network Control Program V7 R4, ACF/SSP V4 R4, EP R12: Resource Definition Guide*, SC31-6223-03
  - *Network Control Program V7 R4, ACF/SSP V4 R4, EP R12: Resource Definition Reference*, SC31-6224-03
  - *Network Control Program V7 R4, ACF/SSP V4 R4, EP R12: Library Directory*, SC30-3785.

The following OS/2 document may be of some help:

IBM Extended Services for OS/2 Programming Services and Advanced Problem Determination for Communications, SO4G-1007.

For the Distributed Console Access Facility (DCAF) Version 1.3 the following documents are needed:

- *DCAF: Installation and Configuration Guide*, SH19-4068
- *DCAF: User's Guide*, SH19-4069
- *DCAF: Target User's Guide*, SH19-6839.

Related Manuals for 3746-9X0 (APPN)

To learn more about the APPN architecture, including high-performance routing (HPR), adaptive rate based flow and congestion control (ARB), dependent LU requesters/servers (DLURs/DLUSs), and other subjects, refer to:

- *Inside APPN - The Essential Guide to the Next-Generation SNA*, SG24-3669.
- *APPN Architecture and Protocol Implementations Tutorial:ecit, SG24-3669.*

The following Enterprise Systems Connection (ESCON), Virtual Telecommunications Access Method (VTAM), and OS/2 documentation may be also helpful:

- Introducing the Enterprise Systems Connection*, GA23-0383*
- Enterprise Systems Connection Migration*, GA23-0383*
- Planning for Enterprise Systems Connection Links*, GA23-0367*
- Introducing Enterprise Systems Connection Directors*, GA23-0363*
- Virtual Telecommunications Access Method V4R3: Resource Definition Reference*, SC31-6438.*
- IBM Extended Services for OS/2 Programming Services and Advanced Problem Determination for Communications*, SO4G-1007.*

For help with TCP/IP, refer to:

- TCP/IP for MVS: Performance Tuning Guide*, SC31-7188.*

To learn more token-ring configurations and the IEEE 802.2 standard, refer to:

- Token-Ring Network Architecture Reference*, SC30-3374.*

Index

Readers' Comments — We'd Like to Hear from You

3745 Communication Controller Models A
3746 Nways Multiprotocol Controller
Models 900 and 950
Planning Guide

(Part 1/3)
Preliminary Second Edition

Publication No. GA33-0457-0A

Please send us your comments concerning this book. We will greatly appreciate them and will consider them for later releases of the present book.

If you prefer sending comments by FAX or electronically, use:

- FAX: 33 4 93 24 77 97
- E-mail: FRIBMQF5 at IBMMAIL
- IBM Internal Use: LGERCF at LGEPROFS
- Internet: rcf_lagaude@vnet.ibm.com

In advance, thank you.

Your comments:

Name	Address
Company or Organization	
Phone No.	

Readers' Comments — We'd Like to Hear from You
GA33-0457-0A

Cut or Fold
Along Line

Fold and Tape

Please do not staple

Fold and Tape

PLACE
POSTAGE
STAMP
HERE

IBM France
Centre d'Etudes et Recherches
Service 0798 - BP 79
06610 La Gaude
France

Fold and Tape

Please do not staple

Fold and Tape

Cut or Fold
Along Line

Printed in United Kingdom