3746, 2210, 2216, and 2220 Interconnectivity: Frame Relay and Related Functions

Brian Dorling Motoaki Nakao Gallus Schlegel Frank Slawitzki

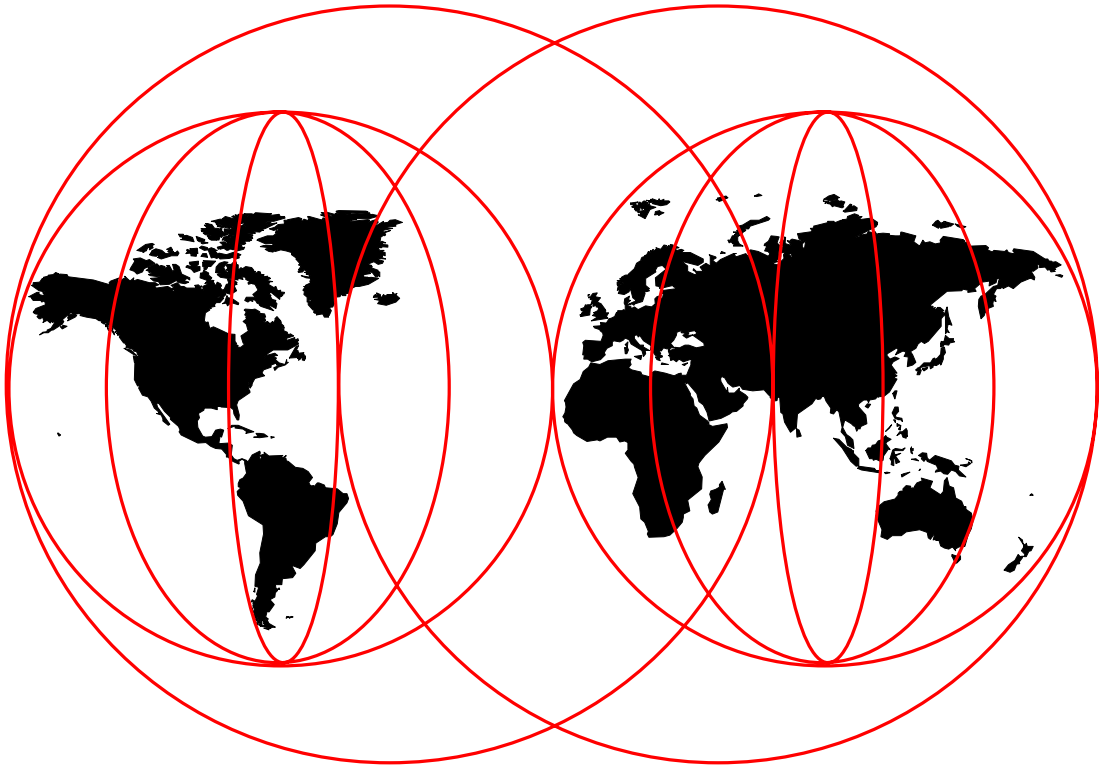

International Technical Support Organization

http://www.redbooks.ibm.com

SG24-2146-00

International Technical Support Organization

SG24-2146-00

3746, 2210, 2216, and 2220 Interconnectivity: Frame Relay and Related Functions

May 1998

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix B, “Special Notices” on page 245.

First Edition (May 1998)

This edition applies to 3746, 2210, 2216 and 2220 at various version and release numbers.

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. HZ8 Building 678
P.O. Box 12195
Research Triangle Park, NC 27709-2195

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

Contents

Figures

Preface

This redbook gives networking professionals guidelines on building and configuring APPN, IP, frame relay, and multiprotocol networks based on, but not limited to, IBM 3746 communications controllers.

The IBM 3746 is shown in configurations together with IBM 2210 Multiprotocol Routers, IBM 2216 Multiaccess Connectors, and workstations running a variety of protocols. Definitions are shown and explained for each node in the configurations, and theoretical information is included to explain how the configurations work.

The configurations cover subarea SNA, APPN, and TCP/IP networking over frame relay backbone networks, using frame relay boundary network node (BNN), frame relay boundary access node (BAN), and data-link switching (DLSw).

In addition this book contains a detailed section on the theory of frame relay networks and congestion control in those networks.

The Team That Wrote This Redbook

This redbook was produced by a team of specialists from around the world working at the Systems Management and Networking ITSO Center, Raleigh.

Brian Dorling is an Advisory ITSO Specialist for Communications Architectures at the Systems Management and Networking ITSO Center, Raleigh. Brian is responsible for a broad range of IBM communication architectures including Advanced Peer-to-Peer Networking (APPN), Multiprotocol Transport Networking (MPTN), Networking Broadband Services (NBBS), Switched Virtual Networking (SVN), and for the 3746 family of multiprotocol communications controllers. Since joining IBM in 1978, Brian has worked as a customer engineer and systems engineer in the networking field in the UK and Germany.

Motoaki Nakao is a Senior Advisory Information/Technology Specialist for Communication Network field in Japan. He has about 20 years of experience in SNA and IP-based networks and network management systems. His areas of expertise include VTAM/NCP, 3745, 3746, NetView, OS/2 and client/server networking. After joining IBM in 1971, he has worked as a systems engineer in the marketing field and technical support department.

Gallus Schlegel is a Senior Network System Engineer in Switzerland. He has over 15 years of experience in networking. Gallus has held various positions in his IBM career including assignments in second level support. He holds a degree in Electronic Engineering. His areas of expertise include most areas of networking, but with a focus on large WAN-based networks.

Frank Slawitzki has been a Support Specialist for telecommunications software in Germany for more than seven years. He supports a range of IBM communication products and architectures including VTAM, NCP/NPSI, 3172 ICP and SNA Communications Program, and Open Systems Adapter in SNA, APPN and TCP/IP environment. He was assigned to the Technical Networking Support Center in La Gaude, France for one year during which he provided EMEA-wide support for the IBM 2220 NBBS Switch and IBM 3745/3746.

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

The Editing and Graphics Team
Systems Management and Networking ITSO Center, Raleigh

Jean-Claude Dispensa
IBM La Gaude

Robert Kraus
IBM Germany

Jean-Pierre Marce
IBM La Gaude

Piet Pieters
IBM Netherlands

Jay Miller
IBM Research Triangle Park

Carsten Seefeldt
IBM Germany

Elmer Valenzuela
IBM Phillipines

Comments Welcome

Your comments are important to us!

We want our redbooks to be as helpful as possible. Please send us your comments about this or other redbooks in one of the following ways:

- Fax the evaluation form found in “ITSO Redbook Evaluation” on page 255 to the fax number shown on the form.
- Use the electronic evaluation form found on the Redbooks Web sites:

  For Internet users `http://www.redbooks.ibm.com/`
  For IBM Intranet users `http://w3.itso.ibm.com/`

- Send us a note at the following address:

  `redbook@us.ibm.com`

Part 1. Introduction and Frame Relay Theory

Chapter 1. Overview

This redbook covers the IBM 3746 Multiprotocol Communications controller Models 900 and 950 base frame, IBM 2210 Multiprotocol Routers, IBM 2216 Multiaccess Connector, IBM 2220 Nways Broadband Switch, and OS/2 workstations running APPN and IP. Although in some cases information is given about the Multiaccess Enclosure (MAE), it is not a part of this book. IBM has announced the 3746 feature code FC #3000, with planned availability in 1998, which will provide direct high-speed connectivity between the MAE and 3746 switch. Software enhancements also available at that time will allow the 3746 and MAE to be configured as a single image, and from the IP routing perspective the 3746 and MAE will be seen as a single IP router. The merging of the 3746 APPN control point and the MAE APPN control point will be supported at a later date.

As the 3746 support for the MAE will change in 1998, it was decided to focus this book on the 3746 base frame. Information about the MAE has been provided in some areas where it will differ from the support in the 3746 base frame. This is included for informational purposes only.

Although the focus of this redbook is building networks using a 3746, configuration information is shown for all machines used in the network. We used an IBM 2220 Broadband switch to simulate a high-speed frame relay backbone network. Although we used a single 2220 instead of a network of switches, it made no difference to our tests.

1.1 Network Test Scenarios

The network scenarios tested were based upon the hardware available to us at the time of testing. The basic configuration was to use a 3746 Model 900 running APPN and IP as an access node to our S/390 servers. This 3746 was connected to the frame relay backbone provided by our 2220 broadband switch, via ESCON to our S/390 servers, and via token-ring to our office LAN. We connected an IBM 2216 multiaccess connector as concentrator node to the frame relay backbone, and had workstations connected to the 2216 via token-ring. An IBM 2210 multiprotocol router was connected over a serial link running point-to-point protocol (PPP) to a serial port on the 2216. The 2210 also had workstations connected to it via token-ring.

1.2 Type 3 Processors

In September 1997 new *type 3* processors became available for the 3746. These processors provide increased performance and double the memory size of type 2 processors. They increase APPN/DLUR performance by up to 100%, and connectivity by up to 200% over type 2 processors.

For 3746 machines running IP, the CBSP3 provides improved dynamic routing protocol performance and more storage for routing tables. The other type 3 processors also provide more space for routing table and ARP caches plus better packet forwarding performance.

Field upgrades are available from type 1 and type 2 processors to type 3 processors.

Table 1 on page 5 shows individual processor connectivity and Table 2 on page 7 shows total 3746 connectivity.

1.2.1 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.

1.2.1.1 Adapter Connectivity

Table 1 on page 5 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 1 on page 5) and total number of PUs (see column Total in Table 1 on page 5) are absolute maximum numbers which cannot be exceeded.

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

The following notes refer to Table 1 on page 5.

Notes:

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 PUs, 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 and X.25).

    For example: TRP2s (without the 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).

1.2.2 Network Node Connectivity

Table 2 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. 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).		

1.3 3746 IP to NCP Internal Connection

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 needed to connect the two routers. To support this function from the 3746 side, the optional feature code #5800 is needed. This is supported by 3746 microcode level ECA 170 and later levels.

The internal connection or connections between the two routers are seen as 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 as a proxy router. 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 1 on page 8 shows a representation of the internal point-to-point IP connection.

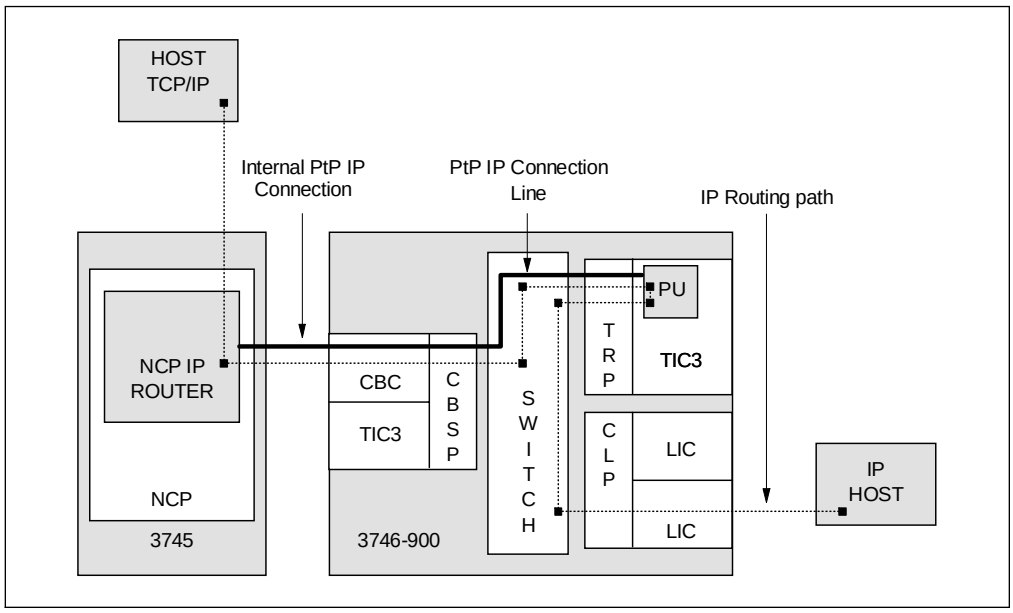

Figure 1. Internal IP PtP Connection

1.3.1 Internal IP Connection Definitions

The following sections explain how to define an internal point-to-point IP connection. Figure 2 on page 9 shows the configuration that is being defined. The definitions shown here include all the information needed for this part of the NCP generation process. The actual PtP link is defined in the LN2240 LINE **1** and IP2240 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** (see Figure 3 on page 10 and Figure 4 on page 11).

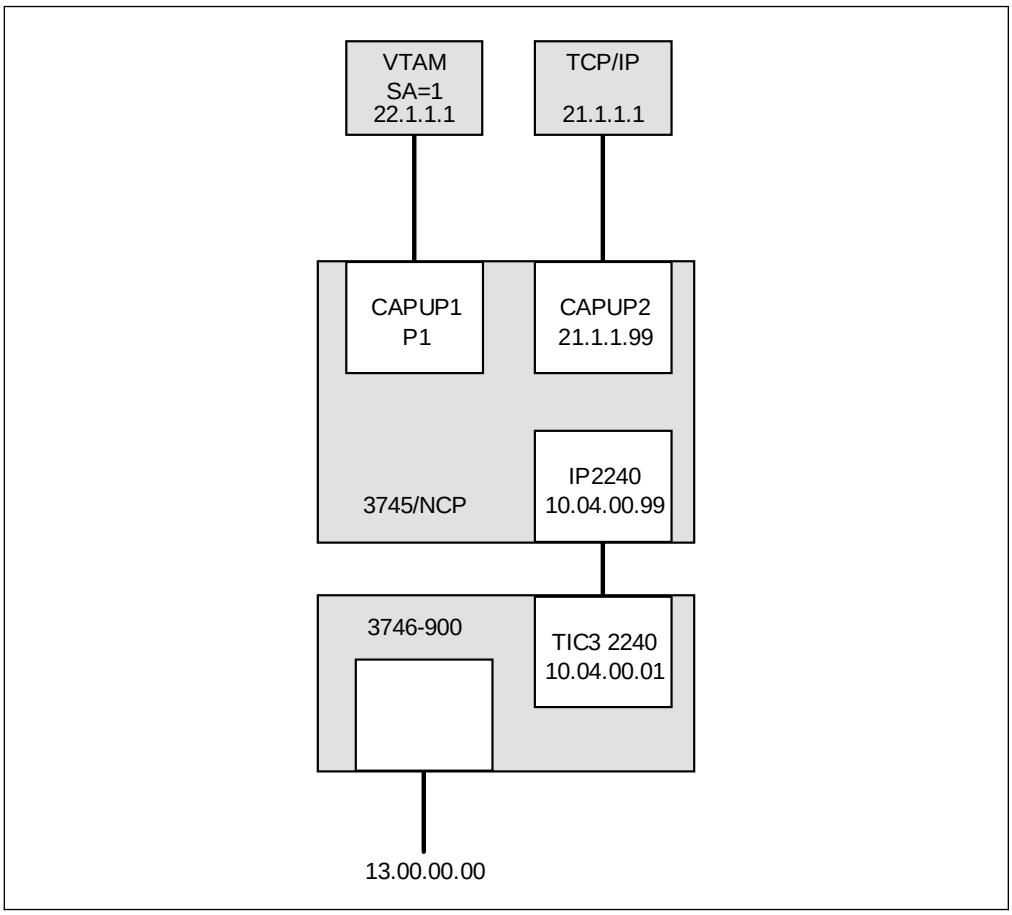

Figure 2. Example Configuration

1.3.1.1 NCP Definitions

Figure 3 on page 10 and Figure 4 on page 11 show the NCP definitions needed to configure the internal IP PtP link from the 3745 side.

```
*************************************************************************
* VTAM HOST DEFINITIONS
*************************************************************************
*
HOSTA01  HOST  BFRPAD=0,DELAY=0.0,INBFRS=6,MAXBFRU=32,
               SUBAREA=1,
               STATMOD=YES,UNITSZ=1024
*
*************************************************************************
*  TOKEN RING LINE 2240
*************************************************************************
*
ODLCGRP1 GROUP ECLTYPE=(PHY,ANY),ADAPTER=TIC3,
               DIAL=NO,LNCTL=SDLC,TYPE=NCP
LN2240   LINE  ADDRESS=(2240,FULL)  1 ,
               LOCADD=400000002240,
               MAXPU=2,
               MAXTSL=16732,
               NPACOLL=YES,
               PORTADD=2,
               SPEED=9600,
               TRSPEED=16
IP2240   PU    ADDR=02,
               INTFACE=(TR2240,2048)  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)
```

Figure 3. NCP Definitions

```
*************************************************************************
*        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=TR2240, 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=TR2240,
                 METRIC=1,DISP=PERM,HOSTRT=NO
*
GENEND   GENEND
         END
```

Figure 4. NCP Definitions

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) for the NCP-IP router.

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.

1.3.1.2 3746 CCM Definitions

The following section describes how to use CCM to create the 3746 definitions for the PtP link.

1. The token-ring port 2240 must be configured first, and IP must be activated on that port.

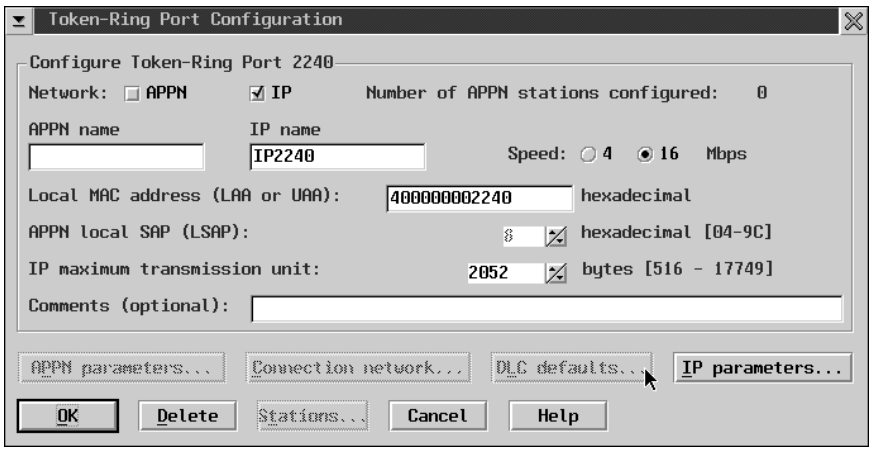

Figure 5. Port Configuration

2. On port 2240, 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 operand of the IPLOCAL **6** statement in NCP.

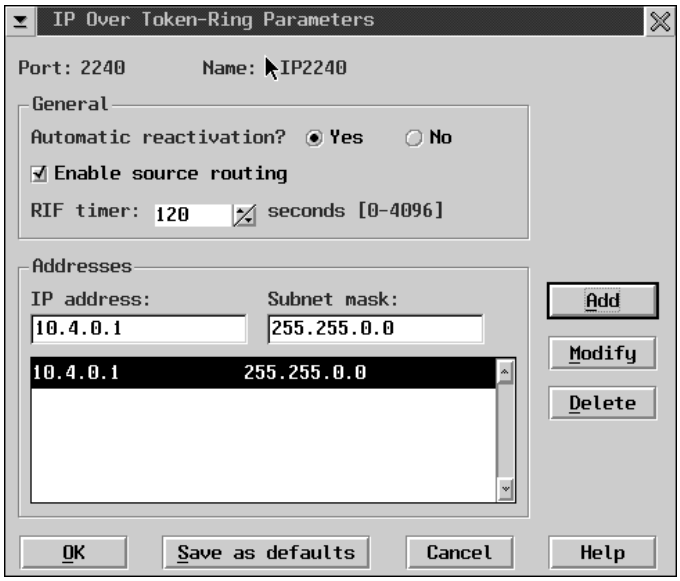

Figure 6. IP over Token-Ring Parameters

3. RIP should be activated as shown on port 2240.

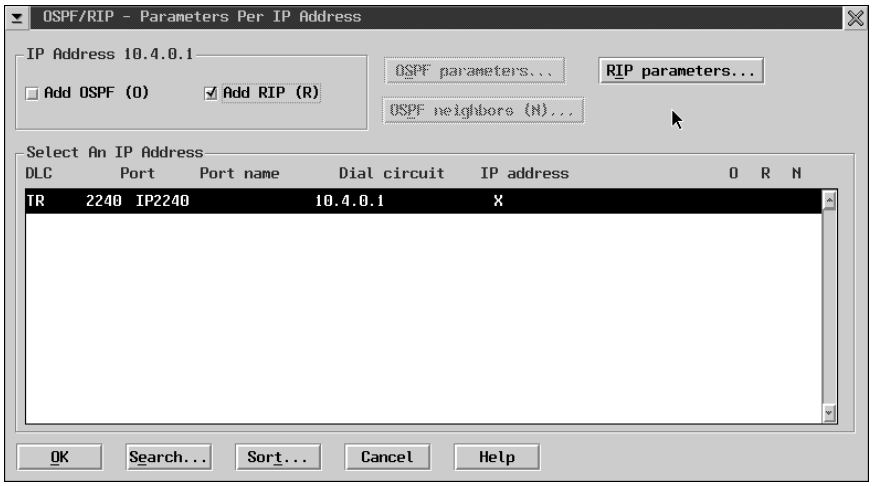

Figure 7. OSPF/RIP Parameters per IP Address

4. The following RIP parameters should be selected for IP address 10.4.0.1.

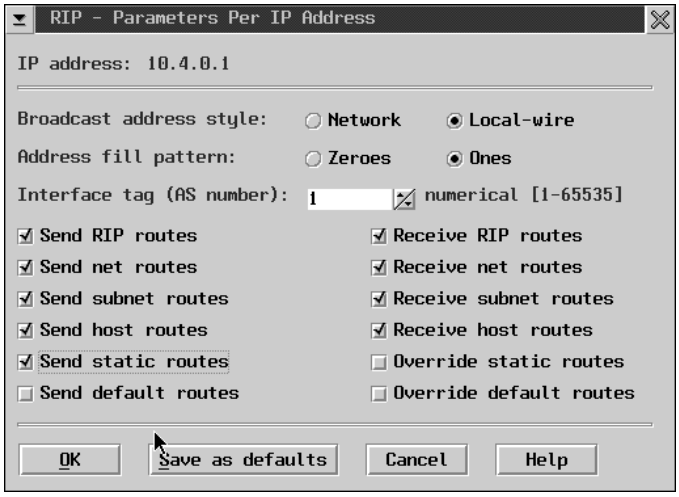

Figure 8. RIP Parameters per IP Address

To use the internal PtP link for NCP-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).

1.4 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 due to the normal completion of all session traffic on a *limited resource* link, the AOD link will automatically be reactivated when it is next 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.

If CP-CP sessions are defined over an AOD limited resource link, the link will be activated as soon as the port becomes active, and never be deactivated as the CP-CP sessions will never terminate.

1.4.1 APPN Processing

The following sections describe how AOD works.

1.4.1.1 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, node type (NN, EN, or LEN), TG number, 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 that reside on an APPN node at the end of an AOD link can be located. Once the AOD link becomes active, the network node can search the adjacent node and sessions can also be activated to undefined LUs.

1.4.1.2 APPN Topology

AOD links are reported 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, which means that topology and route selection services will still calculate routes that cross inactive AOD links. When the AOD link is actually needed again, the AOD link will be reactivated.

1.4.2 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.

1.4.2.1 Defining AOD Links

The first steps are to define the port and link stations on that port. This is shown in Figure 9 and Figure 10.

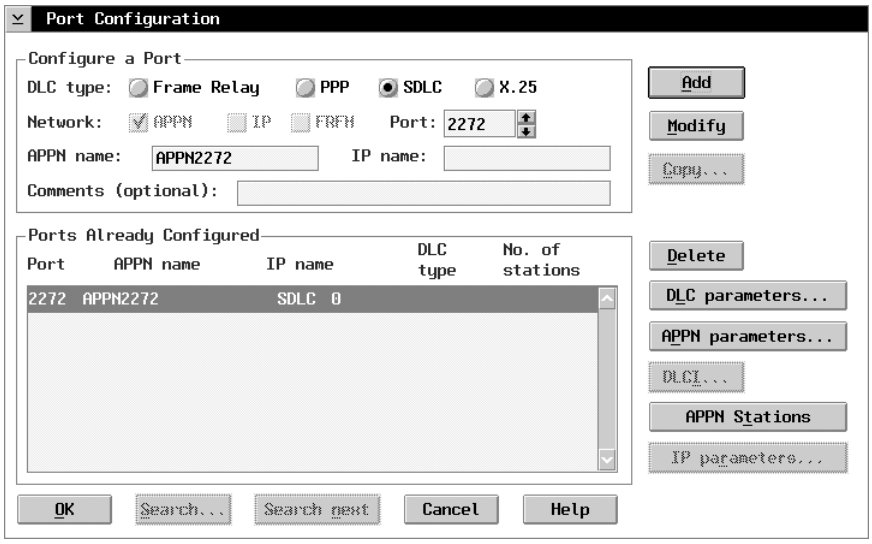

Figure 9. Port Configuration

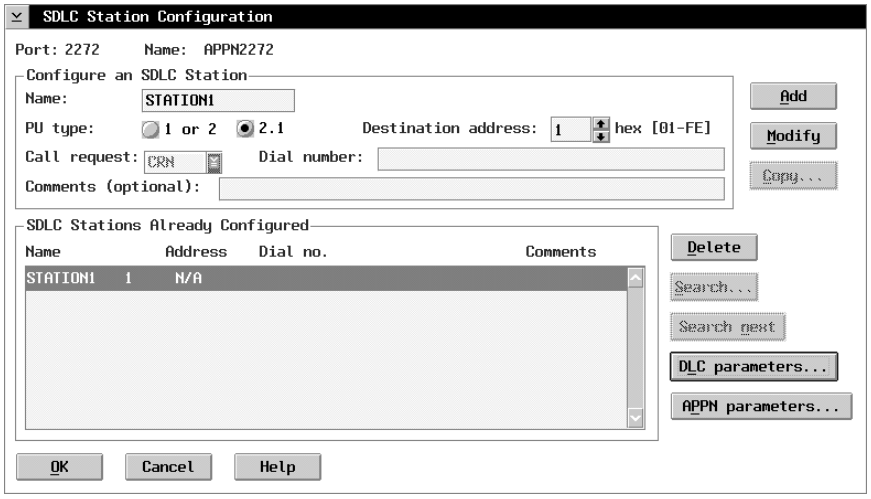

Figure 10. Station Configuration

For the station, we select **APPN parameters**. The resulting dialog is shown in Figure 11 on page 16. 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.

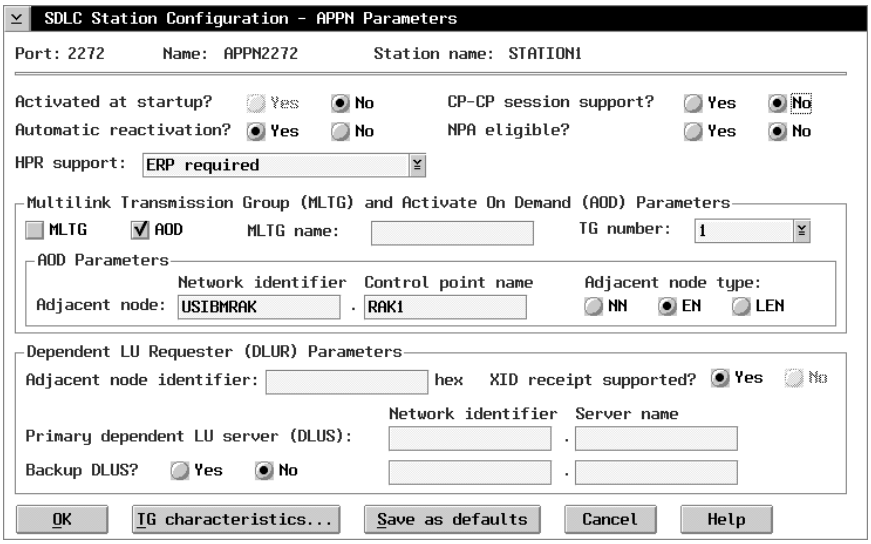

Figure 11. Station Configuration APPN Parameters

1.4.2.2 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 12.

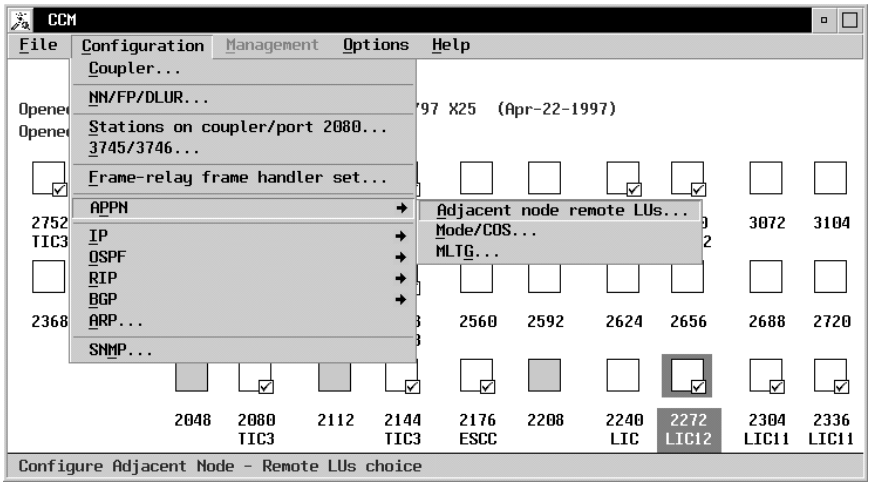

Figure 12. Adjacent Node Remote LUs

The following dialog allows the user to define the LUs on each adjacent node, shown in Figure 13 on page 17. 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 of the definitions shown only one wildcard type is needed, here all three types were shown as an example.

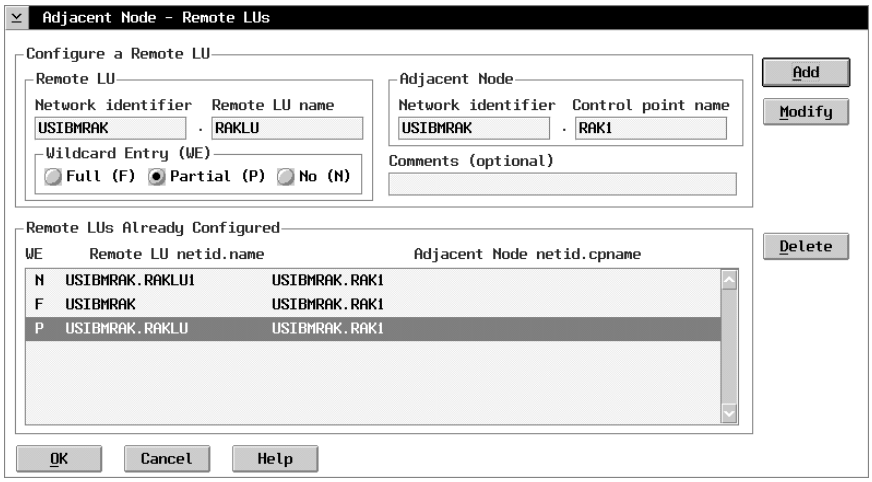

Figure 13. Predefining Remote LUs

1.5 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.

1.5.1 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 that carries this session offers.

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 non-disruptive 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.

1.5.2 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 require the DLC to guarantee delivery of packets, hop-by-hop, and to guarantee delivery of packets in the order that they were transmitted. 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.

In regards to routing and ANR labels, MLTGs are treated the same as single-link TGs. 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.

The following functions are not supported in HPR MLTG:

- Limited resource
- Connection networks
- Nonactivation XID

Much of the HPR MLTG architecture revolves around the handling of the 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.

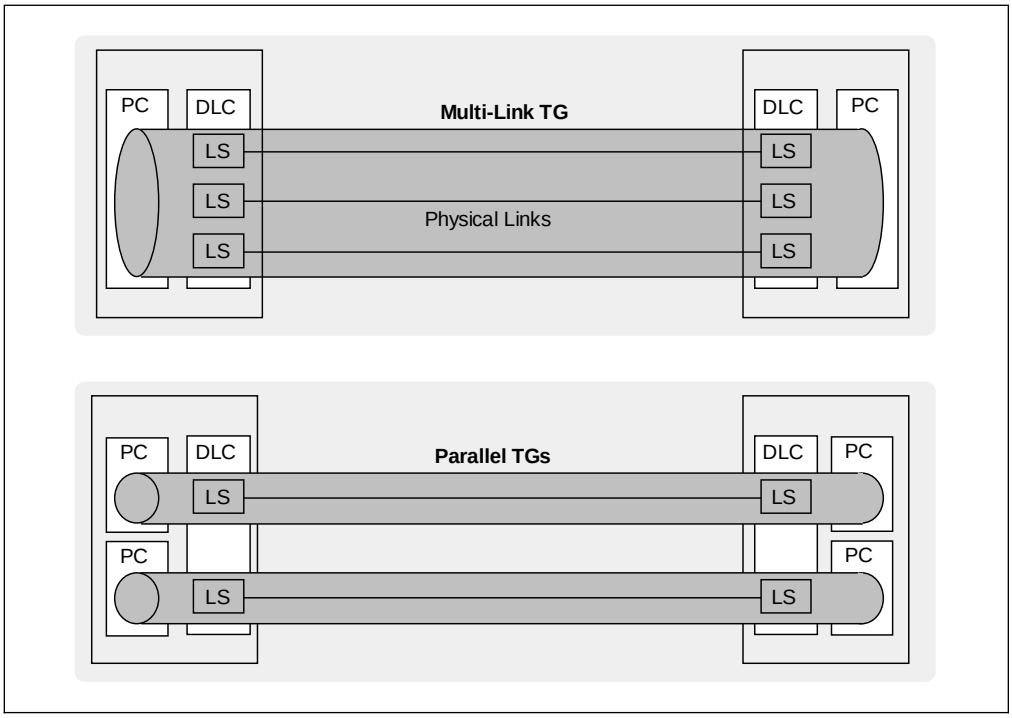

Figure 14. Multilink and Parallel TGs

1.5.3 3746 HPR MLTG Implementation

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, X.25 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 142 on page 172 shows the MLTG definition window.

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 possible to change existing MLTGs and define new MLTGs.

1.6 3746 Extended Functions

From microcode level D46130I (ECA 170) onwards a new system of password usage to activate chargeable features will be implemented in the 3746. This system has the following objectives:

- Allows the user to access new functions included in one or more feature codes, each controlled by a password.
- The password for each function is unique to an individual 3746-9x0.
- The customer pays for the functions used on each 3746-9x0.

1.6.1 Activating Extended Functions

Figure 15 shows how extended functions are activated, and where the passwords are entered.

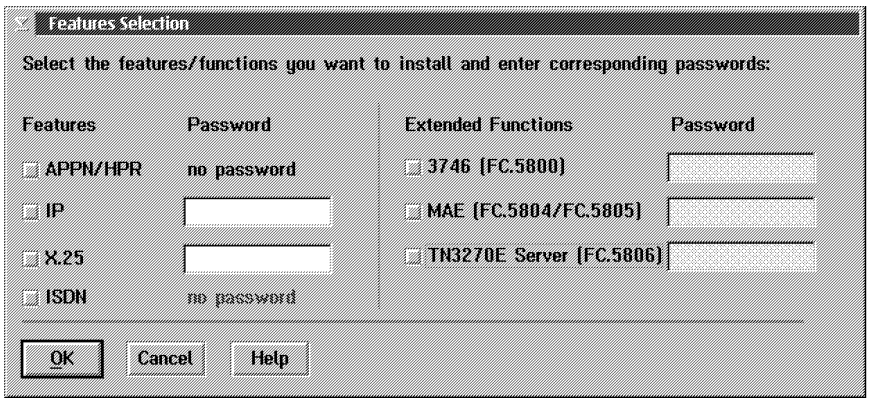

Figure 15. Activating Extended Functions

The following sections give an explanation of how functions have been divided into the new feature codes.

1.6.2 3746 Extended Functions (FC #5800)

Feature code #5800 must be ordered to get the corresponding password and operate any of the following functions.

1.6.2.1 Functions Controlled by the NNP (3746 Models 900 and 950)

The following functions are controlled by the NNP:

- HPR MLTG support over token-ring, Ethernet, SDLC, frame relay, and X.25
- Bandwidth Reservation System (BRS) for IP traffic over PPP lines
- Frame relay enhancements:
  - Frame relay switching (FRFH) - CIR
  - Bandwidth Reservation System (at DLCI level) between IP, APPN, and HPR
- X.25 (requires FC #5030):
  - SNA/DLUR, APPN, HPR, and IP over X.25
  - X.25 port sharing between NNP and NCP traffic (not NPSI)
  - PVCs and SVCs
  - NPM support (requires NPM V2R3 + PTFs)

1.6.2.2 Functions Controlled by NCP (3746 Model 900)

The following functions are controlled by NCP:

- Internal IP coupling to 3746 IP router (requires NCP V7R6)
- Dynamic windowing enhancements for frame relay/ISDN (requires NCP 7.6)
- PRI Euro-ISDN (LIC16/FC.5216) enhancements (requires NCP V7R5+):
  - Automatic backup of frame relay links over ISDN (non-disruptive for NNP to NCP connections)
  - NPM support (requires NPM V2R3 + PTFs)

1.6.3 Multiaccess Enclosure Extended Functions Part 1 (FC #5804)

Feature code #5804 must be ordered in order to get the corresponding password and operate any of the following functions supported via the multiaccess enclosure of the 3746 Models 900 and 950:

- Interactive Network Dispatcher (NetDispatcher)
- RIP V2 (on multiaccess enclosure ports only)
- Native HPR over ATM
- Enhanced ATM adapters: LIC294, LIC295
- Branch Extender
- Enhancements announced in September 1997:
  - ATM

- Frame relay, PPP, ISDN and WAN
- APPN/HPR, DLSw and IP

1.6.3.1 Detailed Description of FC #5804

The MAE Extended Functions provide a set of multiprotocol routing protocols and transport software to enable scalability and load-balancing capabilities for S/390 IP/Web servers connected to the Internet or intranet:

- Interactive Network Dispatcher

  The Network Dispatcher function provides load balancing among a set of IP servers adjacent to the router running this function. The load-balancing mechanism uses technology from IBM's Research Division to determine the most appropriate server to receive each new connection. Subsequent traffic for that connection is then forwarded to the same server. The routing is transparent to users and other applications. The load information is obtained from a set of weights based upon the number of connections active per server, the number of new connections since the last interval, feedback from response time of individual HTTP, FTP, SSL servers, and configurable policy information.

  The Network Dispatcher sees only the incoming packets from the client to the server. It does not need to see the outgoing packets, which significantly reduces the overhead imposed by load balancing. The client's packet is forwarded to the chosen server exactly as it was created. Since Network Dispatcher is also available on AIX, Windows NT, and Sun Solaris, it is useful for many applications such as e-mail servers, Web servers, distributed parallel database queries and other TCP/IP applications. The Interactive Network Dispatcher feature for Multiaccess Enclosure is priced based on the number of servers being balanced.

  Included with the Multiaccess Enclosure Extended functions is a license for Network Dispatcher for IBM Networking (5801-AAR FC #2151). One Network Dispatcher for IBM Networking Use Authorization license (program 5807-AAR FC #1453) must be purchased for each server. The user may not exceed the number of servers for which they are authorized.

- RIP Version 2

  RIP Version 2 adds the following features: route tags to propagate EGP information, subnet masks to support variable subnet masks, next hop addresses to support optimization of routes, authentication for password passing, and multicasting so that multicast can be used instead of broadcast. RIP V2 is available today on multiaccess enclosure (FC #3000) ports only.

- ATM:
  - High-Performance ATM Adapters

    1-port 155 Mbps MMF (FC #3294 - LIC294) and 1-port 155 Mbps SMF (feature number 3295 - LIC295) provide improved performance compared to LIC284 and LIC293 respectively.

  - Native APPN/HPR over ATM

    APPN/HPR supports native ATM so that the router can attach directly to ATM network without LAN emulation or encapsulation. This support includes: ATM signaling of bandwidth, QoS, ATM addressing, connection network support for SVCs, route selection extensions for ATM

characteristics, mapping between ARB and ATM characteristics, and HPR over ATM MIB extensions.

  - Native ATM bridging allows routers to connect frame relay/ATM interworking switches to devices on either PVCs or SVCs that do not support LAN emulation connections.
  - Configurable Quality of Service (QoS) allows LAN emulation networks to take advantage of ATM's QoS capabilities.
  - Next Hop Resolution Protocol (NHRP) enables shortcut routes for IP across ATM networks. NHRP supports zero-hop routing for endstations with NHRP and one-hop routing for stations without NHRP clients.
  - For added failure recovery, a backup gateway for endstations on LAN emulation can now be configured with default gateway IP addresses. If the primary gateway goes down, the backup gateway automatically starts passing packets from the endstation to other subnets. Additionally, the user can configure which ARP server is the primary and backup.
  - Server Cache Synchronization Protocol (SCSP) distributes the SRP servers to eliminate a single point of failure.
  - All the supported routed protocols and native ATM bridging may be multiplexed onto a single ATM permanent virtual circuit.

- Branch Extender

  Branch Extender is an APPN option that can be used to build a large APPN/HPR network that delivers access to customers' existing SNA applications with scalability and cost-effectiveness. It provides a gateway function to interconnect thousands of branch offices to an enterprise wide area network (WAN) and improves efficiency of network flows.

  The network node with Branch Extender addresses the problem of too many network nodes by limiting the gateway node to the topology of the branches it serves on its downstream links and uses a form of default routing for network services on its upstream link. In this way, Branch Extender reduces topology database size and traffic, adds the ability to register resources from a branch to a central directory server in the WAN, and allows for direct connections between subnetworks without using an extended border node. This function is available for any data link type supported by APPN.

- Frame relay, PPP, ISDN, and WAN enhancements:
  - Frame relay

    Frame relay dial circuit interface is configurable on a V.25bis interface type.

    Frame relay data compression (mode-1 FRF.9) is configurable per PVC to run over a frame relay interface.

    Congestion can be reduced with support for congestion management via CLLM messages: SNMP traps are sent on receipt of CLLM, FECN, or BECN frames, the 3746 reduces the transmission rate upon receipt of a FECN, and sends a BECN.
  - PPP

    PPP Bandwidth Allocation Protocol/Bandwidth Allocation Control Protocol (BAP/BACP) adds the ability to dynamically add or drop links over ISDN B channels.

Authentication servers can now be used so that names and passwords need not be configured at each router.

Encryption Control Protocol (ECP) using Data Encryption Standard (DES) Cipher Block Chaining (CBC) mode is now available for PPP.

– ISDN

ISDN I.430 and I.431 is supported to enable interconnecting to the leased-line service from NTT.

– Ethernet

The Ethernet locally administered MAC address can be configured to override the default burned-in address.

– WAN

Bandwidth Reservation (BRS) can assign TCP/IP packets to a BRS class and priority based on the packet's UDP or TCP port number.

A backup frame relay, PPP, or X.25 link can be specified for IP over frame relay when the traffic rate reaches a specified threshold.

Enabling or disabling of adapters can be done from a single operator console without knowing which interface(s) is configured for WAN reroute.

- APPN and DLSw enhancements:

  – APPN/HPR

  Native HPR over ATM (see ATM section)

  Implicit focal point and up to eight backups enable the router to initiate a management session with NetView.

  – DLSw

  A range of source/destination SAPs and MAC addresses can be configured to override circuit priority.

  The DLSw Switch-to-Switch Protocol allows the exchange of a MAC address list between partners.

  NetBIOS session alive spoofing eliminates session alive frames on a dial-on-demand link.

1.6.4 Multiaccess Enclosure Extended Functions Part 2 (FC #5805)

FC #5804 must be ordered to get the corresponding password. FC #5805 is required to operate any of the following functions supported via the Multiaccess Enclosure (MAE) of the 3746 Models 900 and 950:

- FDDI, HSSI, F-Enet, 128 MB (MAE memory)
- ESCON MultiPath Channel (MPC+) support for IP applications
- Enterprise Extender for APPN/HPR over IP backbone
- Channelized T1/E1 and other enhancements

1.6.4.1 Detailed Description of FC #5805

FC #5805 requires FC #5804 and provides the following additional functions for the MAE (FC #3000 and #3001):

- FDDI support

  One FDDI interface per adapter (FC #3286 - LIC286). Operates as either a Dual Attach Station (DAS) or a Single Attach Station (SAS) using multimode fiber (MMF).

- HSSI support

  One HSSI interface per adapter (FC #3289 - LIC289). Supports T3 and E3 speeds.

- 10/100 Mbps Ethernet support

  One interface per adapter (FC #3288 - LIC288) with speeds of either 10 Mbps or 100 Mbps.

- 64 MB Additional system memory (FC #3520)

  An optional second 64 MB DIMM is available on Multiaccess Enclosure for a total of 128 MB of memory for especially demanding environments (TN3270E, DLSw, APPN, DLUR, etc.).

- High-Performance Data Transport (HPDT) for UDP

  High-Performance Data Transfer (HPDT) MultiPath Channel (MPC), also known as MPC+, has been extended to include IP support over the ESCON channel. HPDT UDP extends the efficiencies of HPDT services to applications using the OS/390 UNIX System Services UDP interface. HPDT reduces CPU cycle consumption and achieves a more efficient transfer of data. HPDT UDP is initially targeted for communications between DB2 on OS/390 V2R4 (requires PTF PQ04890) and SAP R/3 application servers. Other UNIX System Services socket applications using UDP, such as NFS and DCE, can also transparently take advantage of HPDT UDP services over the Multiaccess Enclosure ESCON channel.

- High-Performance Data Transport (HPDT) for TCP

  HPDT TCP/IP extends the efficiencies of HPDT services to IP applications using OS/390 V2R5. HPDT reduces CPU cycle consumption and achieves a more efficient transfer of data. It is supported over the ESCON channel of the MAE.

- Enterprise Extender

  Enterprise Extender is a simple set of extensions to APPN High-Performance Routing (HPR) technology to integrate SNA into IP backbones. To the HPR network, the IP backbone is a logical link; to the IP network, the SNA traffic is UDP datagrams.

  Enterprise Extender provides the flexibility for SNA parallel sysplex features that are currently available in HPR networks now to be available to users in networks that have IP backbones, or even IP clients when coupled with TN3270e server support. Enterprise Extender also makes it possible for SNA networks to use IP attachments as alternate and backup routes for the SNA network.

  Enterprise Extender technology can also reduce the demands on the data center routing platforms, and, thus, provide a more cost-effective solution than

other integration technologies. Enterprise Extender seamlessly routes packets through the network protocol edges eliminating the need to perform costly protocol translation and the store-and-forward associated with transport layer functions.

The Enterprise Extender technology also provides many of the traffic control features that SNA users have come to expect. Using Class of Service (COS), SNA applications specify the nature of the services they require from the network (for example, interactive, batch, etc.). Enterprise Extender supports SNA priority in IP environments by mapping the SNA COS priority to UDP port numbers that can be easily prioritized using Bandwidth Reservation System (BRS).

- Channellized T1/E1 support

  This support allows the current ISDN PRI adapter to be configured as a channellized T1 or E1 in lieu of using it for ISDN PRI. Support is provided for frame relay and PPP over individual or groups of DS0s. One or multiple connections are supported on the same physical interface. The bandwidth of each connection will be a multiple of 64 kbps up to the maximum speed of 24*64 for T1 or 31*64 for E1. The time slots for the combined DS0s need not be contiguous.

- Dial-in support for SDLC PU Type 2 devices

  Switched dial-in is the capability for SDLC PU Type 2 devices to dial through a switched data network. The support will be provided through DLSw. It will provide HDX and FDX support as well as NRZ and NRZI. Call answering is supported but a dial-out facility is not.

- X.25 scalability on the Multiaccess Enclosure

  This extends the current X.25 capacity from a limit of 239 VCs to a limit that is memory-dependent and capable of supporting more than 1000 VCs.

- RIP outage-only advertisements:

  Allows RIP advertisements to only be sent on an interface when a route is missing from the IP route table. This will facilitate advertisement on ISDN/V.25bis Dial-on-Demand (DoD) circuits only in circumstances where the DoD circuit also has data traffic to send. Circuits will not be brought up solely for RIP advertisements.

1.6.5 Multiaccess Enclosure TN3270E Server (FC #5806)

This feature must be ordered to get the corresponding password and operate the TN3270E server. This function is supported via the Multiaccess Enclosure of the 3746 Models 900 and 950 and requires the MAE Extended Function FC #5804.

The TN3270E server provides a TN3270 gateway function for TN3270 clients downstream from an SNA/VTAM S/390. The clients connect to the server using a TCP connection which the server then maps to a corresponding SNA LU-LU session that the TN3270 server maintains with the S/390.

The TN3270E server supports the capabilities defined in RFCs 1576, 1646 and 1647.

The connectivity from the TN3270E server to the S/390 uses DLUR in the Multiaccess Enclosure and is supported, locally and remotely, over all the interfaces

that support DLUR. When coupled with Enterprise Extender, the TN3270 servers can be distributed in the network with an IP infrastructure and, therefore, be placed in locations that provide the best scalability and availability without regard to backbone protocol.

1.7 3746 CD-ROM and Optical Disk Support

3746 Licensed Internal Code supporting new functions available December 1997 (or later) will be delivered on CD-ROM only (media FC #9300). Customers with an optical disk reader on their service processor, who wish to use new functions *must* upgrade their service processor or install the latest service processor.

For service processors not equipped with the CD-ROM reader, the Optical Disk media remains available after December 1997 for microcode maintenance up to ECA 170/EC level D46130x. (x = alphabetic letter identifying the technical level of the microcode; x includes possible future levels.) It will also be available for shipment of 3746 MES which does not require any of the new functions available December 1997 or later. For example, an order for processors Type 3 (ESCP3/TRP3/CLP3/CBSP3) is in this category and may include the Optical Disk media (FC #9500).

1.7.1 Required Action

Configurations not yet submitted to AAS must be updated before they are entered in AAS. If you enter an order based on an old CFREPORT file, and your order includes functions available December 1997 or later, make sure that the media feature included in the configuration is FC #9300 (CD-ROM).

Important Reminder

Any order without a service processor *must* include a media feature number (either #9300 or #9500). This requirement applies whether this order is for a new 3746 or for a 3746 MES. (This is applicable to all shipments on or after December 1997.)

1.7.2 Recommendation

CF3745 was enhanced November 1997 to automatically determine the media and service processor features required to support the hardware/microcode functions being configured for a 3746 machine or MES.

It is strongly recommended that existing orders/configurations are validated via CF3745 before entering new orders or altering orders planned to be shipped December 1997 or later.

1.7.3 Service Processor (SP) Prerequisites for CD-ROM Support

The following is a reminder of the SP support for the 3746 functions available starting December 1997:

- Service Processor Type 2 (feature 5052): 3746 functions are supported.
- Service Processors Type 1 (feature 5021): 3746 functions are supported if the SP is equipped with the 64 MB memory expansion (feature 5028) and a 2 GB hard disk drive.

The 3172-based service processors (P/N 41H7520 and P/N 55H7630) are equipped with a 2 GB hard disk drive, but the 9585-based service processor requires the SP hard disk drive Upgrade (feature 5026)

- Oldest Service Processors (feature 5020): 3746 functions are not supported. A Service Processor Type 2 (feature 5052) must be ordered.

Note: The Service Processor Type 2 contains a CD-ROM reader. For the new functions available December 1997 or later, or to benefit from CD-ROM specific functionalities (see below), the Service Processors Type 1 (feature 5021) must be upgraded with the previously mentioned features and with the SP CD-ROM Upgrade (feature 5051). However, it is strongly recommended not to upgrade with these features but to go directly to the Service Processor Type 2 (FC #5052).

1.7.4 The Benefits of CD-ROM Media

Using CD-ROM media provides the following benefits:

Dual Load Module Support

Two levels of 3746 microcode can be installed on the hard disk drive (hard disk drive) of the Service Processor. Benefits of this new support are discussed in the following sections.

Online Microcode Upgrades

The 3746 remains operational during the installation of a new microcode level. While the new microcode is being installed as the non-active level, the 3746 remains in production on the current active level. If the Service Processor (SP) operating system is modified, the SP functions such as CCM support, are not available during this step.

The 3746 is interrupted only during activation of the non-active level (re-IML), which lasts about 5 to 15 minutes, depending on the number of processors in the 3746 configuration. The customer may perform this activation at a later time, when compatible with the network operations. Compared to a microcode upgrade via optical disk media, the 3746 interruption time is reduced from 150 minutes (average) to less than 15 minutes.

Backup Production Microcode Level

After the new microcode level is activated, the previous production level becomes non-active, but remains available for possible re-activation if for any reason the customer needs to switch back to this production level. Should this happen, the 3746 interruption time is again limited to the re-IML time of the 3746.

Testing the Non-Active Microcode Level

The dual level of code also allows the users to test a new level of microcode, maintenance level or functional level, by taking advantage of unused machine time; they can activate a trial level of microcode and then come back to the production level when required.

1.7.5 Online Documentation

The CD-ROM media containing the new Licensed Internal Code of 12/12/97 will also include the following 3746 product information:

- All the customer publications, including the *3745/3746 Overview,* GA33-0180-8, *3745/3746 Planning Guide,* GA33-0457-0, *CCM User's Guide* and service publications (installation and maintenance, etc.).

- Most of the 3746 presentations and documents available today on the IBM intranet home page (http://w3.lagaude.ibm.com/ccp/3746.htm).

Netscape and Acrobat Reader (included in the CD-ROM) allow the user to display the above information at the service processor screen.

1.7.6 Stand-Alone CCM (3746 Controller Configuration and Management)

The following possibility is under study for the near future: all supported levels of the stand-alone CCM would be available on the CD-ROM containing the Licensed Internal Code. This would allow the user to retrieve the CCM level corresponding to the microcode level of his/her 3746(s) and run this CCM on a workstation. Also under study is a simple procedure (no need for diskettes) for installing this CCM on the workstation.

1.7.7 Optical Disk (OD) and CD-ROM Support Details

Data migration from OD to CD-ROM is supported from any microcode level supporting optical disk to any microcode level supporting CD-ROM. Customer parameters (MOSS-E) are saved on diskette and no longer on a second optical disk.

The following are the microcode EC levels:

- D46130x

  D46130 is the *last* microcode EC supporting OD media (and the previous microcode installation methodology). See Table 3 for service processor requirements.

- F12380x - December 1997 GA level of microcode

  F12380 is the *first* microcode EC supporting CD-ROM media and only this media (and the new microcode installation methodology). See Table 3 for required service processor features.

Table 3. Service Processor Requirements

Service Processor Type	**Microcode Delivery Media**	
	CD-ROM	**Optical disk**
Microcode Level	**F12380X**	**up to D46130X**
7585 (Pentium-200Mhz) FC 5052	Yes 1 2 With CD enabled	Yes 2 With OD enabled
3172 (Pentium-90Mhz) FC 5021	Yes With FC 5051 With FC 5028	Yes 3
3172 (486-66Mhz) FC 5021	Yes 4 With FC 5051 With FC 5028	Yes 3 4
9585 (486-66Mhz) FC 5021	Yes 4 With FC 5051 With FC 5028 With FC 5026	Yes 3 4
9577 (486-33Mhz) FC 5020	No	Yes 3 4

CD-ROM operations require the service processor to be equipped with a CD-ROM drive, 2-GB hard disk drive and 96 MB of memory.

Notes:

1 Starting December 1997, SP Type 7585 is equipped with a CD-ROM reader only and no OD drive, 96-MB memory and 2-GB hard disk drive.

2 Between June 1997 and December 1997, SP Type 7585 was equipped with an OD drive and a CD-ROM reader. The OD drive was enabled as default. Starting December 1997, CD-ROM media will be used for shipment of microcode when an MES is ordered for a 3746 using such an SP (CD-ROM media, FC.9300, included in the MES by CF3745). The existing OD drive will be removed and the CD-ROM drive enabled when the MES is installed.

3 See requirements for FC 5026 and FC 5028, depending on functions configured on the 3746 (APPN, MAE, etc.). Refer to CF3745 for detailed information.

4 For better response time at the operator console, it is strongly recommended that this SP is replaced by a 7585 (FC 5052).

1.8 Frame Relay Support History

In addition to being an APPN/SNA node or an IP router, the 3746 is also a frame relay switching node. For IP and APPN/SNA traffic, it also supports connection to a frame relay network as a frame relay terminating equipment (FRTE) node. The following sections give an overview of the history of frame relay support in ACF/NCP, 3745, and 3746.

1.8.1 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 non-A 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 endstation, called Frame Relay Terminal Equipment (FRTE). This is referred to as NCP's INN FRTE function.

Note: A frame relay endstation is known as an FRTE but is sometimes referred to as an FR Data Terminal Equipment (DTE) analogous with X.25 terminology.

1.8.2 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 Handler (FRFH) function.

This allows connection of FRTEs to a network of 3745s and provides end-to-end connections (PVCs) for FRTE pairs. 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 Data Communications Equipment in analogous (DCE) 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.

1.8.3 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).

1.8.4 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 adapters. 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

Frame relay switching between lines on the 3746 Model 900 may be defined in NCP. The information is passed to the Communication Line Adapters (CLAs) on the 3746 Model 900, which provide outboard frame switching totally independent of the 3745.

3746 LMI support

LMI support was added for 3746 Model 900 connected lines.

NCP V7R2 introduces communications rate (CR) support, which allows users to allocate a minimum bandwidth to each virtual circuit, depending on the traffic needs of the corresponding endstations. 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 by other active virtual circuits.

Note: The communication rate must be defined per PVC segment. An end-to-end communication rate requires consistent definitions on each of the PVC segments that comprise a virtual circuit.

1.8.5 ACF/NCP Version 7 Release 3

NCP V7R3 (March, 1995) had the following frame relay support enhancement:

Multiple BNN stations per DLCI for 3746 lines

Support of multiple BNN stations per DLCI on 3746 connected lines was added.

Frame relay over token-ring

Frame relay over token-ring 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 (supported on 3745 base adapters only).

IP over frame relay

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 the permanent dynamic reconfiguration function available with VTAM Version 4 Release 3.

The IP over frame relay capability is RFC 1490 compliant.

Frame relay BAN support

The frame relay boundary access node (BAN), which uses the RFC149 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)

1.8.6 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 (FRFH) 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 be 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.

1.8.7 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 microcode level D46120B 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 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 error 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

3746 microcode level D46120B together with NCP V7R5 supports the following 3746 DLCI sharing options:

- One INN station (routed or bridged frame format) per DLCI
- Multiple BNN/APPN stations (routed or bridged frame format) controlled by NCP or NNP, and 3746 IP per DLCI

Enhanced dynamic windowing algorithm

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 to specify the upper boundary of this delay.
2. To prevent NCP from over-reacting 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.

1.8.8 ACF/NCP Version 7 Release 6

NCP V7R6 (September 1997) had the following frame relay support enhancements:

NCP communications rate enhancements
NCP COMRATE enhancements allow the definition of a relative priority for each protocol that is sharing a DLCI (FRRATES keyword).

NCP dynamic windowing enhancements
The granularity of the enhanced dynamic windowing algorithm frame loss (Dw) and congestion detection (Dwc) suboperands has been improved.

Switched frame relay lines
NCP controlled switched frame relay lines carrying SNA peripheral and subarea traffic are now supported.

Congestion control parameters set via NDF
3746 attached frame relay lines can now have their congestion control parameters (for FECN, BECN and DE) set via NDF parameters. Previously this could only be set from MOSS-E.

1.9 Service Processor (SP) and Network Node Processor (NNP) Support Matrix

The following table gives an overview of the service processor (SP) and network node processor (NNP) features and machine types needed to support each of the available 3746 hardware configurations. All numbers prefixed with a # are feature codes.

Table 4 (Page 1 of 2). Service Processor (SP) / Network Node Processor (NNP) Feature Code and Machine Type

Function	Service Processor					NNP	
	#5020	#5021			#5052 [1]	#5022	#5122 [2]
	9577 [3]	9585 [4]	3172 [5]	3172 [6]	7585 [7]	3172 [8]	7585 [9]
NNP traffic only	Yes	Yes	Yes	Yes	Yes	[11]	[11]
Second expansion enclosure (#5016)	[10]	#5028	#5028	#5028	Yes	[11]	[11]
SP Sharing	[10]	#5028	#5028	#5028	Yes	[11]	[11]
APPN/HPR (NNP) **IP** (#5033)	#5026	Yes	Yes	Yes	Yes	Yes	Yes
APPN/ISR/DLUR/RTP More than: -3000 PUs -9000 LU-LU sessions	#5026	Yes	Yes	Yes	Yes	#5027	#5027
APPN/ISR/DLUR/RTP More than 15000 LUs-LUs sessions (up to 30000) [12]	#5052	#5026 #5028 #5051 [10]	#5028 #5051 [10]	#5028 #5051	Yes	#5211	Yes
3746 -#5802 (SSE) -More than 120 lines controlled by the NNP (up to 240) -CD-ROM support -APING from SP	#5052	#5026 #5028 #5051 [10]	#5028 #5051 [10]	#5028 #5051	Yes	Yes	Yes
MAE -#3000/#3001 -Extended function (#5804) -Extended functions 2 (#5805) -TN3270e server (#5806) -HSSI (T3/E3 speeds) (#3289) -Fast Ethernet (#3288) -FDDI (#3286) -64MB memory expansion (#3520)	#5052	#5026 #5028 #5051 [10]	#5028 #5051 [10]	#5028 #5051	Yes	Yes	Yes

Table 4 (Page 2 of 2). Service Processor (SP) / Network Node Processor (NNP) Feature Code and Machine Type

Function	Service Processor					NNP	
	#5020	#5021			#5052 [1]	#5022	#5122 [2]
	9577 [3]	9585 [4]	3172 [5]	3172 [6]	7585 [7]	3172 [8]	7585 [9]
Previews • EBN	#5052	#5026 #5028 #5051 [10]	#5028 #5051 [10]	#5028 #5051	Yes	Yes	Yes

Notes:

[1] Service Processor Type 2

[2] Network Node Processor Type 2

[3] 486-33Mhz

[4] 486-66Mhz

[5] 486-66Mhz P/N 41H7520

[6] Pentium-90Mhz P/N 55H7630

[7] Pentium-200Mhz

[8] Pentium-90Mhz

[9] Pentium-200Mhz

[10] Recommended alternative: Replace the service processor by a service processor type 2 (#5052) for better response times at the operator console.

[11] No requirement on the NNP.

[12] Requires also a CBSP3 or CBSP3 upgrade.

1.9.1 Service Processor and Network Node Processor Feature Codes

The following is a description of SP and NNP feature codes.

#5020 Service Processor (Type: 9577)

#5021 Service Processor (Type: 9585, 3172 P/N 41H7520 or 3172 P/N 55H7630)

#5022 Network Node Processor (Type: 3172 P/N 41H7522, no longer orderable)

#5026 Service Processor HDD upgrade (HDD 2GB/1GB formatted)

#5027 Network Node Processor Memory Expansion (64MB available on #5022 and #5122)

#5028 Service Processor Memory Expansion (64 MB available on #5021)

#5029 Service Processor Rack Mount Kit (for SP re-installation in the Controller Expansion #5023). #5029 may be required if rack mount equipment is not available. It contains:

- 2 brackets support and 1 plate for the Display Screen
- 2 brackets support for any kind of Service Processor (but 7585)
- 1 drawer kit for Keyboard/Mouse
- 1 plate for the Modem and Optical Disk Drive

#5051 Service Processor CD-ROM drive (Provides a CD-ROM drive for SP #5021)

#5052 Service Processor Type 2 (Type: 7585 with 96MB memory, 2GB HDD, Cd-ROM drive)

#5122 Network Node Processor Type 2 (Type: 7585)

1.10 3746 Microcode EC and ECA Levels

The following describes the functional levels of microcode (ECA), starting with the latest announced ones (possibly not yet available). For each level, only the newly supported functions are indicated, as they are cumulative with the functions of all the other levels down in the list (older levels). Once a new level is available, all the new machines (and hardware MES that ship with microcode) are automatically shipped at this level of microcode. The new level is not distributed to the installed 3746-900s and 3746-950s, except some levels such as the engineering change (EC) level D22510K (ECA number 142) of 12/95.

If the minimum EC level for a requested function is not installed or planned to be installed (for example, as part of an MES) order the most current ECA number via the IBM service representative. When an MES is shipped with microcode (see description below), this microcode is shipped at the most current functional level (ECA number) and technical level (EC number).

The EC number indicated for a given ECA number reflects the technical level of the microcode at the availability date of the ECA. The alphabetic index (A, B, etc.) complementing the EC number is incremented when a new technical level starts shipping. Starting with ECA 167, the operator can display the active ECA and EC numbers of the 3746-9x0(s) connected to the service processor:

1. EC level F12380x (ECA number 175)

   Availability: December 1997

   Automatically shipped with MES containing one of the following feature numbers: 3294, 3295, 5033, 5052, 5122 and 5804. If the minimum EC level D46130D is not installed, the new level is shipped with any MES containing feature numbers 3000, 3287, 5051, 5123, 5203, 5523, 5623, and 5800. If minimum EC level D46120 is not installed, the new level is shipped with any MES containing one of the following feature numbers: 5016, 5027, 5028, 5030, 5216, 5631 or 5032.

   This ECA may be ordered only for *installed* 3746s running at a lower ECA number/EC level and requiring one of the new microcode functions listed hereafter:

   - APING from service processor to APPN/HPR PUs
   - Dual level of code (active and non-active)
   - CD-ROM support (reduced configuration/maintenance time)

   For features 5052 (SP Type 2) and 5051 (SP Upgrade Type 1) shipped after December 1997 this microcode is automatically shipped.

2. EC level D46130I (ECA number 170)

   Availability: 30/11/97 (older levels no longer orderable)

   Shipped on OPTICAL DISK only (for SP with Optical Disk Drive) For service processors configured for CD-ROM support, see ECA 175.

If not installed, automatically shipped with the first MES containing one of the following feature numbers: 3287, 5016, 5022, 5027, 5028, 5030, 5033, 5122, 5123, 5203, 5216, 5523, 5623, 5631, 5632, 5800

This ECA may be ordered only for INSTALLED 3746s running at a lower ECA number/EC level and requiring microcode function(s) which are not orderable via feature code, such as the functional improvements of older ECAs (167 and before).

In particular, feature number 5800 (3746 Extended Functions 1) must be ordered to operate any of the following functions:

- 3746-NCP internal IP coupling (requires NCP V7R6)
- Dynamic windowing enhancements for frame relay/ISDN (NCP 7.6)
- X.25 traffic controlled by the NNP (independent from NCP)
- 3746-controlled traffic:
  - SNA/DLUR, APPN, HPR and IP
  - PVCs and SVCs
  - X.25 port: sharing between 3746 NN, 3746 IP and NCP ODLC traffic
  - NPM support (NPM V2R3 + PTFs)
- HPR MLTG support over TR, Ethernet, SDLC, frame relay and X.25
- Bandwidth Reservation System (BRS): 3746 PPP lines (IP traffic)
- Frame relay switching (FRFH): 3746 lines controlled by the NNP
- CIR for 3746 frame relay lines controlled by the NNP
- Bandwidth Reservation System at DLCI level: IP, APPN, HPR
- PRI Euro-ISDN (LIC16/FC.5216) enhancements (3746-900/NCP V7R5+):
  - Automatic backup of frame relay links over ISDN (nondisruptive for NCP to NCP connections)
  - Support of LIC16 for NPM (NPM V2R3 + PTFs)

3. EC level D46130B (ECA number 167)

   Availability: June 1997

   Automatically shipped with MES containing one of the following feature numbers: 3000, 5022, 5033, 5052 or 5122. If the minimum EC level D46120 is not installed, the new level is shipped with any MES containing one of the following feature numbers: 5016, 5027, 5028, 5030, 5216, 5631 or 5032.

   This ECA may be ordered only for *installed* 3746s running at a lower ECA number/EC level and requiring one of the new microcode functions listed hereafter:

   - HPR/RTP and ARB: TR, Ethernet, SDLC, frame relay and ESCON
   - 3746 NN connectivity increase: More than 4000 PUs (up to 5000) and more than 12000 APPN/DLUR data sessions (up to 15000)
   - CCM: Configuration checking versus installed configuration (CDF-E)
   - Display of the 3746 microcode level (ECA, EC) at MOSS-E console

4. EC level D46120A (ECA number 159)

Availability: December 1996

- IP over leased lines (PPP) and frame relay lines
- NPM support: Ports/stations (APPN/HPR), Processor/TIC3 utilization
- Frame relay line and DLCI sharing: NCP (V7R5), 3746 NN and 3746 IP

5. EC level D46120 (ECA number 157)

   Availability: November 1996

   - 3746 NN connectivity increase: 4000 PUs + 12000 APPN/DLUR data sessions per 3746 NN.
   - Second expansion enclosure (FC.5016): Six processors (total 16).
   - Network node processor memory expansion (FC.5027).
   - Service processor memory expansion (FC.5028).
   - EMEA only: LIC16 (FC.5216) for Euro.ISDN primary support by NCP V7R5 (3746-900).
   - FR BAN for INN traffic with remote 372x connected to 2210/2216 (NCP V7R5).
   - Backup network node processor (FC.5022).
   - Optimization of processor storage utilization: Increased APPN and DLUR connectivity (more PUs/sessions per processor). Up to 100% connectivity increase.
   - Selective load of microcode (APPN/HPR, IP) by category of processors: ESCP2, CLP, TRP2 and (CB)TRP2.
   - NNP no longer required for DCAF TCP/IP consoles on service LANs.

6. EC level D46100A (ECA number 155)

   Availability: September 1996

   - Ethernet interfaces (FC.5631, FC.5632): APPN/DLUR, HPR, IP, etc.
   - 3746 IP Routing (FC.5033) over ESCON, token-ring, Ethernet
   - 3746 HPR/ANR over ESCON, TR, Ethernet, SDLC, frame relay
   - 3746 APPN/DLUR over frame relay (BNN and BAN)
   - Maximum connectivity:
     - 3746-950: 500 PUs per CLP
     - 3746-900: 1000 PUs per CLP, including NCP-controlled frame relay PUs and X.25 PUs
   - APPN/DLUR performance improvement by 30-40% in transactions/sec and data throughput
   - Port sharing:
     - ESCON port: traffic for 3746 NN, 3746 IP, NCP-A and NCP-B
     - TIC3 port: traffic for 3746 NN, 3746 IP and one NCP
   - NetView Topology Manager (3746 APPN/HPR): 3746 local topology
   - Network Management (3746 IP): SNMP (NetView/AIX), Telnet, CCM
   - CCM: 3746 online configuration changes, including ESCON, with:

  - Automatic deactivation/activation of impacted resources
  - Delete/Copy/Search functions
  - 3746-900 with NNP: ESCON Generation Assistant replaced by CCM
- Year 2000 support (applicable also to 3745 Models xxA)
- Network node processor installation option: No loading of the APPN/HPR microcode in the 3746-900 processors

7. EC level D22560D (ECA number 146)

   Availability: 03/29/96

   - 3746 NN support: LIC11, DLUR for LIC11/LIC12, four LICs per CLP, CLP backup, 16 versus four host link stations per ESCP2, 3000 versus 1500 LU sessions per TRP2, 2000 versus 1000 adjacent nodes per 3746 NN and 6000 versus 3000 LU-LU sessions per 3746 NN
   - X.25 Support - FC 5030 (3746-900/NCP V7R4)
   - Non-ERP support over TIC3 for CNN ANR (3746-900/NCP V7R4)
   - 3746-900: 3000 versus 1000 station (aggregate number of PUs over frame relay and virtual circuits over X.25)

8. EC level D22560A (ECA number 144)

   Availability: January 1996

   - 3746 NN support (3746-900, 3746-950): APPN over LIC11, TIC3, and ESCON. DLUR over TIC3. Network node processor (FC.5022).
   - EGA integrated to MOSS-E (3746-900/NCP).
   - 2000 versus 500 PUs on (CB)TRP2 (3746/NCP).

9. EC level D22510K (ECA number 142)

   Availability: December 1995 - automatically shipped to all the installed 3746-900s.

   - V.24 interface (LIC11) support at speeds up to 28.8 kbps (V.35 modems)
   - Native IP routing over ESCON channels (3746-900/NCP V7R3)
   - Frame relay frame switching between 3745 & 3746-900 (NCP V7R4)
   - ISDN Terminal Adapter (BRI/PRI) attachment (3746-900/V7R2)
   - CLP: 1000 stations (aggregate number of PUs over frame relay and virtual circuits over X.25), in addition to 1000 SDLC PUs

10. EC level D22510J (ECA number 138)

    Availability: August 1995

    - CLP: Frame relay BAN and frame relay SAP multiplexing (3746-900/NCP V7R3)

11. EC level D22510I (ECA number 137)

    Availability: June 1995

    - CLP: X.25 (3746-900/NCP V7R3/NPSI V3R8)
    - NPM 6-900/NCP V7R3)
    - ISDN Terminal Adapter (BRI/PRI) attachment (3746-900/V7R2)

- CLP: 1000 stations (aggregate number of PUs over frame relay and virtual circuits over X.25), in addition to 1000 SDLC PUs

12. EC level D22510D (ECA number 134)

    Availability: September 1994

    - CLP: Frame relay (3746-900/NCP V7R2)
    - CLP: Up to 120 versus 100 active lines
    - CLP: Up to 1,000 versus 500 active PUs

Chapter 2. Introduction to Frame Relay

At the end of the eighties and beginning of the nineties important technological changes were taking place which played a major role in the development of frame relay.

The introduction of, usually LAN attached, intelligent workstations changed the data processing paradigm from centralized host computing to distributed processing. The growth of distributed processing, the need for LAN interconnection, and the growing use of graphics and images has lead to exponentially increasing network traffic. In addition to the larger data volumes, traffic patterns are also changing. Traditional networking coped with certain amounts of host interactive traffic that was relatively constant during business hours and used bandwidth offshift for batch file transfer. Due to distributed processing and LAN interconnection today's networks are confronted with traffic patterns which are more unpredictable.

Not only has the demand for connectivity changed, but also the technology to provide networking facilities has been subject to important changes. The introduction of digital and fiber technologies provides faster and more reliable communication but also requires networking technologies which are able to efficiently operate at higher speeds. In order to meet this requirement the concept of fast packet switching has been developed.

2.1 Packet Switching Techniques

Before discussing what fast packet switching, is we first describe packet switching and circuit switching, which is also referred to as time division multiplexing (TDM).

2.1.1 Packet Switching

Packet switching networks allocate capacity only when stations have packets to send, rather than dedicating a portion of their capacity to each active station on the network. Examples of packet switching architectures are X.25, SNA and IP. X.25 and SNA use connection-oriented network protocols meaning that data transfer between endstations requires an end-to-end relationship to be established first. IP uses a connectionless network protocol, and endstations may send data without a prior connection establishment.

Packet switching architectures make packet forwarding decisions based on addresses imbedded in the data. Therefore, the packet switch has to understand the addressing structure and routing information contained within the data it is forwarding. Because the packet switch understands the addressing, it can accept data over a single link and forward it to multiple destinations.

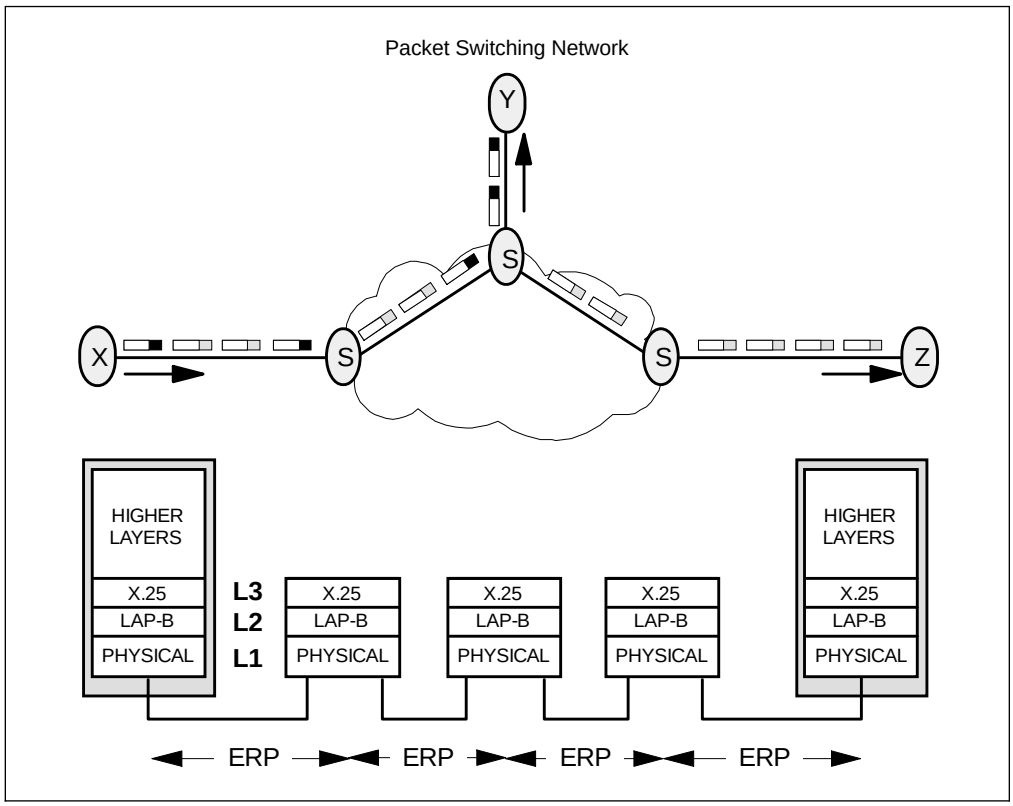

Figure 16. Packet Switching

Figure 16 depicts how the concept of packet switching allows an endstation to multiplex data destined for multiple destinations on a single network attachment. Station X is exchanging data with both Y and Z simultaneously. To allow the network switches *S* to route traffic to the appropriate destination routing identifiers are appended to the data. These routing identifiers either have *global significance*, meaning the same routing identifiers are used at either end of the network or have local significance, meaning that values are assigned independently at either end. If the routing identifiers are not the same then the network packet switches are responsible for changing these identifiers when data packets progress through the network.

As mentioned, X.25 requires a connection, called a *virtual circuit (VC)*, to be established before endstations can communicate. The connections are dynamically established for switched virtual circuits (SVCs) or defined at subscription time for permanent virtual circuits (PVCs). The routing identifiers, or *logical channel numbers (LCNs)*, are assigned during connection establishment. The LCNs for a PVC are fixed numbers. The LCNs for a SVC are dynamically allocated.

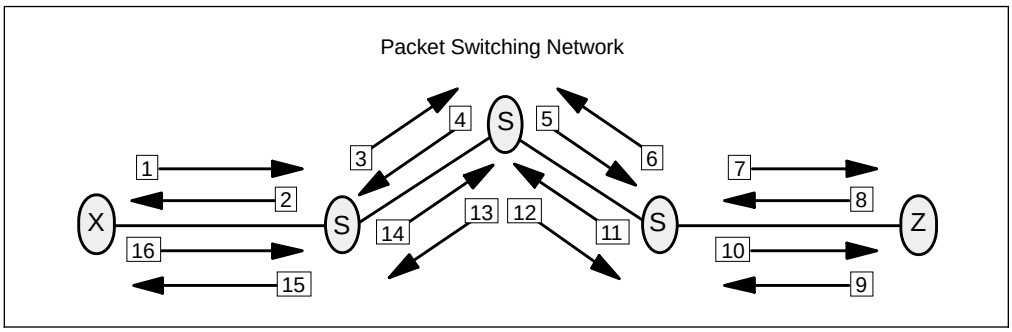

Figure 17. Packet Switching

Traditional packet switches do error checking, acknowledgement and possible retransmission on a hop-by-hop basis rather than end-to-end (see Figure 17); this makes them slower. Therefore, these architectures not only provide efficient utilization of bandwidth, but also provide relatively error-free transmission of data.

Figure 17 shows the order of flows (1 to 15) across a packet switching network to send a single packet and receive an acknowledgement, between nodes X and Y.

2.1.2 Circuit Switching

Circuit switching, unlike packet switching, dedicates a fixed amount of bandwidth to each attached (active) station, regardless of the actual traffic from that station. The circuit switches (for example time division multiplexors) are entirely protocol-independent and do not know or care what the bits they are forwarding represent. Figure 18 on page 46 depicts the concept of circuit switching. Station X is able to transfer data to Y and Z at the same time. When multiplexing techniques are employed X requires only a single physical network attachment which appears as two distinct connections.

Circuit switching does not acknowledge packets and provides no error checking and correcting. These functions are handled by the endstations. Due to the limited processing in intermediate nodes, high throughput rates can be achieved. Circuit switching provides the best choice when traffic is steady and predictable.

The primary weakness of circuit switching is that fixed blocks of bandwidth are reserved for the duration of a connection. The bandwidth allocated to a data session is reserved permanently, even if that data session is idle. This can be very inefficient. In addition, to provide connections to multiple destinations at the same time requires multiple network attachments which can be expensive.

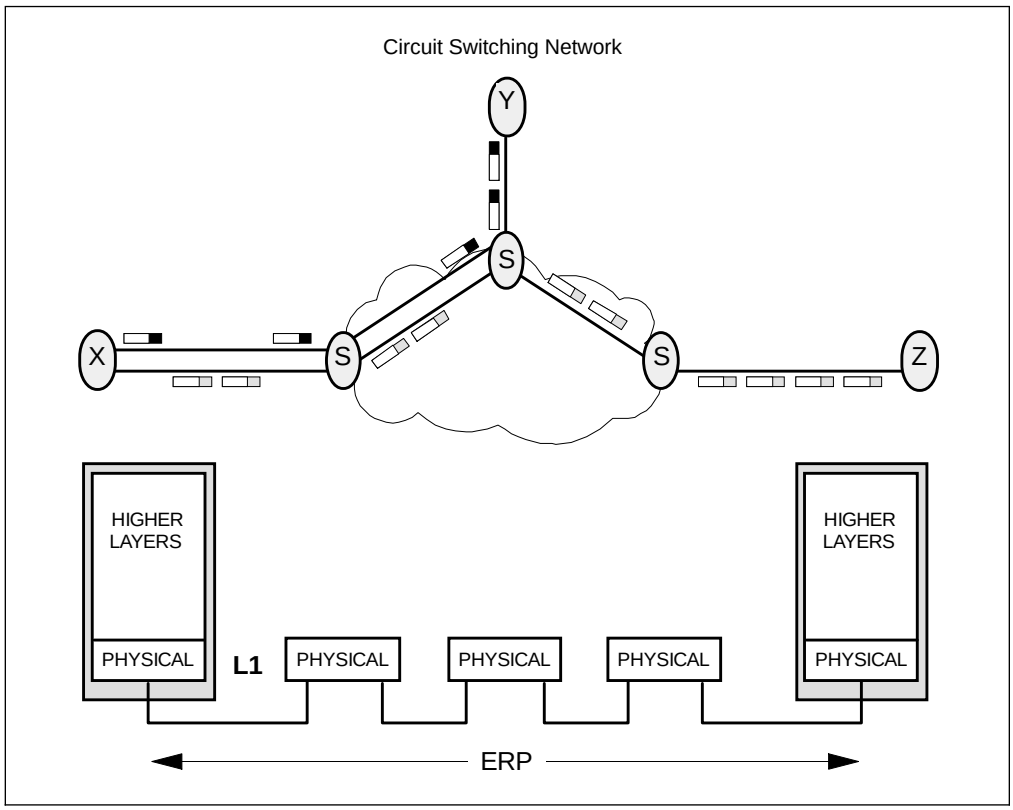

Figure 18. Circuit Switching

2.1.3 Fast Packet Switching

Fast packet switching is a generic term that refers to packet switching technologies that omit most of OSI model[1] layer two processing and all of layer 3-7 processing (in the intermediate nodes) to achieve higher data throughput. The error correction and acknowledgment is performed by the endpoints. Fast packet switching provides the best of both worlds. It uses packet transfer mode to more efficiently allocate available bandwidth to bursty traffic, but it does less per packet processing, which translates into improved throughput. Because it operates below layer 3 of the OSI model it is easy to run multiple protocols over it. Figure 19 on page 47 depicts the concept of fast packet switching using frame relay.

[1] The OSI reference model for networking (or simply OSI model) has been created by the International Standards Organization (ISO). It is assumed that the readers are familiar with the OSI model. If not, readers are referred to one of the many publications on this subject.

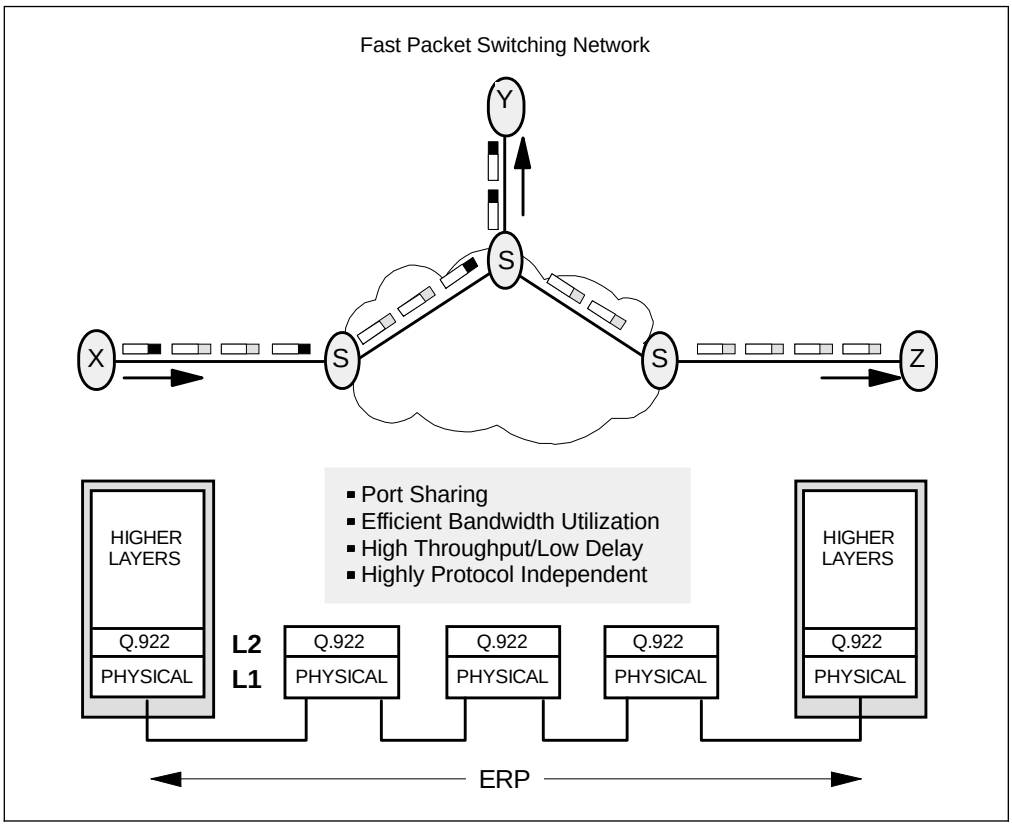

Figure 19. Fast Packet Switching Using Frame Relay

Figure 20 shows the flow of packets (1 to 8) to send a data packet and receive an acknowledgement in a fast packet switching network.

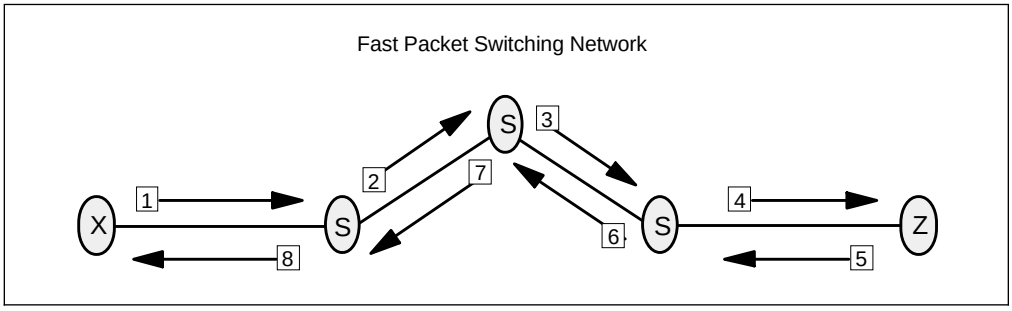

Figure 20. Fast Packet Switching

Fast packet switching is usually used to refer to both frame relay and cell relay. Cell relay, or cell switching, allows the transport of voice, data and full-motion video with a guaranteed quality of service (QoS). Frame relay, although usually limited to data and still images transport, can also transport voice and video, but with less precise, and less predictable QoS. Examples of cell-based transfer mechanisms are Asynchronous Transfer Mode (ATM) and the IEEE 802.6 (SMDS) standard for Metropolitan Area Subnetwork (MAN). For an excellent overview of these, and high-speed networking in general, see *High Speed Networking Technology: An Introductory Survey*, GG24-3816. The technique of frame relay is explained in 2.3, “Frame Relay Technical Description” on page 49.

2.2 Frame Relay, an International Standard

The two organizations which have been the main driving forcing in defining frame relay as an international standard are ITU-T and the American National Standards Institute (ANSI). ITU-T is one of the organizations of the International Telecommunication Union (ITU), a United Nations organization for maintaining and extending international cooperation for the improvement and use of telecommunications. The ITU-T has replaced the former CCITT. ANSI is a non-governmental organization coordinating standards developing activities and representing the USA in international standards bodies.

Table 5. Frame Relay Standards

Description	Date	ITU-T		ANSI	
		#	Status	#	Status
Architectural Framework	1992	I.233	Approved	T1.606	Approved
Frame Mode Bearer Services Interworking	1992	I.555	Approved	-	-
Data Link Layer Signaling [1]	1992	Q.922	Approved	T1.602	Approved
Network Layer Signaling	1992	Q.933	Approved	T1.617	Approved
Congestion Management	1991	I.370	Approved	T1.606 [2]	Approved
Network to Network Interface	1993	I.372	Approved	T1.617	Approved
Framework for Frame-Mode Bearer Services	1993	I.122	Approved	-	-
Core Aspects of Frame Protocol for Use with Frame Relay Bearer Service	1991	-	-	T1.618	Approved

Note:

- [1] The first version (*Blue Books, Melbourne '88*) of the frame relay standard was documented in ITU-T I.122. I.122 still exist (and is often referred to) but merely points to the newer standards.
- [2] T1.606 Addendum

Table 5 gives an overview of the ITU-T and ANSI publications relevant to frame relay. The remainder of this book focuses on the ITU-T standards although it is important to remember that ITU-T and ANSI frame relay standards have been developed in parallel and are closely aligned.

2.2.1 Frame Relay Forum

Several manufacturers of communications equipment started developing their first frame relay implementations before the standards committees had finalized the frame relay standards. At the end of the eighties a number of vendors joined in the *Group of Four*, which later became the *Frame Relay Forum*.

The group of four produced a document in 1990 entitled “Frame Relay Specification with Extensions - Based on T1S1 Standards.” The document reference number is *001-208966*. The document proposed a non-ISDN implementation of frame relay. This document took the proposed ANSI and CCITT standards (I.122) and extended them to suit the needs of the members of the group of four. The document attempts to address issues such as non-ISDN physical level connections, network management and service extensions such as broadcasting.

Incompatibilities (see for example 2.3.5.5, “LMI and the Group of Four” on page 68) exist between the early implementations based on this document and more recent implementations based on the ANSI and ITU-T standards.

The Group of Four invited “Any vendors with frame relay products or plans to implement frame relay.... to join in this effort.” This invitation led to another 17 companies joining them in October 1990 and formed the basis of the *Frame Relay Forum*. The number of members had reached 50 by 1992 and a European Chapter was formed in 1992.

Members of the Frame Relay Forum have agreed on the following implementation agreements:

Table 6. Frame Relay Forum Implementation Agreements

Description	FRF Document	Date
Frame Relay User-to-Network Interface	FRF.1.1	1/19/1996
Frame Relay Network-to-Network Interface	FRF.2.1	7/10/1995
Frame RelayMultiprotocol Encapsulation	FRF.3.1	6/22/1995
Frame Relay Switched Virtual Circuit	FRF.4	1/5/1994
Frame Relay / ATM Network Interworking	FRF.5	12/20/1994
Frame Relay Service Customer Network Management	FRF.6	3/1/1994
Frame Relay PVC Multicast Service and Protocol Description	FRF.7	10/21/1995
Frame Relay / ATM PVC Service Interworking	FRF.8	4/14/1995
Data Compression over Frame Relay	FRF.9	1/22/1996

2.3 Frame Relay Technical Description

The ITU-T frame relay standards originated within the framework of ISDN. However, there is nothing in the standard that precludes it from being implemented in a non-ISDN environment. Few, if any, frame relay implementations are based on ISDN physical level interfaces.

2.3.1 Frame Relay Reference Model

An ISDN *Basic Rate Interface* (BRI) user device has one 16 kbps D channel for call signalling and up to two 64 kbps B channels for data communications over the ISDN interface. In order to take full advantage of this high-speed, low error-rate B channel connection, the CCITT (now the ITU-T) standards committees defined the *Additional Packet Mode Bearer Service* now known as *frame relay*. The title for this recommendation changed in 1993 to *Framework for Frame Mode Bearer Services* (I.122), which is simply a list of other recommendations related to frame mode bearer services.

2.3.2 Frame-Mode Bearer Services I.233

The architectural framework defined in ITU-T I.233 makes a distinction between the two types of *Frame-Mode Bearer Services*. ITU-T I.233.1 describes the frame *relaying* bearer services, and ITU-T I.233.2 describes the frame *switching* bearer services.

Frame Relaying Bearer Services

This is a basic network service for the transfer of data-link frames over a D, B, or H, channel. This service has the following characteristics:

- The UNI allows for the establishment of multiple virtual calls and/or multiple virtual circuits to multiple destinations.
- For virtual calls, the control signalling is performed in a logically separate manner via a D channel signalling protocol.
- User data is transmitted in *Link Access Procedure for Frame-Mode Bearer Services* (LAPF) frames.
- The network preserves the order of frames transmitted when they are delivered.
- The network detects errors and discards the affected frames.

Frame Switching Bearer Services

This has all the characteristics of the relaying bearer service described above, plus these additional characteristics:

- Frames are transmitted with acknowledgements returned to the transmitting user.
- Flow control is supported across the UNI in both directions.
- The network detects and recovers from transmission, format, and operational errors.
- The network detects and recovers from lost or duplicated frames. are delivered.

This means that the frame relaying bearer service is an unreliable multiplexed service where it is possible that frames may be lost or duplicated by the network, and there is no flow-control mechanism across the UNI. It does provide that frames are delivered in the order in which they were transmitted. The frame switching bearer service is a reliable multiplexed service which provide flow- and error-control, more like an X.25 network.

2.3.2.1 User-Network Interface Protocol Architecture

Figure 21 on page 51 depicts the protocol architecture that supports the frame-mode bearer services. As in other areas of ISDN, functions are divided between two modes of operation:

Control Plane (C plane)

This is involved in the establishment and termination of logical connections. The C plane protocols are between the subscriber and the network.

User Plane (U plane)

This is responsible for the transfer of data between subscribers. The U plane protocols are end-to-end.

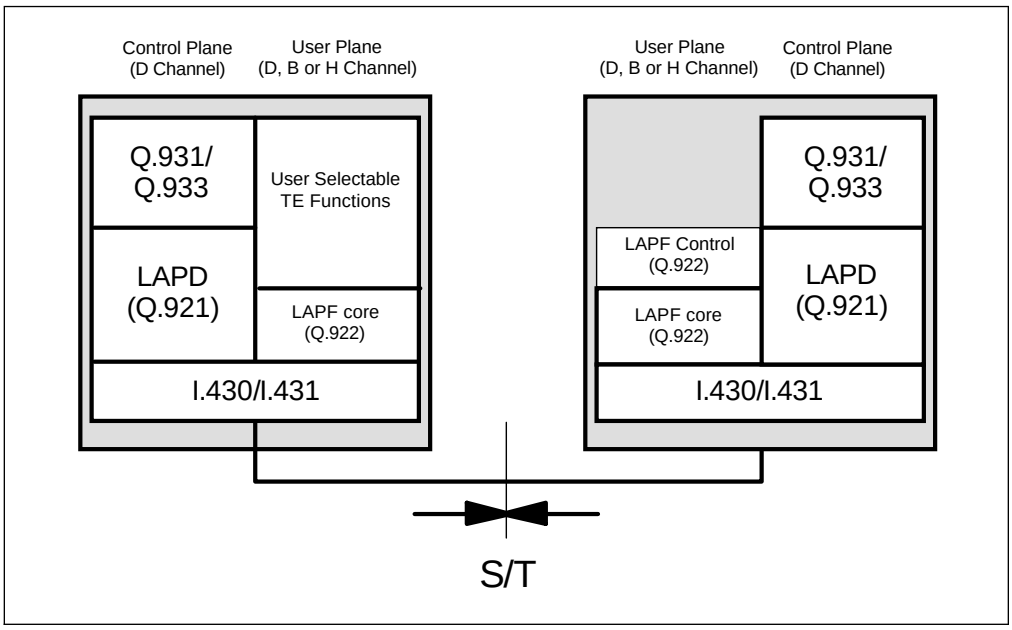

Figure 21. User-Network Interface Protocol Architecture

Note: The ISDN reference points S and T refer to points where protocols are defined. Often, when describing the frame relay architecture, the S/T reference points are shown as the points across which the frame relay protocol is implemented.

The link access procedure for the D channel (LAPD) (Q.921) protocol is used to pass signalling information over the D channel between the user and the ISDN exchange. The *link access procedure for the frame-mode bearer services* (LAPF) (Q.922) is derived from LAPD and defines the frame relay communications protocol used over the B channel. Only the *core* functions of LAPF are used for frame relay. See 2.3.3.2, “The Data Link Control Layer” on page 53 for more information on LAPF.

The core functions of LAPF in the user plane constitute a sublayer of the data link layer. These core functions provide a bare service for transferring data link frames from one subscriber to another, with no flow- or error-control.

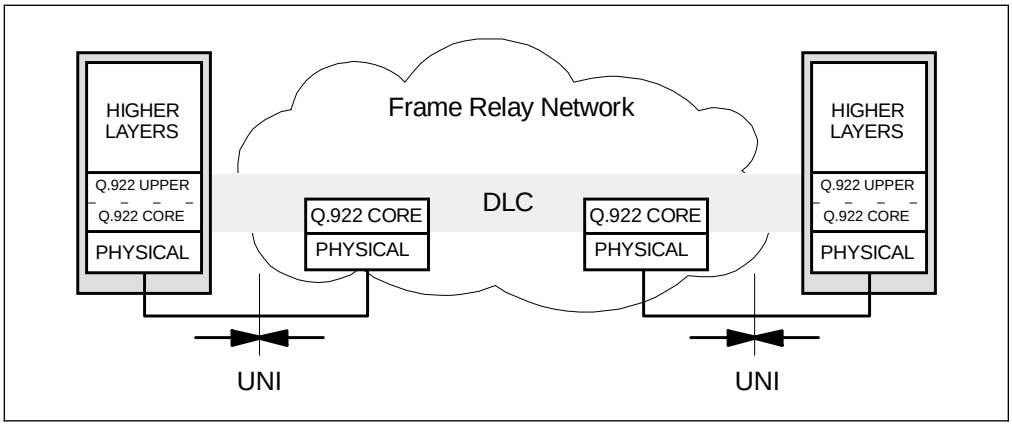

Figure 22. Frame Relay Reference Model

2.3.3 Frame Relay

The ITU-T frame relay standard encompasses both the frame relaying service description and the user-network interface protocol architecture. To understand these concepts, the reference model shown in Figure 22 on page 51 is used. Depicted is user equipment attached via a user-network interface to a network providing the frame relaying function.

Note: In IBM Frame Relay documentation the terms *Frame Relay Frame Handler (FRFH)* and *Frame Relay Terminating Equipment (FRTE)* are often used to distinguish between the networking and the user function respectively.

The frame relay standards describe only the lowest two layers, data link control and physical layer, of the OSI model. Although not depicted in Figure 22 on page 51 it is assumed that to allow communication between network-attached equipment, they have also implemented the procedures and functions described in the higher layers.

The network frame relaying service is described in general terms only. The service should provide an unacknowledged transfer of frames from user-network to user-network interface and:

1. It preserves their order as given at one user-network interface point if and when they are delivered at the other end.
2. Protocol capabilities are available so that the network may discard erroneous frame if it elects to do so.
3. It detects transmission, format and operational errors (for example, frames of unknown DLCI).
4. Frames are transported transparently; only the address and FCS field may be modified.
5. Frames are not acknowledged (within the network).
6. Throughput may be monitored and enforced.
7. Congestion control procedures may be put in place.

In the case of permanent virtual circuits (PVC) no real-time call establishment is necessary and parameters are agreed upon at subscription time.

The recommendation limits itself to a service description with an exception of item 5 which specifies some of the internal operation. It is felt that a network offering a fast packet relaying service should not acknowledge the frames internally. The only reason for acknowledging is to enable retransmissions when data is lost, as the service description does not imply a guaranteed end-to-end delivery of frames and endstations should take provisions to protect against lost frames anyway.

2.3.3.1 The Physical Layer

As mentioned in the previous sections, the ITU-T and ANSI standards have been defined within the context of ISDN. There are no international standards for providing frame relay over conventional physical circuits (for example, V.35, X.21 or G.702). However, most (if not all) implementations are based on these.

2.3.3.2 The Data Link Control Layer

Frame relay layer two functions are based on the use of LAPF data link control (DLC) protocol. LAPF is divided into *DL-CORE* and *DL-CONTROL* functions, also known as Q.922 UPPER and Q.922 CORE. The use of DL-CORE (see Figure 22 on page 51), the frame relay data link control core functions, is mandatory; the use of DL-CONTROL is optional.

Layer two core functions to be implemented at both ends of the user-network interface are:

- Frame delimiting, alignment and transparency
- Frame multiplexing and demultiplexing using the address field
- Inspection of the frame to ensure that it consists of an integer number of octets (prior to zero bit insertion or following zero bit extraction)
- Inspection of the frame to ensure that it is neither too long nor to short
- Detection of transmission errors

ITU-T Q.922 describes the format, depicted in Figure 23, of the layer two frames exchanged on the user-network interface.

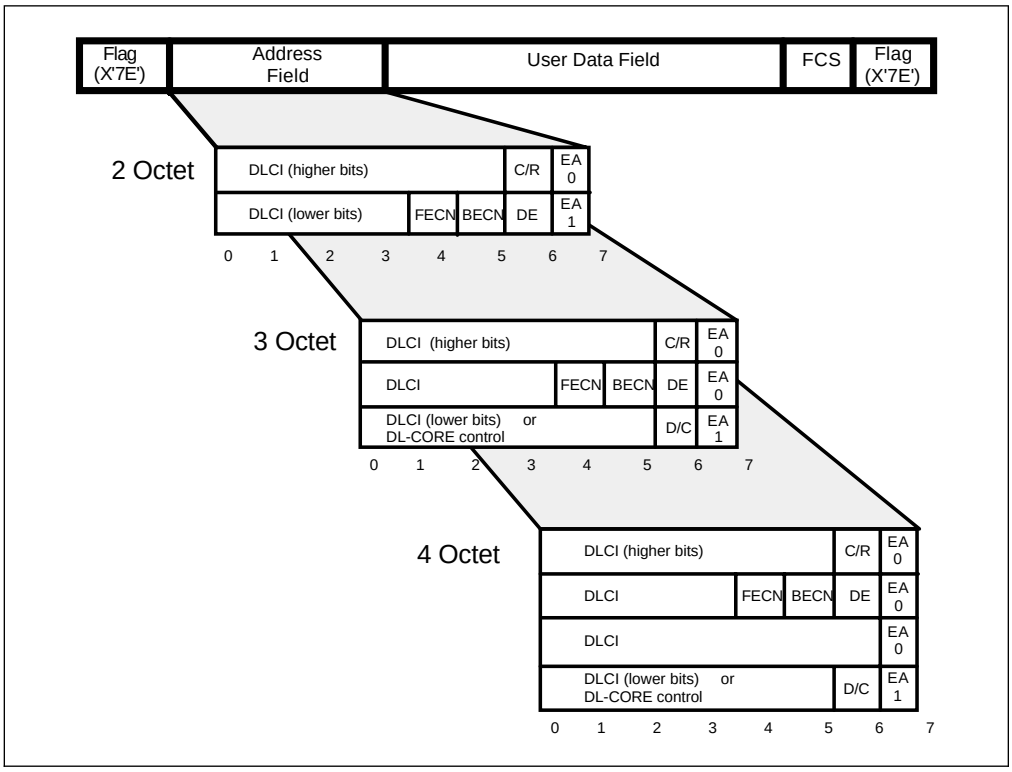

Figure 23. Frame Format

Frames exchanged consist of a pair of flags delimiting the frame, an address field, the user data field and the frame check sequence (FCS) field. The fields are explained in detail in the following section:

Flag

The flag delimiting the frame is a single octet with value X'7E' (B'01111110'). All frames must start and end with one or multiple flags. The closing flag may be used as an opening flag for a consecutive frame.

The above bit sequence allows the receiver to synchronize on start and end of a frame. To avoid the occurrence of this bit stream within the user data, the originator inserts a 0 bit after each sequence of five consecutive 1 bits. The receiver removes each 0 bit after five consecutive 1 bits. This technique is called *bit (de)stuffing*.

Address Field

The address field is either two, three or four octets long. The default frame relay address field used by most, if not all, implementations, is a two-octet field.

The address field contains the following information:

DLCI

Data Link Connection Identifier: A multiple bit field that uniquely identifies a virtual circuit.

Frame relay is a packet switching technique (see also 2.1, "Packet Switching Techniques" on page 43). The concept of virtual circuits allows an endstation to maintain multiple connections, to multiple endstations on a single user-network interface at the same time. Each virtual circuit is identified by a different DLCI.

The DLCIs used along the path of a virtual circuit may be different for each hop. DLCIs only have local significance.

Traffic destined for or originating from each of the partner endstations is multiplexed, carrying different DLCIs, on the same user-network interface. The DLCI is used by the network to associate a frame with a specific virtual circuit.

The DLCI is a concatenation of the DLCI bits contained within the address field octets. The length of the DLCI (see Table 7) depends on the length of the address field and the setting of the D/C field (see remainder of this section).

Table 7. Length of Data Link Connection Identifier (DLCI)

Address Field Size	D/C = 1	D/C = 0
2 Octets	10[1]	10[1]
3 Octets	10	16
4 Octets	17	23

Note:

1. The DLCI has a fixed length within the two-octet address field. The use of the D/C field is not applicable.

Table 8 on page 55 depicts the DLCI ranges and the allocation to different networking functions.

Table 8. DLCI Values

10-Bit DLCI		16-Bit DLCI		17-Bit DLCI		23-Bit DLCI		Function
Begin	End	Begin	End	Begin	End	Begin	End	
0		0		0		0		LMI Channel[1]
1	15	1	1023	1	2047	1	131,071	Reserved
16	991	1024	63,487	2048	126,975	131,072	4,194,303	Available to the user
992	1007	63,488	64,511	126,976	129,023			Layer two management[2]
1008	1022	64,512	65,534	129,024	131,070			Reserved
1023		65,535		131,071				In channel layer two management[3]

Notes:

1. For details see 2.3.5, “Local Management Interface (LMI)” on page 60.
2. Used for information related to the network, such as the consolidated link layer management message (see 2.3.4.2, “Consolidated Link Layer Management” on page 59).
3. User to pass management interface messages that relate to the higher layers across the connection.

DLCI/DL-CORE

DLCI or DL-CORE field: A six-bit field contained within the last octet of a three- or four-octet address field. See the explanation of the D/C field.

D/C

DLCI/DL-CORE Indication: A one-bit field contained within the last octet of a three- and four-octet address field indicating whether the first six bits of the last addressing octet should be interpreted as the least significant bits of the DLCI (D/C = '0'), or is providing (D/C = '1') additional, still to be defined, control functions.

CR

Command/Response Indication: One-bit field indicating whether the frame is a command or response frame. Its use is not defined within the frame relay protocol and values are passed transparently across the network.

EA

Extended Address Bits: One-bit field that will be set to 1 in the last octet and 0 in the preceding octet(s) of the addressing field.

FECN

Forward Explicit Congestion Notification Bit (see 2.3.4, “Congestion Control” on page 56): One-bit field that notifies the user that the network is experiencing congestion in the direction the frame was sent. However, it is assumed no obligation exists, that users will take action to relieve the congestion.

BECN

Backward Explicit Congestion Notification Bit (see 2.3.4, “Congestion Control” on page 56): One-bit field that notifies the user that the network is experiencing congestion in the reverse direction the frame was sent. However, it is assumed that if no obligation exists, users will take action to relieve the congestion.

DE

Discard Eligibility Bit (see 2.3.4, “Congestion Control” on page 56): One-bit field indicating whether (DE = '1') or not this frame should be discarded by the network in preference to other frames (for example, during congestion).

User Data Field

The user data field contains the actual data to be transported. It should consist of an integral number, before/after bit (de)stuffing, of octets. The maximum size for this field is network-dependent; the ITU-T defines a minimum maximum of 260 bytes. A minimum maximum of 1600 bytes is strongly recommended. The contents of this field are passed unchanged across the network and not interpreted by the frame relay protocol.

Frame Check Sequence

The frame check sequence (FCS) is used to check that the data has been received without error. It consists of two octets containing a cyclic redundancy check using the ITU-T error checking polynomial ($x^{16}+x^{12}+x^5+1$). The FCS operates on all bits in the frame excluding the flags and the FCS itself. Frames with an incorrect FCS will be dropped. The FCS has to be recalculated when bits in the address field change. An example is the DLCI which has local significance only.

In addition to the layer two core functions, additional functions such as frame acknowledgment, flow control and recovery from detected transmission, format and operational errors need to be implemented. These functions are a user responsibility. ITU-T I.922 DL-CONTROL or other, either standard or proprietary protocols may be used. Many implementations divert, usually only slightly, from ITU-T I.922 DL-CONTROL because the standard was not finished at the time of implementation.

This delineation of the data link control function is very similar to the delineation of the IEEE LAN standards between the Logical Link Control (LLC) and the Media Access Control (MAC) layers. For a description of how the DLC layer user functions have been implemented on several IBM networking products refer to Appendix A, “IEEE Logical Link Control 802.2” on page 239.

2.3.4 Congestion Control

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.

As described in 2.2, “Frame Relay, an International Standard” on page 48, the service offered by a frame relaying service is the transparent and unacknowledged transfer of frames. The user data received will be as 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 endstations and LMI messages sent by the network, endstations may also receive Consolidated Link Layer Management (CLLM) messages generated by the network reporting congestion. The use of CLLM is not

widespread; for details see 2.3.4.2, “Consolidated Link Layer Management” on page 59.

Frame relay networks using *in-band* congestion signaling report congestion by using bits in the frame address field. The endstation will receive no other frames, (the exception being the LMI message as described in 2.3.5, “Local Management Interface (LMI)” on page 60) than frames sent by another endstation.

The network is able to inform endstations 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.

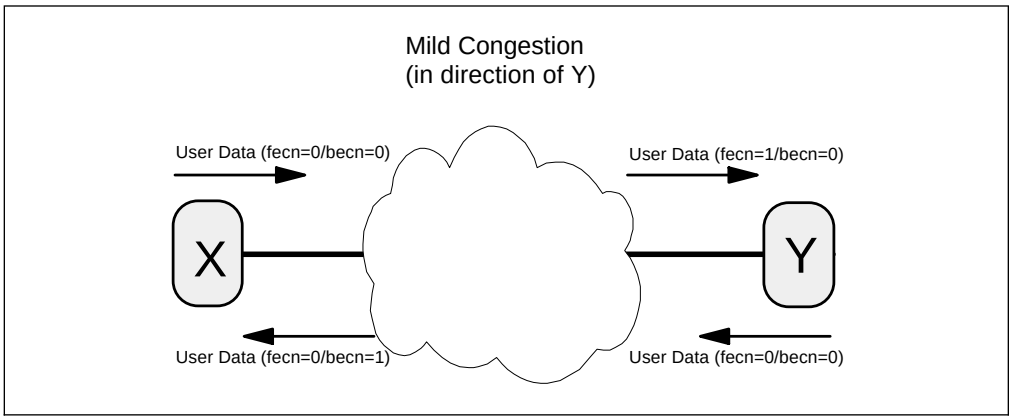

Figure 24. In-Band Congestion Signaling

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 endstations 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 relaying network will usually start discarding frames during severe congestion. Endstations 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 endstations 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 endstations, has therefore been considered inadequate to networks providing a frame relaying service and additional provisions have been defined.

2.3.4.1 Committed Information Rate and Burst Sizes

The maximum number of bits per seconds that an endstation 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)*.

Endstations 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. However, the network 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 a increment of time T_c. The CIR is also referred to as *minimum acceptable throughput*.

B_c and B_e are expressed in bits, T_c in seconds, the access rate and CIR in bits per second. B_c, B_e, T_c, CIR are defined per DLCI, the access rate is valid for per 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 9.

Table 9. 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	$T_c = (B_e/\text{Access Rate})^2$

Note:

1. Table depicts the valid parameter configurations. Other configurations are for further study.
2. When the two communicating endstations 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 circuits 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. Only a number of considerations are mentioned and careful study is required.

Networks may mark frames above B_c with discard eligible (DE) but have plenty of spare capacity to transport the frame, or the reverse, have limited capacity and discard excessive frames immediately. Networks may mark frames above B_c+B_e with discard eligible (DE), and possibly transport them, or just drop the frames as suggested by ITU-T I.370.

Network managers always try to balance costs and performance and have to examine the frame relay service provider charging schemes. Networks may implement fixed charging dependent 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 flow control mechanisms implemented on top of the layer two 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:

$$B_c+B_e = \text{Access rate}*T_c$$

Depending on functions implemented on top of the layer two core functions, lost frames may be quickly detected and recovered from. This may be a time-consuming activity severely impacting performance. In the latter case subscribing to high values of CIR, B_c and B_e is important.

2.3.4.2 Consolidated Link Layer Management

In general it can be said that the in-band congestion signaling designed within the frame relay standards limits the networks ability to react properly on congestion and prevents the implementation of more advanced congestion control schemes known for example, within SNA. In-band frame relay congestion control relies upon setting of the FECN bit in the data being sent to the destination, which then subsequently should request the originator to slow down, thus in fact relying upon end-to-end flow control procedures, which may very well be absent.

This limitation (described within recommendation Q.922) has led the ITU-T to develop an optional signaling mechanism called Consolidated Link Layer Management (CLLM). CLLM is used on a separate DLCI, therefore out of band, allowing the network to pass control messages to the user. When a network becomes congested the network may use the FECN/BECN bits within the user data frames on the user's data DLCI, use the CLLM messages or use both.

The CLLM message contains a list of DLCIs that are likely to cause the congestion. The endstation is expected to relieve the congestion by limiting the data transfer on the DLCIs identified. For two-octet address field frames the CLLM messages are sent on DLCI 1007. CLLMs are sent within LAPF exchange identification (XID)

frames. The C/R field within the frame header (see Figure 23 on page 53) is set to 1 indicating a response. (The receiver is not expected to reply.) For three- and four-octet address fields the use of CLLM is still under study.

Due to the limited (if any) number of networks that have implemented CLLM, format and additional functions of the CLLM are not detailed further.

2.3.5 Local Management Interface (LMI)

The Local Management Interface (LMI) as described in this section is a set of procedures and messages specified in ANSI T1.617 and ITU-T Q.933, defined to operate between a user device and a frame relay network, which provide status and outage notification for frame relay permanent virtual connections (PVCs).

LMI procedures identified are:

Unidirectional LMI support

Unidirectional LMI support supports status queries from the *user* side to the *network* side. The network side responds with the requested information.

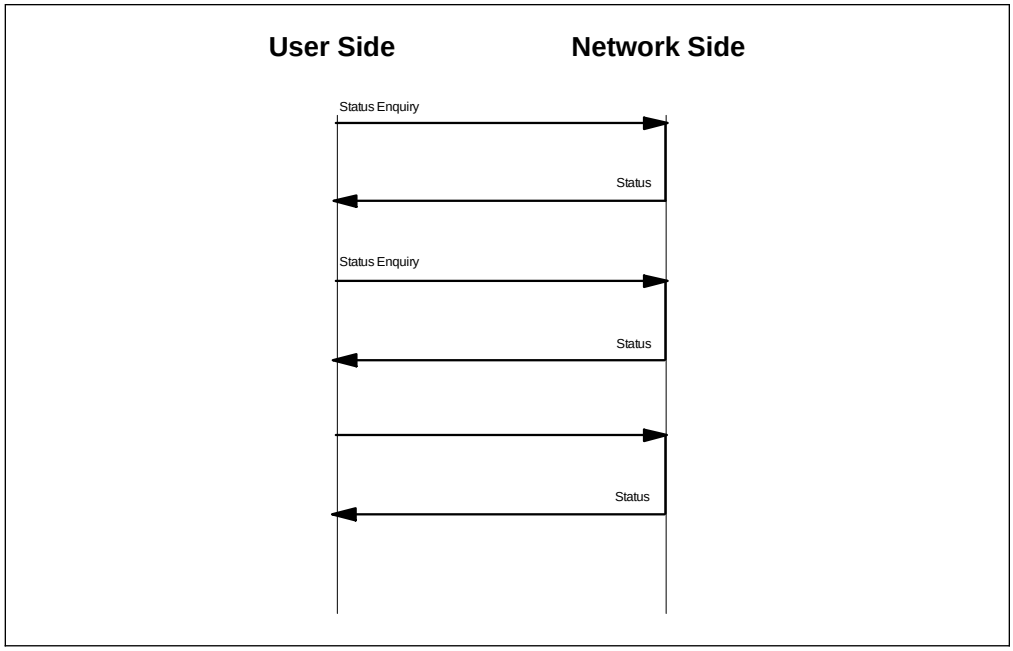

Figure 25. Unidirectional LMI Support

Note: Frame relay can be used to connect two stations using a direct point-to-point connection without going through a frame relay network. This can be a local connection (connect two systems back-to-back) or a high speed point-to-point communications link (such as T1 or E1).

When using the unidirectional LMI support, due to the non-symmetrical nature of the procedures, one of stations has to function as the user side, and the adjacent station as the network side.

Bidirectional LMI support

The (optional) bidirectional support allows both sides of a physical link connection to support user-side and network-side LMI procedures concurrently. Separate sequence numbers, error counters and error

thresholds are used and network side LMI functions operate on both sides.

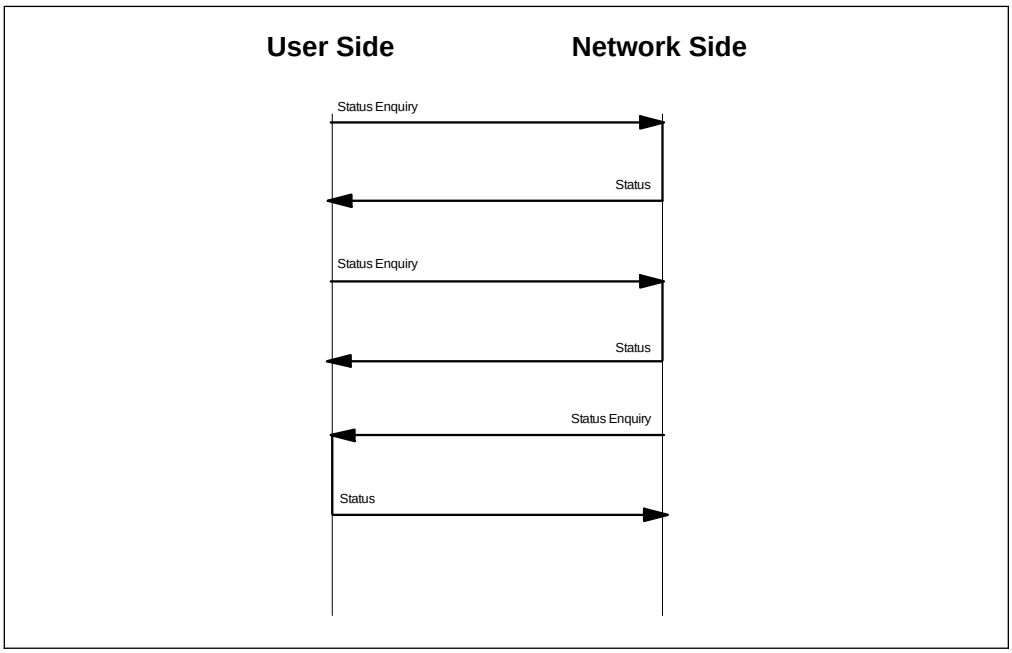

Figure 26. Bidirectional LMI Support

When both ends of a connection support both user and network-side LMI procedures a balanced connection results. Bidirectional LMI support becomes important when connecting multiple networks; see also 2.3.6, “Network-to-Network Interface (NNI)” on page 69.

Asynchronous LMI support

Using either unidirectional or bidirectional update procedures could cause significant delay in informing that changes in PVC status have occurred. For example, using the default poll timer (T391=10) and poll timer counter (N391=6) it may take up to 60 seconds before a user is informed about the inactive status of a PVC. During this period considerable amounts of data, not able to reach its destination, can be generated.

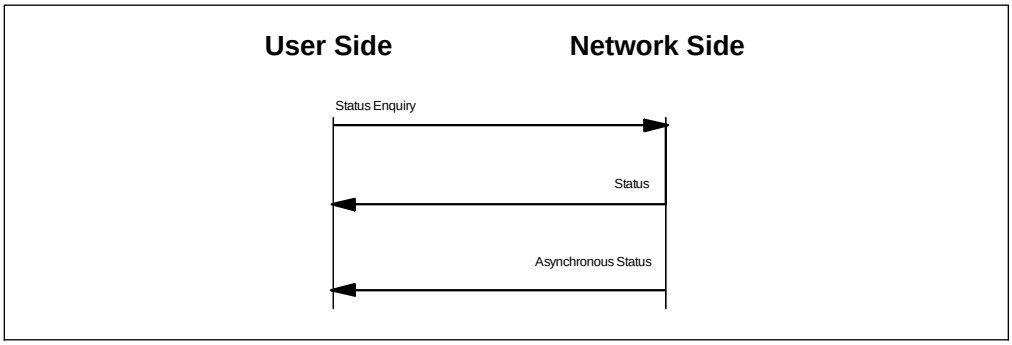

Figure 27. Asynchronous LMI Support

For this reason asynchronous support has been created. The asynchronous procedure allows the network side to send unsolicited PVC status using the Status message with a special report type of

asynchronous status. The Status (asynchronous) message includes PVC status for a single PVC.

The Status (asynchronous) is not part of the periodic polling process (the link integrity verification information entity (IE) is not included) and it does not support the *new* status. (The new status bit has no significance.) For details see “Notification of PVC Status” on page 64.

The unidirectional support is mandatory, but the bidirectional and asynchronous support are optional LMI procedures.

2.3.5.1 LMI Timers

The aforementioned LMI procedures use the following timers (Table 10) and counters (Table 11).

Table 10. LMI System Parameter - Timers

Timer	Description	Range (sec.)	Default (sec.)	Started	Stopped	Actions Taken When Expired
T391	Link Integrity Verification polling timer	5-30	10	Transmit status enquiry		Transmit status enquiry. Record error if status message not received.
T392	Polling verification timer	5-30	15	Transmit Status	Receive Status Enquiry	Record error by incrementing N392. Restart.

Note:

1. T392 > T391

Table 11. LMI System Parameter - Counters

Counter	Description	Range	Default	Usage	User or Network Side
N391	Status (full) polling counter	1-255	6	Polling cycles	User side
N392	Error Threshold	1-10	3	Errors	Both
N393	Monitored events count	1-10	4	Events	Both

Note:

1. N392 ≤ N393

As the LMI describes an interface that is not balanced[2] a user and a network side are distinguished; see also 2.3.5, “Local Management Interface (LMI)” on page 60. DLCI 0 is used for transferring LMI messages between adjacent frame relay devices. For details about the LMI message formats refer to 2.3.5.3, “LMI Message Formats” on page 65.

[2] A balanced (or symmetrical) interface assumes the same kind of role at either end of a connection.

2.3.5.2 LMI Message Types

There are two messages supported:

- Status Enquiry (X'75')
- Status (X'7D')

These messages are used to perform three procedures:

- Link integrity verification (LIV)
- Notification of PVC status
- Notification of added or deleted PVCs

The three procedures are discussed in more detail later.

Link Integrity Verification (LIV): The link integrity verification (LIV) procedures require the user side of a physical link to exchange sequence numbers periodically on a defined polling interval (T391, see Table 10 on page 62), so each side can determine whether the physical connection to the adjacent node is still functional.

The LIV procedure is depicted in Figure 28.

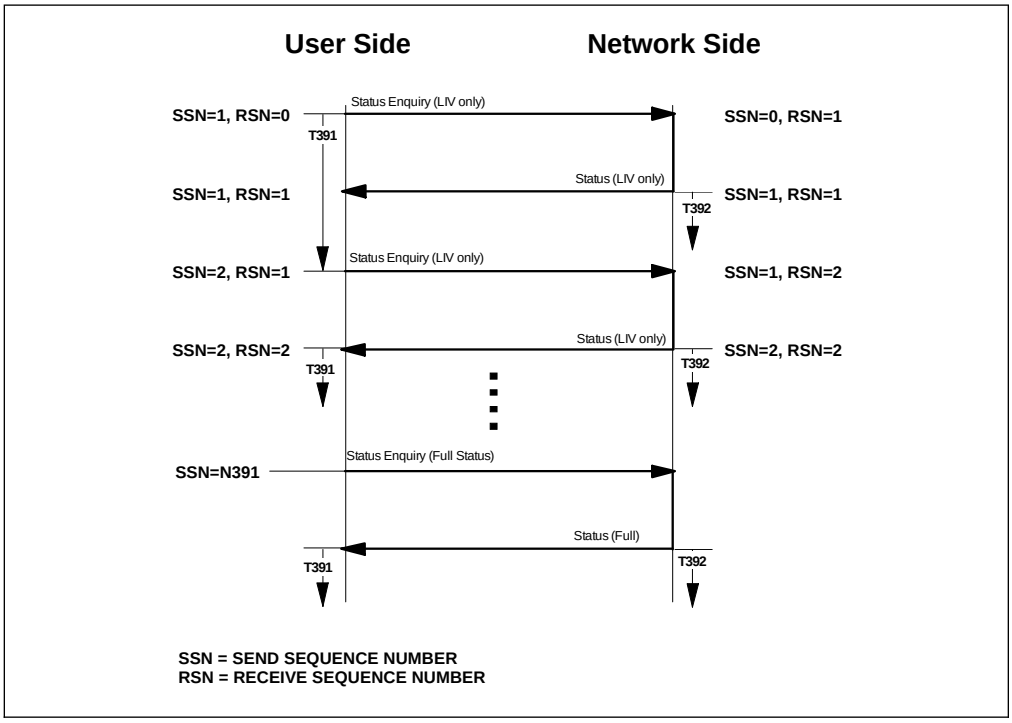

Figure 28. Link Integrity Verification Procedure. Status Enquiry(full) is used for link integrity verification as well.

Periodically the user side transmits a Status Enquiry message. The network side responds with a Status message containing the requested link integrity verification IE. Both sides maintain a send sequence number and a receive sequence number. The link integrity is checked by a sequence number generation and checking process.

The user side sends a Status message containing a link integrity verification IE. The report type may either be full status or link integrity only. The link integrity IE contains the send sequence number of the sender (user side), which is

incremented each time a Status Enquiry message is sent, and a receive sequence number that equals the send sequence number contained within the link integrity information element of the last Status message received from the the network side.

The network side must respond by sending a Status message containing a link integrity verification with the send sequence number of the sender (network side), which is incremented each time a Status message is sent, and a receive sequence number that equals the send sequence number contained within the link integrity information element of the last Status Enquiry message received from the the user side. The network side maintains a timer (T392, see Table 10 on page 62) that is restarted each time a Status message is sent.

Link integrity errors are recorded:

- By the user side - if not receiving a valid Status message containing a receive sequence number matching the last send sequence number sent in a Status Enquiry message within T391 seconds
- By the network side - if not receiving a valid Status Enquiry message containing a receive sequence number matching the last send sequence number sent in a Status message within T392 seconds

As an optional LMI feature, stations could declare all PVCs over a given physical link inactive when detecting that N392 (Table 11 on page 62) errors have occurred out of the last N393 intervals (T391 for the user side, T392 for the network side).

Notes:

1. The above procedure impacts all PVCs across a frame relay interface.
2. There are no details in the frame relay specification about how an interface determines that the error situation is cleared; therefore, recovery will be implementation-dependent.

Notification of PVC Status: This polling procedure, depicted in Figure 28 on page 63 requires the user side to request PVC status at least once per N391 (see Table 11 on page 62) polling intervals. The user side can request PVC status at any and all polling intervals.

The PVC status tells the user side whether existing PVCs are active or inactive and have been added or have been deleted. Procedures for added or deleted PVCs are discussed in “Notification of Added or Deleted PVCs” on page 65. The same Status Enquiry and Status messages used for link integrity verification are also used by the PVC status procedure. A report type of *full* status in the Status Enquiry indicates a user-side status request. The network side must respond with a Status message with report type of full status along with PVC status IEs for each existing PVC.

A PVC status of active indicates to the user side that the PVC is available for use. An inactive status indicates that the PVC is unavailable and that the user side should not use it. Any data sent by the user over an inactive PVC most likely will not reach its destination. A status of active does not guarantee that data will reach its destination since the network can discard frames when needed due to congestion or errors.

Notification of Added or Deleted PVCs: This procedure is part of the Notification of PVC Status procedure.

The network side sets the New bit on in a PVC status IE in addition to the active/inactive status indicator if the PVC has been added since the last Status message containing full PVC status was sent. The network side resets the New bit if the next Status Enquiry received indicates that the user side received the Status messages (that is, last received sequence number in Status Enquiry matches current sequence number sent in the Status message). If the Status message has not been acknowledged, the Status bit will be repeated the next time PVC status is requested by the user side.

The network side indicates that a PVC was deleted by not including a PVC status IE for the deleted PVC in any more Status messages. If the user side receives a PVC status IE with the New bit on for a previously existing PVC, then this indicates that the network side deleted and added the same PVC since the last status was sent to the user side. By not allowing the New bit to be reset until the user side acknowledges receiving the Status message, the network side ensures the user side is able to detect deleted PVCs.

An error situation occurs when PVC status information is received for PVCs that do not exist (New has never been set), or the receipt of a full status message which omits a PVC currently being used. In these cases the user side should mark the PVC as *active* and *unavailable*, respectively.

2.3.5.3 LMI Message Formats

The basic format of the LMI message is depicted in Figure 29 on page 66.

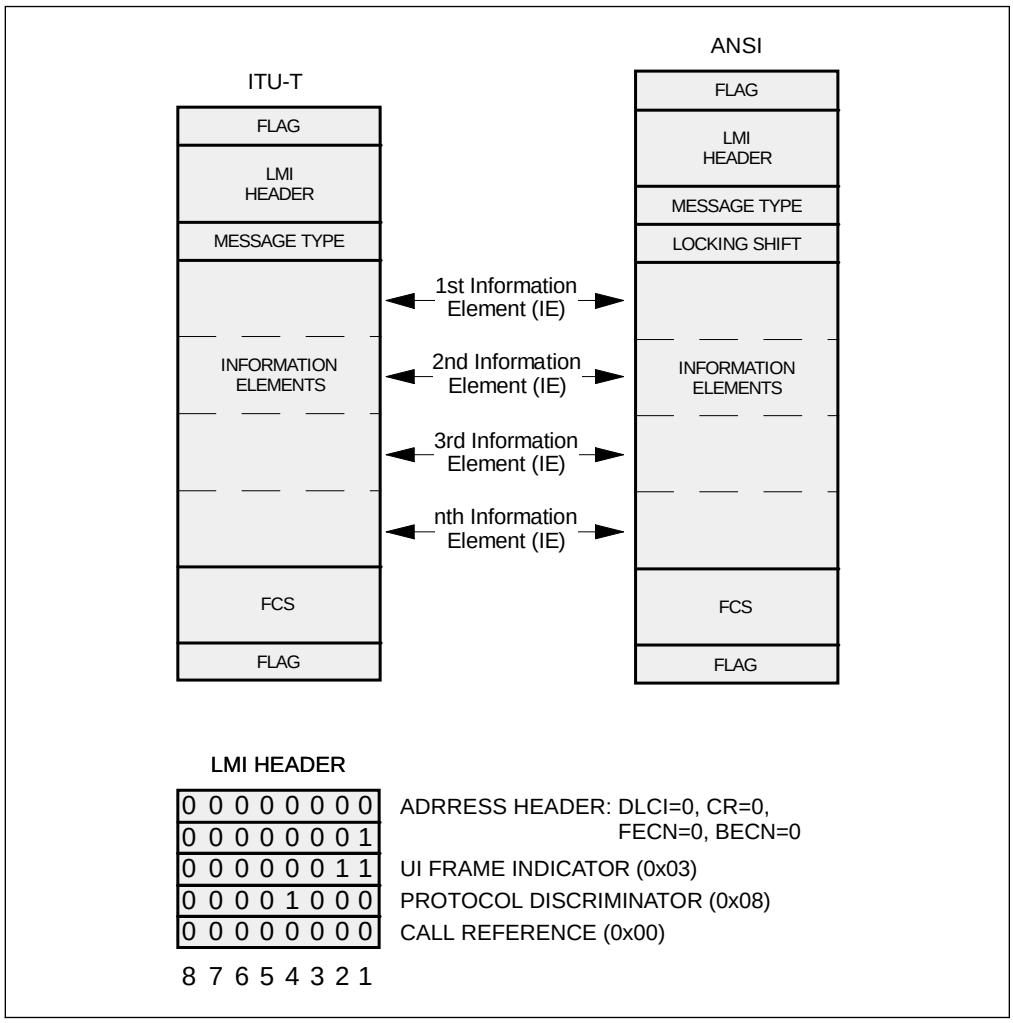

Figure 29. ANSI and ITU-T LMI Message Format

The management services field identifies the type of management message. There are two messages supported:

- Status Enquiry (X'75')
- Status (X'7D')

Note: Currently ITU-T and ANSI have defined three types of management messages:

- Call establishment
- Call clearing
- Miscellaneous messages

Status Enquiry and Status are both miscellaneous messages. Call establishment and call clearing messages refer to the use of switched virtual circuits (SVCs).

Information Element Formats: An LMI message contains one to *n* variable length information elements (IE). The three types of IEs identified are shown in Figure 30 on page 67.

The values and usage of the information element identifiers defined by ITU-T and ANSI are shown in Table 12 on page 68.

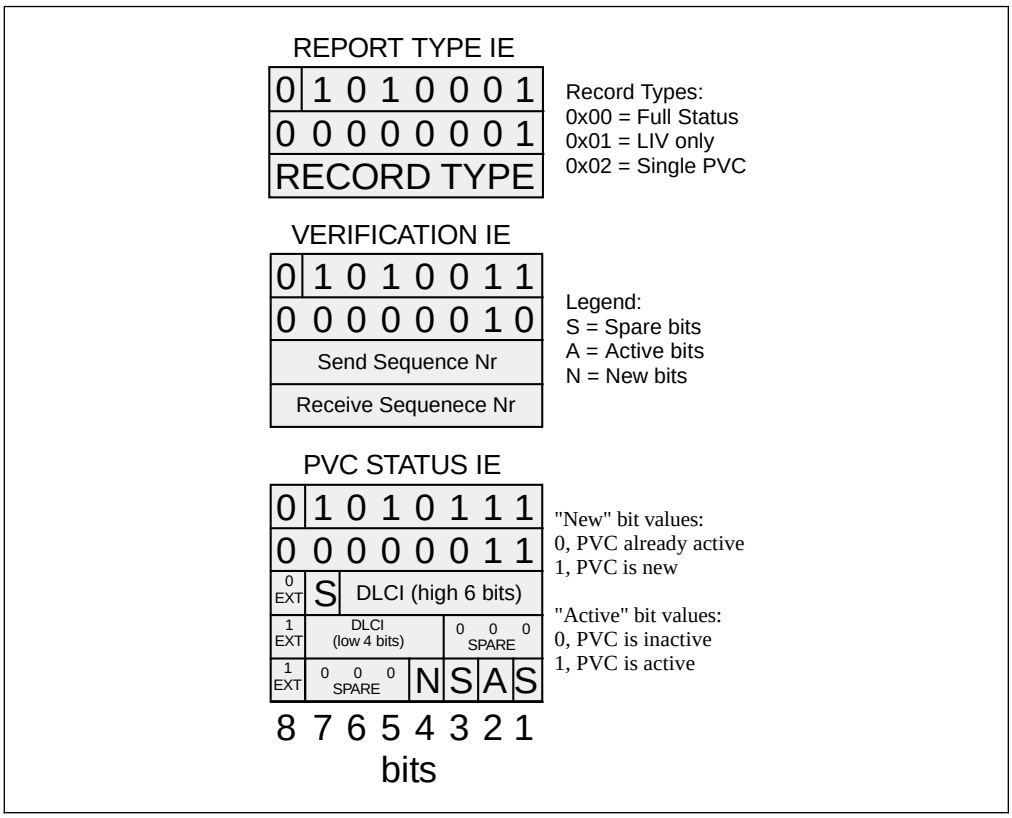

Figure 30. ITU-T Information Element Format

Status Enquiry and Status Message Formats: Status Enquiry messages are periodically sent for link integrity verification or to request notification of (full) PVC status. Figure 31 depicts the format of the Status Enquiry message, and the status message formats returned.

A Status message is either sent in reply to a Status Enquiry message or sent unsolicited informing the adjacent station about the status of PVC changes on the frame relay connection. Figure 31 shows the different format of the Status messages.

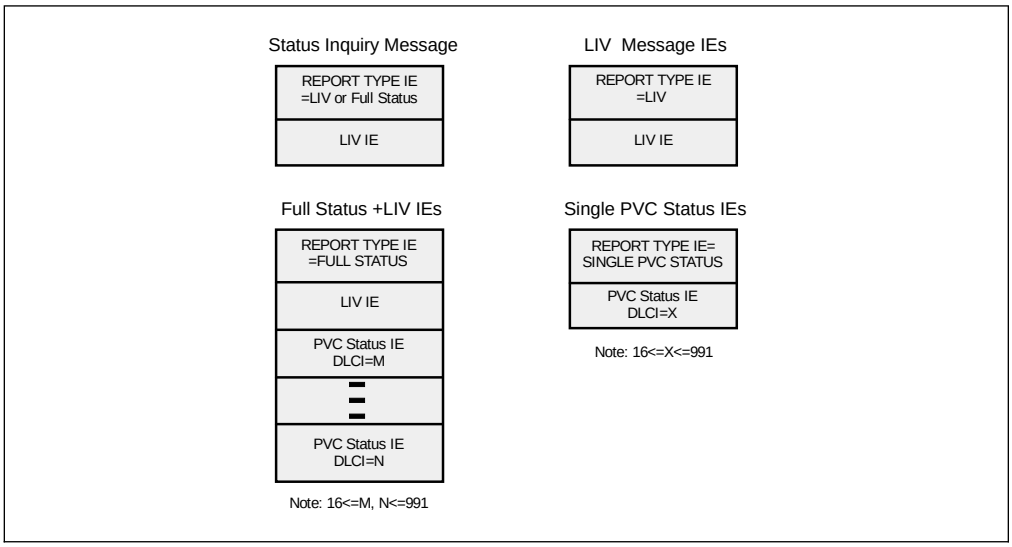

Figure 31. LMI Status Enquiry and Status Report IE Formats

The Record Type field in the Report Type IE specifies one of the record types listed below:

- X'00' - Full Status Request
- X'01' - Link Integrity Verification Only
- X'02' - Single PVC Asynchronous Status

Note: The status enquiry message format also includes an LIV IE. As described in “Link Integrity Verification (LIV)” on page 63 the link integrity is checked by the sequence number generation and checking process. Every time the user side issues a status enquiry message, the send sequence number (SSN) is incremented and placed in the SSN field in the LIV IE that is included in the status enquiry message. This allows the network side to place this SSN in the receive sequence number (RSN) field of the LIV IE it sends back to the user side.

2.3.5.4 ANSI and ITU-T LMI Standards

As mentioned ANSI T1.617 (US standard) and ITU-T Q.933 (international standard) cover LMI for frame relay PVC support. The operating procedures are the same for both but there are a few frame format differences:

- ITU-T used code set 0 and ANSI uses code set 5. There is no locking shift byte after the message type in LMI messages for the ITU-T version; therefore, all LMI messages in the ITU-T format will be one-byte shorter than in the ANSI format.
- The report type, the link integrity verification and the PVC status IE identifier are different between the two standards; see Table 12. The format of the data within the IEs is the same.

Table 12. IE Value Differences

Information Element (IE)	ITU-T Value	ANSI Value
Report Type	X'51'	X'01'
Link Integrity Verification	X'53'	X'03'
PVC Status	X'57'	X'07'

It is implementation dependent if an ITU-T implementation on one end of a link will interoperate with an ANSI implementation at the other end. It is advised to always use either ITU-T or ANSI at both ends of a link.

2.3.5.5 LMI and the Group of Four

LMI specifications defined by the Group of Four (see page 48) within *Frame Relay Specifications with Extensions* differ from and are incompatible with both ITU-T and ANSI standards. This is the basis for the LMI standard commonly known as REV.1.

The channel used for LMI is 1023 under the Group of Four while 0 under ITU-T/ANSI. This will prevent LMI from functioning with a Group of Four implementation on one side and an ITU-T or ANSI implementation at the other end of a frame relay link. In addition, the Group of Four have defined the unidirectional LMI procedures only, referred to as the *heartbeat process*.

While running LMI on DLCI 1023 at one end of a circuit, and DLCI 0 at the other end will not stop the frame relay circuit from operating, this will prevent the LMI from functioning.

2.3.6 Network-to-Network Interface (NNI)

The frame relay standards describe both the interface between an end user station, a frame relay network, called User-to-Network Interface (UNI), and the interface between adjacent frame relay networks, called Network-to-Network Interface (NNI). Figure 32 depicts the concept of UNI and NNI.

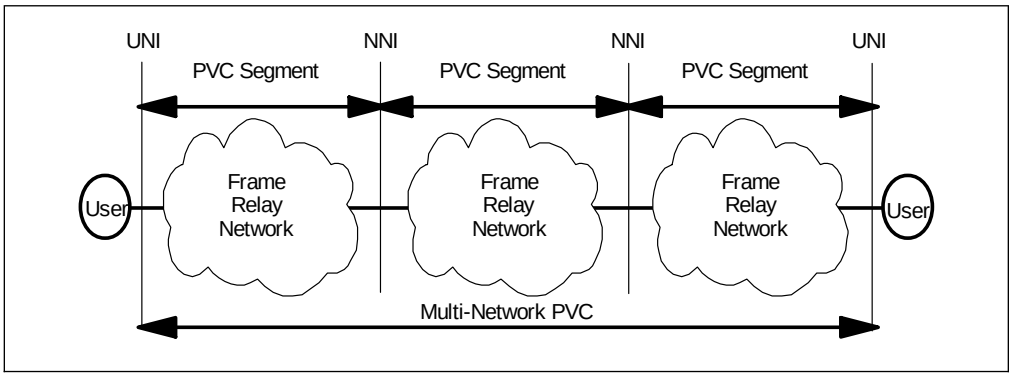

Figure 32. User and Network-to-Network Interface

An end-to-end PVC connection is composed of multiple PVC segments. PVC segments are bound by either an UNI and an NNI or two NNIs. For a PVC to be active it is required that all PVC segments are active.

Multiple, either public or private, networks can be concatenated. The network to which an end user station attaches is called an access network. Intermediate networks are called transit networks. Each network is considered to use local addressing, meaning independent DLCI numbers on either end of a PVC. The DLCI numbers used on a link between two adjacent networks, however, must be the same at either end.

The UNI describes the interface between the network and the end user station, the NNI the interface between adjacent networks. Some devices, for example ACF/NCP, are able to provide both UNI and NNI interfaces at the same time.

The NNI is based on the same standards as the UNI. To provide a balanced interface both network and user LMI procedures must be provided at both ends; see Figure 26 on page 61. To allow immediate notification of changing PVC status the use of the asynchronous LMI update procedure by either network is recommended. (see Figure 27 on page 61)

2.3.6.1 NNI Consideration

On an NNI the connecting link will most likely be shared by a high number of DLCIs. To offer end users acceptable response times, adequate link and switching capacity must be available between adjacent networks. Provisions, for example by enforcing agreed committed information rates, must be put in place to avoid end users monopolizing resources.

Special consideration is required to the maximum frame size allowed between two networks. Bidirectional LMI procedures make sure that the stations on either end of the link regularly receive LMI Status (full) messages. Within the full report five-octets are required to report status for each DLCI. When using the two-octet addressing scheme (see Table 8 on page 55) a maximum frame size of approximately 4896 octets is required to accommodate all 976 DLCIs available.

One should make sure the maximum frame sizes are sufficient at both ends of a link.

The DLCI addressing scheme employed should be carefully examined. Although ANSI/ITU-T have defined a two-, three- or four-octet addressing scheme most vendors have only implemented the two-octet scheme, thereby limiting the available DLCIs to endstations on an NNI to 976. This number may not be sufficient for specific installations. For more details on addressing, see 2.3.3.2, “The Data Link Control Layer” on page 53.

The aforementioned problems, limited switching capacity, limited frame size and limited number of DLCIs, may be alleviated by using multiple links between two adjacent networks as each link provides an independent NNI. The use of multiple links, however, requires careful planning. For example, although one can simply split up DLCIs, the traffic generated per DLCI may vary considerably and an equal traffic distribution on each of the available links may be difficult to accomplish.

2.3.7 Multiprotocol Interconnect over Frame Relay

The frame relay standards defined by ANSI and ITU-T as described in the previous sections, provide an unacknowledged and transparent, protocol-independent, data transfer. The frame relay network provides virtual circuits that allow data exchange between stations attached to the network. DLCIs, one on either end, uniquely identify the virtual circuit.

The Internet Engineering Task Force (IETF) has proposed a change, within Request For Comment (RFC) 1490, to provide multiprotocol interconnect over (not restricted to) a single DLCI. RFC 1490 obsoletes RFC 1294. Without this change a separate DLCI has to be assigned to each protocol supported between the two end users. This quickly uses up DLCIs and increases network costs, since carriers usually increase charges per DLCI.

RFC 1490 describes an encapsulation method for carrying LAN interconnection traffic over a frame relay backbone. Additionally, it describes a simple fragmentation procedure for carrying large messages over a frame relay network with a smaller frame size and procedures to obtain a higher level protocol address dynamically: Address Resolution Protocol (ARP), Reverse Address Resolution Protocol (RARP) and Inverse ARP. The remainder of this section summarizes RFC 1490.

No dynamic procedure has been defined to identify the use of RFC 1490. Endstations are assumed to have prior knowledge of the virtual circuits on which RFC 1490 encapsulation is employed and must configure the virtual circuits explicitly for this purpose. Protocol data units (PDUs) must be encapsulated within the frame format depicted in Figure 33 on page 71.

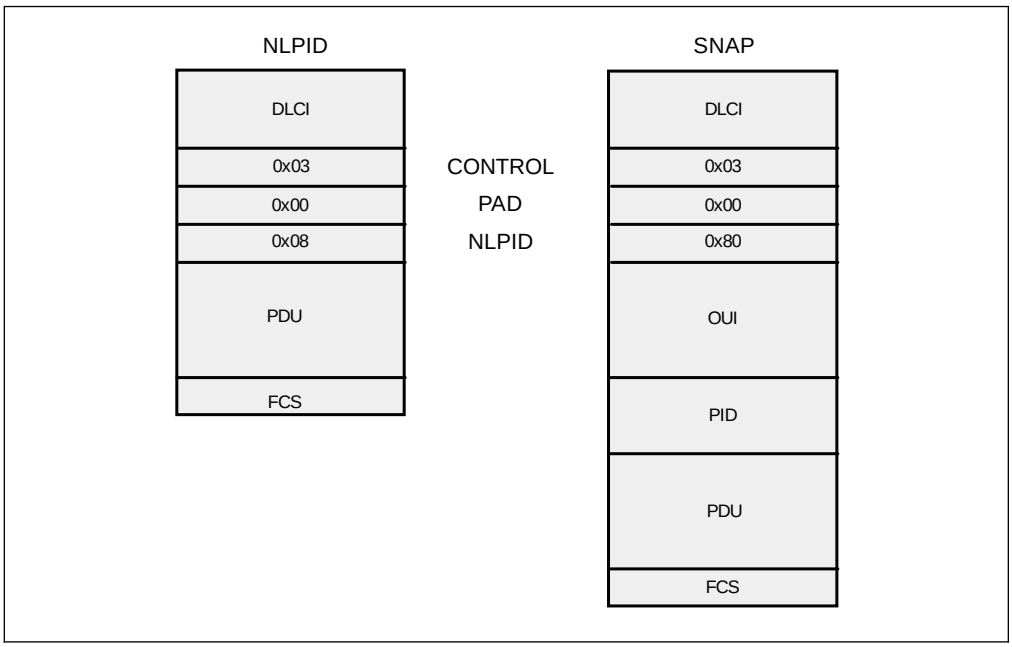

Figure 33. RFC1490 Encapsulation

The following section describes the fields shown in Figure 33:

Control Field

The control field is a one-byte field either indicating X'03' (UI, Unnumbered Information) or X'AF'/X'BF' (XID, eXchange IDentifier, request and response respectively).

The use of XID is optional and used for data link layer parameter negotiation. Parameters negotiated are maximum frame size, retransmission timer and maximum window size. If a station supporting XID receives an XID frame, an XID response will be returned indicating the required parameter settings. If XID exchange is not supported, either values bilaterally agreed on or default values should be used. The default values (as defined in ITU-T Q.922) are a frame size of 262, assuming a two-octet address field, a frame retransmission time of 1.5 seconds and a maximum window of 3, 7, 32 or 40 for 16 kbps, 64 kbps, 384 kbps and 1.535 Mbps or higher, links respectively.

Padding Field (PAD)

The pad is an (optional) field to align the protocol data unit (PDU) to a two-octet boundary. This improves PDU processing time. If available, it must be X'00'.

Network Level Protocol ID (NLPID)

The NLPID is a one-byte field registered by ISO and ITU-T. It contains values for many different protocols. A few of the values are listed in Table 13 on page 72.

Table 13. Network Level Protocol ID (NLPID)	
NLPID	**Function**
X'00'	Not Used[1]
X'08'	CCITT Q.933[2]
X'80'	Subnetwork Access Protocol (SNAP)
X'81'	ISO CLNP
X'82'	ISO ESIS
X'83'	ISO ISIS
X'CC'	Internet IP
X'CE'	Ethertype

Notes:

1. A value of X'00' would not allow the NLPID to be distinguished from a PAD.
2. See 3.1.1, “NLPID and SNAP Encapsulation” on page 87.

IEEE Subnetwork Access Protocol (SNAP)

The SNAP is used to indicate the presence of a PDU for protocols for which no NLPID has been defined.

Note: For some protocols, for example IP, both NLPID and SNAP values have been defined. RFC 1490 recommends PDU encapsulation should use NLPID whenever possible.

If the NLPID equals SNAP (X'80'), the encapsulated protocol data is preceded by a five-octet SNAP header; see Figure 33 on page 71.

The SNAP header contains two fields:

Organization Unique Identifier (OUI)
The SNAP header contains a three-octet OUI that identifies the organization which administers the meaning of the two-byte *Protocol Identifier (PID)* following it.

Protocol Identifier (PID)
The PID identifies the protocol data unit following the RFC1490 header.

Some of the PIDs and OUI combinations are depicted in Table 14 on page 73.

Table 14. IEEE Protocol Identifiers (PID)		
OUI	**PID**	**Protocol**
X'000000'	X'0800'	IP
X'000000'	X'0806'	ARP
X'000000'	X'0807'	XNS
X'000000'	X'6003'	DECnet
X'000000'	X'8035'	RARP
X'000000'	X'809B'	Appletalk
X'000000'	X'8137'	IPX
X'0080C2'	X'0001'/X'0007'	IEEE 802.3[1]
X'0080C2'	X'0002'/X'0008'	IEEE 802.4[1]
X'0080C2'	X'0003'/X'0009'	IEEE 802.5[1]
X'0080C2'	X'0004'/X'000A'	FDDI[1]
X'0080C2'	X'000B'	IEEE 802.6[2]
X'0080C2'	X'000D'	Fragmented PDUs

Notes:

1. These OUI/PID values are used when encapsulating frames in bridged format. Two PID values are given, differentiating between the FCS of the LAN PDU transported being included or not.
2. For IEEE 802.6 BPDUs one value has been defined, as the header of the MAC frame indicates the presence of the cyclic redundancy check field.

Fragmented PDUs will be encapsulated using the SNAP format with an OUI of X'0080C2' and a PID of X'000D'.

2.3.7.1 Bridged and Routed Protocol Data Units

The main function of RFC 1490 is to detail how data should be encapsulated when sent across a frame relay network. Two types of encapsulated protocol data units (PDUs) are distinguished: routed and bridged PDUs.

Within frame relay packets carrying bridged frames the RFC 1490 header does not contain any reference to the layer three data transported; that is, the PDU is networking layer independent. The format of the bridged frame is the same as a frame sent on a local area network (LAN). For example a medium access (MAC) header and, optionally, a routing information field (RIF) are included. An NLPID value of X'80', indicating SNAP, and an OUI value of X'0080C2', the 802.1 organization code, should be used. Table 14 depicts the PID values for the various LAN protocols.

When sending bridged frames the frame relay connection is used as a (virtual) LAN. On this LAN, MAC addresses must be assigned to both ends of the DLCI. Figure 34 on page 74 depicts the concept of virtual LAN.

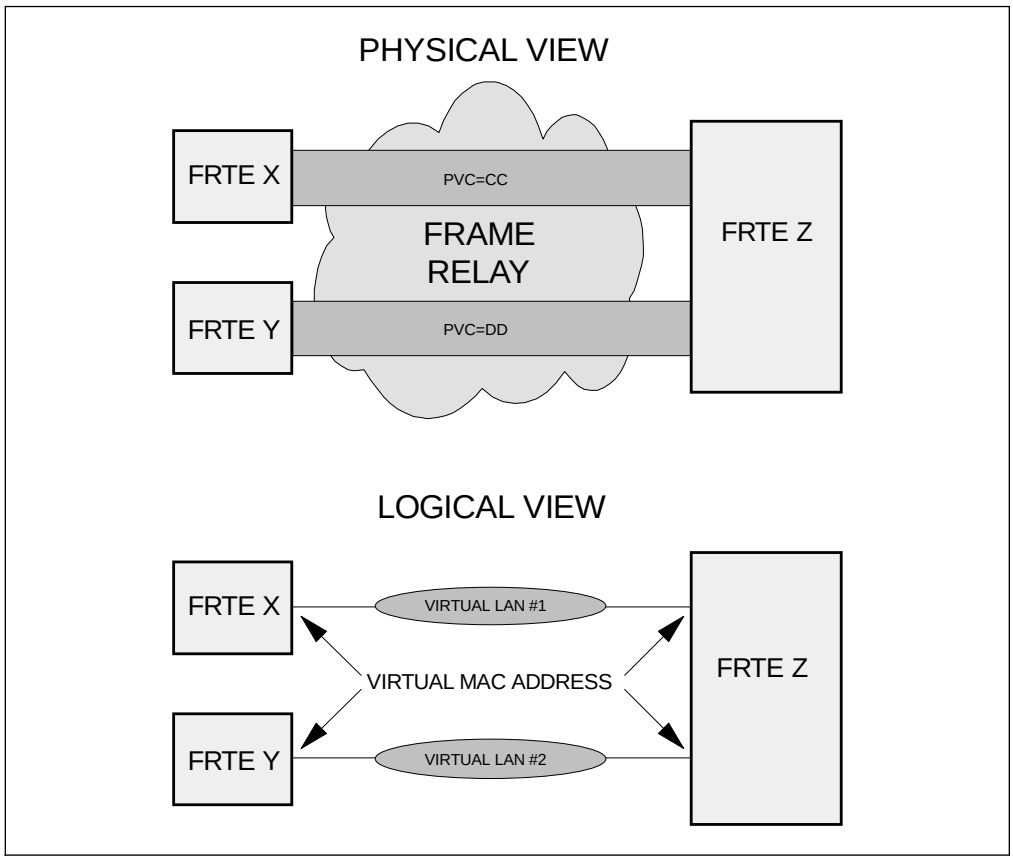

Figure 34. Bridged Frames Using a Virtual LAN

Note: It should be pointed out that the term *bridged* frame does not imply a device is bridging. It only identifies the encapsulation method employed.

In the above figure bridged PDUs are sent between X and Y without any actual LAN bridging taking place in FRTE Z. Instead they are relayed by the frame relaying function in FRTE Z.

The RFC 1490 header of routed PDUs does identify the layer three, or networking layer protocol (IP, SNA, XNS, DECnet, etc.) of the imbedded message. The layer three protocol is indicated by either the NLPID or, if NLPID equals SNAP, the PID. Also data encapsulated according RFC 1490 plus the IBM frame relay extensions is transported as routed PDUs. For routable[3] protocols, such as IP, XNS, DECnet, etc., a layer three protocol address must be assigned to either end of each DLCI.

Note: It should be realized that the term *routed* PDUs does not imply a device is routing; it only identifies the encapsulation method employed.

3 Despite the fact that SNA should be considered as a non-routable protocol, which means no layer three SNA addresses can be identified, SNA data can be transported within routed PDUs.

2.3.8 Frame Relay Access Device (FRAD)

As mentioned before frame relay offers capabilities that makes its use very attractive. It is efficient, allows bandwidth sharing and is protocol-independent. No network manager however, will decide to migrate to a network infrastructure consisting of frame relay only. Not only because time is needed for frame relay to "prove" itself and for network managers to understand and appreciate its concepts but also because it would require considerable investment to equip all endstations with the necessary frame relay hardware and software. Therefore, protocols such as SDLC, X.25, ASYNC, BSC and others will continue to exist for years to come.

Frame relay access devices (FRADs), also known as frame relay assemblers and disassemblers, enable endstations to communicate using the service of a frame relay network without native frame relay support on the endstations themselves. The FRAD concept is depicted in Figure 35.

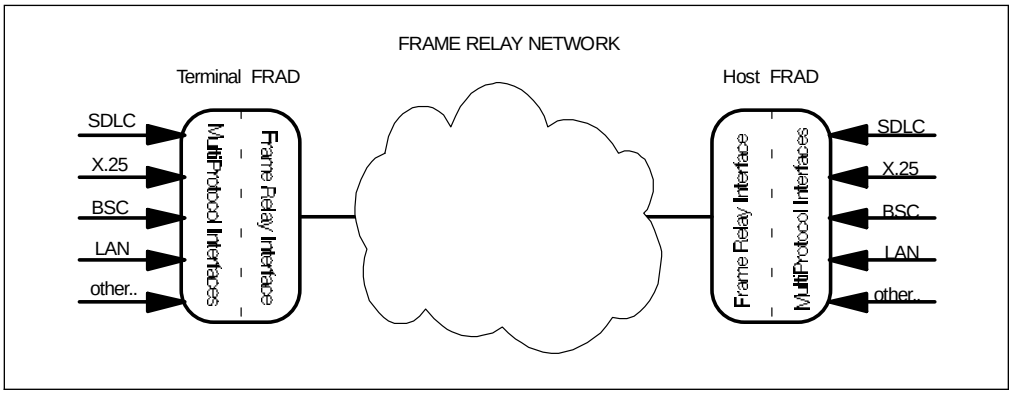

Figure 35. Frame Relay Access Devices

Endstations connect to FRADs using, for example, SDLC, ASYNC, BSC or LAN protocols, which then, via their frame relay network connection enable end-to-end communication. The number of active DLCIs between two FRADs is implementation-dependent. One might think of a single DLCI for all FRAD-to-FRAD traffic, separate DLCIs for each end-to-end connection or hybrid solutions.

FRADs operate in pairs, one on either end of the frame relay network. In a situation where the functions of the FRADs are dissimilar one often distinguishes *terminal* and *host* FRADs. The terminal FRAD is responsible for attaching to remote equipment, and the host FRAD for connecting to wide area networking or host application facilities.

The term FRAD originates from the packet switching world. X.25 ASYNCH PADs (packet assembler/disassembler) have been known for many years and allow asynchronous traffic to be transported over an X.25 network. In addition many X.25 PAD manufacturers provide support for SDLC and BSC. The operation of an X.25 ASYNCH PAD has been standardized, within ITU-T recommendations X.3, X.28 and X.29. For the other X.25 PAD functions proprietary protocols are used. Similar functions have been developed for IP networks. IBM's data link switching (DLSw) has become the de facto standard for transporting SNA and NetBIOS traffic over IP.

Due to the lack of standards no generally accepted definition of a FRAD exists. Essential FRAD functions as defined within this publication, are:

- Frame relay (de)encapsulation
- Fragmentation and assembly of long messages
- Supplementary functions, such as:
  1. Protocol conversion
  2. Protocol optimization

  An example of protocol conversion functions are facilities to convert SDLC data into IEEE LLC 802.2. Examples of protocol optimization functions are: performing the polling of SDLC's peripheral equipment locally instead of end-to-end polling, suppression of IBM LLC 802.2 keep-alive messages flowing over the frame relay network, alternative routing in case of DLCI failures, data compression and others.

One of the most common types of FRAD implementations, the SDLC FRAD, is depicted in Figure 36.

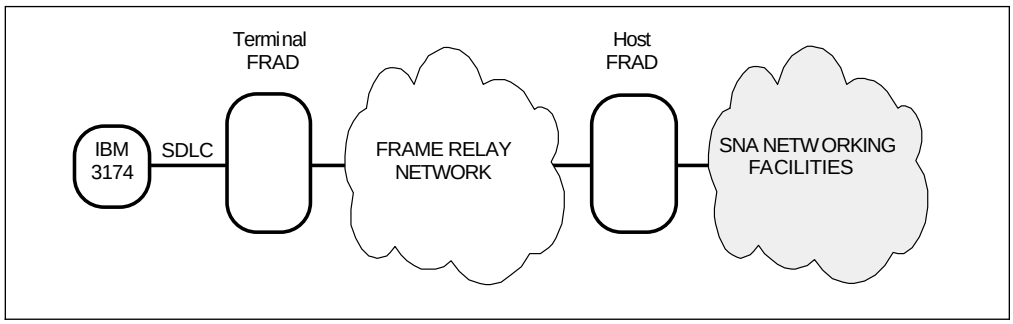

Figure 36. SDLC Frame Relay Access Device

In this example an IBM 3174 without native frame relay support requires access to SNA networking facilities via a frame relay network. The IBM 3174 attaches to a terminal FRAD using its SDLC and SNA peripheral node protocols. The terminal FRAD converts the SDLC peripheral data into a format understood by the host FRAD, encapsulates and possibly fragments the message and forwards it to the host FRAD. This host FRAD de-encapsulates, assembles, and converts the data into the appropriate SNA formats and forwards the traffic to its final destination using SNA networking facilities.

Careful study of vendor product implementation is required to decide on FRAD interoperability. De facto standard for encapsulation and fragmentation is RFC 1490; see 2.3.7, “Multiprotocol Interconnect over Frame Relay” on page 70. The implementation of RFC 1490 is itself no guarantee that FRADs on either end of the frame relay network are compatible and able to interoperate. This mainly depends on the supplementary functions performed by the FRAD.

2.3.8.1 Frame Relay Access Devices and SNA

IBM networking products offering FRAD functions for SNA stations use one of the three implementations shown in Figure 37 on page 77. The implementations are incompatible and two FRADs are only able to interoperate when both FRADs comply to the same standard.

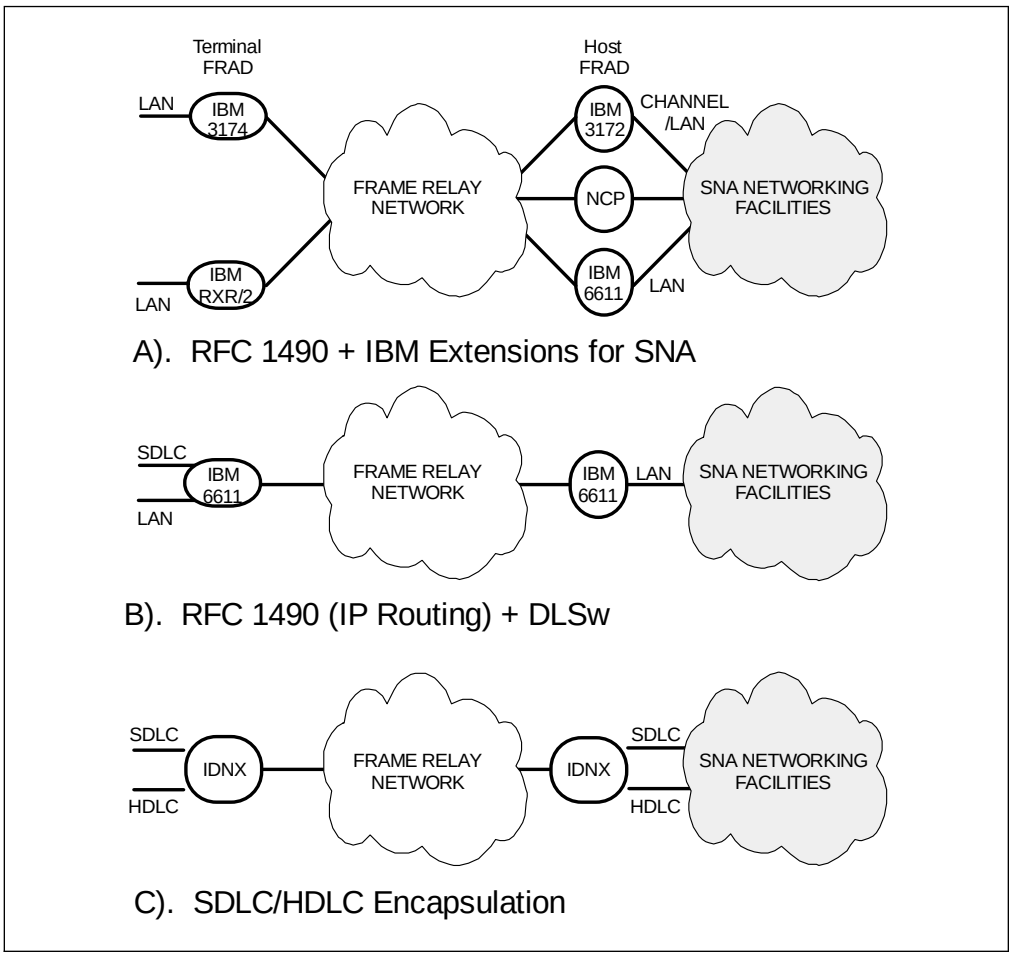

Figure 37. IBM Frame Relay Access Devices

Example A

In Figure 37, example A shows FRAD implementations based on RFC 1490 plus the IBM extensions for SNA devices. IBM networking products providing host FRAD functions based on this standard are the IBM 6611, IBM 3172, 2218, and ACF/NCP. The IBM 6611 attaches to the SNA networking facilities via a LAN connection, and the IBM 3172 via a channel link or LAN connection. ACF/NCP provides, offering SNA and FRAD functions at the same time, immediate access to SNA networking facilities.

Any terminal FRAD that converts the traffic of its attached stations to a data stream complying to RFC 1490 plus the IBM extensions for SNA stations, is able to interoperate with either of the three aforementioned host FRADs. Equipment attached to the terminal FRAD appears to the SNA network as LAN-attached SNA devices. IBM networking products providing terminal FRAD functions based on this standard are the IBM 3174 and RouteXpander/2 (RXR/2). The IBM 3174 and RXR/2 support LAN, Ethernet and token-ring attachment equipment only.

IEEE LLC 802.2 keep-alive messages (for details see the *inactivity timer* on page 241), will be periodically sent across the frame relay network between host and terminal FRAD.

Example B

Example B in Figure 37, depicts an alternate implementation based upon IP routing over frame relay, as defined within RFC 1490, and the IBM data link switching (DLSw) function. Although usually not identified as such, any device that has implemented DLSw, allowing the transport of SNA data within IP datagrams, and supports IP over frame relay is able to perform FRAD functions for SNA equipment.

The IBM 6611 is able to perform terminal FRAD functions for LAN and SDLC-attached stations. Connectivity to the SNA networking facilities requires a LAN interface. The equipment attached to the terminal FRAD appears to the SNA network as LAN-attached SNA devices.

Because DLSw is used, supplementary FRAD functions are offered; remote equipment is locally polled instead of end-to-end, limiting frame relay traffic and making connections less vulnerable to timing dependencies, and FRAD-to-FRAD message delivery is guaranteed, with automatic rerouting if a DLCI fails. For details about the IBM DLSw function see the IBM 6611 product documentation.

Example C

Example C in Figure 37 on page 77, shows a third implementation that is used by IDNX. IDNX encapsulates the SDLC or HDLC traffic of its attached devices, fragments long messages and sends the data across the frame relay network. Encapsulation is done using proprietary protocols and IDNXs are required at either end of the frame relay network.

Any type of device is supported as long as either SDLC or HDLC is used as the layer two protocol. For the SNA network the remote device appears as locally attached. The layer two protocol remains unchanged, within the SNA network. The remote SDLC devices need to be defined as SDLC devices, and the remote HDLC as HDLC devices.

IDNX does not offer supplementary FRAD functions. All data, including polling messages, will traverse the frame relay network unchanged.

2.3.9 Private versus Public Frame Relay Network Service

The prime movers to develop frame relay have been the demand of accommodating the growing requirement for connectivity of LAN-based protocols across the enterprise and the demand to make more efficient use of new networking technologies (in particular, faster and more reliable lines, made available). Also as mentioned in the *IBM Network Blueprint* frame relay is setting the stage for a future migration to ATM. It is no longer a question that frame relay, due to its protocol-independence and high efficiency, offers an adequate solution for these demands.

Many network managers might have considered the migration to frame relay but are reluctant to migrate overnight. Most private networks comprise a significant investment in both equipment and skill geared towards efficient transport and management of traffic and while the propensity and importance of LAN-based applications is growing, mission critical applications are often host-oriented and will continue to be so for years to come. Therefore, while LAN connectivity is becoming increasingly important, migrations should not impair the ability and performance of the mission-critical traffic.

Frame relaying services can be obtained from either public service providers or by using facilities with additional or already available in-house equipment. Public services are offered by a wide variety of carriers, although coverage outside the USA tends to be less pervasive. Solutions to enable the implementation of a private frame relay network offered by IBM, are ACF/NCP, Nways 2220 and IDNX.

The question of private versus public service is not new, and difficult, if not impossible, to answer in general as it depends greatly on the conditions pertinent to the specific environment one is dealing with. It is not unlikely that the optimal combination of functionality and cost reduction and other factors playing a role, will be achieved through hybrid networks, a combination of both public and private networks. The following provides some considerations.

A major justification for public frame relay service is cost reduction as compared to traditional leased line networks. Customers connect to the network via dedicated local access links at a specified bandwidth (usually up to T1 or E1 speeds). Virtual circuits are then defined to provide connectivity to multiple sites. Key to cost reduction is that network providers run very high-speed (T3, E3) links through their network. These links are relatively cheap, in terms of cost against capacity, and used by multiple customers. An additional benefit is that the high-speed lines provide the ability to improve response times.

Another possible advantage is that the service provider has implemented additional availability features, such as automatic re-routing around failed lines or nodes. Important also is the outsourcing aspect. Much of the network administration can be relegated to the carrier, reducing network administration and management burden.

However, private implementation provides a much greater degree of control to the network. For example, organizations may have concern over data security in utilizing a public facility or require a higher degree of availability and performance. If only the wide area lines are leased from a carrier, the amount of the network that is out of the organization's control is limited.

Private frame relay may sometimes be the only option available in certain geographical areas as public service may be absent or not pervasive enough. Also, it is possible that depending on network configuration and carrier pricing structure, private frame relay is a more cost-effective solution.

Upgrading an existing private network to a private frame relay implementation generally involves less network restructuring. In particular, as with the ACF/NCP implementation, all that is required is a software upgrade to existing hardware. Finally some private networks provide value-added functionality which may be desirable in the user environment.

2.3.10 Voice over Frame Relay

Note: This section is based on the *White Paper, A Discussion of Voice over Frame Relay, October 1996*, issued by the Frame Relay Forum.

Voice over frame relay technology offers the possibility to consolidate voice and voice-band data (fax and analog modems) with data services over the frame relay network. Although implementation agreement work for carrying voice over frame relay is progressing within organizations such as the Frame Relay Forum's Technical Committee, there is currently not a uniform standard or implementation

agreement defined for vendor equipment interoperability or for the transport of voice across a carrier's public frame relay network.

Voice must be transmitted in *near-real-time*. This means that transmission and network delays must be very small, otherwise speech cannot be recognized by the listener. Various technical advances (for example, speech digitization and compression) have now made the implementation of voice transmission over packet networks viable.

Still it should be understood that frame relay networks have inherent limitations, it was these limitations, among others, that lead to the development of ATM networks, networks that are designed to transport voice, video, and data traffic. ATM networks offer a guaranteed quality of service (QoS), which allows speech to be transmitted in near real time.

The first step in transmitting voice over a digital medium is speech digitization. Advances in this field allow the amount of *raw* data that needs to be transmitted, to be greatly reduced. This allows more speech to be transmitted on lower speed lines, and produces less load on the networking equipment.

2.3.10.1 Speech Digitization

Analysis of voice samples have shown that less than a quarter of a normal dialog contains the essential information that needs to be transmitted. The rest of it consists of pauses, repetitive patterns and background noise. The following techniques are used for voice digitization:

Removal of repetitive speech sounds

Repetitive sounds (such as the s in the word snake or long o in the word loan) are caused by vibration of the vocal cords. Transmission of these identical sounds is not necessary and their removal can decrease network load.

Removal of pauses (silence suppression)

Pauses between words and sentences, and gaps that come at the end of one person talking but before the other begins, also can be removed. The pauses may be represented in compressed form and they can be re-created at the destination side of the call to maintain the natural quality of the spoken communication. The suppression and removal of silent periods can also significantly decrease the bandwidth needed for voice.

Voice frame formation

The above two steps make it possible for voice to be efficiently compressed. The remaining speech information may be digitized and placed into packets suitable for transmission over frame relay networks. These packets tend to be smaller than average data frames. This helps to reduce transmission delay across the frame relay network.

Voice Compression

Various algorithms are used to sample the speech pattern and reduce the information sent, at the same time retaining the highest voice quality level possible. One way of doing this is using the Adaptive Delta Pulse Code Modulation (ADPCM) algorithm. ADPCM can reduce the speech data rate to half that of the Pulse Code Modulation (PCM), which is an ITU standard for digital voice coding. ADPCM may be used in place of PCM while maintaining approximately the same voice quality.

Another method of doing voice compression is to use Digital Signal Processors (DSPs). These are microprocessors that are designed specifically to process digitized signals such as those found in voice and video applications.

As the available bit rate for voice is reduced from 64 kbps to 32, 16, 8, and 4 kbps or below, the strategies for redundancy removal and bit allocation need to be ever more sophisticated. The typical bit rate that is used at this moment is 8 kbps per voice channel.

Echo cancellation

Echo occurs when the transmitted voice is reflected back to the point from which it was transmitted. In voice networks, echo cancellation devices are used within a carrier's network. The longer the distance, the more delay, and the more likely echo will result.

Voice transmitted over a frame relay network will also face propagation delays. As the end-to-end delay increases the echo would become noticeable to the end user if it is not canceled. Since frame relay carriers do not use echo cancellation equipment in their frame relay networks, it is up to the equipment vendor to address echo cancellation in the equipment.

2.3.10.2 Technical Challenges

The following basic problems that are inherent in frame relay networks must be overcome to make voice over frame relay usable and affordable:

Delay and delay variation

The bursty nature and variable frame sizes of frame relay may result in variable delays between consecutive packets. The variation in the time difference between each arriving packet is called jitter.

Jitter can make it difficult for the receiving end side to smoothly regenerate voice. As voice is a continuous wave form, a large gap between the regenerated voice packets will result in a distorted sound.

To avoid dropping frames, data can be buffered at the speech decoder sufficiently to account for the worst case delay jitter through the network.

Frame loss

Voice over frame relay can usually withstand infrequent packet loss better than data. If a voice over frame relay packet is lost, the user will most likely not notice. If excessive frame loss occurs, it is equally unacceptable for voice over frame relay and data traffic.

Traffic integration - fax and modem support

Apart from voice over frame relay, it is possible to support fax and data modem services as well. This ability may prove to be beneficial to end users that have high fax traffic volumes between branches and headquarters.

Voice band fax and data modem signals can be demodulated and transmitted as digital data in packet format. However, it is difficult to reliably compress fax and data modem signals to achieve the low bandwidth utilization often necessary for the most efficient integration over frame relay. Sometimes a scheme is implemented where voice is compressed to a low bit rate, but upon detection of a fax tone, the

bandwidth is reallocated to a higher bit rate to allow for faster fax transmission.

Prioritization

To minimize voice traffic delay and thus a service degradation, a prioritization mechanism may be used that assigns a higher priority to delay sensitive traffic. For example, it is possible to configure different priority queues for each input traffic type. Voice and fax traffic can be placed in the highest-priority queue, for faster delivery to the network. Lower priority data traffic can be buffered until the higher priority voice and fax packets have been sent.

This prioritization can be implemented by VFRADs, but frame switching equipment inside the network treats all arriving packets with the same priority as there is currently no frame relay prioritization mechanism. To achieve expedited transmission of voice packets, all nodes in a frame relay network would need to implement a prioritization scheme.

Fragmentation

Fragmentation is used to break up larger blocks of data into smaller, less delay-creating frames. This is another means used to ensure the highest voice quality level possible. Fragmentation attempts to ensure an even flow of voice frames into the network, minimizing delay.

As with prioritization, VFRADs may implement a fragmentation scheme, but voice frames could still be delayed inside the network as the fragmentation scheme currently being discussed is limited to the FRTEs at the edge of the network and not to frame relay switches.

Fragmentation also reduces jitter because voice packets can be sent and received more regularly. When used in combination with a prioritization mechanism, a better service level for voice frames may be achieved.

Multiplexing techniques

Sometimes bandwidth optimization multiplexing techniques such as logical link multiplexing and subchannel multiplexing are used. Logical link multiplexing allows voice and data frames to share the same PVC, thus saving money for extra PVCs. Subchannel multiplexing combines multiple voice conversations within the same frame. This reduces packet overhead, which can be useful on low-speed links.

2.3.10.3 Single or Dual Access

Figure 38 on page 83 depicts a typical scenario for voice over frame relay. A frame relay network is being used to carry data traffic and voice traffic. PBXs are interconnected as well as smaller *key systems*. (These are small PBX systems with few subscribers connected to them.)

Imagine the installation shown in Figure 38 on page 83 without any voice traffic. Large frames can be sent between all FRADs. This keeps the utilization of the FRADs low. Adding voice from the key system in location 3, will cause some fragmentation of data packets that are sent on the same access link in the frame relay backbone. As this location only has a key system, the increased utilization of the VFRAD/router will probably be low. In a typical network of one or two central data centers and many remote branches (such as location 3), the effect on the central FRADs will be much greater.

A *single access* configuration as shown for location 2 will receive small voice packets, plus fragmented data packets on a single link from the frame relay network. Increased voice traffic will severely impact the data traffic due to the increased utilization on the VFRAD, even if the frame relay network is capable of handling the traffic.

A *dual access* configuration as shown for location 1 can reduce the impact of voice traffic on the existing data traffic by separating the two types of traffic each onto their own access link into the frame relay network.

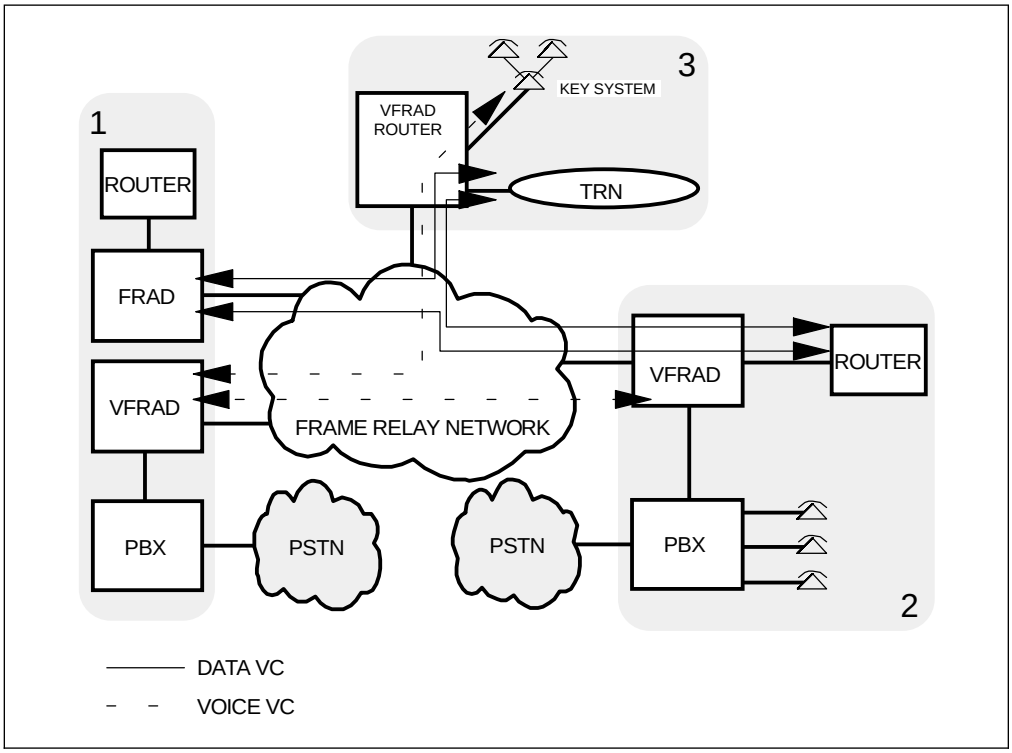

Figure 38. Voice over Frame Relay - Single and Dual Access

2.3.11 Fragmentation

Fragmentation provides frame relay DTEs and DCEs with the ability to fragment long frames into a sequence of shorter frames, which are then reassembled into the original frame by a receiving DTE or DCE. Fragmentation is discussed in the frame relay forum *Voice over Frame Relay Implementation Agreement FRF.12*. Fragmentation is needed to control delay and delay variation when real-time traffic such as voice is carried across the same interfaces as data. The FRF.12 implementation agreement:

- Allows real-time and non-real-time data to share the same frame relay UNI or NNI link.
- Allows the fragmentation of frames of all formats.
- Defines the fragmentation procedure that can be used by other protocols or implementation agreements (such as FRF.11 Voice over Frame Relay).
- Defines three fragmentation models:

Locally across a UNI

This is used to allow real-time and data frames to share the same UNI interface between a DTE and the frame relay network. Since UNI fragmentation is local to the interface, it allows the network to take advantage of higher internal trunk speeds by transporting larger frames, which is more efficient than transporting a larger number of small frames.

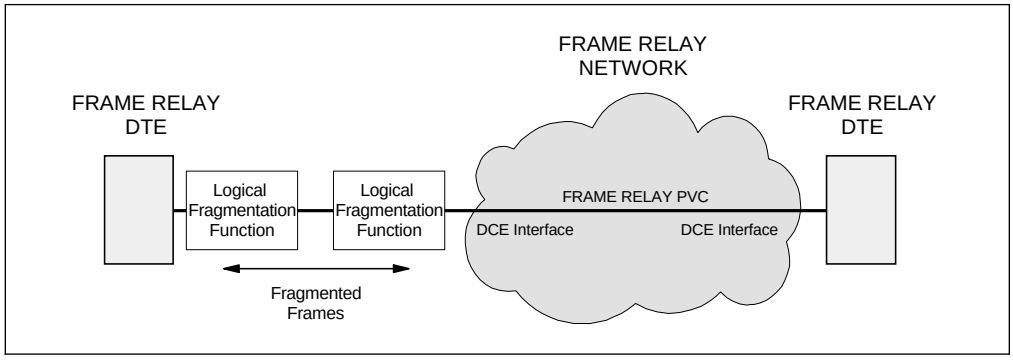

Figure 39. UNI Fragmentation

Locally across a NNI

This is used to allow real-time and data frames to share the same NNI VC. It allows delay-sensitive traffic on one NNI VC to be interleaved with fragments of long data frames on another VC using the same NNI.

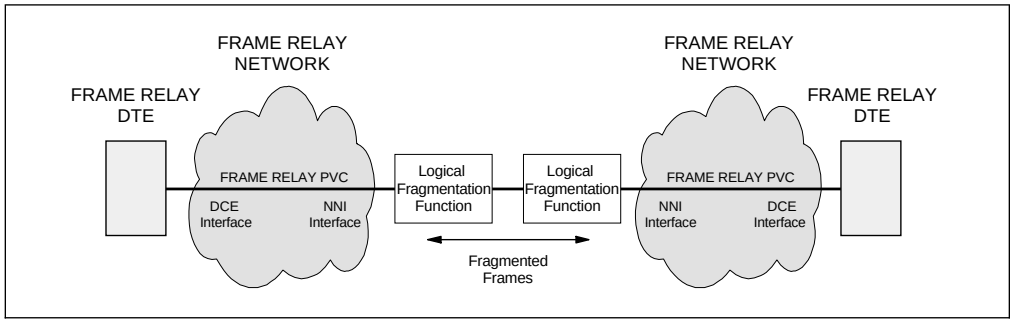

Figure 40. NNI-to-NNI Fragmentation

End-to-end

End-to-end fragmentation is used between peer DTEs, and is restricted to use on PVCs only. It is most useful when peer DTEs wish to exchange real-time and non-real-time data using slower interfaces, but either one or both UNI interfaces does not support UNI fragmentation.

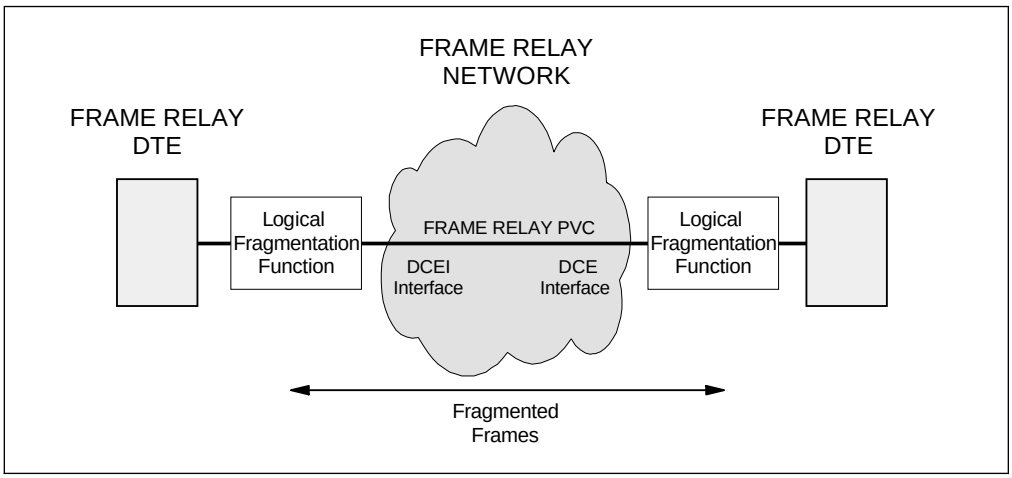

Figure 41. End-to-End Fragmentation

Note: The functionality specified in the implementation agreement has been illustrated as a stand-alone *logical fragmentation function*. It is expected that this function will be implemented in the DTE, DCE, and NNI interfaces.

2.3.11.1 Fragmentation Formats

For interface (NNI and UNI) fragmentation, a two-octet fragmentation header precedes the frame relay header. Figure 42 shows this format. The beginning fragment bit (B) is set to one on the first data fragment derived from a frame and to zero on all following fragments from the same frame. The ending fragment bit (E) is set to one on the last data fragment and to zero on all other fragments from the same frame. B and E may be set together. The control bit (C) is set to zero in all fragments and is reserved for future use. The sequence number is a 12-bit binary number that is incremented modulo 2**12 for every data fragment transmitted on a DLCI. There is a separate sequence number for each DLCI. The *Seq #1* is the high 4 bits, and the *Seq #2* is the lower 8 bits of the sequence number.

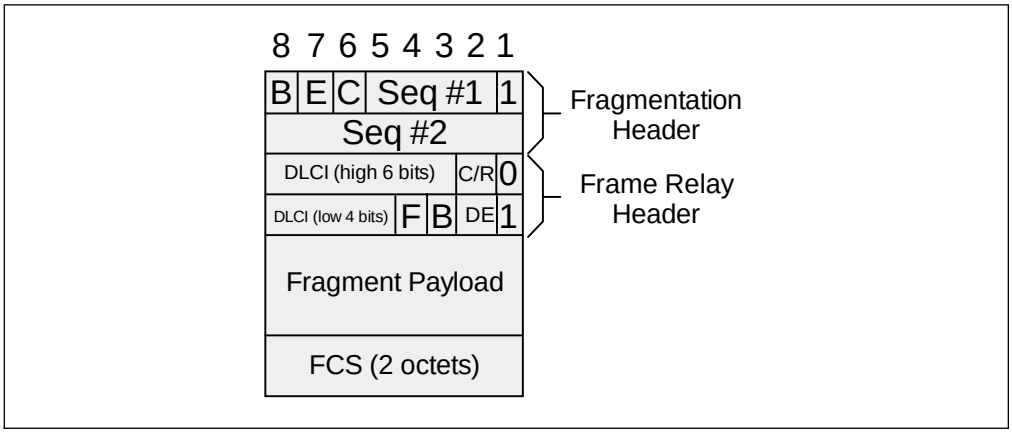

Figure 42. UNI and NNI Fragmentation Format

For end-to-end fragmentation, a two-octet fragmentation header follows the frame relay header. Figure 43 on page 86 shows this format. This allows for transparent transmission of the frames by frame relay nodes which do not support fragmentation.

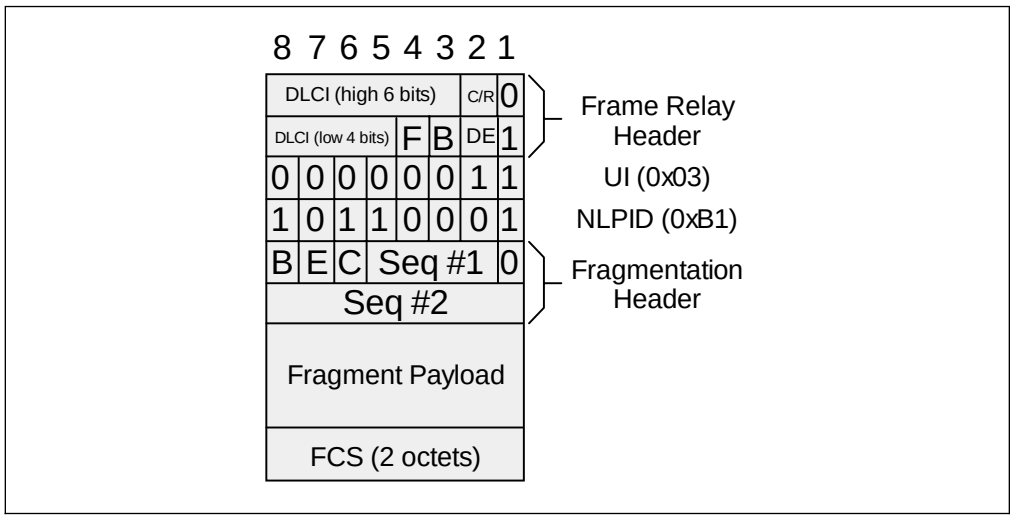

Figure 43. End-to-End Fragmentation Format

Chapter 3. IBM Frame Relay Extensions

IBM introduced SNA support of frame relay networks in 1992, *NCP V6R1* which was available in August 1992. It allowed NCPs to be connected to each other over a frame relay network. NCP acted as a frame relay endstation, called Frame Relay Terminating Equipment (FRTE).

Since 1992 IBM has steadily expanded frame relay support in its product offerings, most SNA-capable products now support frame relay.

Frame relay provides an unacknowledged (that is basically unreliable) transparent, protocol-independent data transfer. Additional format descriptions and procedures are required to allow endpoints to differentiate between the protocols being transported, and to provide a reliable transport.

3.1.1 NLPID and SNAP Encapsulation

2.3.7, “Multiprotocol Interconnect over Frame Relay” on page 70 describes Multiprotocol Interconnect over Frame Relay, which is defined in RFC1490. Frames can have unique Network Level Protocol IDs (NLPIDs) assigned, or they can be coded as IEEE Subnetwork Access Protocol (SNAP) frames.

When encapsulating protocols do not have an NLPID or SNAP assigned, the NLPID value of X'08' can be used to imply encoding according to rules defined in ITU-T Q.933. According to this recommendation the four octets following this NLPID identify the layer two (L2) and layer three (L3) protocol of the imbedded protocol data unit. Figure 33 on page 71 depicts the frame format, and Table 15 shows the layer two and three fields IBM has reserved for SNA and NetBIOS traffic identification.

The two-octet L2 protocol field is either X'4C80' or X'5081'. X'4C80' indicating that IEEE 802.2 LLC is used as the L2 protocol; X'5081' meaning no L2 protocol at all. When L2 = X'4C80', the first octet of the L3 protocol field contains X'70' to indicate the protocol is user-specified and defined in the second octet. The latter will be set to one of the values shown in the following table:

Table 15. Q.933 Traffic Encodings

Q.933	L2	L3	Protocol
X'08'	X'4C80'	X'7081'	Subarea (PU4 to PU4) Traffic
X'08'	X'4C80'	X'7082'	Peripheral (PU4 to PU2, FID2) Traffic
X'08'	X'4C80'	X'7083'	APPN (FID2) Traffic
X'08'	X'4C80'	X'7084'	NetBIOS Traffic
X'08'	X'4C80'	X'7085'	HPR (NLP) Traffic with 802.2 Headers
X'08'	X'5081'	X'7085'	HPR (NLP) Traffic without 802.2 Headers

Note: Although L2 identifiers are defined for peripheral (non-APPN) traffic and APPN traffic, in fact for frame relay no distinction is made between these two types of traffic.

APPN high performance routing (HPR) traffic transported over frame relay links does not use IEEE 802.2 LLC. It uses the L2 protocol ID for NONE and the L3

protocol ID for HPR. HPR transported using the routed frame format does not use a SAP as there is no L2 protocol. When using the bridged frame format, a SAP is used.

3.1.2 Bridged and Routed Frame Formats

The main function of RFC 1490 is to detail how data should be encapsulated when sent across a frame relay network. Two types of encapsulated protocol data units (PDUs) are distinguished: routed and bridged PDUs. The terms *routed* and *bridged* do not imply that the frames are being routed or bridged. It is more accurate to say that the bridged frame format contains the neccesary L2 data (MAC header) to allow it to be bridged at the other side of the frame relay network. The routed frame format does not include such information and therefore, decisions on what to do with the frame must be based on L3 information within the user data.

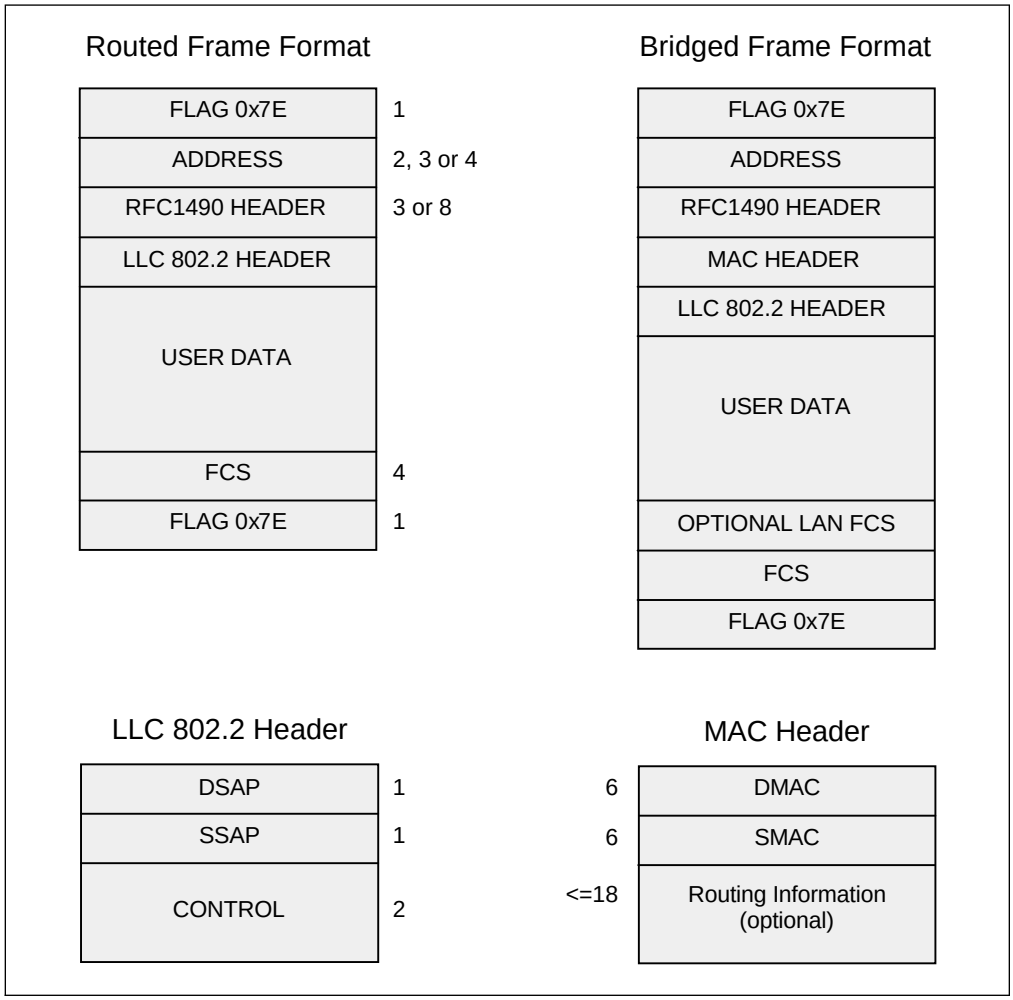

Figure 44. Routed and Bridged Frame Formats

Figure 44 depicts the two methods used to transport SNA over frame relay.

In addition to the frame relay and RFC 1490 control fields:

- The routed frame format contains an IEEE 802.2 LLC data unit, including source and destination service access points (DSAP and SSAP).
- IBM's FR BAN implementation of the bridged frame format optionally contains an IEEE 802.2 LLC LLC data unit encapsulated in an IEEE 802.5 MAC

(token-ring) frame, including source and destination MAC addresses, and a routing information field (RIF).

Due to its smaller header field, the use of the routed frame format is more efficient. However, the bridged format is more versatile and enables the multiplexing of a virtually unlimited number of SNA connections on a single DLCI.

Figure 45 shows routed frames for SNA, APPN and HPR traffic with error recovery (ERP), and without error recovery.

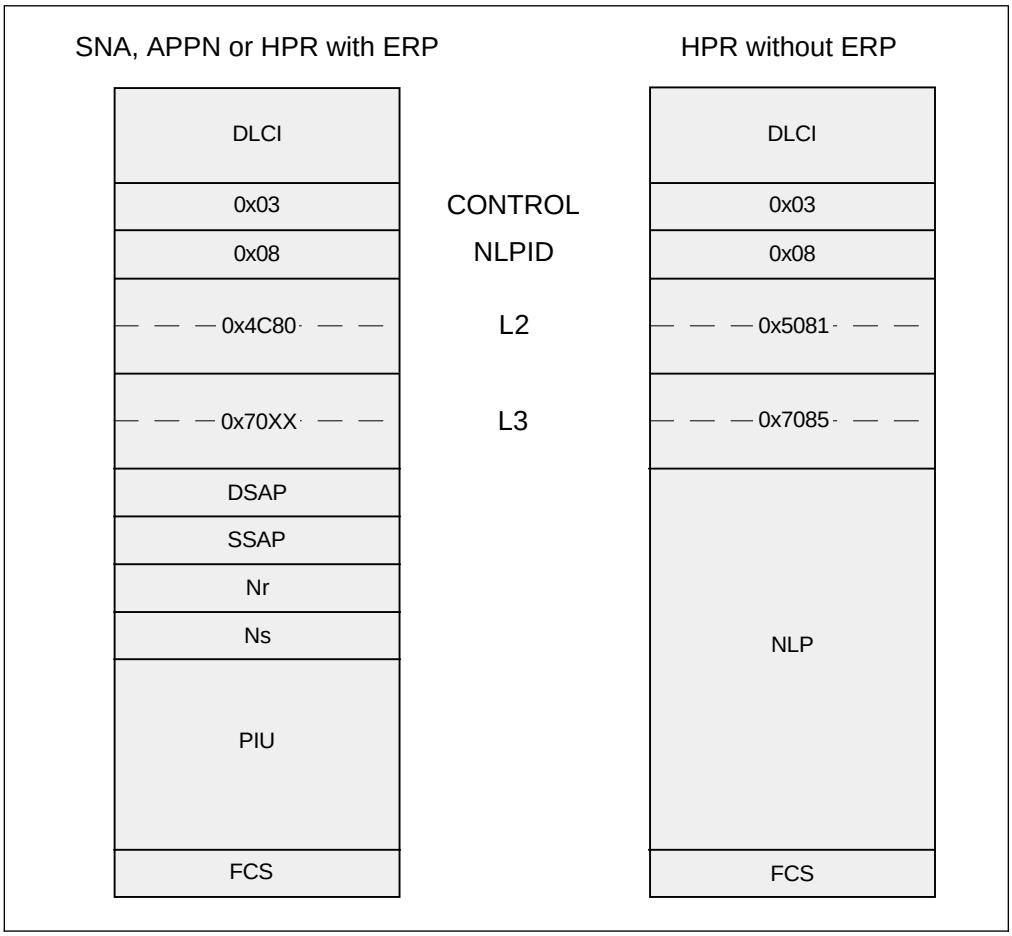

Figure 45. SNA Routed Frame Formats

Figure 46 on page 90 shows bridged frames for SNA, APPN and HPR traffic. The 3745 and 3746 support sending and receiving frames without the optional LAN FCS; the optional LAN FCS is never sent by the 3745 and 3746.

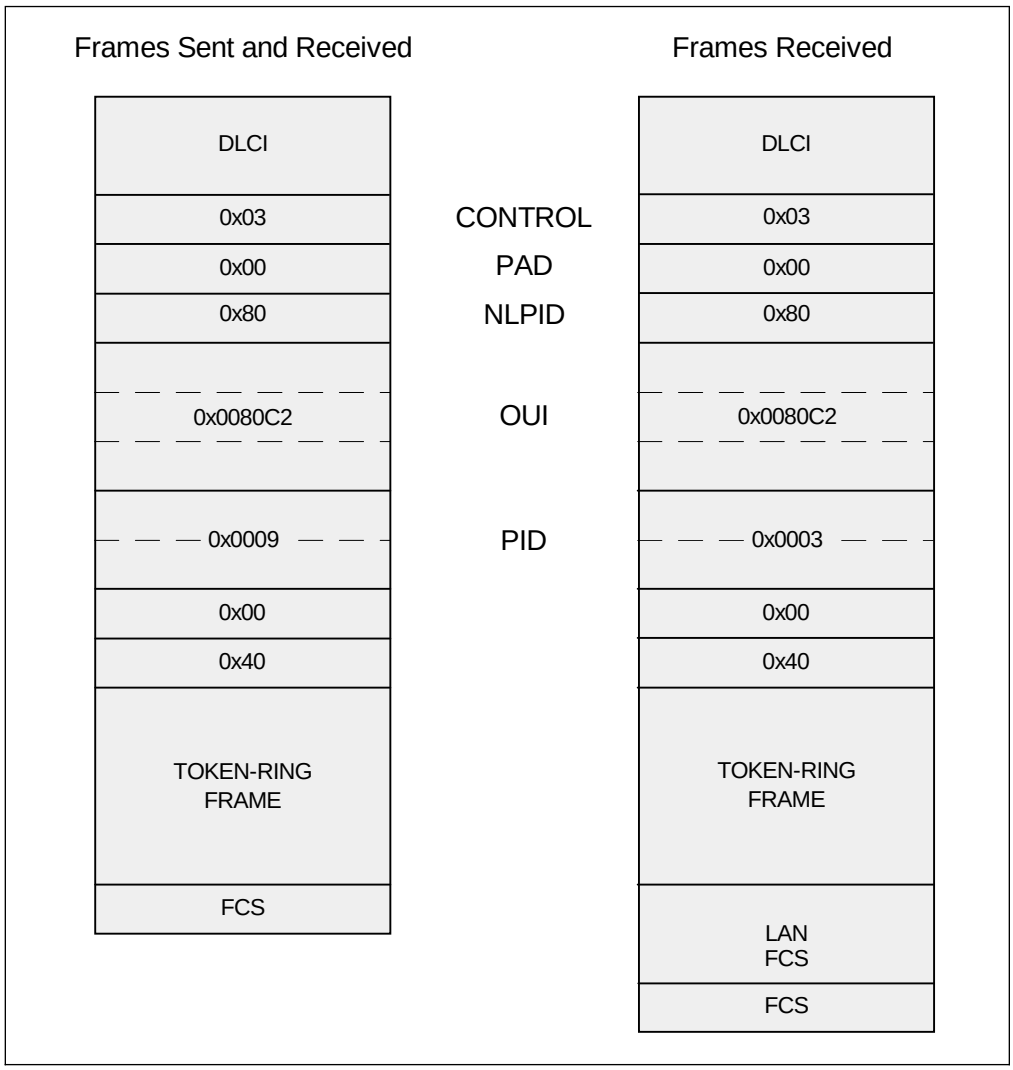

Figure 46. SNA Bridged Frame Formats

IBM supports frame relay *boundary network node (BNN)* and frame relay *boundary access node (BAN)* traffic from peripheral SNA devices primarily to IBM 3745/3746 Nways Communications Controllers. BAN and BNN traffic use the bridged and routed frame formats of RFC1490 to transport SNA traffic over frame relay networks. The following sections describe BNN and BAN.

3.2 SNA over Frame Relay

RFC1490 defines how multiple protocols should be encoded for transport by frame relay networks. SNA and base APPN (without HPR) require that a reliable transport protocol is used. IBM uses IEEE 802.2 LLC functions and procedures to provide a reliable (from the higher layers point of view) transport of SNA data over frame relay (see Appendix A, “IEEE Logical Link Control 802.2” on page 239 for a description of IEEE 802.2 LLC).

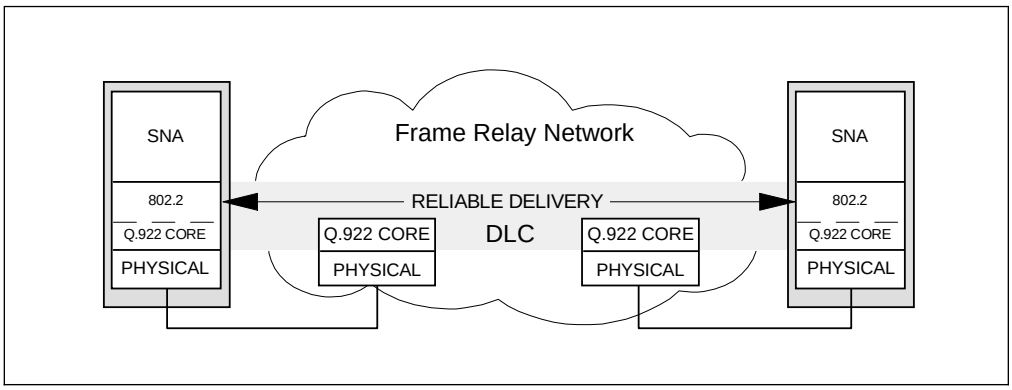

Figure 47. IBM SNA over Frame Relay

IBM frame relay BNN and BAN, which use standard encapsulation techniques, are IBM's own solution for the efficient transport of SNA traffic between peripheral routers and IBM 3745 and 3746 Nways communications controllers.

3.2.1 SNA PU Multiplexing over Frame Relay

There are a number of choices for multiplexing local and downstream PUs over a frame relay link:

One DLCI per PU

For a private frame relay network this may not be problem as far as cost is concerned, but as system definitions are required to define the PVC-PU relationship, it soon becomes a problem as the number of PUs increases. Public frame relay networks charge by the PVC, therefore, it is essential for a customer to use as few PVCs as possible. This method leads to a significant cost disadvantage over solutions where multiple PUs are multiplexed onto a single PVC.

SAP multiplexing

There are 127 valid SAPs available within the RFC1490 routing header. Through system definitions, it is possible to associate each individual PU with a SAP.

MAC multiplexing

This uses the bridged frame format. The MAC address and SAPs are used to differentiate between frames destined for different PUs, either in the BAN router or downstream from it.

DSPU Support

Several router vendors have announced DSPU support. In this approach, the router contains a single SNA PU. All LUs belonging to downstream PUs are made to look like LUs under this PU. This approach is also definition-intensive and it hides the true network topology from management application.

3.2.2 Frame Relay Boundary Network Node (BNN)

A frame relay boundary network node (BNN) uses the *RFC1490 Routed Frame Format* for encapsulation of SNA traffic over frame relay networks. A BNN provides support for SNA functions on the frame relay access device (FRAD) and for downstream PUs (DSPUs) to access the frame relay WAN link. A drawback of this encapsulation method is that traffic on a single DLCI arriving at a peripheral frame relay node, contains no other information in the standard frame relay header that

can be used to determine if or how PU traffic has been multiplexed onto a single DLCI. APPN nodes communicating over frame relay links use this type of encapsulation, routing of frames to nodes downstream is done via the APPN function.

In the case of SNA, the RFC1490 header specifies SNA, and therefore, the frame is passed to the DSAP specified in the IEEE 802.2 LLC header. Figure 48 depicts BNN local SNA support. A single IEEE 802.2 LLC connection (A to A) is required across the PVC.

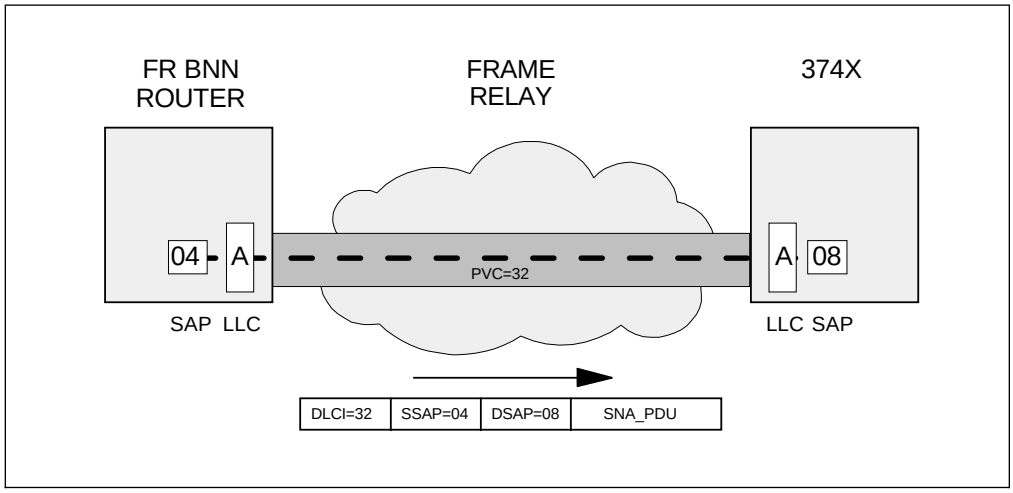

Figure 48. IBM 374X BNN Local SNA Support

Figure 49 on page 93 depicts BNN DSPU support. Some kind of L3 router function is need to route the traffic from DSPUs onto the upstream PVC. An SNA-capable BNN router is shown at the top, SAP multiplexing is shown at the bottom. Example frame formats are also shown. Downstream LLC (D to D, for example) is terminated at the BNN router. In each case a single LLC connection is used over each of the upstream PVCs.

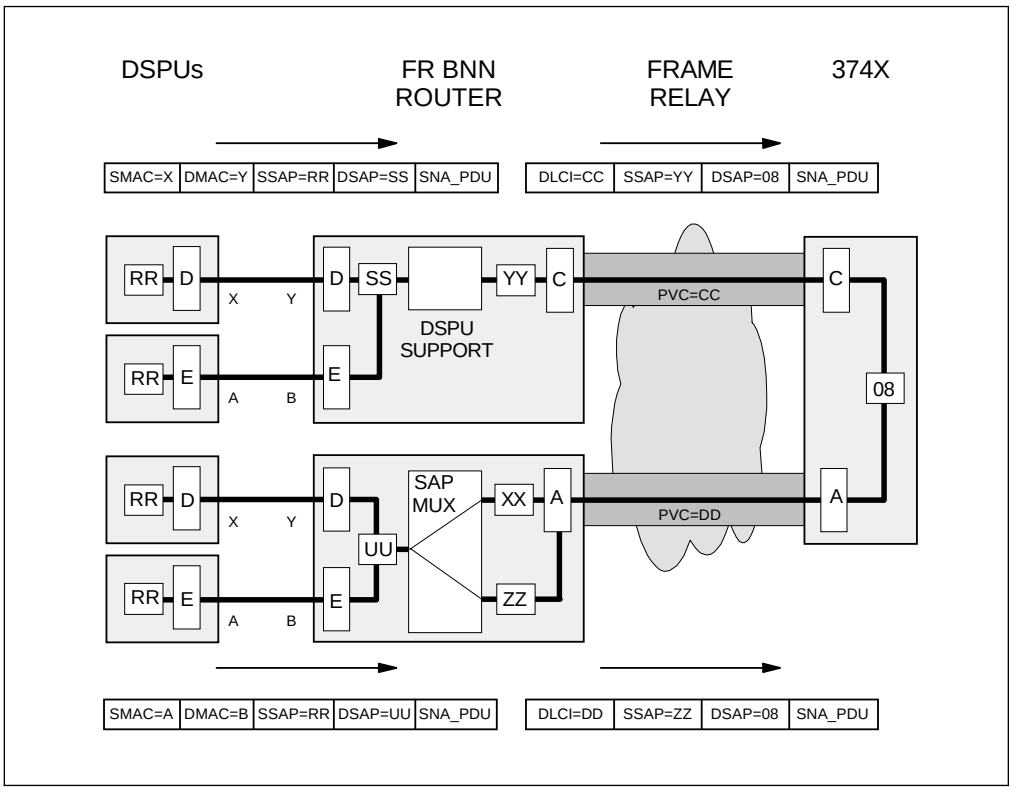

Figure 49. BNN DSPU Support

FR BNN supports multiple LLC connections over a single DLCI. Multiple LLC connections provide simultaneous SNA communication with AS/400 or 3174, and their downstream PUs (DSPUs), over a single DLCI.

3.2.3 Frame Relay Boundary Access Node (BAN)

A frame relay boundary access node (BAN) uses the *RFC1490 Bridged Frame Format* for encapsulation of SNA traffic over frame relay networks. The BAN support can be divided into different functions:

Support of Source Route Bridging (SRB)

Support for the RFC1490 Bridged Frame format, and the ability to bridge traffic between LAN and frame relay WAN links.

Local SNA use of SRB

The ability to bridge traffic from the SNA components in the BAN router internally onto the frame relay WAN link.

LLC Termination

Nodes that support BAN-2 can terminate LLC connections in the BAN router and multiplex that traffic onto LLC connections across the frame relay WAN.

Before the introduction of FR BAN, either TCP/IP encapsulation or FR BNN were used to transport SNA over frame relay. FR BAN was designed to specifically address the limitations of FR BNN described in 3.2.2, “Frame Relay Boundary Network Node (BNN)” on page 91 when used to support multiple downstream LAN devices. Use of the RFC1490 bridged frame format removes the architectural limit on the number of stations supported on a particular PVC. LAN stations can be added to the downstream network with no configuration required at the BAN router.

Alerts sent from downstream stations contain both PVC and MAC address information allowing full identification of the source.

The bridged frame format also results in less overhead at the BAN router. Source Route Bridging (SRB) information in the frame allows each frame to be bridged through the BAN router (rather than routed by it), leading to significantly higher throughput. As a disadvantage, FR BAN frames have a larger header leading to lower throughput on frame relay links.

Frame relay BAN was designed specifically to allow LAN and SDLC peripheral nodes (SNA T2.0 and T2.1) direct frame relay access to the NCP boundary function of the 3745 and 3746 Model 900. This support is a combination of the FR BAN router and support in the 3745/3746 which supported the sending and receipt, and internal routing within the 3745/3746 of bridged format frames.

3.2.3.1 Frame Relay BAN-1

When using BAN-1, the frame relay access node or BAN router is functioning as a LAN bridge. Through MAC level bridging, the remote SNA nodes can access the destination MAC address directly. In the case of the 3746 NN, this is by using the virtual LAN address assigned by the 3746 NN to its frame relay interface. The SNA function residing on the BAN router node, and DSPUs (for example, see Figure 50 and Figure 51 on page 95), can access the BAN function. For DSPUs the routing information field contained within the RFC 1490 bridged SNA frame will indicate its LAN connectivity.

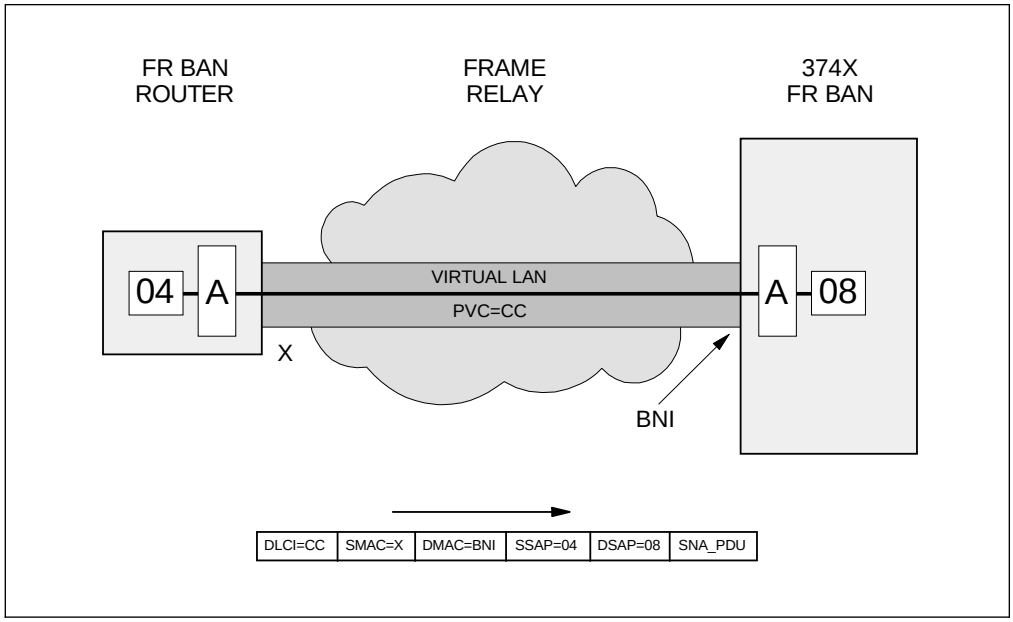

Figure 50. IBM 374X BAN-1 Local Support

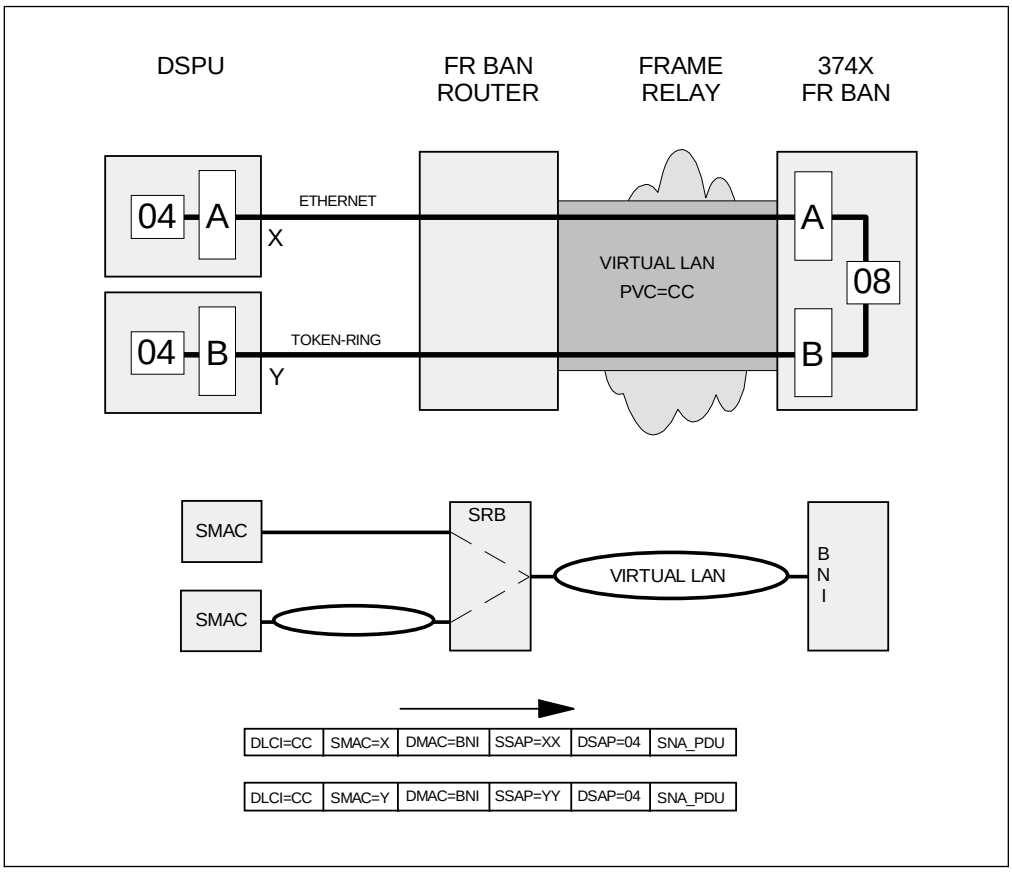

Figure 51. IBM 374X BAN-1 DSPU Support

As with FR BAN-1 the remote access node is only bridging, and the LLC connections between DSPUs and remote node are end-to-end.

When using the FR BAN-1 feature, the BAN router is responsible for filtering unwanted traffic. It has to examine each frame, and only those with the destination MAC address of the remote node should be allowed to pass.

Note: Although the IBM 374X (for example) will simply discard all non-SNA data, filtering unnecessary data will improve line utilization and improve response times.

3.2.3.2 Frame Relay BAN-2

When using FR BAN-2, the DSPUs no longer access the remote node using its virtual LAN address. Instead they use a BAN MAC address defined on the BAN router providing the frame relay connection to the remote node. If the FR BAN router has multiple PVCs to the remote node, multiple BAN MAC addresses can be assigned. For example, the BAN router depicted in Figure 52 on page 96 has assigned two separate BAN MAC addresses. Each BAN address maps onto a separate PVC.

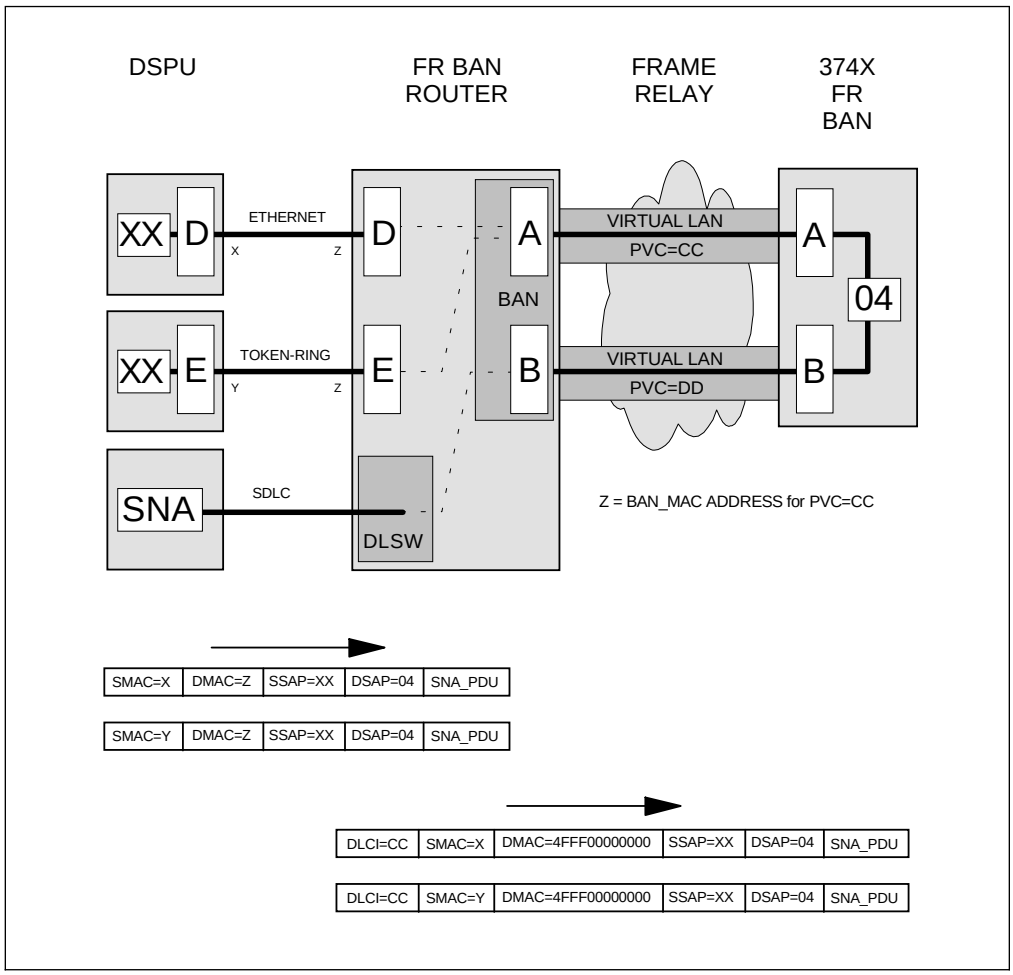

Figure 52. IBM 374X BAN-2 Support

With FR BAN-2, the BAN router does not function as a LAN bridge. Instead, it provides a very effective way of preventing unnecessary traffic passing the frame relay network. Only SNA traffic destined for the BAN MAC address will be forwarded.

The router terminates the LLC traffic received from attached end stations. At the same time, the router establishes a new LLC connection to the NCP over the frame relay network, and converts the BAN MAC address into the 3746 NN virtual MAC address vice versa. Thus, though two LLC connections exist within the connection, the break between them is transparent both to the NCP and to the endstation.

Note: In both cases (BAN-1 and BAN-2), once traffic has left the frame relay BAN router, the traffic is in the standard RFC1490 Bridged Frame Format. There is no difference.

3.2.3.3 Frame Relay BAN-2 and Data Link Switching

One side benefit that BAN-2 brings is the support of SDLC attached devices. Normal DLSw support encapsulates SNA traffic in TCP and then transports the data over IP links. The DLSw function also terminates the SDLC or LAN LLC connections, preventing timeouts from occurring due to delays in the IP part of the network. BAN-2 support terminates the SDLC connection, but the SDLC traffic is then internally passed to one of the outbound frame relay LLC connections. There is no encapsulation of the SNA data in TCP (see Figure 53 on page 97).

In addition, DLSw SNA traffic coming into the BAN-2 router over IP connections is removed from its TCP encapsulation and internally routed onto outbound LLC connections.

In both DLSw methods described previously, SNA data is transported in IEEE 802.2 LLC frames across the frame relay network.

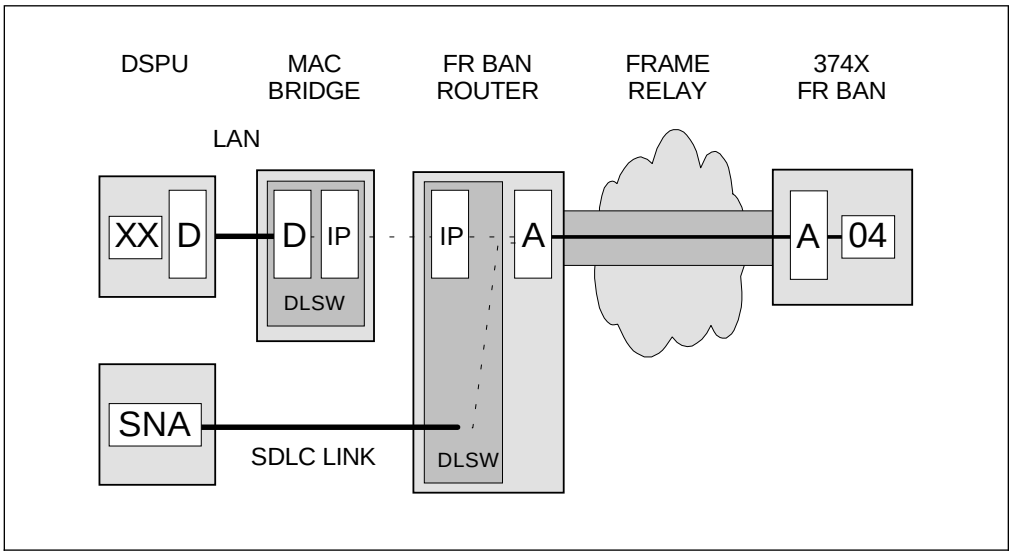

Figure 53. IBM 374X BAN-2 and DLSw

3.2.3.4 IBM 3745 and 3746 Model 900 BAN Support

Note: Although the 3746 NN supports the bridged frame format, it is not a LAN bridge; no LAN bridging is possible between the virtual LANs on the 3746 NN.

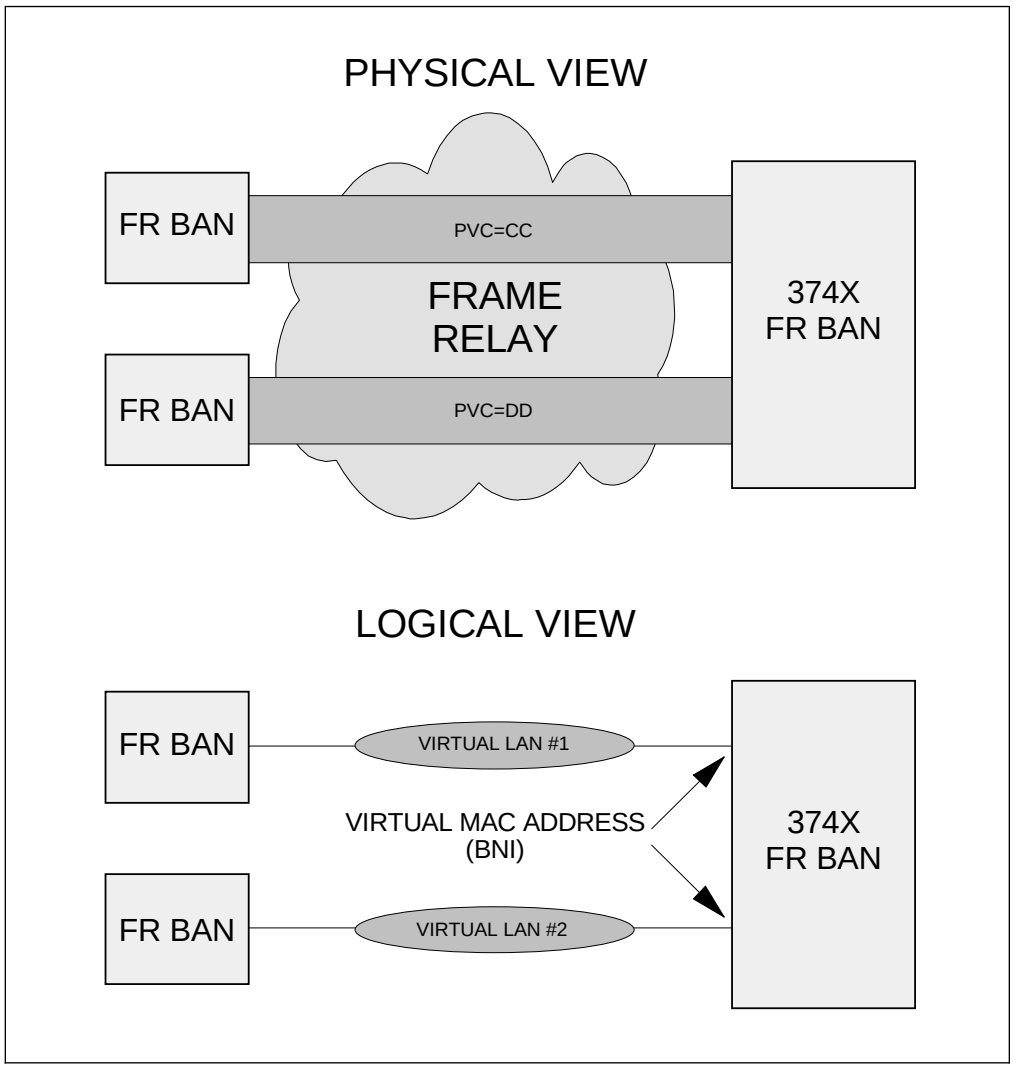

Figure 54. FR BAN Virtual LANs

The MAC address assigned to the 3745 (we will use 3745 to refer to 3745 and 3746 Model 900 in this section) is known as the *Boundary Node Identifier (BNI)*. This BNI is the destination address for all traffic origination at downstream stations. Whether the BAN router uses BAN-1 or BAN-2, the 3745 always receives frames with a destination address that corresponds to its BNI.

Frame relay BAN support in the 3745 allows for the internal transfer of the bridged frame format frames between the adapters and SNA functions in the 3745 and 3746. Initial support allowed BAN frames to be internally transferred to the SNA boundary function (peripheral traffic) and the APPN network node in the 3746 NN (APPN traffic). Later the capability to transfer BAN frames to the subarea function (INN traffic) was also added.

With FR BAN, only one DLCI is normally needed. However, FR BAN may use many DLCI connections between the remote node and the IBM 374X. In some cases, you may want to set up more than one DLCI to handle FR BAN traffic. If multiple LLC connections are used on a single DLCI, the IBM 374X differentiates between the DSPUs using the different source and destination SAP pairs, or different source MAC addresses, within each individual LLC connection.

3.2.4 Comparison of SNA Transport over Frame Relay Networks

There are many methods available for the transport of SNA, APPN and HPR traffic over frame relay networks. This section compares these methods and describes where each method is best suited.

3.2.4.1 SNA Encapsulation in IP

Where a router network already exists, and a small percentage of the total traffic is SNA, IP encapsulation through DLSw can be a cost-effective solution. IP traffic over frame relay is supported by NLPID encapsulation (NLPID=X'CC') (see Table 13 on page 72). The TCP layer provides reliable delivery over the frame relay network as shown in Figure 55.

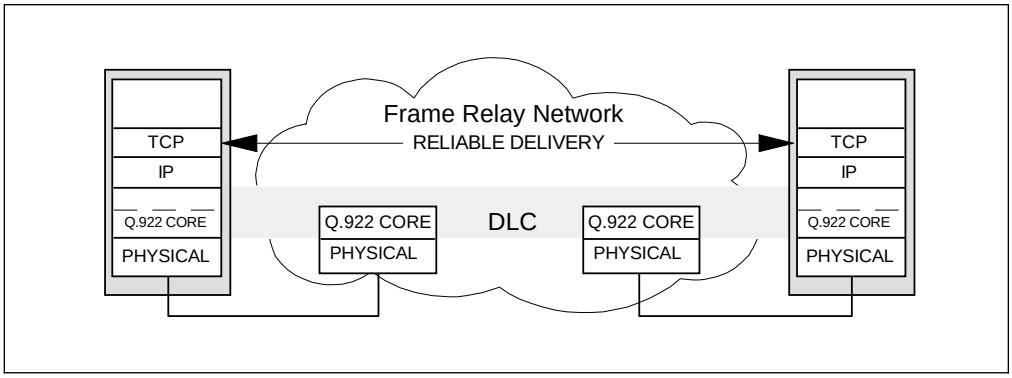

Figure 55. IP over Frame Relay

As previously discussed, BAN-2 can terminate local SDLC and LLC connections providing some protection against timeouts caused by delays in the IP network. The downside of the IP encapsulation approach is the decreased effective throughput of the network. Frame relay headers, IP headers, TCP headers, DLSw information and SNA headers are all transported across the WAN links. In addition, intermediate routers have a lower throughput than frame relay switches.

DLSw requires a terminating DLSw router at both ends of the IP network. This means a router must be between a 374X and the frame relay network (as shown in Figure 56 on page 100).

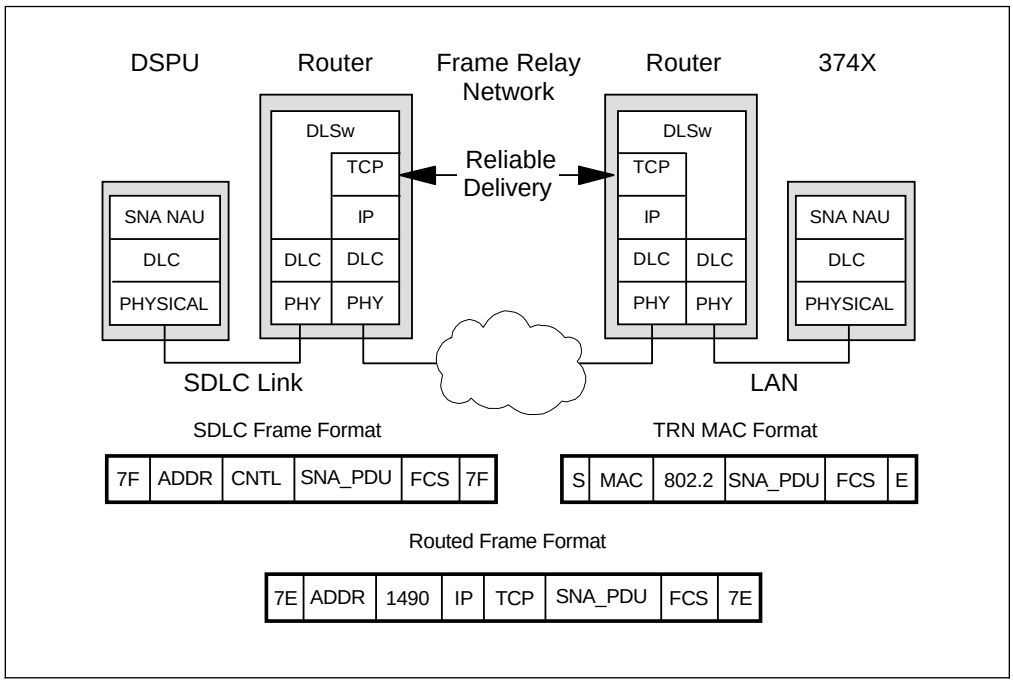

Figure 56. DLSw Transport of SNA

3.2.4.2 APPN or BAN

For base APPN traffic (non HPR), should the FRAD be an APPN network node (NN) or a BAN/BNN router?

Certainly, for full APPN management of all intermediate hops, the optimal solution is to use an APPN NN as a FRAD. In this case, the APPN traffic is routed by the NN as shown in Figure 57 on page 101. Downstream APPN nodes may be connected by a various DLC types; the upstream traffic may be multiplexed onto a single upstream LLC connection, or onto multiple connections. Using full APPN routing, no definitions are required; route discovery and setup is dynamic. One drawback to this approach may be the processing power and storage needed by the FRAD.

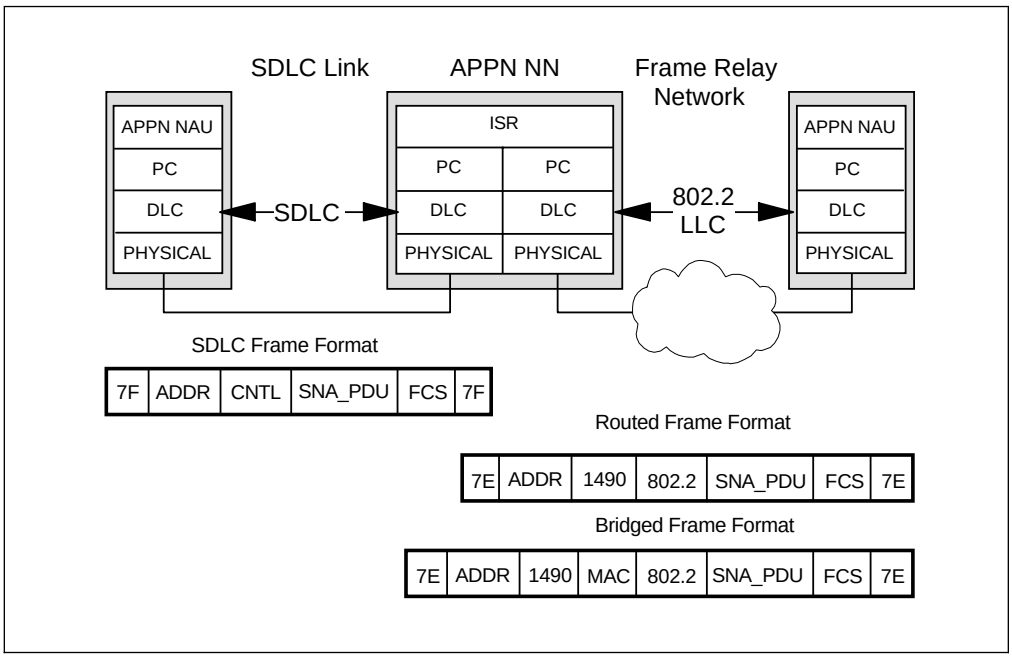

Figure 57. APPN Capable FRAD

For peripheral SNA, BAN bridging provides an optimal solution where downstream PUs must be supported. For APPN traffic, (where an APPN capable FRAD is too costly), then a simple BAN capable FRAD would provide fast bridging between the downstream links and the frame relay WAN. The BAN bridge is transparent to SNA and APPN network management.

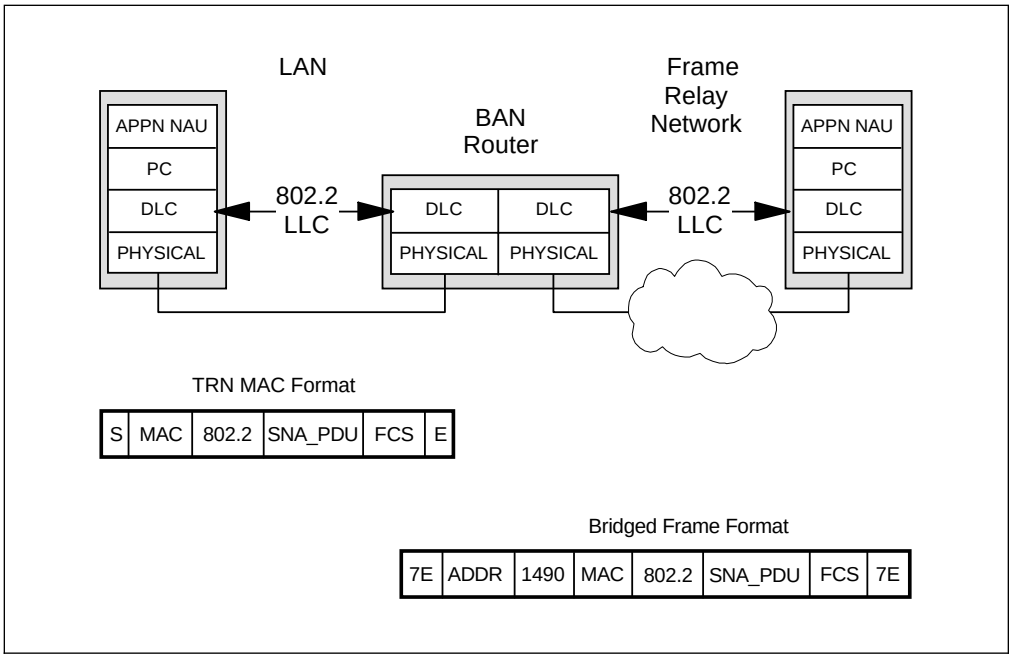

Figure 58. Frame Relay BAN Router

In both these solutions, the frame relay network connects directly to the 374X controller.

3.2.4.3 HPR Traffic over Frame Relay

APPN high performance routing (HPR) supports transport of SNA traffic over links without error recovery procedures (ERP). The RTP component provides end-to-end reliable data transport for the APPN sessions using HPR. Although HPR works without ERP on most link types, not all implementations support non-ERP traffic on all link types. Where non-ERP is supported, Figure 59 shows how APPN traffic can be transported across LANs and frame relay networks without LLC being used. This allows very efficient use of the WAN links due to the low overhead in each packet. The ANR routing function in intermediate nodes is also very fast and produces very low utilization in intermediate nodes such as the APPN FRAD shown.

Note: The bridged and routed frame formats can be used to transport HPR traffic across a frame relay network.

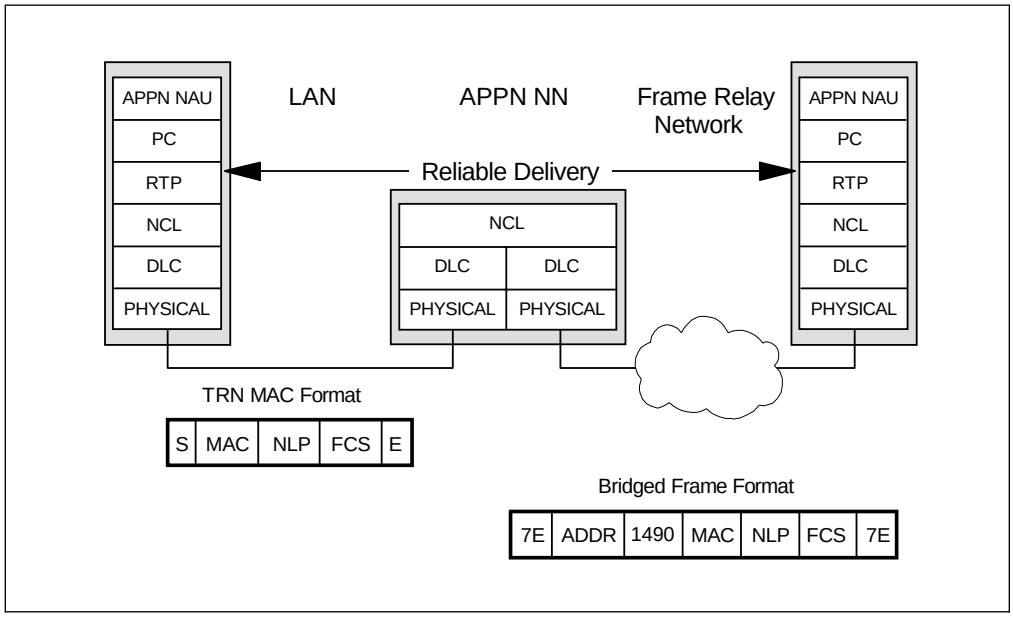

Figure 59. HPR over Frame Relay Example

3.2.4.4 Optional Intermediate Node Interaction with ARB

The adaptive rate-based (ARB) congestion and flow control algorithm is designed to let RTP connections make efficient use of network resources by providing a congestion avoidance and control mechanism.

The basic approach used in this algorithm, as the name implies, is to regulate the input traffic (offered load) of an RTP connection based on conditions in the network and the partner RTP endpoint. When the algorithm detects that the network or the partner endpoint is approaching congestion and the path becomes saturated, resulting in increased delays and decreased throughput, it reduces the rate at which traffic on an RTP connection is allowed to enter the network until these indications go away. When the network or partner endpoint is sensed to have enough capacity to handle the offered load, the algorithm allows more traffic to enter the network without exceeding the rate that the slowest link on the path of an RTP connection or the receiver can handle.

For a mechanism such as ARB to function correctly in frame relay networks, there must be interaction between the frame relay congestion notification mechanisms

(FECN/BECN) and the routing transport protocol. This interaction is achieved by mapping the FECN indication into the network layer header (NHDR) Slowdown indicator as described below. When transporting HPR traffic over DLSw over frame relay, this interaction is not possible.

In Figure 60, an RTP connection exists between nodes A and D, which is carried over a virtual circuit through a frame-relay network between nodes B and C. The ARB sender in node A has increased its (allowed) send rate to a value that is greater than the CIR defined for the frame-relay virtual circuit. When the actual send rate exceeds the CIR, the frame relay network sets the forward error congestion notification (FECN) bit, but as long as the frame-relay network has enough bandwidth available, data will arrive in node C with no additional delays. Because the ARB receiver in node D does not measure an increase in delays, it will notify the ARB sender in node A to further increase the send rate. The ARB sender in node A will continue increasing its send rate until it exceeds the excessive information rate (EIR), which is the maximum rate allowed over the frame-relay virtual circuit. The frame-relay network discards packets that are in excess of EIR, which in turn (because packets are lost) causes the ARB sender in node A to cut its send rate drastically (by one half). Because still there are no measured delay increases, the ARB sender will again increase its send rate beyond EIR and will have to cut back after a number of measurement intervals. This wide oscillation in the send rate is undesirable because it reduces overall throughput.

The preferred method of operation in this scenario is to notify the ARB sender when congestion starts to occur (that is, when CIR is exceeded). This then causes the ARB sender to moderate its send rate such it never exceeds EIR (thus avoiding packet losses). This minimizes oscillation and increases overall throughput. To achieve this, the frame-relay DTE in node C, when receiving the FECN indication, maps the FECN bit into the network layer header (NHDR) Slowdown indicator. This eventually causes the ARB sender to reduce its send rate.

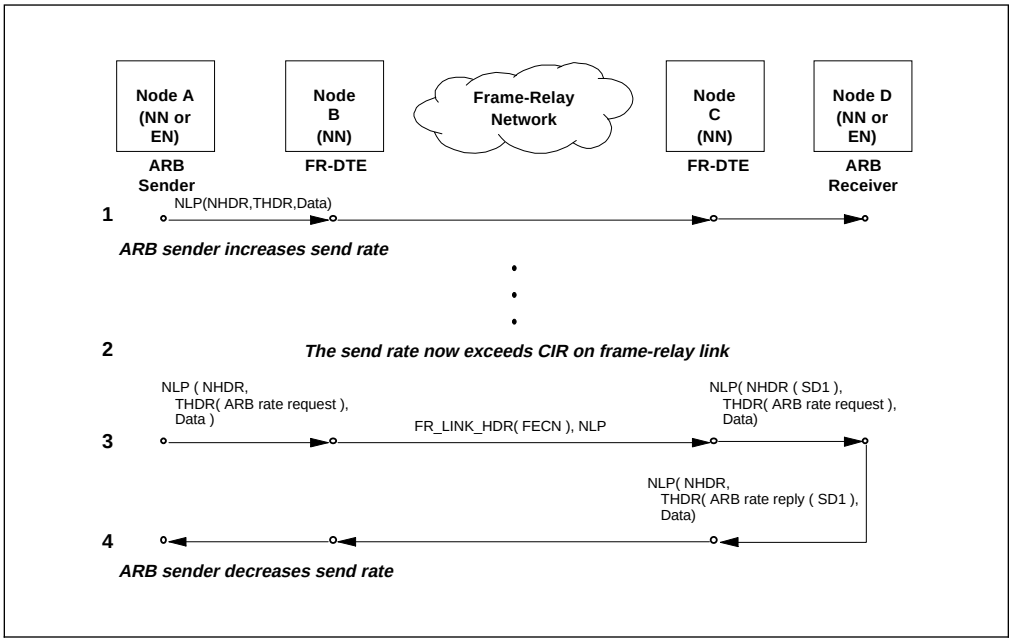

Figure 60. Intermediate Subnet Causing ARB Send Rate Reduction

Figure 60 explains how this works in detail.

1. The ARB sender in node A is sending network layer packets (NLPs) and keeps increasing its send rate because there are no delays in the network.
2. At this point, the send rate exceeds the CIR over the frame relay virtual circuit.
3. The frame relay network sets the FECN indicator on the frame relay link header (FR_LINK_HDR) in every packet that exceeds the CIR. When node C receives a packet with the FECN indicator, it sets the Slowdown1 (SD1) indicator in the NHDR. (That is, it maps the FECN indicator into the slowdown1 indicator.) The ARB receiver in node D observes SD1 when it is set in an NLP that contains an ARB rate request. (The SD1 indicator is ignored in NLPs that do not contain an ARB rate request.)
4. The ARB receiver reflects the SD1 condition by setting the Slowdown1 indication in the ARB rate reply. This causes the ARB sender in node A to reduce its send rate and prevent packet loss by not allowing EIR to be exceeded on the frame-relay virtual circuit.

3.2.5 Summary of BNN and BAN

The following section is a short summary and comparison of frame relay BAN and BNN.

Boundary Network Node (BNN) has the following features:

- Boundary network node (BNN) function uses the routed frame format.

  As the routed frame format provides a more efficient data transport in the WAN, it is recommended that FR BAN support only be used when necessary.
- PU multiplexing using the following techniques:
  - Each SNA node uses a different PVC
  - By SAP multiplexing (only 127 concurrent upstream LLC connections per DLCI are possible)
  - DSPU support by hiding the downstream PUs
- FR BNN is system definition intensive if DSPUs are used.
- FR BNN is not supported on the IBM 6611.

Boundary Access Node (BAN) has the following features:

- BAN uses the Bridged Frame Format
- BAN frames includes layer 1 addresses (MAC addresses)
- BAN supports SDLC devices via DLSw
- No system definitions of LAN stations are needed
- There is no architectural limit to number of stations on a PVC
- BAN produces less router overhead, therefore better throughput than BNN
- BAN supports the multiplexing of SNA and IP traffic on the same or different PVCs
- When using BAN for Ethernet-attached DSPUs, the access node must support translational bridging
- BAN supports a maximum of 127 LLC connections to a single remote MAC address

There are two versions of FR BAN:

FR BAN bridging (FR BAN-1) The access node bridges end-to-end LLC connections between DSPUs and the NCP. FR BAN-1 is supported on 6611, 2210, RXR/2, 2217, 3172, AS/400, 3174 and frame relay token-ring bridge program/DOS.

FR BAN routing (FR BAN-2) DSPUs use BAN MAC address of FR BAN router at the destination MAC. The router terminates LLC traffic from endstations and establishes new LLC connections to the NCP. The router performs conversion between BAN MAC address and NCP boundary node identifier address. FR BAN-2 is supported on 6611 and 2210. FR BAN-2 also supports SNA equipment on remote 6611s or 2210s, using DLSw

In general, FR BAN-1 is preferred as it provides a fast delivery of data with minimal overhead. However, one drawback is that if traffic on a DLCI is high, session timeouts may occur.

With FR BAN-2, session timeouts rarely occur since LLC connections are terminated at the router, and new LLC connections are used in the WAN.

When using FR BAN to connect to a 3174, 3172, 2217 and RXR/2, SNA communication is possible with these nodes, and with token-ring attached endstations.

When using FR BAN to connect to AS/400, SNA communication is possible with the AS/400 itself. DSPUs are not supported.

When using FR BAN to connect to the frame relay token-ring bridge program/DOS, SNA communication is possible with token-ring attached endstations.

When using FR BAN to connect to a 6611 and 2210, SNA communication is supported on SDLC, token-ring, Ethernet, and IP-connected SNA stations. To transport SNA data over IP requires the data link switching (DLSw) function.

3.3 ATM/FR Interworking for APPN HPR Traffic

Frame relay is an ideal *feeder* network for higher speed ATM networks. IBM is currently defining a native ATM DLC for APPN networks. This allows existing APPN applications to gain access to ATM QoS and traffic contracts without changes being made to the applications themselves. Native access to ATM networks will allow existing APPN to use and gain the full benefits of ATM without the use of an enabling protocol, for example, muliprotocol over ATM (MPOA).

Frame relay service interworking can be used between an HPR node with a native ATM DLC and a frame relay connected APPN node only if they have compatible LLC functions. Logical data link control (LDLC) which is the base for ATM, and LLC2 (the base for frame relay) are not compatible since they use different mechanisms to determine when an activation XID exchange is complete, to deactivate a TG, and to monitor link availability.

For service interworking to work for HPR traffic, either optional frame relay LDLC support and the control flows over RTP option set must be implemented by an HPR node on the frame relay side, or optional ATM LLC2 support must be implemented

by the HPR node on the ATM side. For service interworking to work for FID2 traffic, ATM LLC2 support must be implemented by the HPR node on the ATM side. Otherwise, the two nodes perform XID negotiation, and the node that discovers they have no common DLC capability sends an XID with an XID Negotiation Error (X' 22') control vector to reject the link connection.

For more information on the native ATM DLC for HPR refer to the redbooks *Interworking over ATM*, SG24-4699, or *Inside APPN and HPR: The Essential Guide to New SNA*, SG24-3669.

3.3.1 Frame Relay 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.

3.3.2 Frame Relay PVC Sharing

To accomplish 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.

Internal routing of incoming frames is done as follows:

1. Check DCLI number. (This may be an FRFH DLCI.)
2. Is the traffic IP (NLPID=X‘CC’)?
3. Routed frame format (NLPID=X‘08’):
   - L2 field indicates ERP or non-ERP.
   - L3 field indicates INN, BNN or APPN, or HPR.
4. Bridged frame format (NLPID=X‘80’).

INN traffic always goes to the NCP; BNN/APPN traffic is routed by the DSAP. HPR traffic is routed by ANR label, therefore HPR packets can be switched at the adapter level in the 3746 without being routed to software components.

3.3.2.1 Frame Relay Frame Handler Functions

The NCP frame relay frame handler functions (FRFH) 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 NCP FRFH function can switch VCs between ports on the 3745 and ports on the 3746.

The 3746 frame relay frame handler function is controlled and defined from CCM, and can switch frame relay traffic between VCs on 3746 adapters.

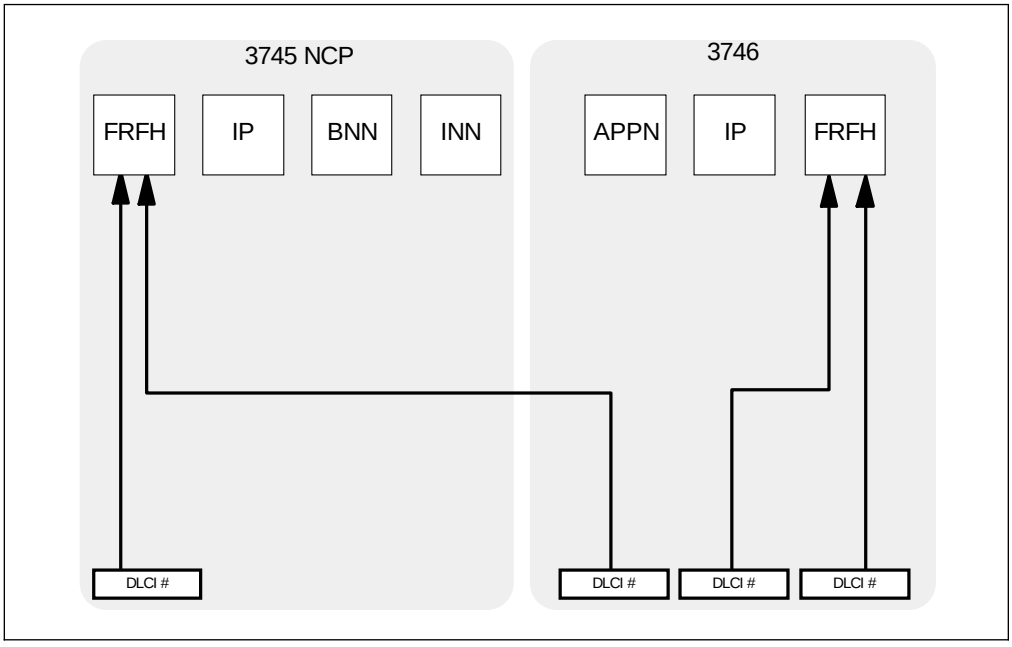

Figure 61. 3745 Frame Handler Function

3.3.2.2 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; 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 62 on page 108 shows the IP support.

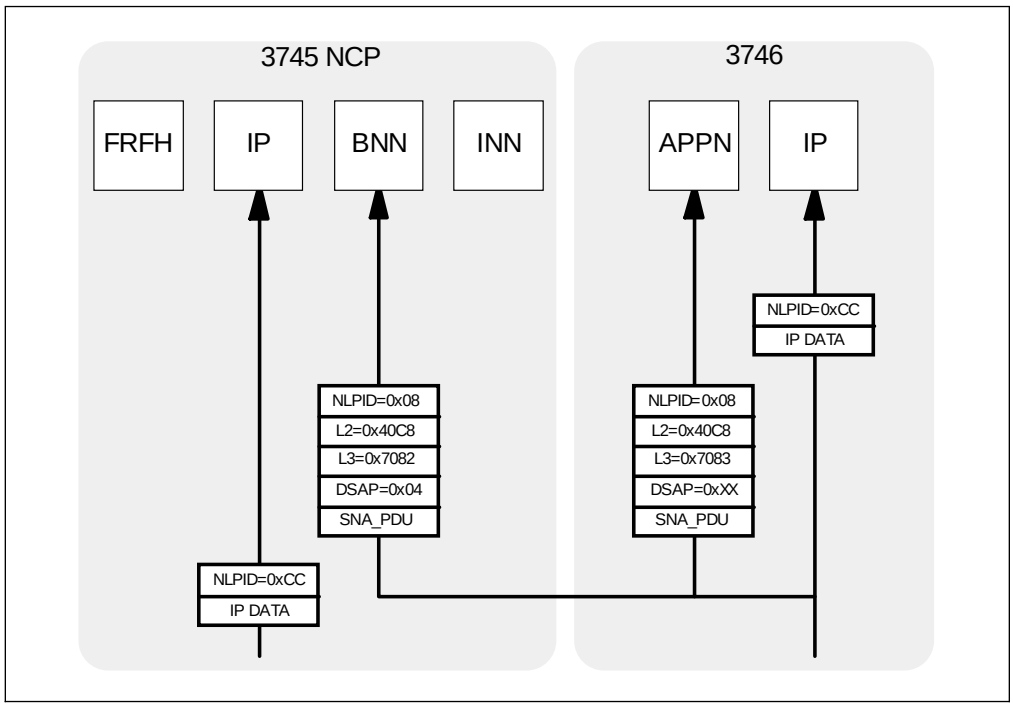

Figure 62. 3745 and 3746 IP over Frame Relay

3.3.2.3 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 63 on page 109). 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. A single INN station per DLCI is supported whether on 3745 or 3746 adapters.

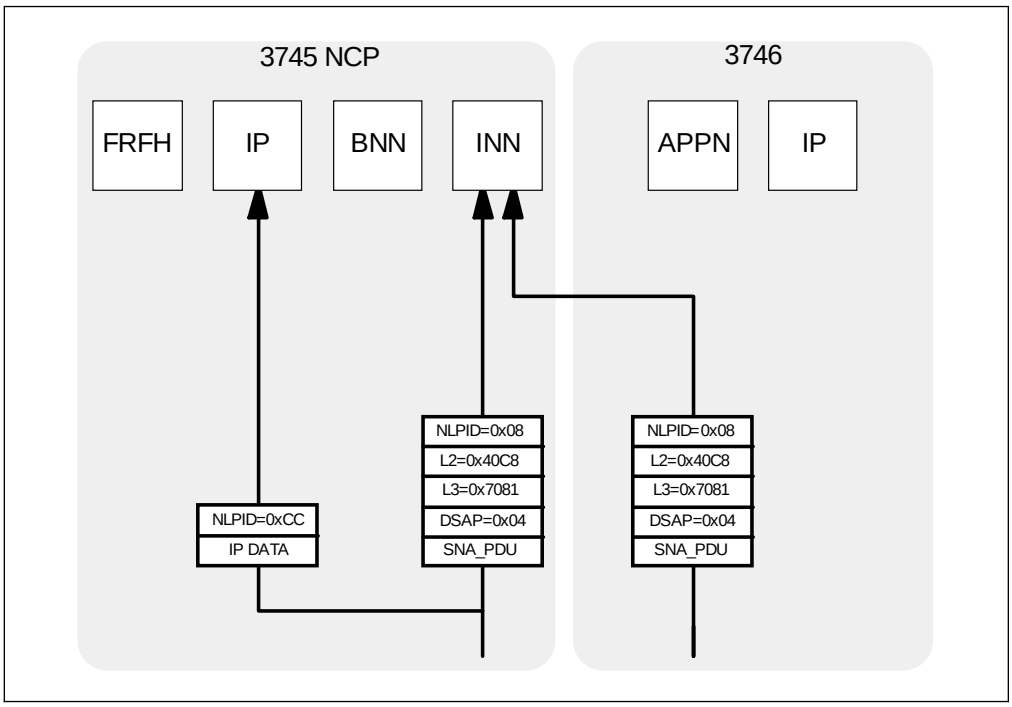

Figure 63. 3745 INN Traffic - Routed Frame Format

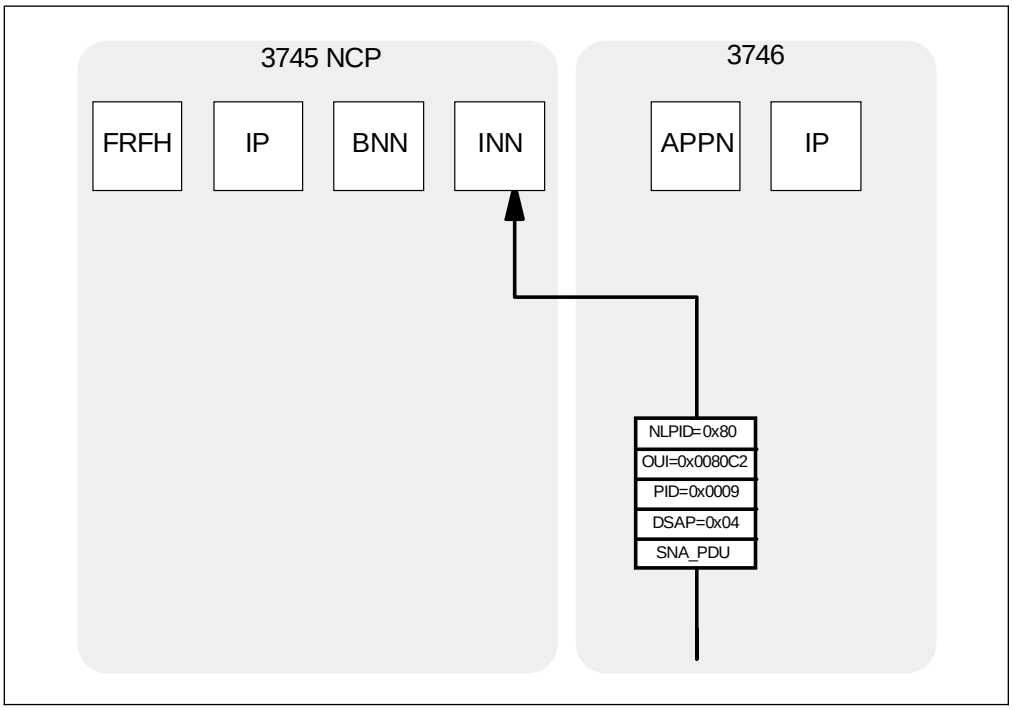

Figure 64. 3745 INN Traffic - Bridged Frame Format

3.3.2.4 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 65 and Figure 66).

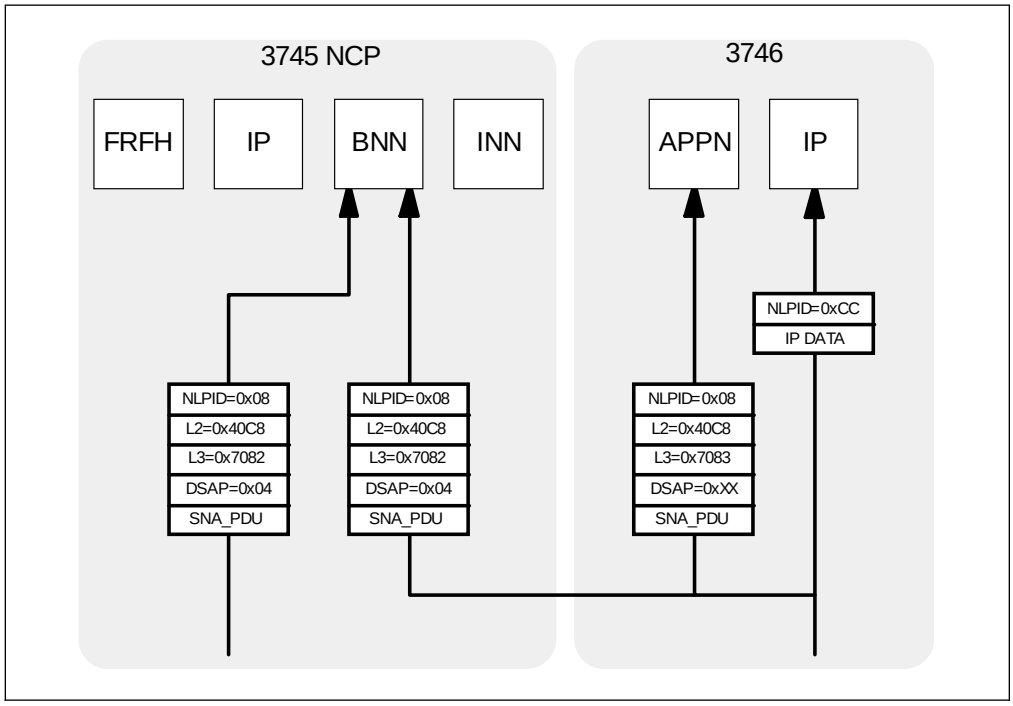

Figure 65. BNN/APPN Traffic - Routed Frame Format

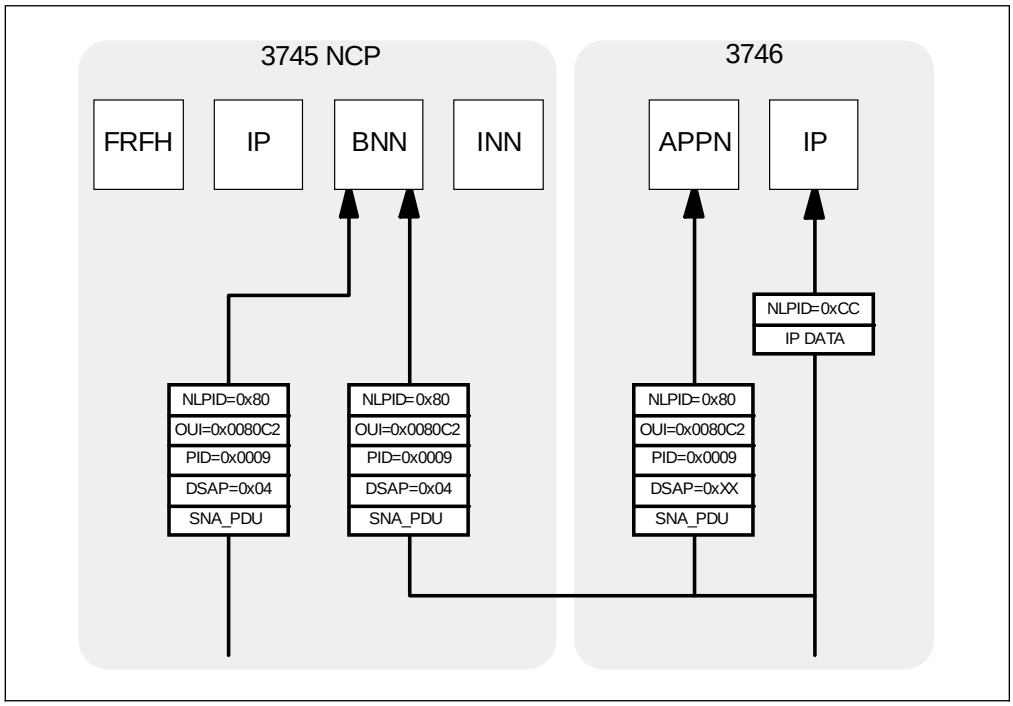

Figure 66. BNN/APPN Traffic - Bridged Frame Format

3.3.3 Frame Relay Scheduling

For frame relay there are two different levels where scheduling occurs. The first is at the DLCI level; the second is at the protocols level. For the DLCI level, we can use COMRATE or CIR to control the transmission of data into the network. At the protocol level, we can use BRS to assign portions of a PVC's bandwidth to different protocols.

3.3.3.1 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 that 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 done via the CCM at 3746 configuration time. Note that CR and CIR can *not* be shared on the same physical line. Also, line sharing is *not* possible with NCP controlled traffic when CCM has defined the link as having a given CIR. The CIR 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. COMRATE is still active under CIR so if the physical pipe throughput is reached, the COMRATE process will create BECNs by using B_c as a COMRATE definition. This means that each DLCI's CIR will be reduced so

that the sum of all CIRs will equal the access rate, and the relationship between the new CIRs will be the same as the relationship between the B_cs.
- 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 T_c 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:

$MIR = 0.25 * B_c / T_c$, or

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,

If $B_c = 0$, then $EIR = B_e / T_c$.

This is different from the 2216 or 2210, which do not define T_c and therefore commit the whole physical bandwidth as EIR.

FRFHs do not implement A-CIR. They set:

$CIR = (B_c + B_e) / T_c$

If the first hop is congested, they set:

$CIR = B_c / T_c$

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 67 on page 113 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.

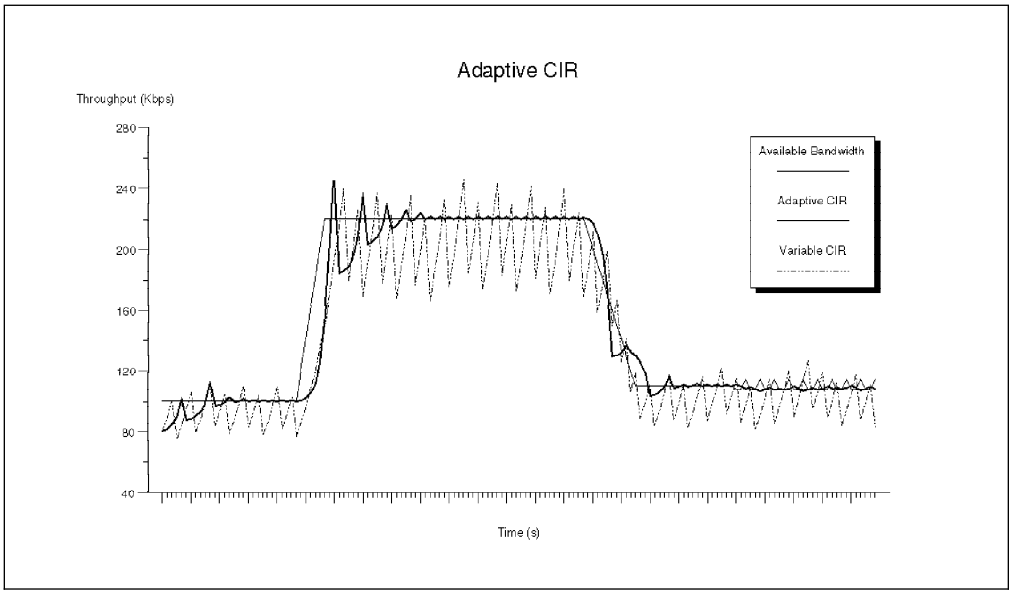

Figure 67. Adaptive-CIR vs Available Bandwidth

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 run-time 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.

3.3.3.2 Bandwidth Reservation System (BRS)

BRS works in addition to the CIR assigned to a DLCI. It allows the user to define how the CIR of a DLCI should be divided between the different protocols that are using that DLCI. BRS is only used when there is more traffic to send than the available bandwidth, and BRS is not used for FRFH DLCIs.

BRS supports three different traffic types:

- SNA (APPN/DLUR, and HPR-ERP) with link-level error recovery

- HPR¬ERP (without link-level error recovery)
- IP

IP traffic can further be differentiated in that five IP sockets can also be defined. The total bandwidth assigned may not exceed 100% of the CIR.

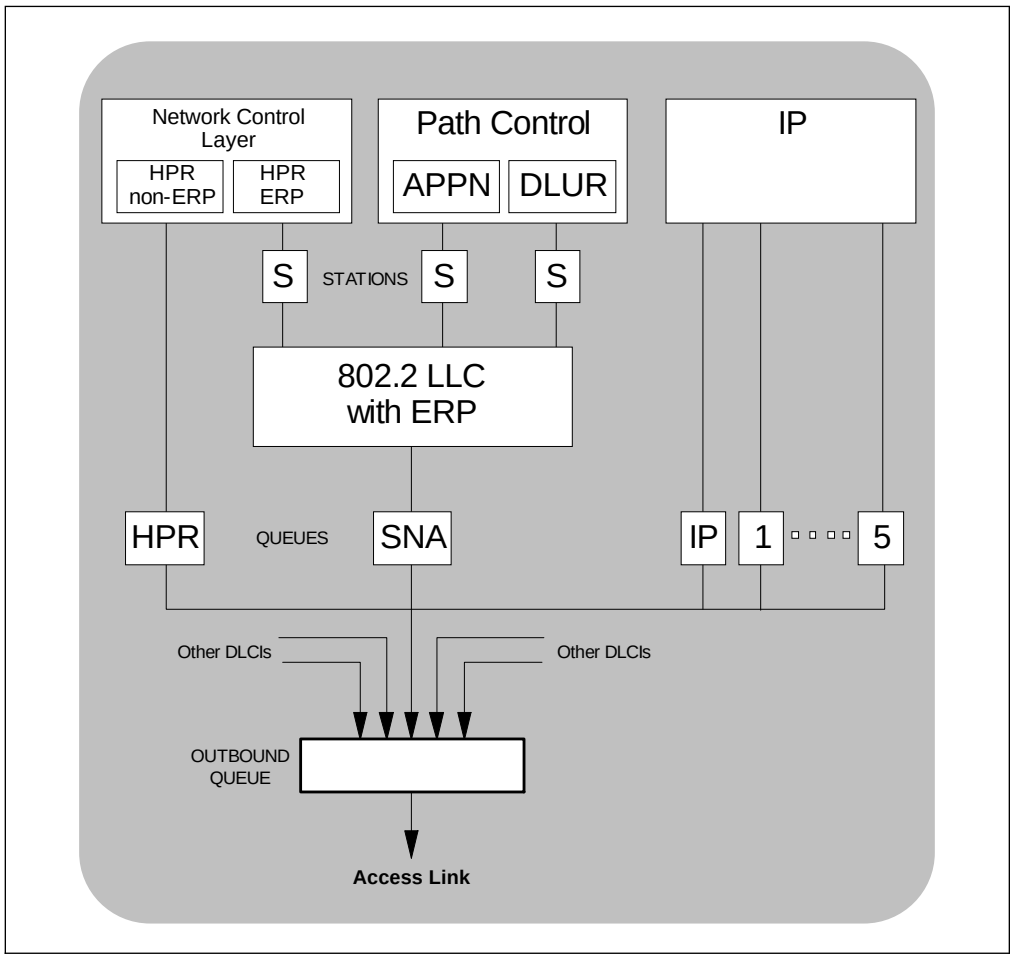

Figure 68. 3746 Frame Relay BRS

When CIR is enabled, T_c, B_c, and B_e can be defined for each DLCI. B_c, or 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 DLCIs:

T_c The measurement interval (T_c) is defined at the DLCI level. The default value is 0.1 second. 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 may be assigned a percentage of those CIR values:

```
BRS
DLUR/APPN/HPR (ERP):      T%
HPR (non-ERP):            U%
IP:                       V%
IP socket #1:             W%
.
.
IP socket #5:             Z%
```

The percentages can be defined between 0% and 100%. A zero value means that no bandwidth is reserved for that protocol. This means that all this traffic will be discarded when congestion occurs. If a protocol has no BRS defined, then that protocol's traffic does not participate in BRS; all that protocol's traffic is transmitted whatever the level of congestion.

3.3.4 Consolidated Link Layer Management (CLLM)

CLLM is an optional FR management function that is not widely supported by the industry but it has been adopted by some frame relay switch manufacturers. CLLM provides some of the same management information provided by LMI, in particular, outage notification. CLLM's main use is to provide asynchronous congestion notification to attaching devices. A single CLLM message may indicate outage or congestion for multiple PVCs. The frame relay protocol supports the following standards for CLLM:

1. ANSI T1.618
2. ITU-T (CCITT) Q.922 Annex A
3. ITU-T (CCITT) X.36 Annex C

On the 3746 from microcode level D46133 onwards, CLLM messages are received and acted on by the 3746.

3.3.4.1 Using CLLM

If your network provider supports CLLM, you can configure 3746 Frame Relay to throttle down its transmit rate for PVCs contained in a CLLM message. CLLM messages contain a cause code that indicates the type and severity of the problem being reported. The device reacts differently depending on the cause code and the CIR configured for each PVC contained in the CLLM message. When the device receives a CLLM message that indicates:

- A short-term condition, and the configured CIR for the PVC is nonzero, the frame relay protocol will throttle the transmit rate for the affected PVCs by the configured information rate (IR) decrement percentage.
- A long-term condition, the frame relay protocol will set the transmit rate for the affected PVCs to the calculated minimum IR.
- Facility or equipment failure or maintenance action, or if the CIR was configured as zero, the FR protocol will continue to transmit any queued data for the affected PVCs but will not accept any more outgoing packets from the upper layer protocols until the congestion condition is cleared.

Once a CLLM message for a PVC has been received, if the device does not receive any CLLM messages or BECNs within the Ty timer period or if a frame

without a BECN is received, the device will consider the congestion condition cleared and gradually return the PVC to its configured transmission rates.

If you are using CLLM to control congestion, you must not configure DLCI 1007 for any other use.

3.3.4.2 Configuring CLLM

CLLM support can be configured from CCM and the CONFIG line command interface.

Part 2. Basic Test Network

Chapter 4. Transport Network

The transport network we used for our test scenarios was built using a wide area frame relay backbone network to connect the 3746 to the 2216, ESCON to connect the 3746 to the S/390 host system, PPP to connect the 2216 and 2210, and token-ring as LAN transport. This section describes the hardware configuration and the definitions process to configure the interfaces used.

The FR network we used in our test scenarios is provided by an IBM 2220 Nways broadband switch. For our scenarios the frame relay backbone provided looks the same as the attached equipment as a frame relay network provided by a frame relay service provider. For attached equipment it makes no difference whether we use a single 2220 or a network of 2220s. As all our test equipment was installed in a single test room, we connected the 3746 and 2216 to the IBM 2220 by direct cables instead of using modem connections.

This chapter shows the how to configure all frame relay connected hardware, 2220, 3746, and 2216. In the case of the 3746 and 2216, IP addressing is directly coupled to the frame relay connected ports, so we have also included the screens that show how to define the local IP address associated with the frame relay port. The detailed discussion of IP addressing is contained in Chapter 5, “TCP/IP Test Scenario” on page 143.

4.1 Frame Relay Interfaces

In our test scenarios we use V.35 interfaces to connect the 3746-900 and the router to the frame relay network. At the 3745-900 side, an LIC12 at T1 speed (1.5 Mbps) serves as the FR access link. At the router site, a V.35 serial link is connected to the FR network at a speed of 512 kbps. This emulates a real life scenario where a single high-speed access link from the 3746 into the frame relay backbone connects to peripheral nodes on the frame relay backbone that are connected with lower speed access links.

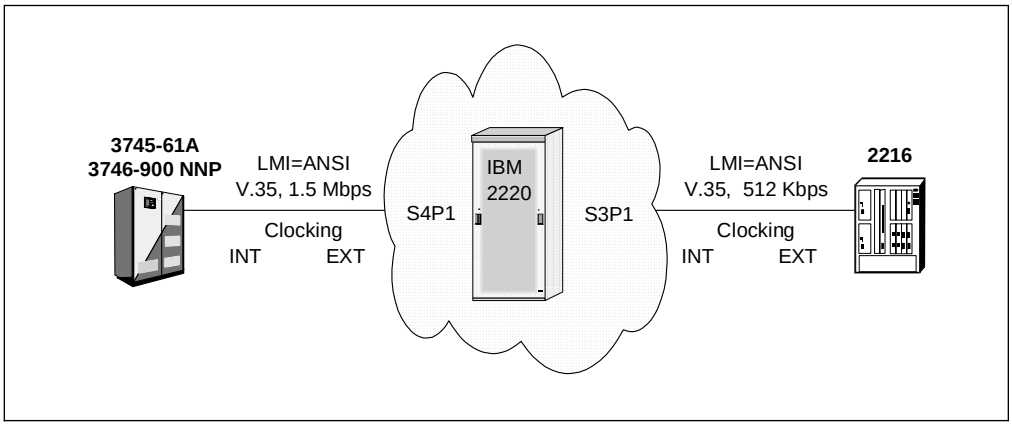

Figure 69. The Frame Relay Network

4.2 Frame Relay DLCIs

The frame relay backbone network provides three permanent virtual circuits (PVCs) and each PVC is identified by its data link connection identifier (DLCI). The frame relay backbone supplies DLCIs 21, 22, and 23 at the peripheral side, corresponding to DLCIs 31, 32, and 33 at the 3746 side (see Figure 70 on page 121).

The PVCs are defined with the following characteristics:

Table 16. Frame Relay Bandwidth Adaptation Definitions

DLCI #	Parameters	DLCI #	Parameters
31	ACC=1.5 Mbps	21	ACC=512 kbps
	CIR=16 kbps		CIR=16 kbps
	BWA=1 kbps-1.5 Mbps		BWA=1 kbps-512 kbps
32	ACC=1.5 Mbps	22	ACC=512 kbps
	CIR=16 kbps		CIR=16 kbps
	BWA=1 kbps-1.5 Mbps		BWA=1 kbps-512 kbps
33	ACC=1.5 Mbps	23	ACC=512 kbps
	CIR=16 kbps		CIR=16 kbps
	BWA=1 kbps-1.5 Mbps		BWA=1 kbps-512 kbps

Bandwidth adaptation (BWA) is a unique feature of the IBM 2220 Nways broadband switch; it is a type of dynamic CIR. The 2220 continually measures the traffic on a particular PVC. If the traffic offered exceeds the current traffic contract, the bandwidth for that PVC is increased dynamically within the specified limits. Likewise, if less traffic is offered, the bandwidth is decreased. If a PVC is defined with **bandwidth adaptation=yes**, the CIR value has significance only when the PVC is set up. Then, the CIR is increased or decreased according to the offered traffic and of course according to the bandwidth available on the trunks. See *Installing and Implementing 2220 Networks*, SG24-2589 for details.

The 2220 can be configured using either the *network management station* (NEM) or the stand-alone configuration tool called *node configuration tool* (NCT). The screens are identical for both methods. The NCT is available for the AIX, OS/2 and Windows platform.

4.3 2220 Frame Relay Configuration

We must perform the following steps to configure the 2220:

1. Define the frame relay ports on the line interface couplers (LIC)
2. Define LMI options on each frame relay port
3. Define potential connections (forward path DLCI-DLCI of the PVC)
4. Define the traffic parameters on the PVC
5. Define virtual connections (return path DLCI-DLCI of the PVC)

Figure 70 on page 121 shows how the the PVCs are defined through the network.

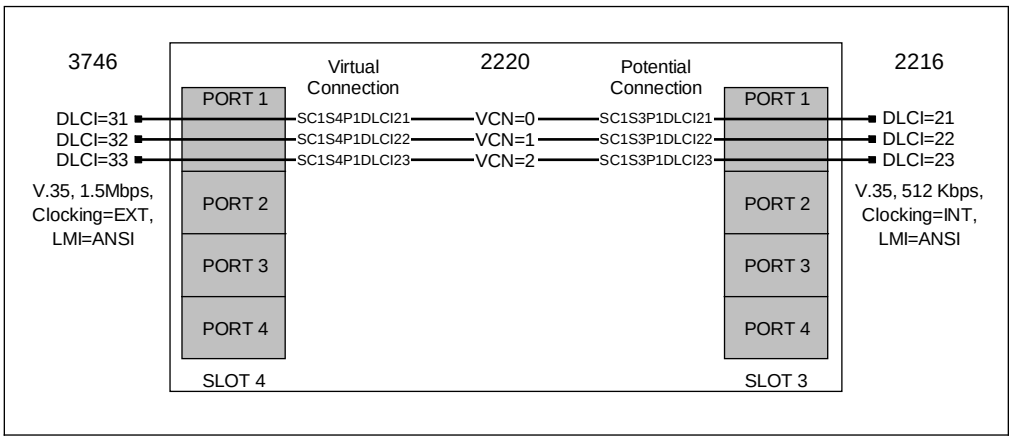

Figure 70. Frame Relay DLCIs

The following screens were captured from the 2220 configuration tool and show the details of the frame relay configuration process.

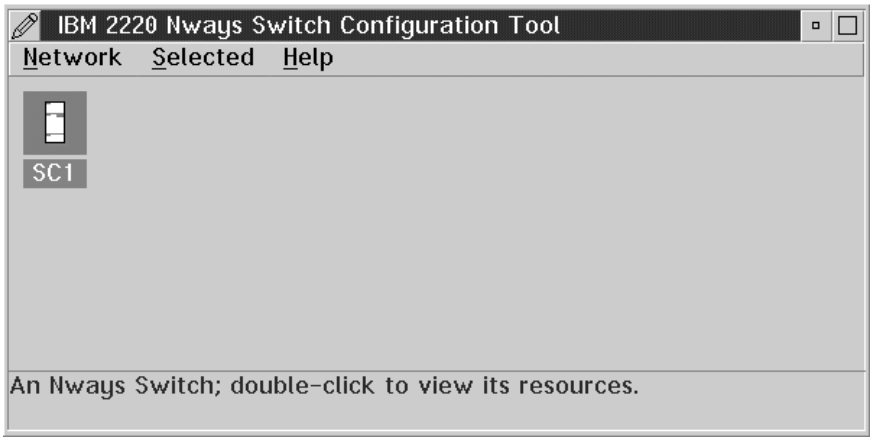

Figure 71. 2220 NCT Main Screen

When the NCT is started, the screen shown in Figure 71 is displayed. This screen shows all 2220 nodes currently configured at this NCT. At this time we have configured just one node, node SC1. To open this node's configuration, double-click the icon. This will display a screen showing the resources that may be configured for this node.

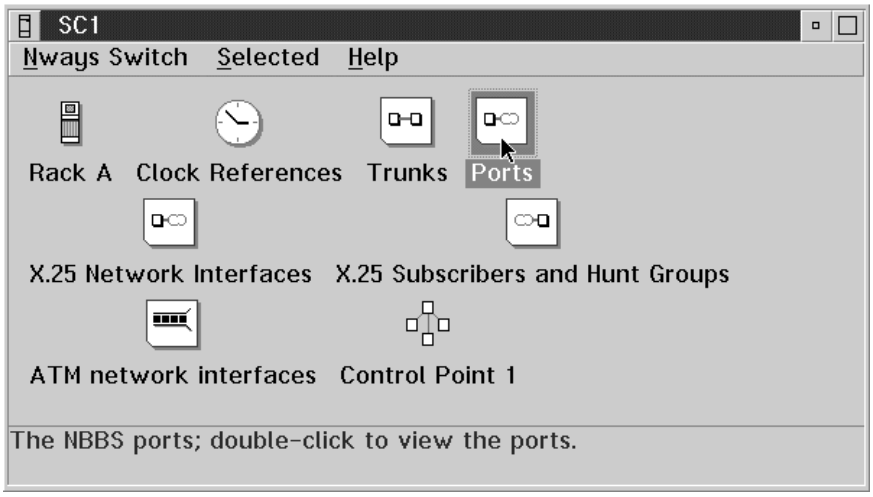

Figure 72. Resources to Configure

Figure 72 shows the icons of all resources on the node SC1. In order to configure PVCs, we need to define the FR logical ports first. Click on the port icon to list all logical ports defined.

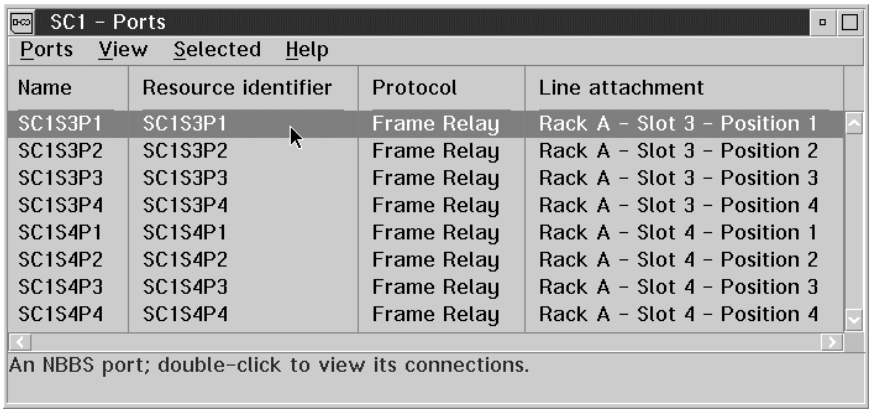

Figure 73. Logical Ports List

In slots 3 and slot 4 of the 2220, we have V.35 interface cards, each interface card supports four high-speed V.35 interfaces. We have defined four logical ports on each card and named them S3P1 to S3P4 (slot 3) and S4P1 to S4P4 (slot4). The following screens show how to define these logical ports. Click on **Ports** and then **Create** from the screen shown in Figure 73. The logical port configuration screen will be displayed.

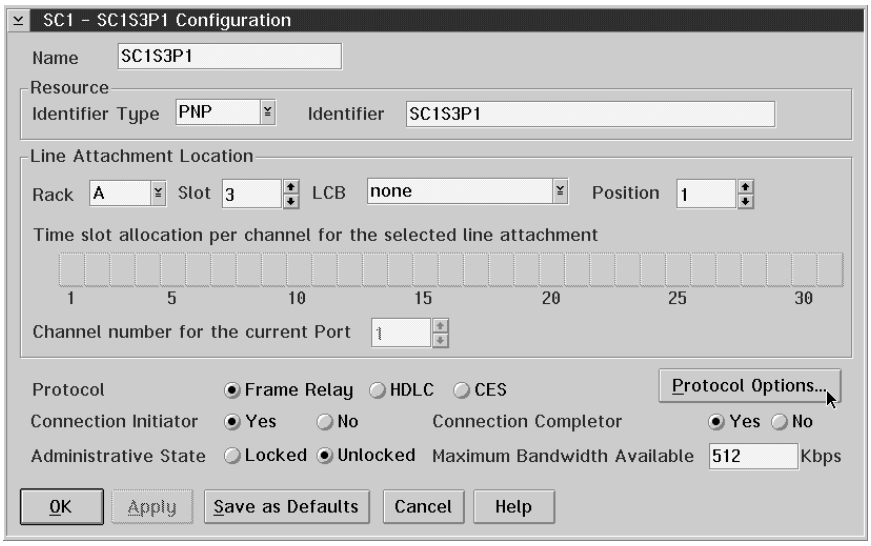

Figure 74. Logical Port Screen

Here we give this port a name (SC1S3P1) and an identifier (SC1S3P1). They can be identical. Then we need to specify to which hardware location this logical port is assigned. It is rack A, slot 3, position 1. If the hardware interface is a channelized E1 or T1, we could assign one or multiple channels to a logical port. The next section down the screen is the protocol. Here we select **Frame Relay**. The field Maximum Bandwidth Available is filled in automatically according to the physical port speed. Leave all other fields at their default. Then to specify the LMI options, click the **Protocol Options...** button.

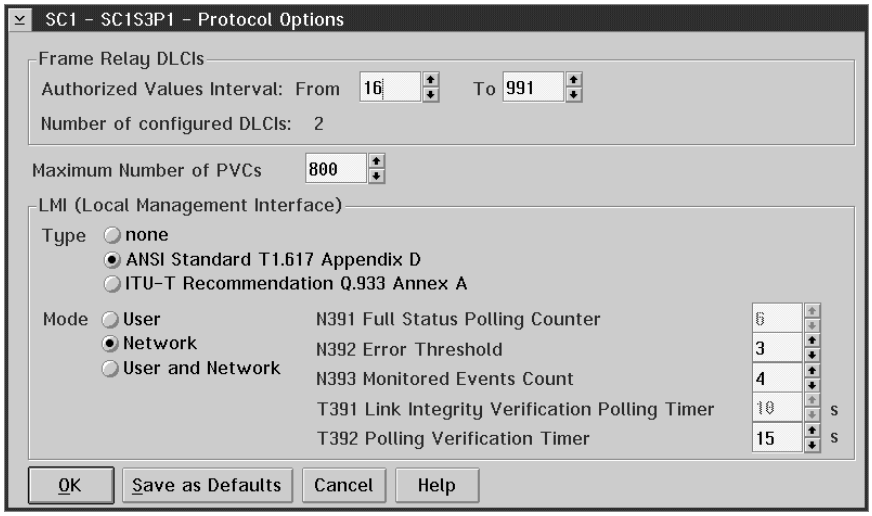

Figure 75. Protocol Options

At the top part of the screen, the range of the DLCIs and the maximum number of PVCs for this port can be specified. The LMI type has to match the definition of the attached frame relay equipment, that is the 3746 and the 2216. To see the 3746 LMI definitions see Figure 84 on page 129; to see the 2216 LMI definitions see Figure 96 on page 136. We selected LMI type **ANSI**, and LMI mode **Network**. Click **OK** to save the data and go back to the previous screen. Click **OK** again, this will get you back to the Logical ports list screen.

Now we need to configure the PVCs on the logical ports we previously defined. To do this, select a port in the list and double-click it. We selected **SC1S3P1**. This displays all connections on that logical port.

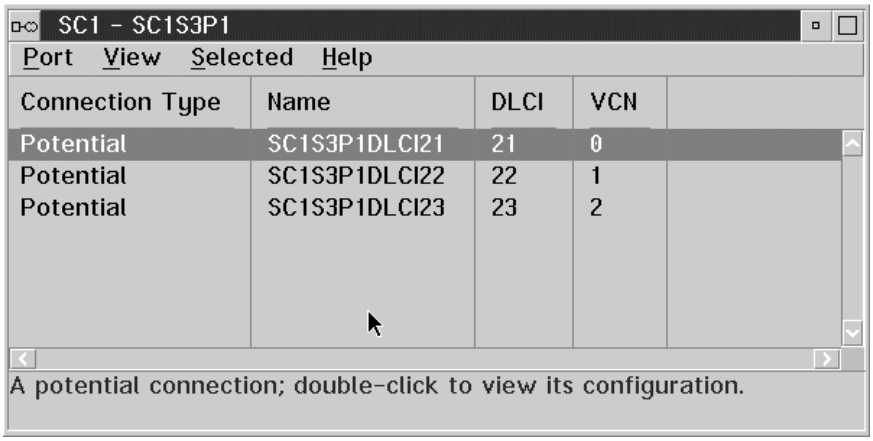

Figure 76. SC1S3P1 Potential Connections

In the 2220, the forward and the return path of a PVC must be specified individually and the paths may be given different names. The forward path is called a *potential connection*, the return path is called a *virtual connection*. The potential connection originates at the source port; the virtual connection originates at the destination port. The potential and virtual connections are identified as belonging together by the *virtual circuit number* (VCN), which is a unique identifier used throughout the network.

To open a potential connection, just double-click the appropriate line on the screen. To create a potential connection, click **Port**, then **Create Potential Connection**.

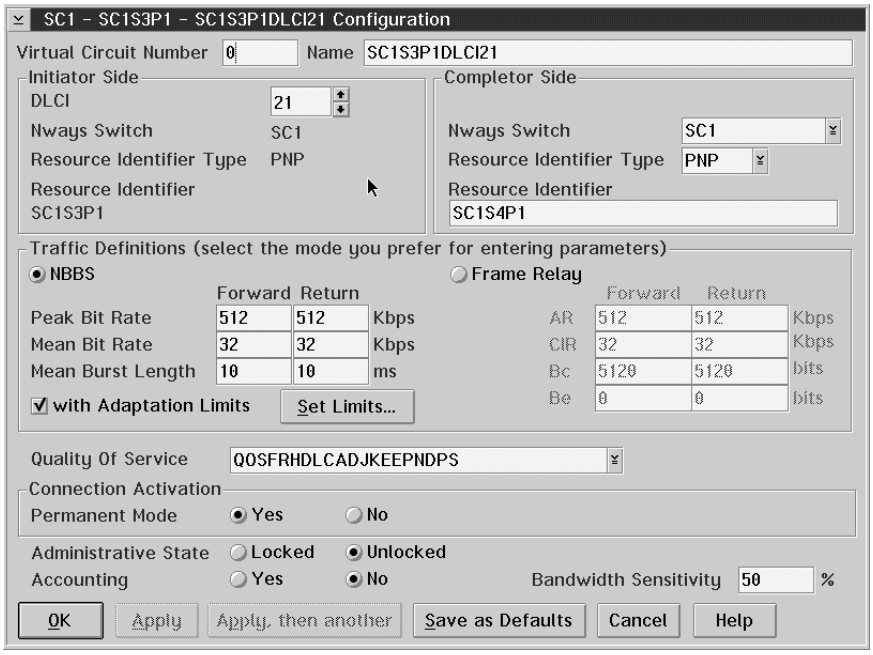

Figure 77. Potential Connection

This is the definition screen for a potential connection. All definitions for a PVC, except the DLCI at the destination side, are done at this screen. As shown in Figure 70 on page 121 the DLCI used at the originating side of the PVC can be different to the DLCI used at the destination side of the PVC.

Virtual Circuit Number
VCN is an internal number and assigned automatically. It cannot be changed.

Name
The name of the PVC.

Initiator Side
Specifies the DLCI number used at the originating side of the PVC.

Completer Side
Specify the Nways Switch name and the Resource Identifier (hardware port) at the destination side of the PVC.

Traffic Definitions
Here we can specify the frame relay PVC parameters in native frame relay mode as ACC, CIR, Bc and Be values, or in NBBS mode as peak, mean and burst length. The NBBS mode is an added value to the frame relay mode as it allows the user to run a PVC with dynamic bandwidth adaptation. Bandwidth adaptation can be configured to adapt within certain pre-specified limits. To activate these limits, click the button **with Adaptation limits**.

Quality of Service
Quality of service (QoS) is a term used in ATM networks. A QoS table specifies a connection in terms of end-to-end delay, cell delay variation, cell loss priority, etc. The 2220 offers the same mechanisms, not only for ATM, but for all protocols. Here we can choose a QoS table for each connection. There are various predefined tables for each protocol. There is also a possibility to add your own tables to the list. For details please refer to *IBM 2220 Nways Switch Configuration and Implementation Guide*, SG24-2589.

Connection Activation
Select **Yes** for permanent mode.

Administrative State
Select **Unlocked**.

Accounting Select
Select **No**.

Bandwidth Sensitivity
When the bandwidth used changes by this amount, then accounting information will be recorded.

Click the button **Set Limits..** to define the bandwidth adaptation limits. Figure 78 on page 126 shows the adaptation limit values set by the configuration tool.

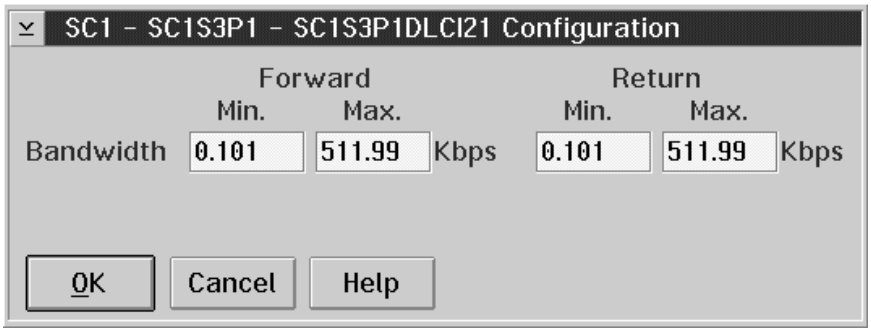

Figure 78. Bandwidth Adaptation Limits

We changed the values shown to an adaptation range from 1 kbps up to the access rate for the forward and the return paths. Click **Ok**.

We have now configured all PVC parameters except the DLCI used at the destination port. To do that, go back to the logical port list screen and click on the destination port, **SC1S4P1**.

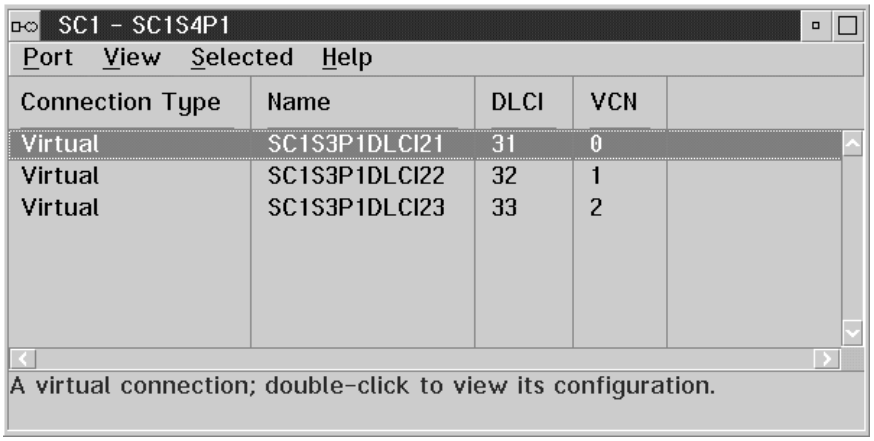

Figure 79. SC1S4P1 Virtual Connections

The screen shows the virtual connections that were created automatically by the configuration tool. The only parameter that must be defined manually is the DLCI number. Double-click the line to open the configuration screen for a virtual connection.

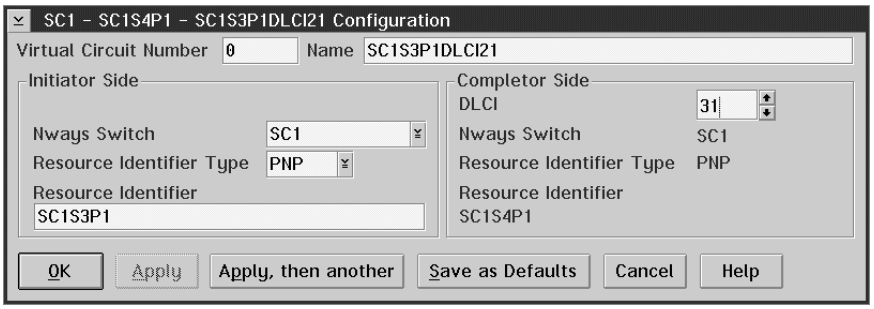

Figure 80. SC1S4P1 Virtual Connection configuration

Enter the DLCI number at the Completer Side. Click **OK** to save this data and repeat this step for the other virtual connections.

Close all open windows. This concludes the configuration of the frame relay PVCs. If you used the offline configuration tool (NCT) to do this configuration, you need to export this configuration to a file and then import this file in the IBM 2220 Nways switch and activate it. If you used the Network Management station to do this configuration, then everything is dynamic and the PVCs just defined will already be active and ready to use.

4.4 3746 Frame Relay Configuration

Figure 81 shows the 3746 CCM ports configuration screen. Each port represents a 3746 coupler. The three kinds of couplers listed below have been configured. We will be using port 2336 for the 3746 to 2220 link.

- LIC12 (port address 2336): interface for the frame relay network
- ESCC (port address 2176): interface for the S/390 ESCON channel
- TIC3 (port address 2144): interface for the token-ring network

Although the LIC11 (port address 2304) is checked in this figure, this port is not referred to by our test scenarios.

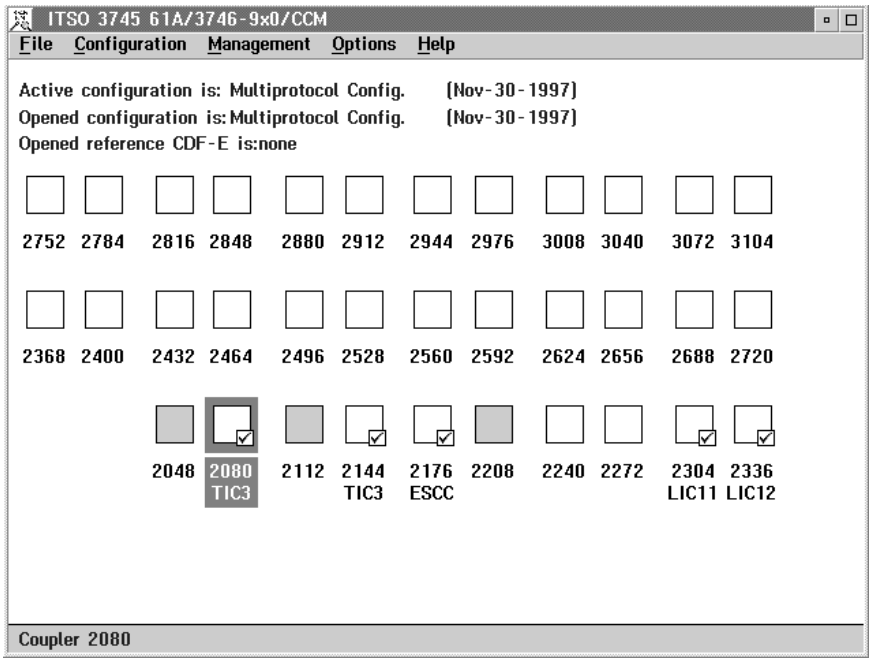

Figure 81. 3746 Ports Configuration

Figure 82 on page 128 shows the basic LIC12 port configuration dialog. This dialog can be reached by selecting the icon indicating **2336/LIC12** in Figure 81. We specified APPN and IP networking over a single frame relay port for our multiprotocol networking environment. Both APPN and IP must be selected on this screen to allow us to make APPN and IP definitions for this port later on in the configuration process.

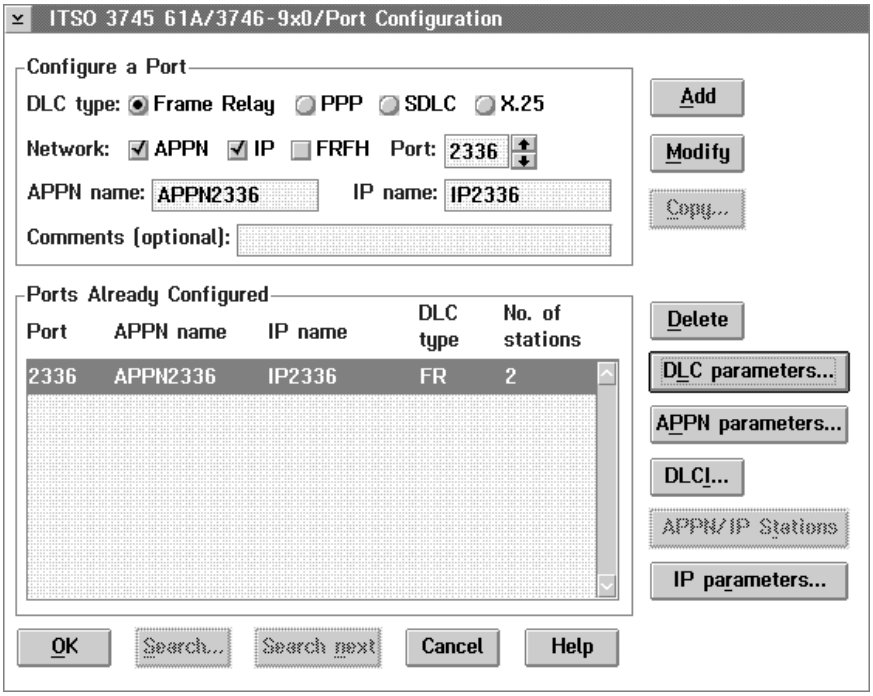

Figure 82. LIC12 Port Configuration

Figure 83 on page 129 shows the data link control parameters defined for the frame relay port. This dialog is reached by selecting the button **DLC parameters...** in Figure 82.

Communications rate (COMRATE) or committed information rate (CIR) must be selected as the bandwidth management method for this port. CIR must be selected if you wish to use the *bandwidth reservation system* (BRS). If the CIR is selected, you will later need to specify CIR parameters and optionally BRS parameters for each DLCI defined on this port. The data block size (DATABLK) parameter in this figure defines the minimum bandwidth that is allocated to this port when there is more data to send than the access rate allows. DATABLK * 8 is the value in bits that is entered in the COMRATE spin button (see Figure 89 on page 132).

The parameter boundary node identifier (BNI) defines the local virtual MAC address of this port. It is used by remote boundary access node (BAN) routers as the destination MAC address when making peripheral connections to this port. This value is identical for all the DLCIs present on this physical line. The default is '4FFF00000000'. The default value for the 3746 local SAP is 8. This is the LSAP value used in frames transmitted by the 3746 on this port.

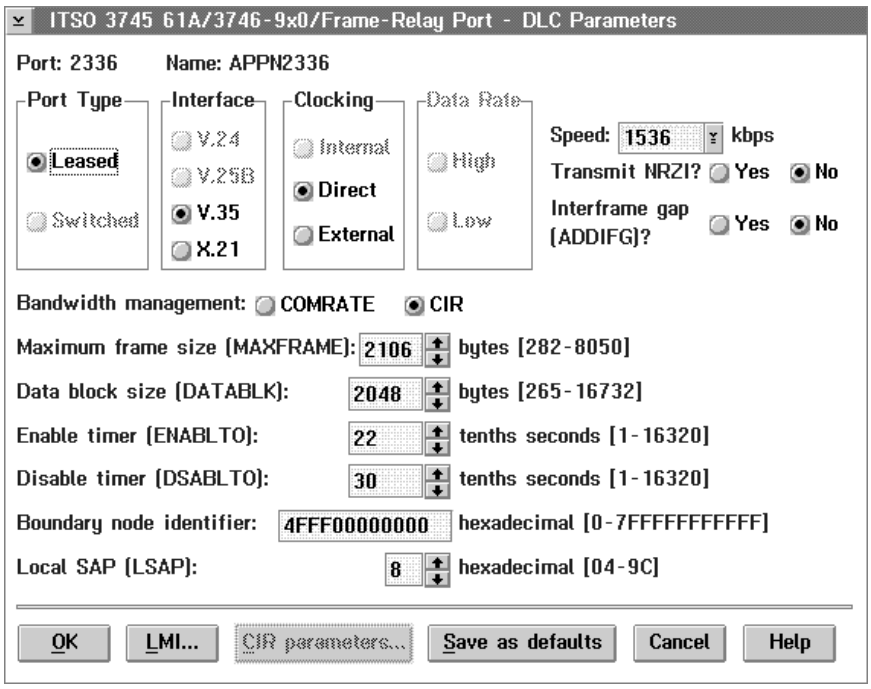

Figure 83. DLC Parameters for the Frame Relay Port

Figure 84 shows the Local Management Interface (LMI) parameters for the LIC12 port. This dialog is reached by selecting the **LMI** button in Figure 83. This parameter must match the LMI category supported by the frame relay network, in our case ANSI.

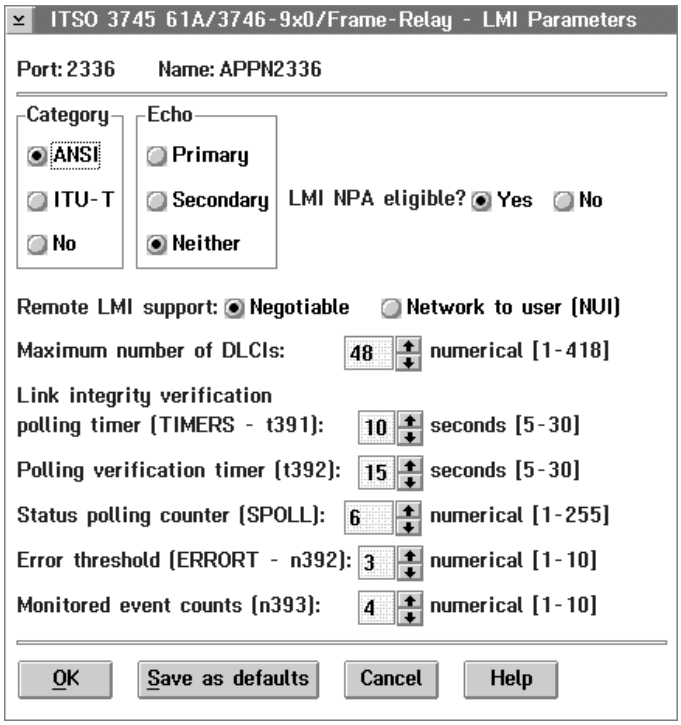

Figure 84. LMI Parameters for the Frame Relay Port

Figure 85 on page 130 shows IP parameters for the LIC12 port. This dialog is reached by selecting the button **IP parameters...** in Figure 82 on page 128. In

our test scenarios there is one IP subnet in the frame relay network. On the 3746 we can define multiple local IP addresses on a single frame relay port, but all DLCIs on this port use the same local IP address. Other specifications are set to the defaults except we have automatic reactivation enabled.

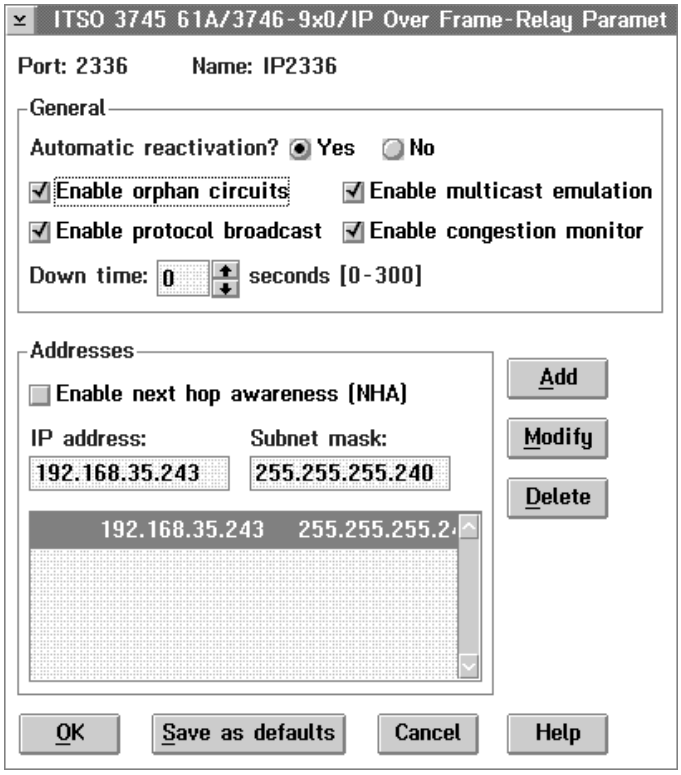

Figure 85. IP Parameters for the Frame Relay Port

On port 2336, two DLCIs are defined at this time. Figure 86 on page 131 shows the IP parameters for DLCI 33. This dialog is reached by selecting the button **DLCI...** in Figure 82 on page 128. The DLCI number must be between 16 and 991 and is unique on this port. The default is 32.

The remote IP address must also be specified if the check boxes Enable multicast emulation and Enable protocol broadcast in Figure 85 have not been selected. This is because the inverse ARP cannot be used to determine the remote IP address on each DLCI when those functions are disabled. Therefore the remote IP address must be predefined.

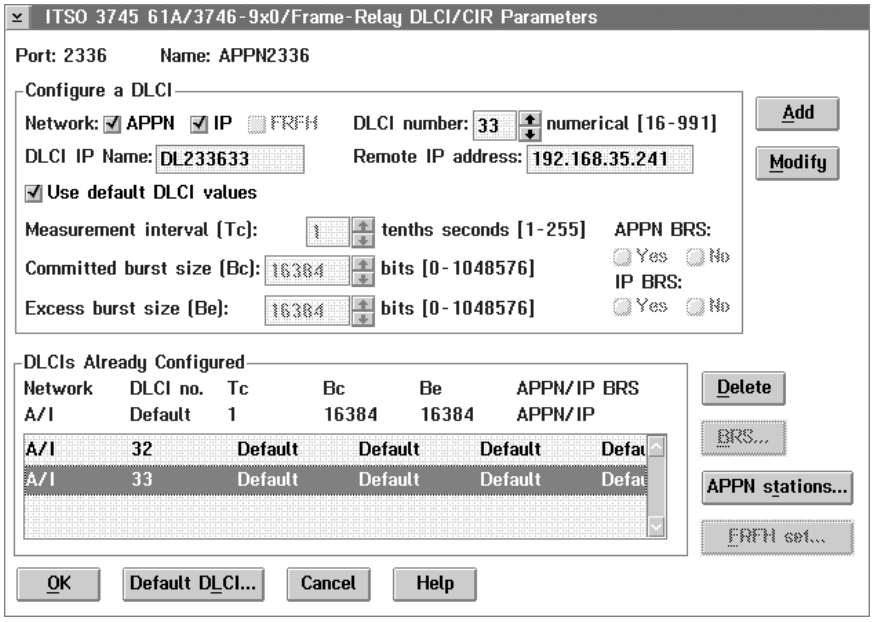

Figure 86. IP Parameters for the Frame Relay DLCI 33

Depending on whether COMRATE or CIR was specified on the port definition, you can set bandwidth management parameters for each DLCI. Figure 86 shows that the Tc, Bc, and Be parameters can not be defined. This is because COMRATE was selected for this port. Figure 87 and Figure 88 on page 132 show CIR and BRS defaults for the DLCI. The dialog in Figure 87 is reached by selecting the button **Default DLCI...** in Figure 86. The dialog in Figure 88 on page 132 is reached by selecting the button **BRS...** in Figure 87.

As Figure 87 shows, we can define default values of Tc, Bc and Be and we can choose whether Bandwidth Reservation System (BRS) is used or not used for the DLCI. The CIR value is calculated by Bc/Tc. The default value for Tc is 1/10th of a second and the default value for Bc is (8 * DATABLK in bits), the DATABLK value comes from Figure 83 on page 129.

Figure 88 on page 132 shows how the CIR is divided between the protocols running over the DLCI by the BRS function. In this case, default values are shown. The default prioritizes APPN/HPR ERP, HPR non-ERP and IP data flows almost equally. We can define up to five IP sockets in addition to the above three protocol types.

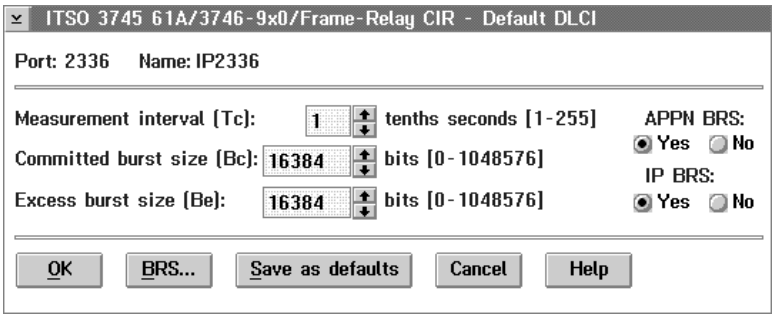

Figure 87. CIR Defaults for the DLCI

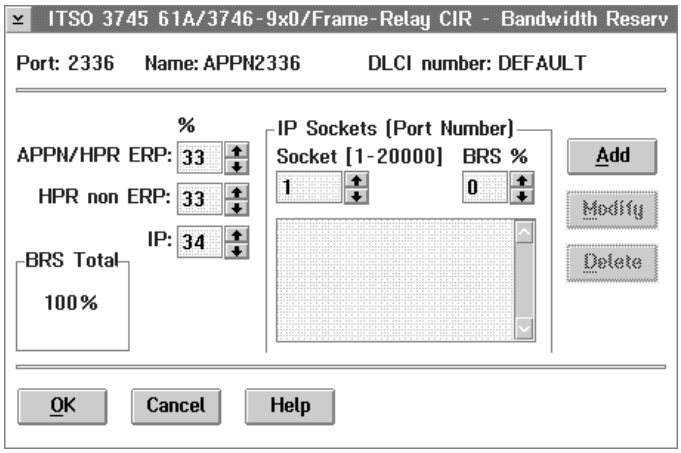

Figure 88. BRS Defaults for the DLCI

Figure 89 shows the COMRATE defaults for the DLCI. This dialog is reached by selecting the button **Default DLCI...** in Figure 86 on page 131 (if COMRATE was selected as the bandwidth management method on the port definitions). In this dialog we can define default values for COMRATE for APPN and IP.

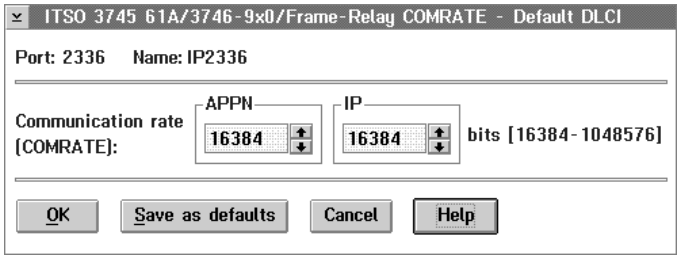

Figure 89. COMRATE Defaults for the DLCI

4.5 2216 and 2210 Frame Relay Definitions

This section deals with the hardware and interface configuration of the IBM 2216 multiaccess connector.

The easiest way to configure the IBM 2216 and the IBM 2210 is to use their configuration tools. Although the two tools are separate programs, the look and feel of both versions is exactly the same. For an in-depth description of the IBM 2210 and IBM 2216 please refer to the redbook *IBM2210 Nways Multiprotocol Router Description and Configuration Scenarios*, SG24-4446.

After starting the 2216 configuration tool, the following navigation window pops up.

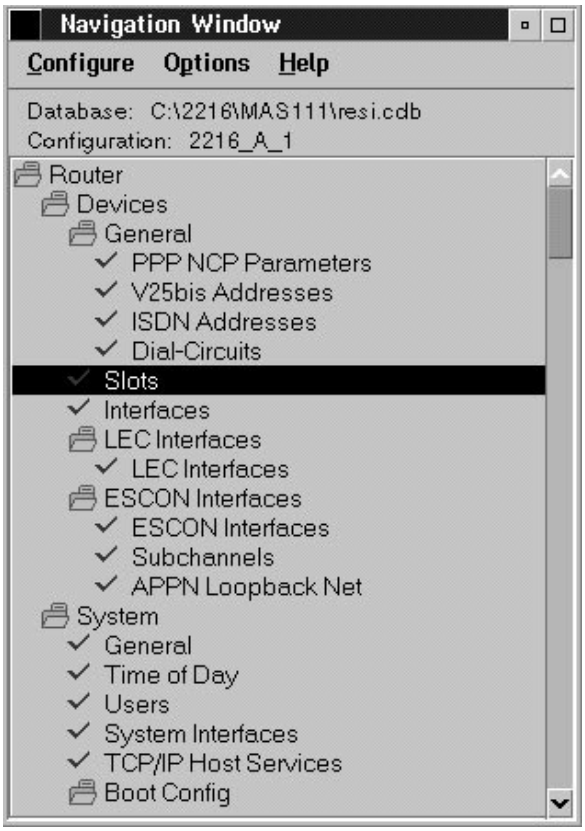

Figure 90. The Navigation Window

The navigation window is designed to make it as easy as possible to configure the router from top to bottom. The first thing to do is to configure the slots. Click under Router/Devices/General the item **Slots**, This opens the Slot Browser window.

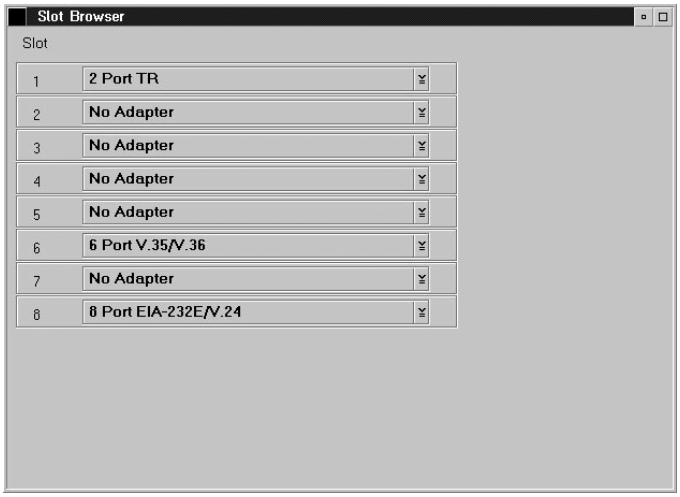

Figure 91. Slot Browser

In the Slot Browser window (see Figure 91) the card type in each of the 8 feature slot can be selected. In our configuration, slots 1, 6, and 8 contain token-ring, V.35, and and V.24 cards. New or changed cards may be specified here. Now go back to the navigation window and click on the menu item **Interfaces**. The Device Interfaces window appears (see Figure 92 on page 134).

Here each interface is configured. Click on the **Configure** button for the interface 0 (TR Slot 1 port 1).

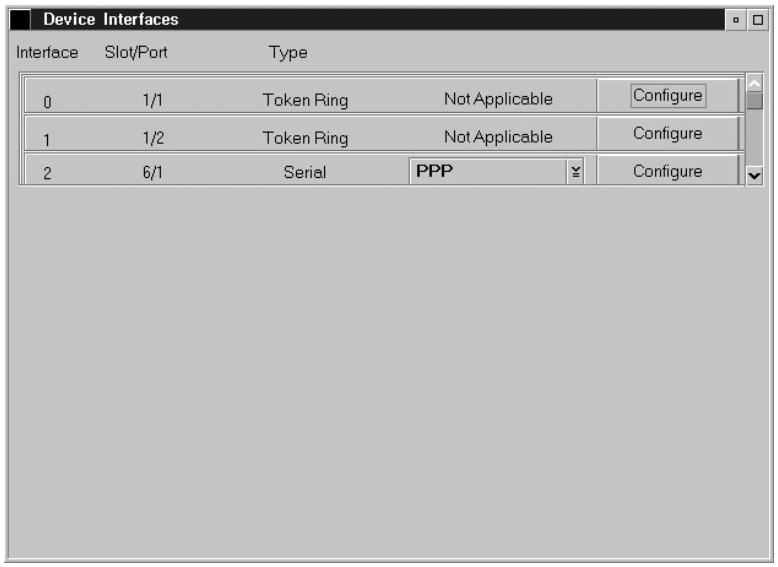

Figure 92. Device Interfaces

In the Device Interfaces window for the token-ring port (see Figure 93), we configured the parameters as shown. That concludes the token-ring interface configuration.

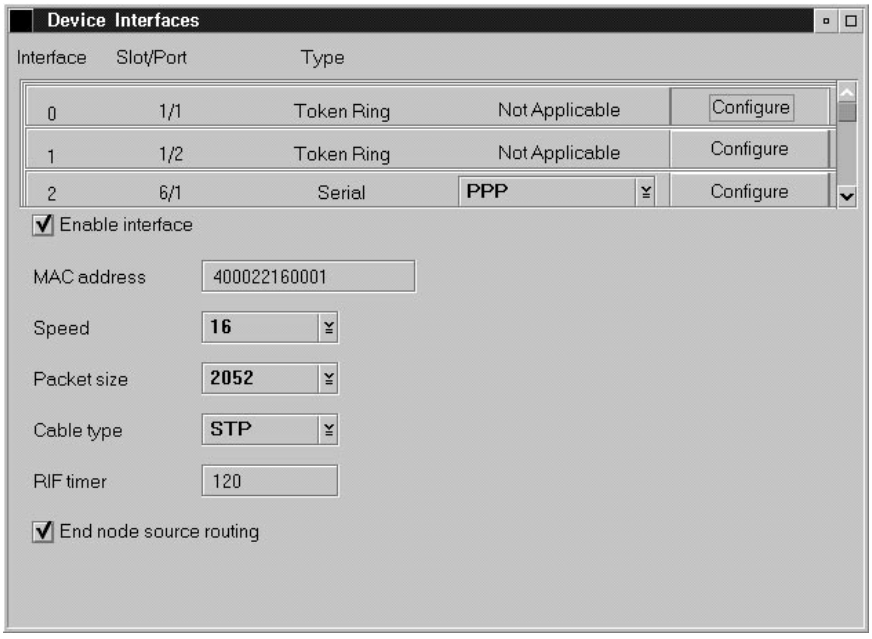

Figure 93. Device Interfaces: TR

Next the PPP link to the 2210 will be configured (see Figure 94 on page 135). This is on interface 2 (Slot 6 port 1). Click on the **Configure** button for this interface.

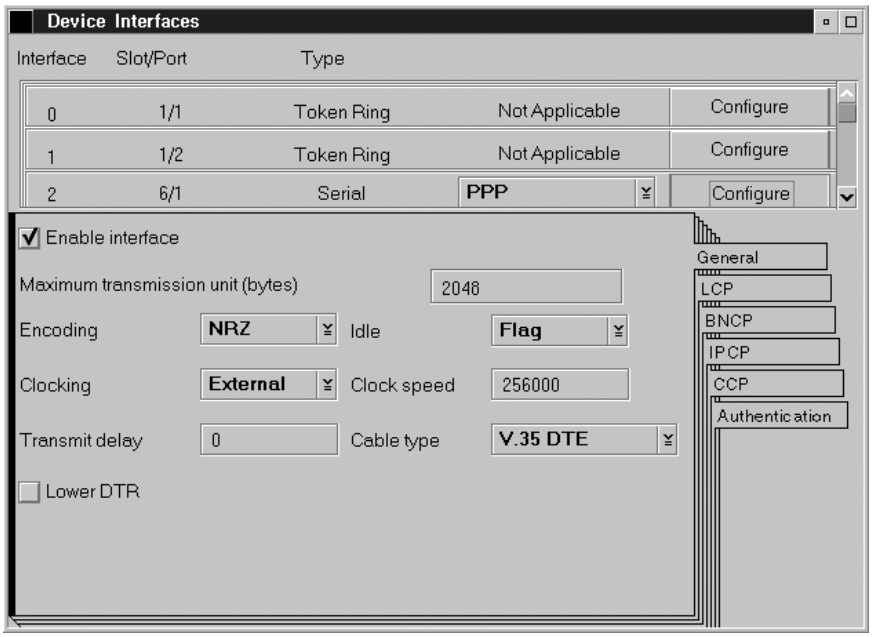

Figure 94. Device Interfaces: Serial PPP

On the General page (see Figure 94), enable the interface and fill in the other parameters as shown. The other pages can be left at their defaults.

Interface 3 (slot 6 port 2) is connected to the frame relay network. We changed the PPP protocol to frame relay, then clicked on the **Configure** button for this interface.

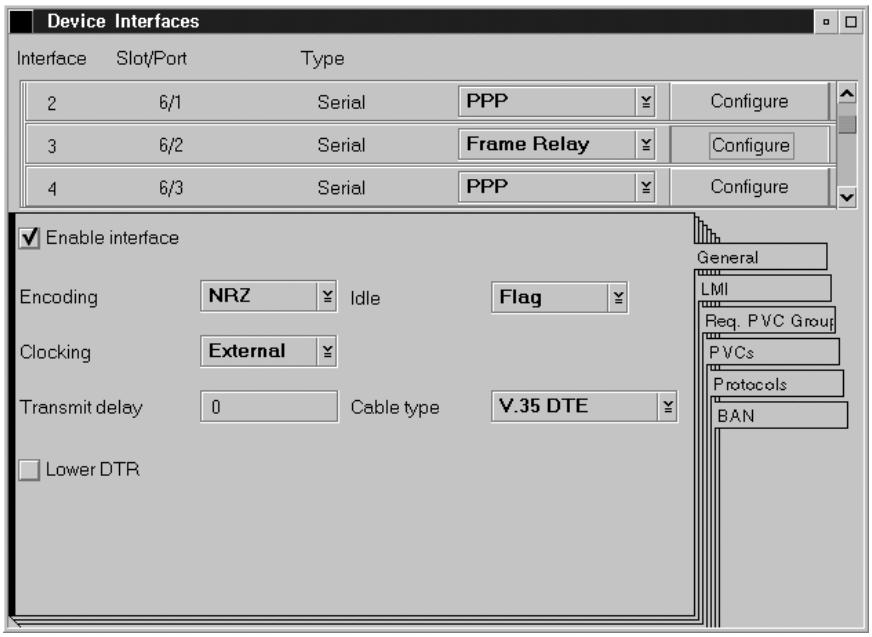

Figure 95. Device Interfaces: Frame Relay General

On the General page (see Figure 95), enable the interface and fill in the other parameters as shown. Then select the **LMI** page.

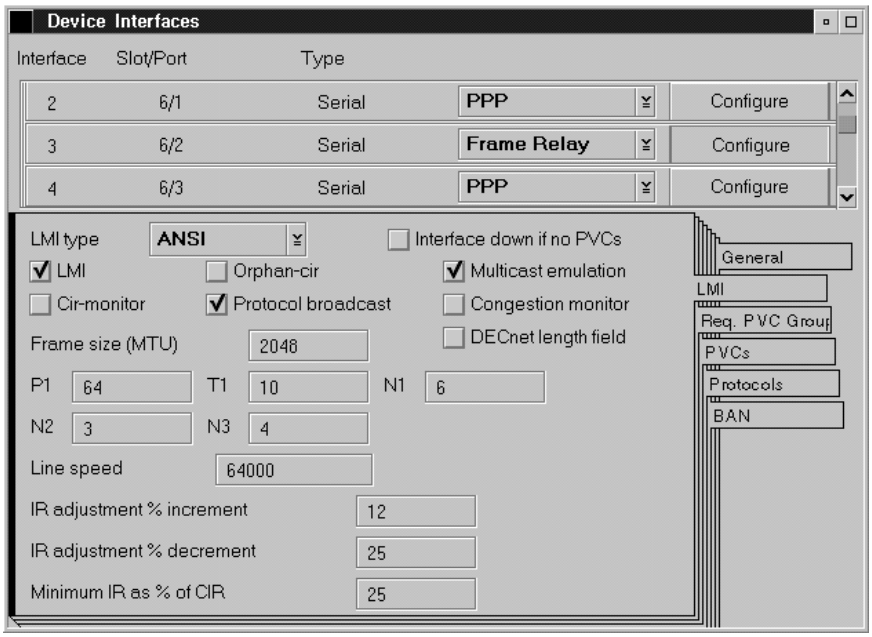

Figure 96. Device Interfaces: Frame Relay LMI

On the LMI page (see Figure 96), select the LMI parameters as shown. LMI type ANSI corresponds to the LMI type supported by the frame relay network. Then select the **PVC** page.

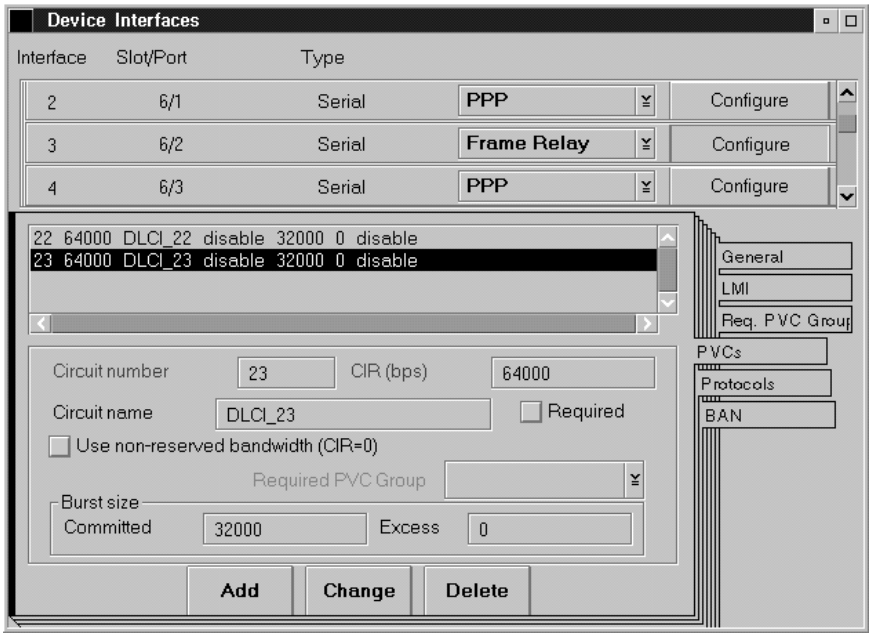

Figure 97. Device Interfaces: Frame Relay PVC

On the PVC page (see Figure 97), we specified the parameters for DLCIs 22 and 23 as shown.

4.5.1 2216A Router IP Configuration Activation

For clarity, we have divided the descriptions of the frame relay and TCP/IP configuration into two separate chapters. From the 2216 point-of-view, the IP configuration steps described in 5.3.4, “2216A Router IP Definitions” on page 150 should be completed before downloading and activating the 2216 configuration. The download and activation process is described in the following text.

To download this configuration to the router, click on the menu option **Configure/Communication/Single Router** at the top of the navigation window. This activates the dialog shown in Figure 98.

Figure 98. Communicate...

This dialog is self-explanatory. One feature is the Date and Time field for the restart option. The current date and time is the default, which means the router restarts immediately after the new configuration is downloaded. To schedule the restart for a later date or time, these fields can be set accordingly.

4.6 3746 ESCON Port 2176 Definitions

Figure 99 on page 138 shows the basic ESCC port configuration. This figure appears when you select the icon **2176/ESCC** in Figure 81 on page 127. The ESCC is a coupler for the ESCON channel connection. We can define the APPN and IP network in one ESCC port for the multiprotocol network environment.

The ESCON port must be enabled for desired network types. These are APPN, IP and/or SNA/Subarea network. A dedicated port name must be specified for APPN and IP networks.

The ESCON director works as an intermediate path switch between the 3746-9x0's ESCON port and the S/390 I/O subsystem. The address of ESCON director's port

that is connected to the 3746-9x0 is referenced in CCM as Control Unit Link Address, which is E0 in our case.

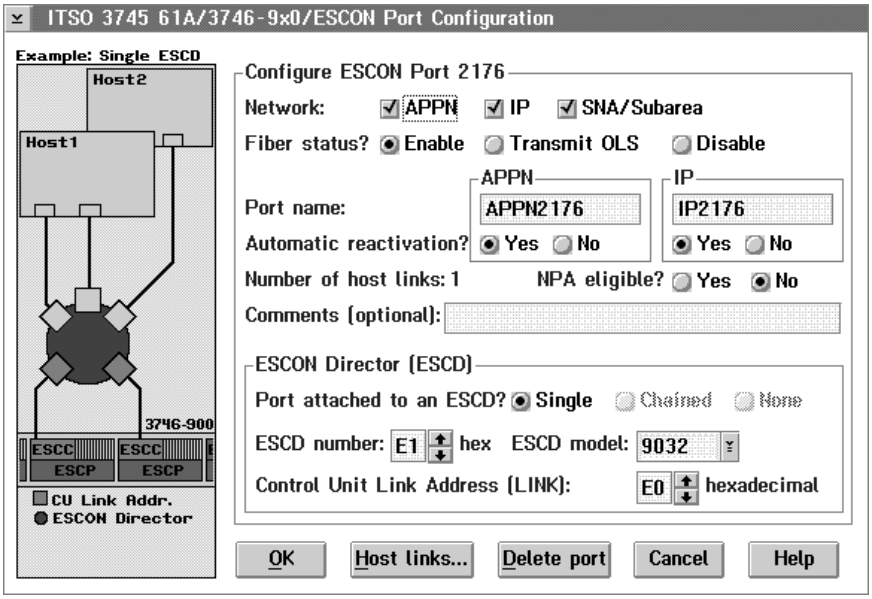

Figure 99. ESCON Port Configuration

Figure 100 on page 139 shows the ESCON host link configuration. This figure appears when you select the icon **Host links...** in Figure 99. Definitions in both Figure 99 and Figure 100 on page 139 build up an end-to-end ESCON host link between the 3746-9x0 and the S/390.

One host link (referenced by No.1) is defined in basic host mode. Host link names for the APPN and IP network over this ESCON port are HL2176A and HL2176I respectively.

In our scenario, the S/390 operating system is one VM system named SYS6, which hosts multiple guest MVS/ESA and OS/390 systems.

The address of ESCON director's port that is connected to the S/390's CHPID is referenced in CCM as Host Link Address, which is E1 in our case. You can find that the end-to-end host link is made up by the following address parameter sequence directed from S/390 to 3746-9x0.

Table 17. ESCON Host Link Parameters	
S/390 Channel Path ID	CHPID: 18 (hexadecimal)
ESCD Host Link Address	HLA: E1 (hexadecimal)
ESCD Control Unit Link Address	CULA : E0 (hexadecimal)
3746 ESCON port number	ESCC : 2176 (decimal)

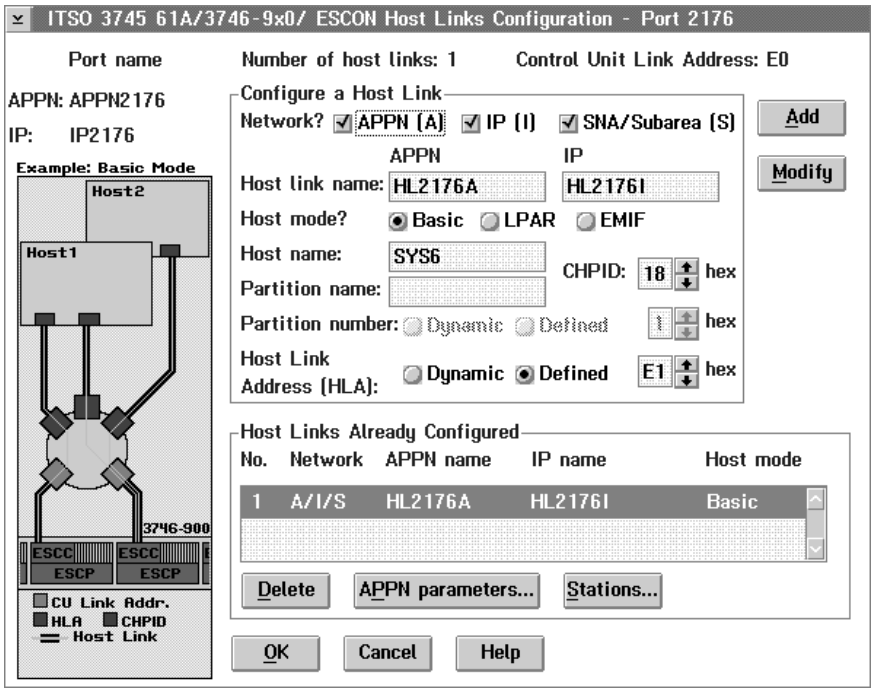

Figure 100. ESCON Host Link Configuration

Figure 101 on page 140 shows the IP parameter definition for one IP station named ST92F. This figure appears when you select the icon **Stations...** in Figure 100.

The IP station named ST92F is one of the stations on the ESCON host link No.1 that is defined in Figure 100. The last three letters of this station name, that is '92F', represent the unit address on the MVS I/O subchannel to which the station is connected. The station is identified by a unit address between x'01' and x'10'. For example, the station ST92F is identified by the unit address X'10'. So, up to 16 stations can be defined on an ESCON host link. These stations are attached to various S/390 systems and subsystems by different I/O subchannels. Refer to Figure 105 on page 145 to see the overall stations defined on this host link.

You can define the IP address for the IP station even if the other IP station's address is defined in the same IP subnet.

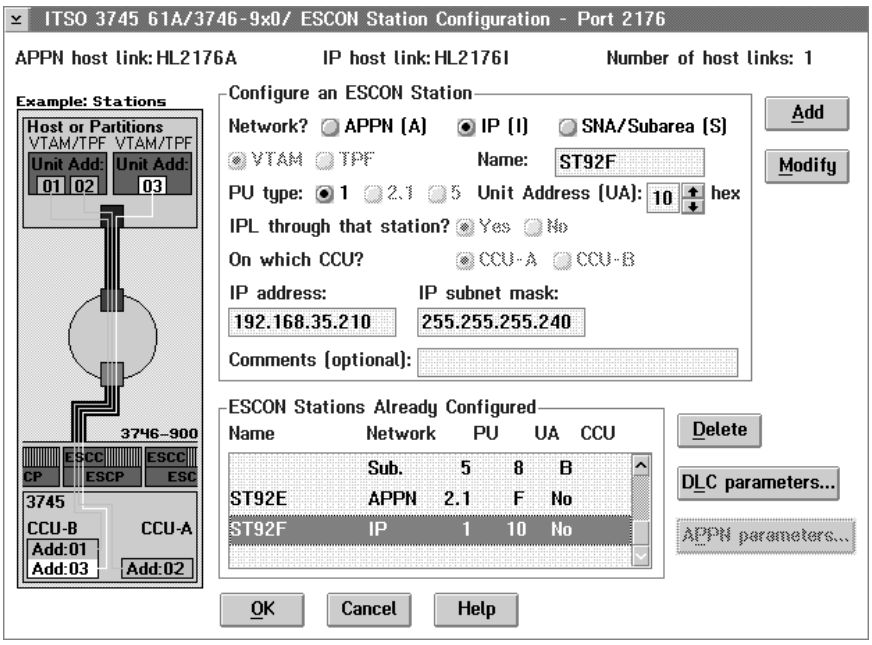

Figure 101. IP Parameters for IP Station on ESCON Host Link

4.7 3746 Token-Ring Port 2144 Definitions

Figure 102 shows the basic TIC3 port configuration. This figure appears when you select the icon **2144/TIC3** in Figure 81 on page 127. We can define an APPN and IP network in one token-ring port for the multiprotocol network environment.

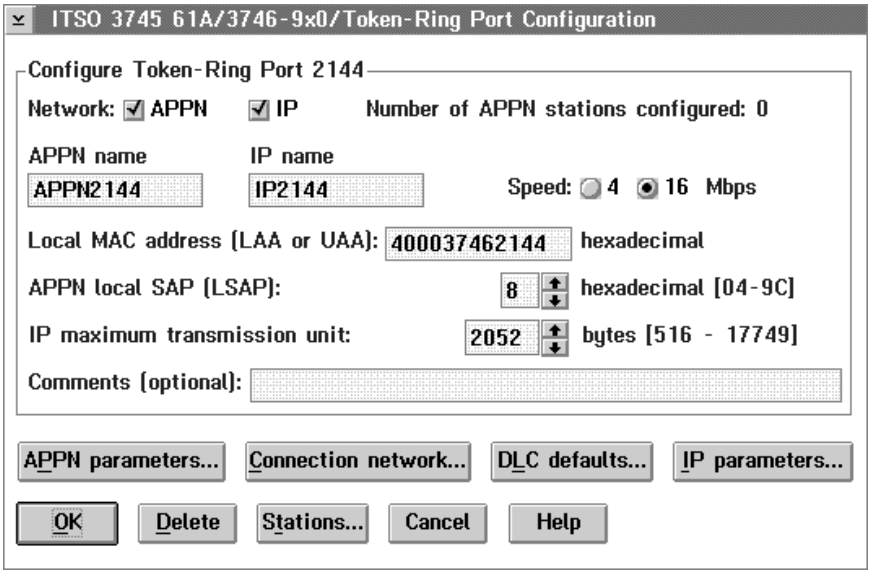

Figure 102. TIC3 Port Configuration

Figure 103 on page 141 shows IP parameters for the TIC3 token-ring port. This figure appears when you select the icon **IP parameters...** in Figure 102.

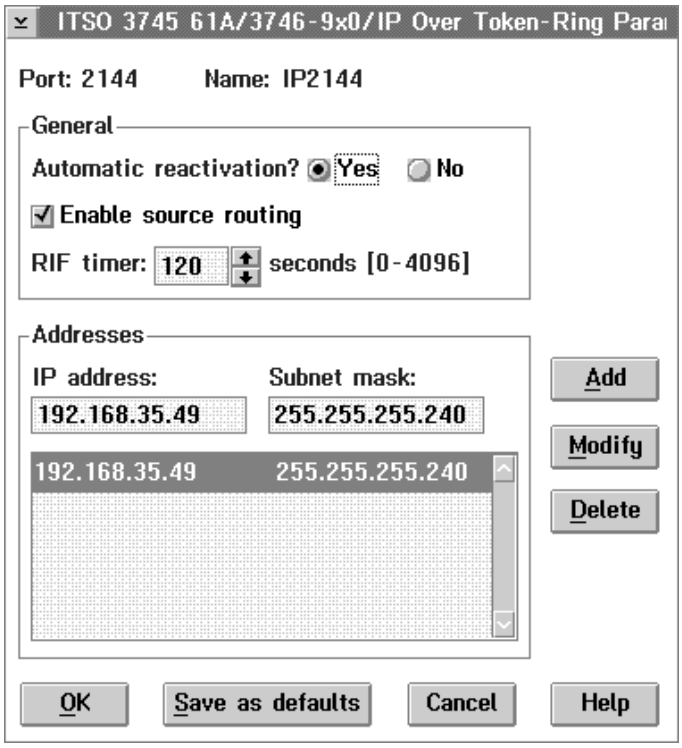

Figure 103. IP Parameters for the Token-Ring Port

Chapter 5. TCP/IP Test Scenario

In this scenario we focus on overlaying a TCP/IP network on to the physical network defined in Chapter 4, “Transport Network” on page 119. The configuration reflects a typical network that is found at many customer sites.

As a dynamic routing protocol we selected OSPF. All the routers used apart from TCP/IP for MVS support OSPF. The IP addressing scheme we used in this scenario is the same for all subsequent scenarios.

Figure 104 shows the base IP configuration that was used on the physical network. IP routing was enabled on the 3746 for the ports used in this configuration; connectivity was established over ESCON to TCP/IP for MVS running on the host S/390 systems, over frame relay to the 2216, and over token-ring to the workstation.

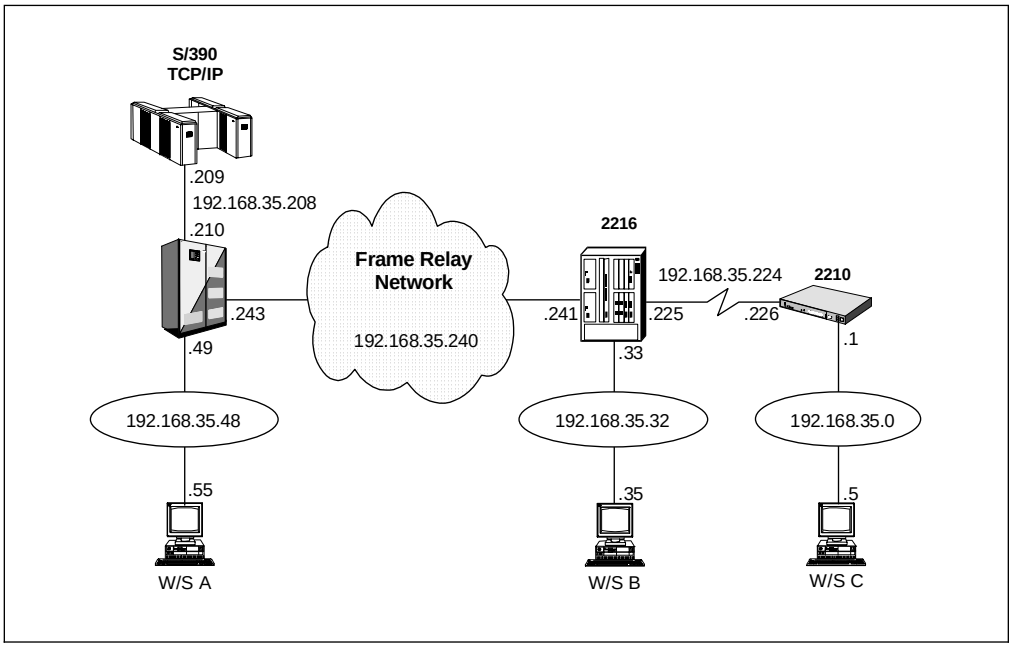

Figure 104. The Base IP Network

5.1 IP Subnets and Addresses

We used an IP network address of 192.168.35.0 with a mask of 255.255.255.240. This gives us 16 subnets with 14 hosts in each subnet. The router is always given the first address in each subnet. The following example shows the available subnets and how they were used in our scenarios.

Table 18. IP Addresses Used

Subnet(s)	Hosts: From->To	Used by:
192.168.35.0	192.168.35.001->014	TR remote LAN (2210-B)
192.168.35.16	192.168.35.017->030	reserved
192.168.35.32	192.168.35.033->046	TR remote LAN (2216-A)
192.168.35.48	192.168.35.049->062	TR local LAN (3746-900)
192.168.35.64	192.168.35.065->078	
192.168.35.80	192.168.35.081->094	
192.168.35.96	192.168.35.097->110	
192.168.35.112	192.168.35.113->126	
192.168.35.128	192.168.35.129->142	
192.168.35.144	192.168.35.145->158	
192.168.35.160	192.168.35.161->174	
192.168.35.176	192.168.35.177->190	
192.168.35.192	192.168.35.193->206	Frame Relay Network 3746-3746
192.168.35.208	192.168.35.209->222	ESCON Channel subnet
192.168.35.224	192.168.35.225->238	PPP Link 2210-A to 2210-B
192.168.35.240	192.168.35.241->254	Frame Relay Network

5.2 Definitions

The following section shows the TCP/IP definitions made for the test network.

5.2.1 TCP/IP MVS Definitions

This section deals with the host TCP/IP for MVS configuration. The host TCP/IP is connected to the 3746-900 via an ESCON channel link. Please refer to 5.3, “3746-900 Definitions” on page 146 for more information on the 3746-900 configuration. Figure 105 on page 145 shows the actual configuration used. As the 3746 and 3745 are being used for other traffic too, there are more stations defined than we actually used in our test scenarios.

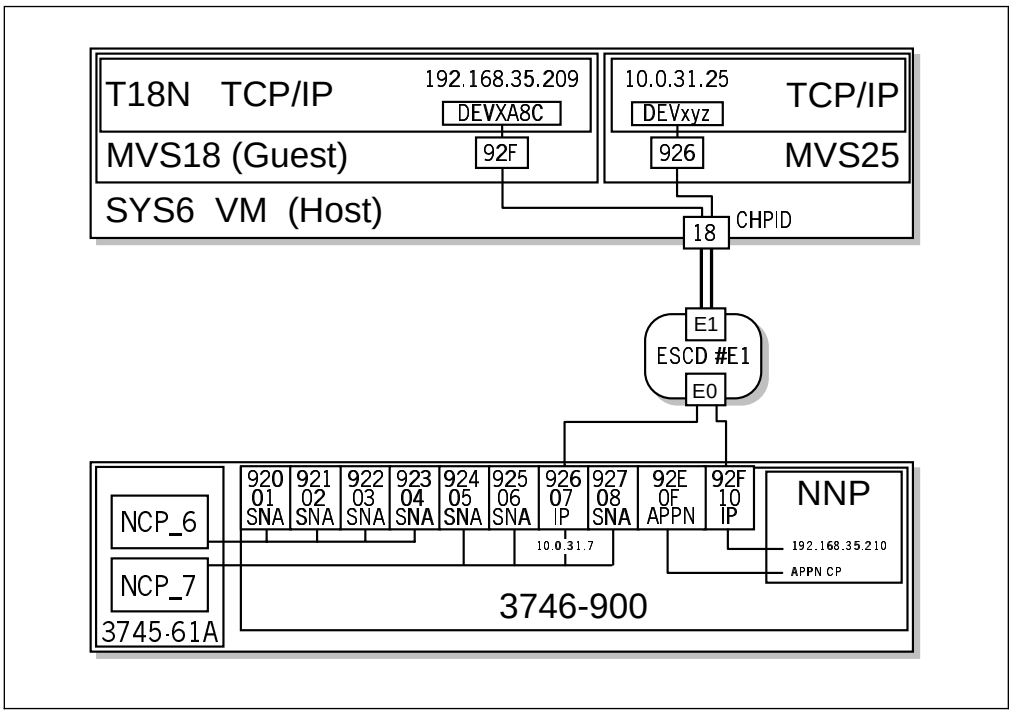

Figure 105. TCP/IP MVS ESCON Connection to the 3746-900 IP Router

For the MVS IOCDF/HCD setup, there are no IP specifics to take into account for an IP over ESCON connection. Over an existing ESCON fiber, we must define a host link to the host system. On this host link up to 16 link stations may be defined. For the 3746 IP router we used link station 0x10. This is the 16th link station at the 3746 side. At the host side this corresponds to address 0x92F, which is the 16th control unit address. The host link is shared with link stations assigned to NCP_7, NCP_6, and the 3746 APPN control point. The other link stations are not discussed further in this book.

The routing protocol between TCP/IP MVS and the 3746-900 is limited to RIP by the host TCP/IP for MVS; OSPF is not supported. This must be taken into account when defining general IP parameters on the 3746. This means that if the outer OSPF network should know about, and advertise, routes to the host TCP/IP (only known by RIP), then the import option must be enabled on the 3746 as shown in Figure 111 on page 150. This option allows a two-way exchange of information between RIP and OSPF.

In order to make the 3746-900 ESCON IP link station and IP address known to the TCP/IP MVS program the definitions shown in Figure 106 on page 146 are needed in the tcpparms profile member of TCP/IP for MVS:

```
; ****************************************************************
; *                                                              *
; * PROFILE.TCPIP for V3R2 Stack T18N on MVS18                   *
; * TCPIP.T18N.TCPPARMS(PROFILE)                                 *
; *                                                              *
; ****************************************************************
; *           3746-9x0 IP over Channel attachment (ESCON)        *
; *           To NN061A                                          *
; ****************************************************************
;
DEVICE DEVXA8C  CDLC 92F 200 200 2060 2060
LINK   LINKXA8C CDLC 0   DEVXA8C
;
; ****************************************************************
; * HOME IP addresses.                                           *
; * The first IP address in the HOME list is the default         *
; * local address for the net or subnet.  For links to other     *
; * networks or subnetworks, the first address in each network   *
; * or subnetwork in the HOME list is also a default local       *
; * address.                                                     *
; ****************************************************************
;
HOME
   .
   192.168.35.209  LINKXA8C     ; ESCON CDLC / IP TO 3746-900 NNP
   .
;
; ********************************************************************
; * RouteD Routing statements for dynamic routing
; * Static definitions are in SYS1.TCPPARMS(GATEWAYS)
; ********************************************************************
;
BSDROUTINGPARMS false    ;  Default max mtu size of 576 bytes is used
;                           for packets to non-adjacent networks.
;    LINK     maxmtu   METRIC   SUBNET MASK     DESTINATION ADDRESS
;                                                (point to point links)
   .
   .
     LINKXA8C  2040       0     255.255.255.240 192.168.35.210
   .
ENDBSDROUTINGPARMS
;
; ****************************************************************
; * Start devices                                                *
; ****************************************************************
;
   .
  START DEVXA8C              ; ESCON CDLC / IP TO 3746-900 NNP
   .
;
; end of profile
```

Figure 106. TCP/IP for MVS TCPPARMS PROFILE Member

5.3 3746-900 Definitions

This section describes the IP definition process for the IBM 3746. The same IP configuration is also used in later test scenarios. The 3746 IP protocol stack can use all hardware ports on the 3746. It can also share ports with the 3746 APPN protocol stack, and NCP protocol stacks.

The 3746 is configured to work as an IP router. For a detailed description of the 3746 IP support, refer to *3746 Models 900 and 950 IP Configuration Guide*, SG24-4845. As well as routing IP over the frame relay network interface, IP is also

routed over the ESCON channel interface to TCP/IP on an S/390 host system, and over a token-ring network interface as a LAN interface.

3746 controller configuration and management (CCM) is a tool that runs on the 3746 service processor (SP) or a stand-alone PC, and is used for defining 3746 interfaces, protcols, and managing the 3746's resources. CCM definitions can be moved between a stand-alone PC and the SP via the import/export function. The management function is only available on the service processor, which is attached to the 3746 and NNP via the service LAN.

5.3.1 3746 Interface IP Definitions

4.4, “3746 Frame Relay Configuration” on page 127 shows how IP addresses are defined for the frame relay interfaces we are using. The IP addresses assigned to the hardware interfaces were defined as the interfaces were defined.

5.3.2 General 3746 IP Definitions

The following general IP definitions for dynamic routing protocols were also needed.

Figure 107 shows the 3746 general IP definition options menu. The option parameters were all left at their default values, which means no static routing was defined.

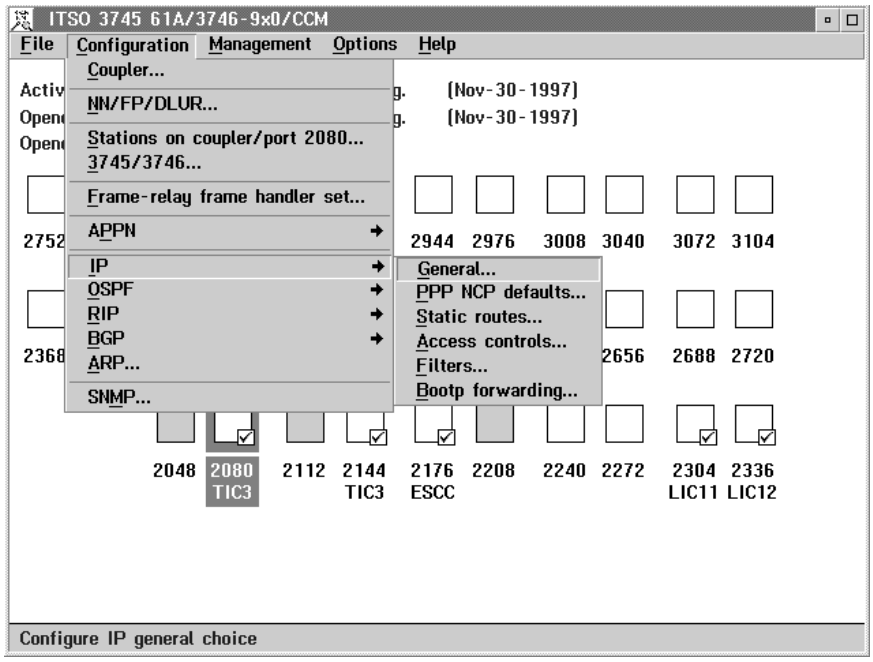

Figure 107. General IP Menu Options

5.3.3 General 3746 OSPF Definitions

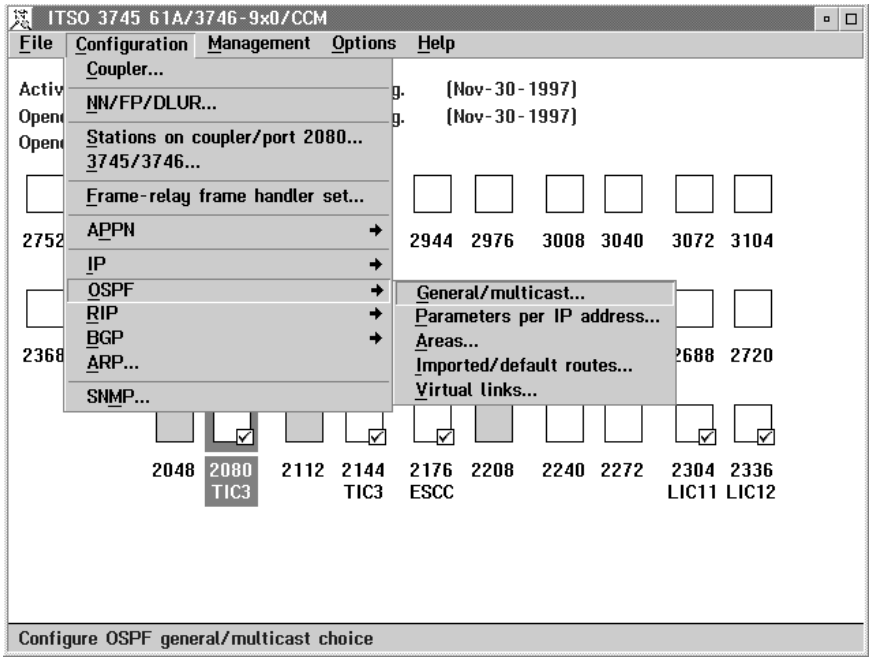

Figure 108. General OSPF Menu Options

Figure 108 shows the menu options for general OSPF options. OSPF routing was enabled in the General/multicast option. The default area number 0.0.0.0 was used.

Figure 109 on page 149 shows an overview of dynamic routing parameters for all defined interfaces. This dialog provides a central overview of all your dynamic routing protocol definitions. This dialog is reached by selecting the **Parameters per IP address...** menu option in Figure 108. Each interface has its IP address listed here with the routing protocols defined (OSPF or RIP) and the OSPF neighbors, if any, are specified.

These values shown are a summary of individual definitions that are explained in the following sections. Please do not pay attention to the IP address 10.0.31.7 in our scenario.

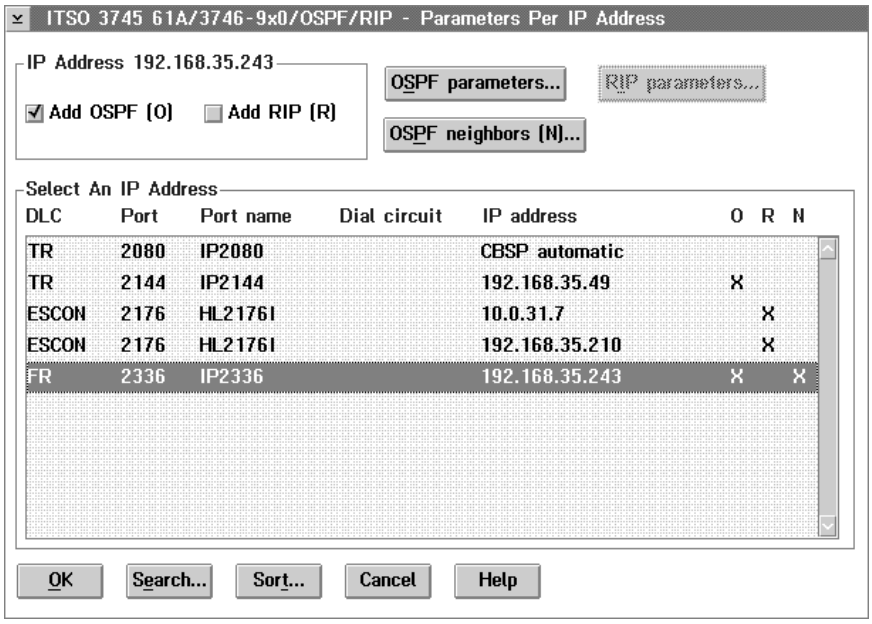

Figure 109. Interfaces and Their IP Addresses

Figure 110 shows the IP addresses of OSPF neighbors that are reachable via the frame relay network. The local IP address on the 3746 frame relay port, 192.168.35.243 is also shown. This dialog is reached by selecting the **OSPF neighbors...** button in Figure 109. The addresses shown are the remote IP addresses that can be reached over the frame relay DLCIs.

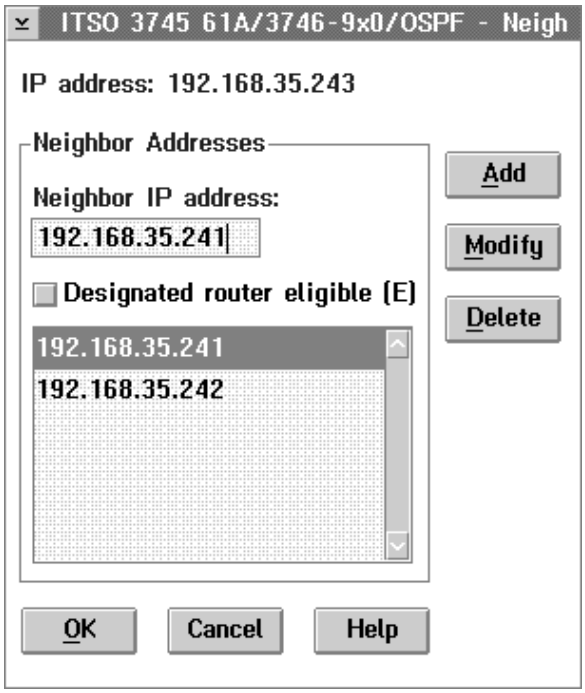

Figure 110. IP Addresses of OSPF Neighbors

Figure 111 on page 150 shows OSPF imported and the default routes specification dialog. This dialog is reached by selecting the **Imported/default routes...** menu option in Figure 108 on page 148. In our scenarios, the dynamic routing protocol used between the 3746 and the 2216 is OSPF. RIP is used between the 3746 and

TCP/IP for MVS as it is the only dynamic routing protocol supported by TCP/IP for MVS. Therefore, we defined the option to import RIP routes (routing information) into OSPF as shown. This allows OSPF to advertise and use the routes that are known to RIP.

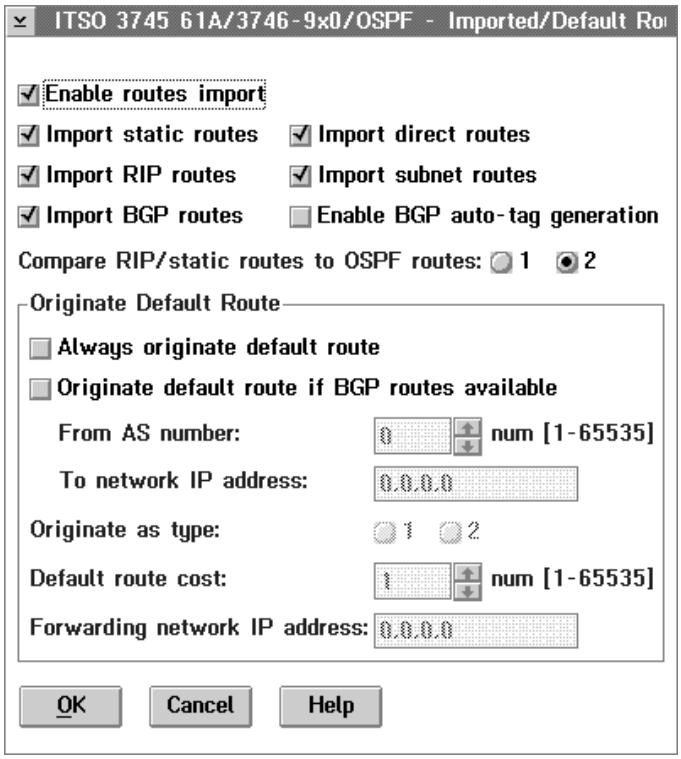

Figure 111. OSPF Imported Routes

5.3.4 2216A Router IP Definitions

The 2216 hardware interfaces were defined in 4.5, “2216 and 2210 Frame Relay Definitions” on page 132. On each DLCI that we plan to use for IP, we must define the remote IP address reachable over that DLCI. On the frame relay port (port 6, interface 2) configuration notebook we must select the **Protocols** tab and configure as shown in Figure 112 on page 151. We defined IP to run on DLCI 23.

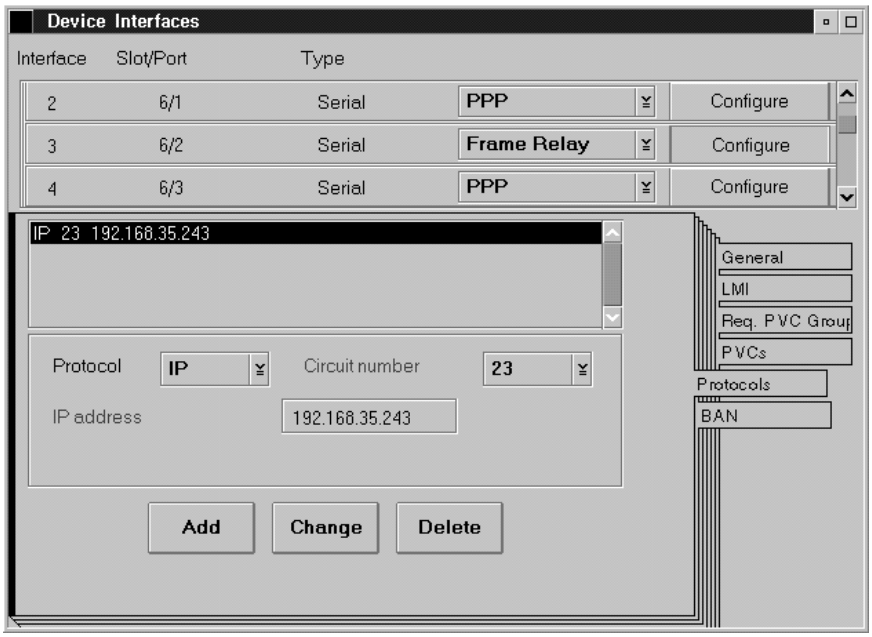

Figure 112. Device Interfaces: Frame Relay Protocols

Now we go back to the navigation window and select **System/General**. We can give the router a name here (see Figure 113). This name will appear as a prompt when you are logged in to the router.

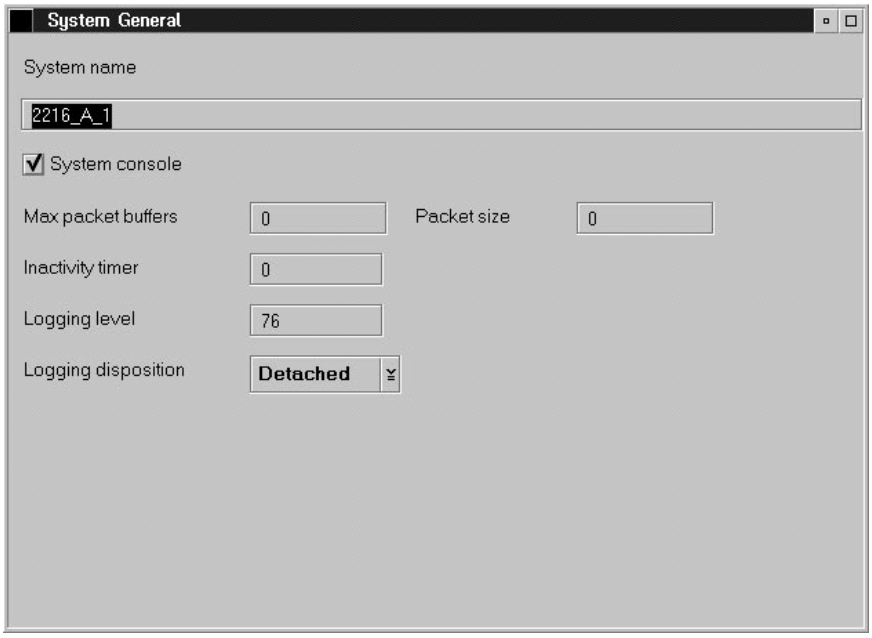

Figure 113. System General

Further down in the navigation window, there is the SNMP menu option. Select **General** under **Communities**. On this dialog we defined two communities: one having read-write and trap access, one having only read and trap access (see Figure 114 on page 152). In order to download the configuration from the 2216 configuration tool, a community name with read-write access must be defined.

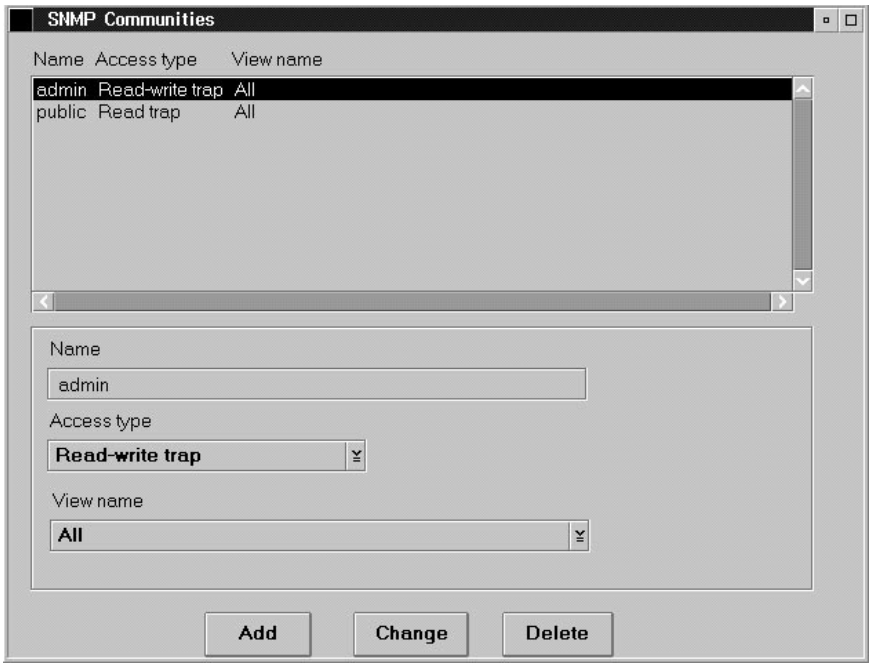

Figure 114. SNMP Communities

Further down in the navigation window, there is the IP menu option. Select **Interfaces** under **IP**. Here (see Figure 115), we can define an IP address and mask for each interface.

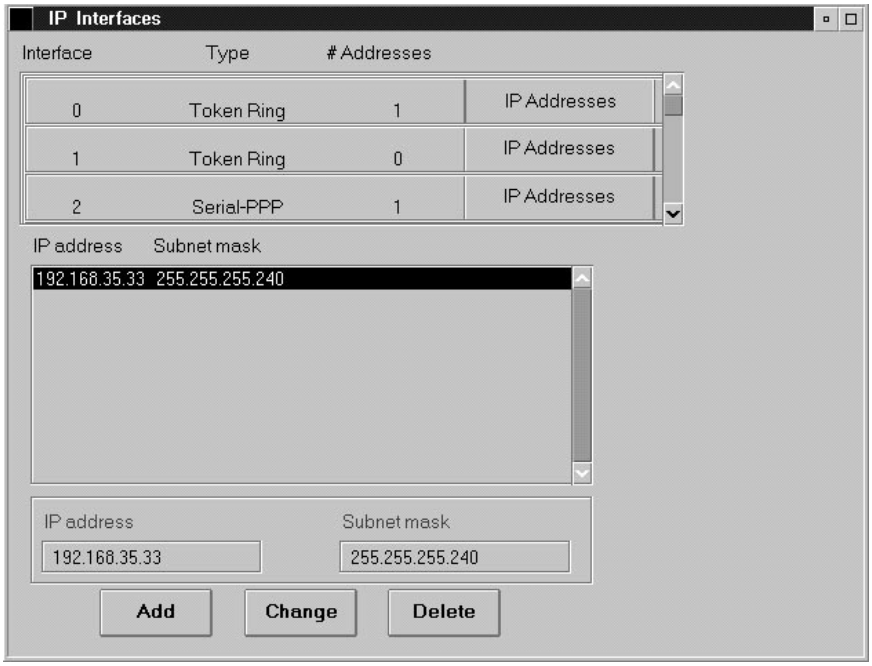

Figure 115. IP Interfaces

Having done that, a routing protocol needs to be enabled. We chose OSPF for the whole network. In the navigator window select **OSPF/General**.

Figure 116. OSPF General

The dialog shown in Figure 117 appears. Check the **Enabled** box. Close this dialog and select **OSPF/Area** from the navigator window.

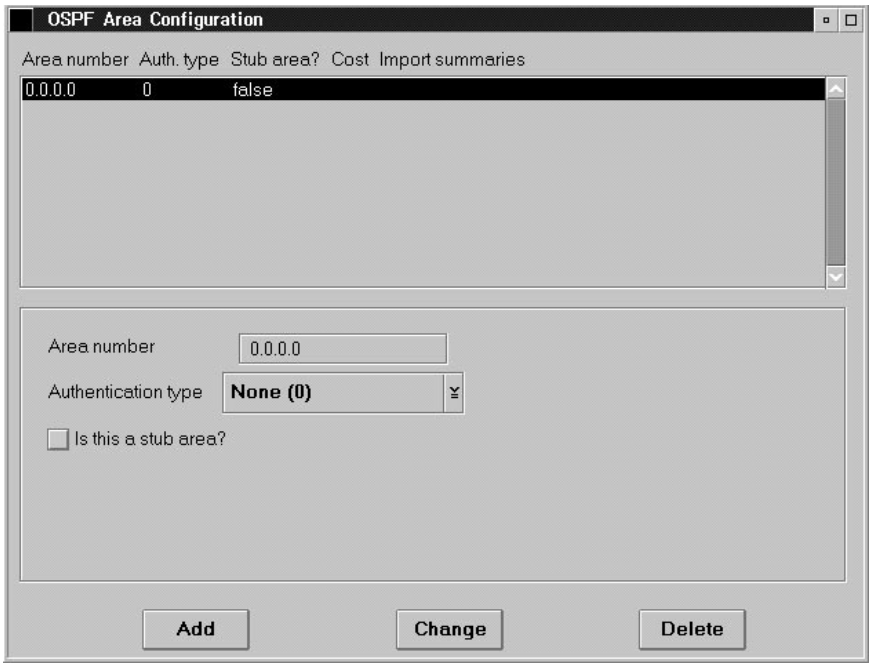

Figure 117. OSPF Area Configuration

The dialog shown in Figure 118 on page 154 appears. Since the network we are configuring is a single OSPF, it must be area 0.0.0.0. Close this dialog and select **OSPF/Interfaces** from the navigator window.

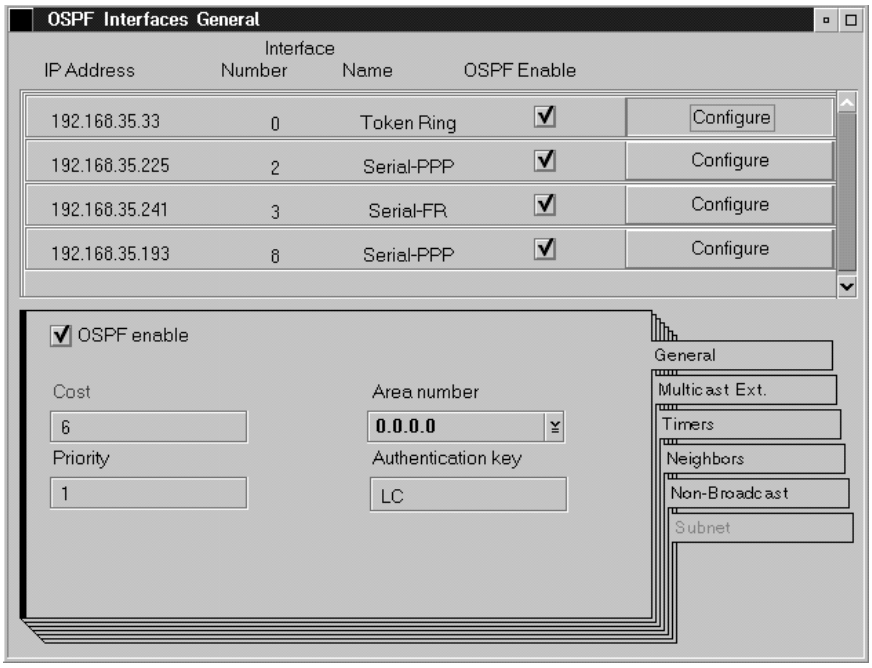

Figure 118. OSPF Interfaces General

The dialog shown in Figure 119 appears. By default, all interfaces are enabled for OSPF. Cost factors are calculated automatically according to the speed of the interfaces.

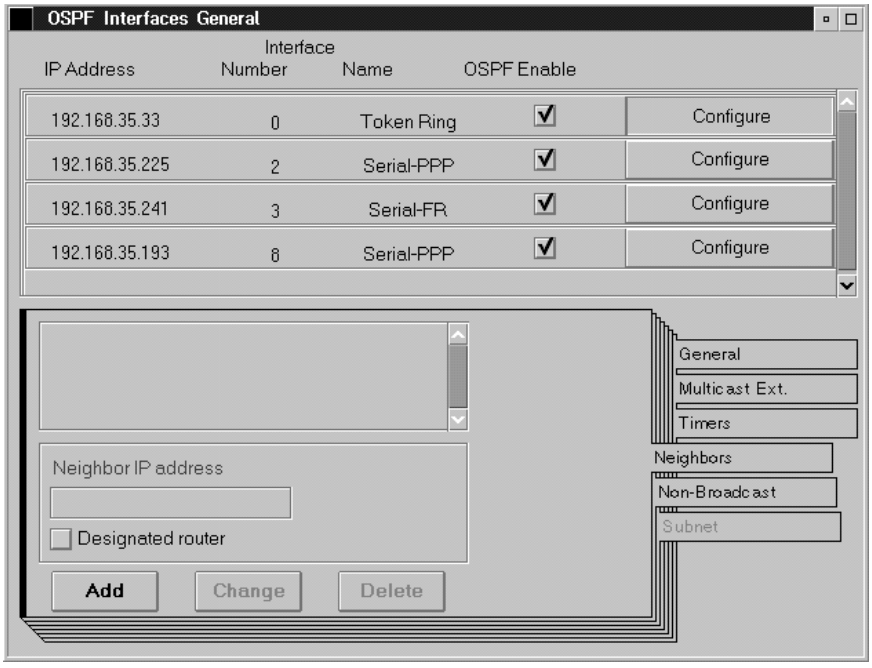

Figure 119. OSPF Interfaces General Neighbors

But what about neighbors? Do we need to define neighbors? Select the **Neighbor IP address** field in the above window and press PF1. This will open a context-sensitive help window. Figure 120 on page 155, and Figure 121 on page 155 show the help screens for OSPF neighbors.

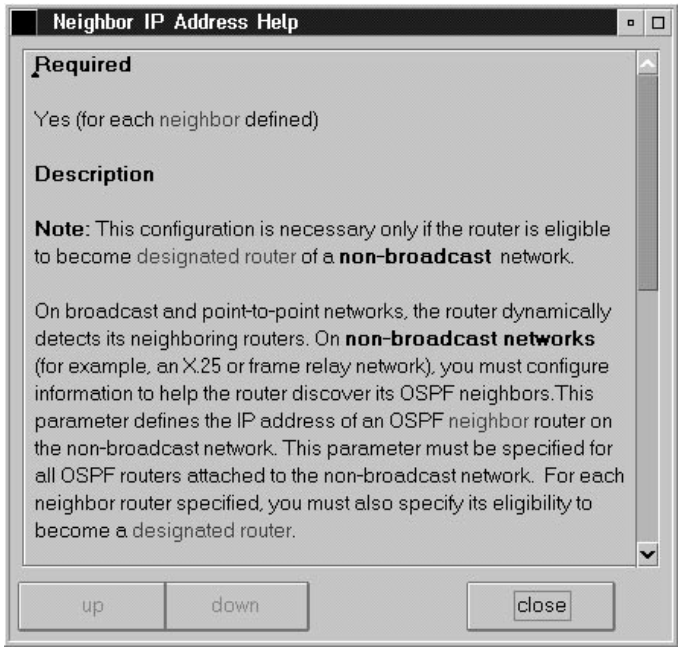

Figure 120. Neighbor IP Address Help (Part 1)

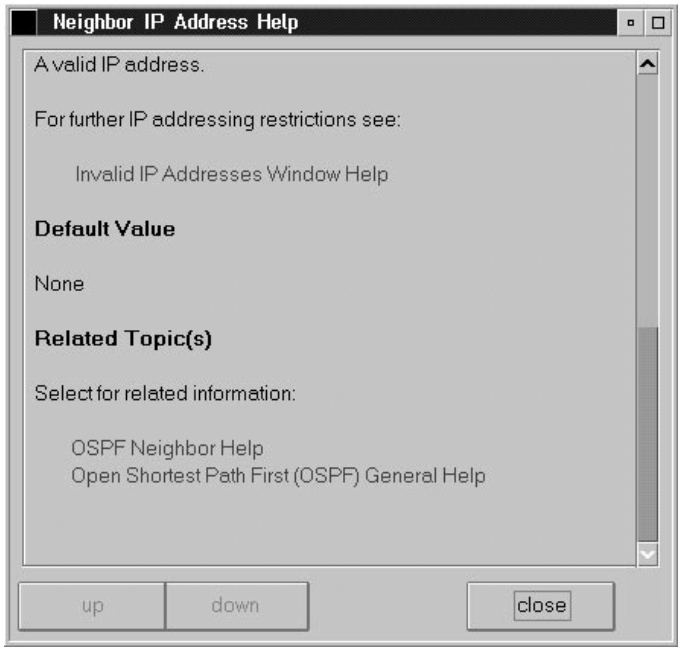

Figure 121. Neighbor IP Address Help (Part 2)

Neighbors need to be defined only on frame relay links, because they are nonbroadcast networks. Close the help window and define the neighbor IP address on the frame relay interface as shown in Figure 122 on page 156.

That concludes the 2216 router configuration. This configuration must now be loaded into the 2216 and activated (see 4.5.1, “2216A Router IP Configuration Activation” on page 137).

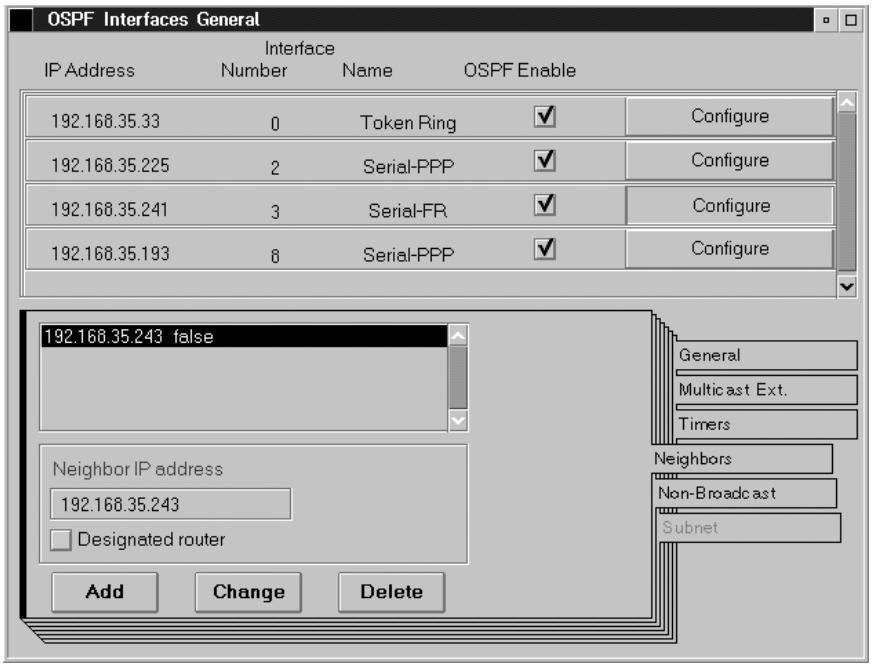

Figure 122. OSPF Interfaces General Neighbors

5.3.5 Router 2210-B Definitions

For the IP scenario, the configuration of the router 2210-B is very similar to the 2216-A. Since the configuration windows for the 2210 are exactly the same as for the 2216, there is no need to repeat them here. The differences are:

- There is no frame relay link.
- The IP address of the serial link is 192.168.35.226.
- The IP address of the token-ring is 192.168.35.1.
- The router name is 2210-B.

5.3.6 Workstation TCP/IP Definitions

The workstations we used were PS/VPs with OS/2 Warp installed. For the IP scenario, the configuration of the workstations is straight forward. In fact, only the IP address, the subnetmask and a default gateway needs to be configured. All other parameters are optional. The following screens show the TCP/IP configuration for workstation A. For workstation B and workstation C, configuration is similar, just the IP addresses, subnet masks, and default gateways are different. The addresses used are listed in Table 19.

Table 19. Workstation IP Addresses

Workstation	IP Address	IP Mask	Gateway
A	192.168.35.50	255.255.255.240	192.168.35.49
B	192.168.35.35	255.255.255.240	192.168.35.33
C	192.168.35.5	255.255.255.240	192.168.35.1

Although the complete configuration of the workstation IP stack is outside the scope of this book, various screens have been included to give an overview of the

process. The workstation TCP/IP configuration notebook is opened by typing TCPCFG from an OS/2 command line. Figure 123 on page 157 shows the configuration notebook. On the Network page we specify the interface we wish to configure and then enter its IP address and subnet mask.

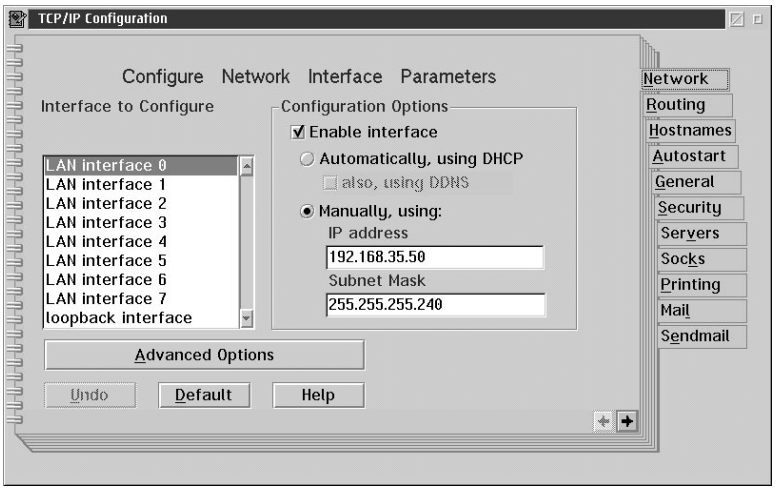

Figure 123. TCP/IP Configuration Notebook: Network

Select the **Routing** tab of the notebook and specify the default gateway for this workstation (see Figure 124).

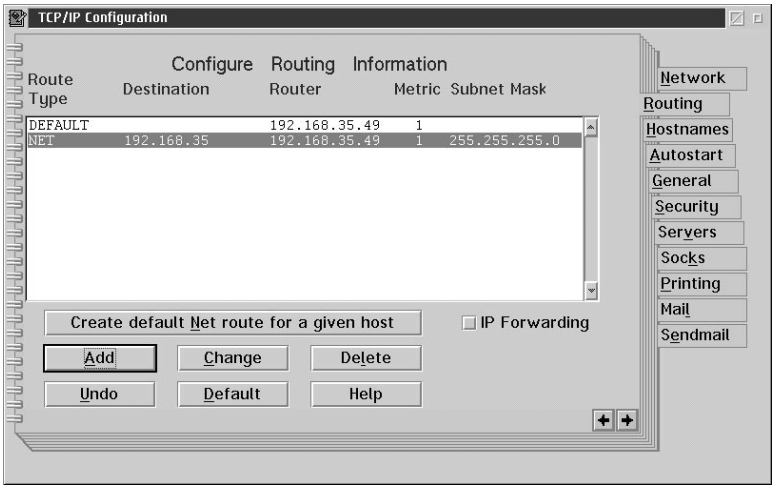

Figure 124. TCP/IP Configuration Notebook: Routing

Select the **Hostnames** tab of the notebook and specify the host name of this workstation and the local domain name (see Figure 124).

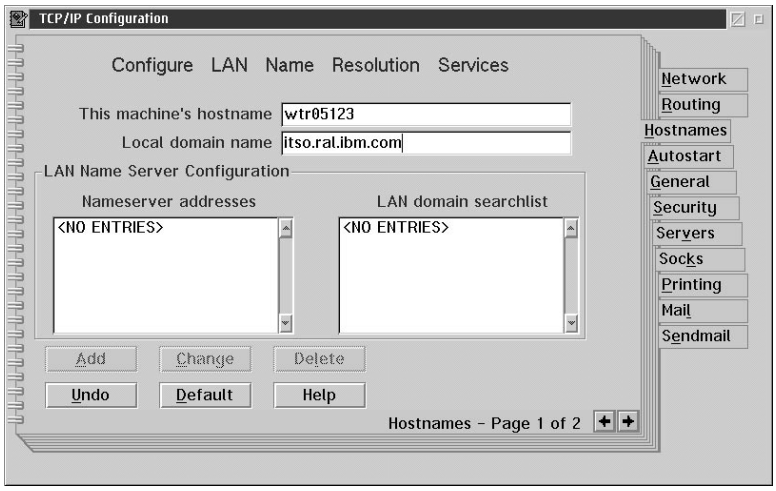

Figure 125. TCP/IP Configuration Notebook: Hostname

5.3.7 3746 CCM Management Displays

Figure 126 shows the IP protocol Management menu options available from CCM. In addition to these displays, it is possible to TELNET to the 3746 NNP and issue line commands to display information about the 3746 IP router and components. For a detailed description see *3746 Nways Controller Models 900 and 950: IP Implementation Guide*, SG24-4845.

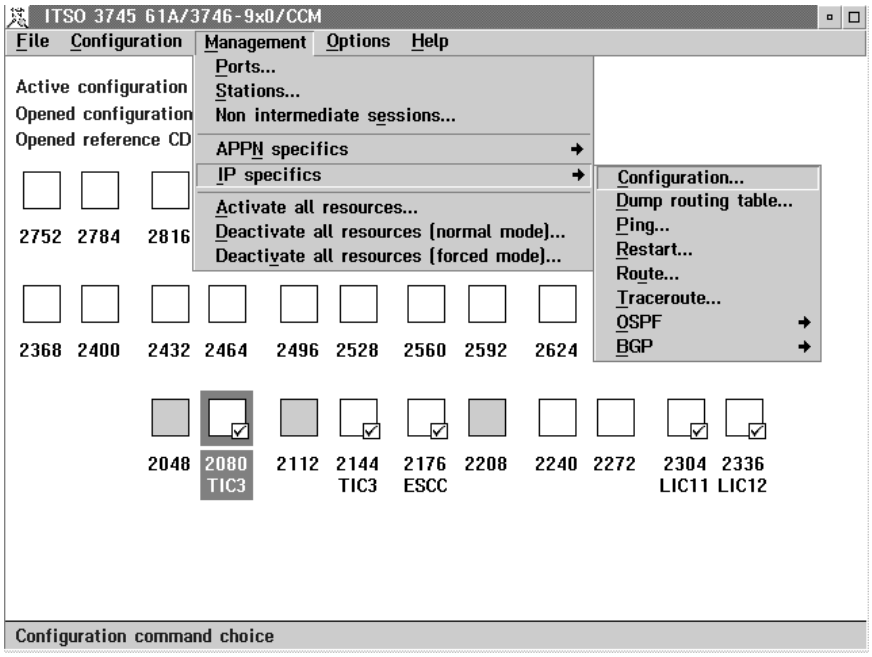

Figure 126. CCM Menus for IP Protocol Management

Figure 127 on page 159 shows the IP results dump display. This display is reached by selecting the **Dump routing table...** menu option in Figure 126. In this figure the current routing table is displayed. The IP interfaces are shown with internal labels (for example, TKR/1 for token-ring port #1). To see which 3746 port address corresponds to this label, see Figure 128 on page 159.

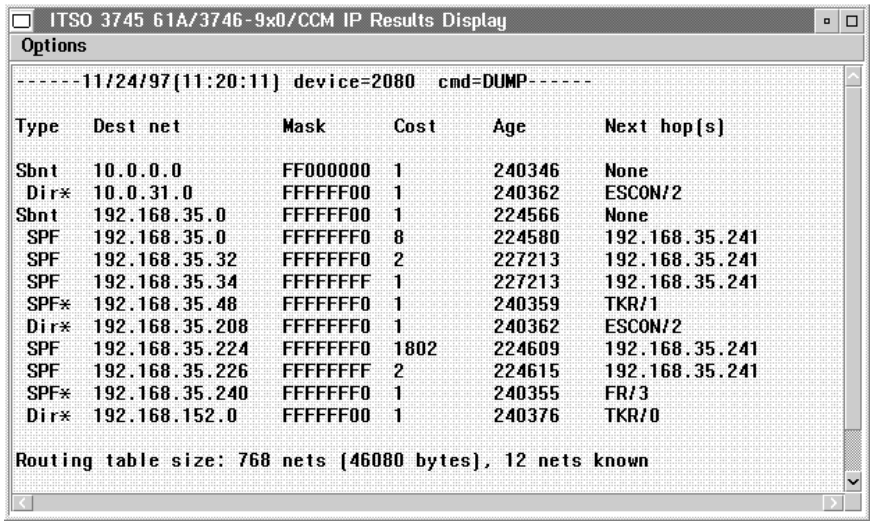

Figure 127. CCM IP Results Dump Display

Figure 128 shows the IP results configuration display. This display is reached by selecting the **Configuration...** menu option in Figure 126 on page 158. In this display the current state for the interfaces is also indicated.

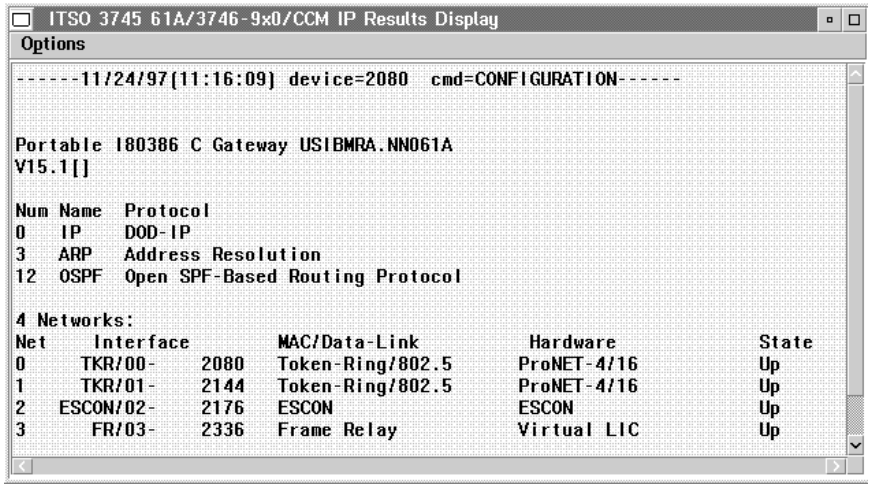

Figure 128. CCM IP Results Configuration Display

Figure 129 on page 160 shows the ports management display. This is not an IP-specific display, but it shows the status of the 3746 hardware ports. This display is reached by selecting the **Ports...** menu option in Figure 126 on page 158. We can see the following ports which are routing IP: port 2080, which is the service LAN token-ring port, port 2144, which is the 3746 token-ring port that is used for connectivity to workstation A, the ESCON port 2176, port 2336 the frame relay port, and port HL2176I, which is the host link to the host system.

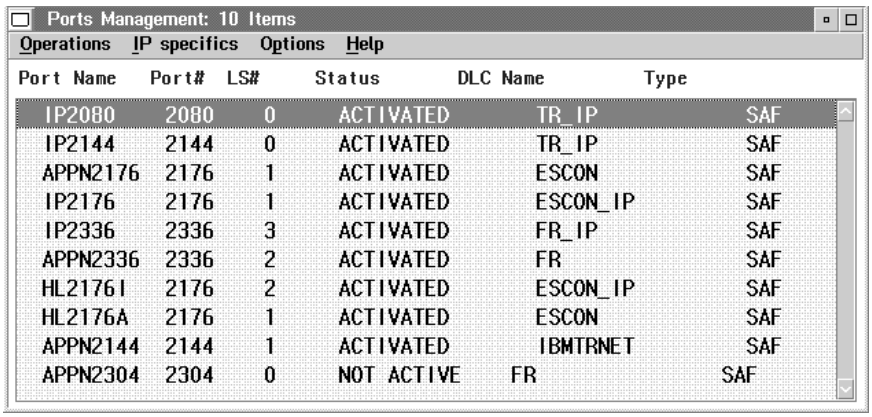

Figure 129. CCM Ports Management Display

Figure 130 shows the stations management display. This is also not IP-specific. This display is reached by selecting the **Stations...** menu option in Figure 126 on page 158. In this figure the current state for the station is also indicated. The link station ST92F is our IP link station to TCP/IP for MVS. This was defined in Figure 101 on page 140.

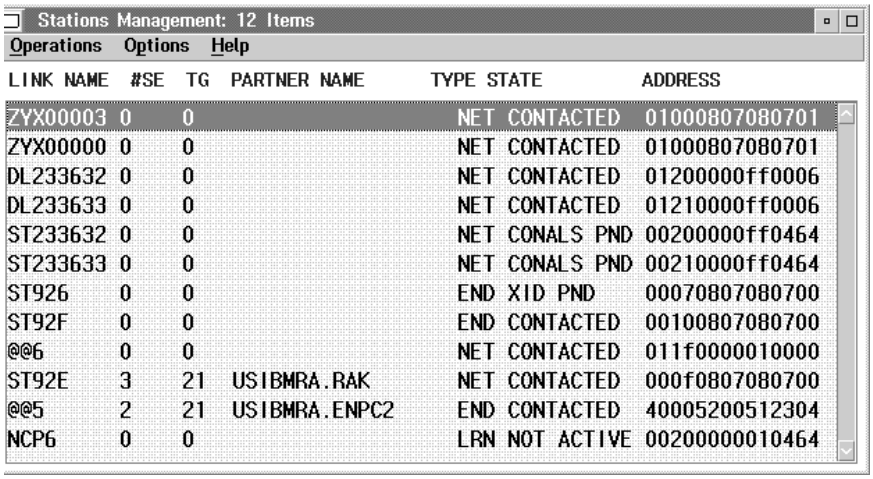

Figure 130. CCM Stations Management Display

Part 3. Test Scenarios

Chapter 6. Frame Relay Boundary Access Node

In this scenario, we focus on the frame relay boundary access node (BAN). The router 2216A is configured for BAN and is called the BAN router. BAN uses the RFC1490 bridged frame format for encapsulation of SNA over frame relay networks. There are two types of BAN implementation: BAN1 and BAN2.

In BAN1, the frame relay access router works as a LAN bridge. This bridging provides a transparent data path between the endstations and the 3746. The implementation of BAN in the 2216 and 2210 routers provide what is called a BAN PVC MAC address. This is a virtual MAC address to which the end stations send the data bound for the 3746.

In BAN2, the router terminates the LLC traffic received from the end stations. At the same time, the router establishes a new LLC connection to the 3746 over the frame relay network and maps the traffic from the other LLC connections to this connection. That means, that there are two LLC connections:

- From the endstation to the frame relay BAN router via LAN/WAN
- From frame relay BAN router to the 3746 via frame relay

It is possible to map connections onto a single LLC connection over the BAN PVC.

From the 3746-900 point of view, there is no difference between BAN1 and BAN2. In both cases the 3746 receives frames in the bridged frame format. This encapsulation is used for SNA traffic. IP traffic may also be sharing the same DLCI as SNA traffic, and it uses the RFC1490 routed frame format. The 3746 supports receiving bridged and routed frame format frames on the same PVC.

The BAN PVC MAC address can be the same for all BAN routers in the network. The advantage of this method is that all workstations in the network can be configured with the same destination MAC address.

The router 2210B is serving a remote location and is connected to the 2216A over a PPP link. The SNA data are transported to the BAN router using DLSw.

Figure 131 on page 164 shows the frame relay BAN test configuration.

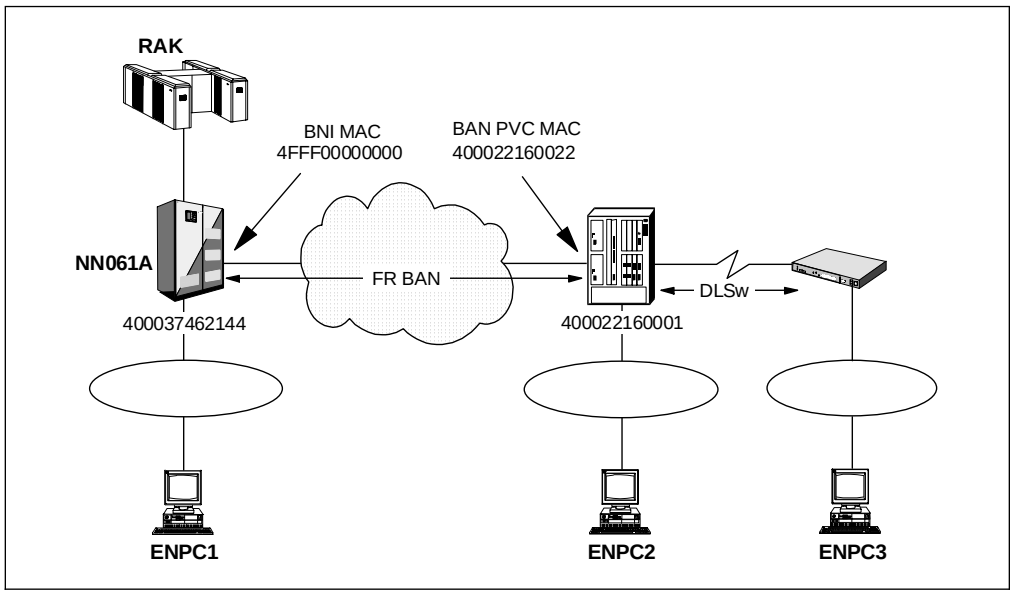

Figure 131. Frame Relay BAN Scenario

6.1 Definitions

The following sections show the definitions made for this test scenario.

6.1.1 VTAM Definitions

For this scenario we added a VTAM V4R4 under MVS/ESA as an APPN network node (NN). This NN has physical connectivity to the rest of our network over an ESCON connection to the 3746 network node. The VTAM NN is called USIBMRA.RAK.

The VTAM start options specified in ATCSTRxx member are shown in Figure 132 on page 165.

```
SSCPID=20,HOSTSA=20,XNETALS=YES,BN=YES,                                X
CDSERVR=YES,                                                           X
ENCRYPTN=NO,                                                           X
IOPURGE=45,                                                            X
GWSSCP=YES,                                                            X
CMPVTAM=4,                                                             X
SORDER=SUBAREA,                                                        X
SRCHRED=ON,SRTIMER=120,SRCOUNT=60,                                     X
CONFIG=K0,SUPP=NOSUP,                                                  X
NETID=USIBMRA 1 ,HOSTPU=ISTPUS20,HOSTSA=20,SSCPNAME=RAK 2 ,            X
SSCPDYN=YES,SSCPORD=PRIORITY,                                          X
SSEARCH=YES,                                                           X
ASYDE=TERM,                                                            X
NOTRACE,TYPE=VTAM,IOINT=0,                                             X
NOTRACE,TYPE=SMS,ID=VTAMBUF,                                           X
PPOLOG=YES,                                                            X
NODETYPE=NN 3 ,                                                        X
INITDB=NONE,                                                           X
CPCP=YES,                                                              X
VRTG=YES,                                                              X
CSALIMIT=0,                                                            X
OSITOPO=LLINES,                                                        X
OSIMGMT=YES,                                                           X
TNSTAT,CNSL,TIME=20,DYNLU=YES
```

Figure 132. VTAM ATCSTR Member

Notes:

- 1 NETID is set to USIBMRA.
- 2 SSCP name is RAK.
- 3 VTAM is started as a network node.

VTAM defines the direct connection with the 3746-900 NNP network node in the local SNA major node. This is shown in Figure 133. The link control unit address used for this link is 92E (4). See Figure 105 on page 145 and Figure 101 on page 140 to see how this corresponds to the 3746 definitions.

```
*********************************************************************
*   LOCAL SNA MAJOR NODE FOR 3746-900  (CPNAME=NN061A)              *
*********************************************************************
         VBUILD TYPE=LOCAL
CP90061A PU    CUADDR=92E 4 ,                                          X
               CONNTYPE=APPN,                                          X
               CPCP=YES,                                               X
               DYNADJCP=YES,DYNLU=YES,                                 X
               MAXBFRU=15,                                             X
               PUTYPE=2,                                               X
               XID=YES,                                                X
               ISTATUS=ACTIVE
```

Figure 133. VTAM Local SNA Major Node

6.2 3746 APPN Definitions

In this scenario the 3746 will be used as an APPN network node. The 3746 is the only network node in the network except for the VTAM network node. Therefore, we made the 3746 the network node server for all of the workstation APPN end nodes. The 2216 and 2210 have no APPN functionality in this test scenario.

Figure 134 shows the logical APPN view of the frame relay BAN test scenario.

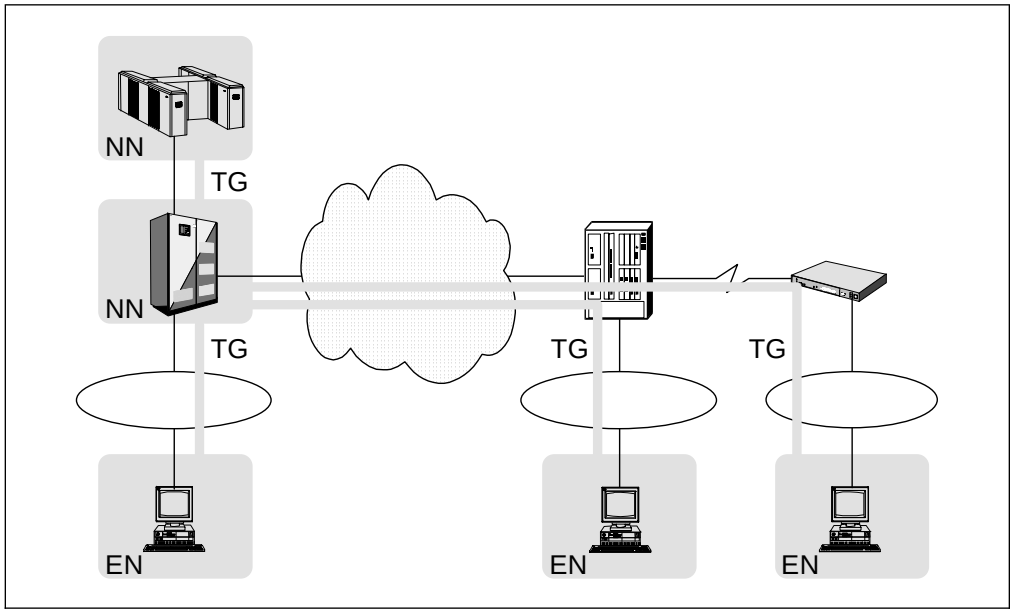

Figure 134. Frame Relay BAN APPN Logical View

6.2.1 3746 General APPN Definitions

Figure 135 on page 167 shows the menu options for the main NN definitions (NN/FP/DLUR) and general APPN options (APPN). All parameters not discussed here were left at their default values.

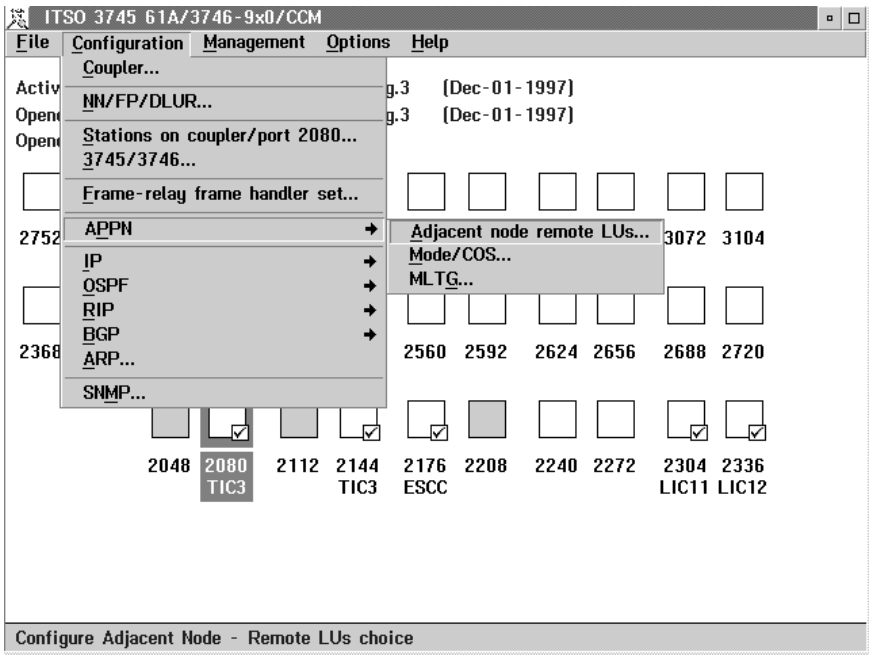

Figure 135. Main APPN Configuration Menu

Figure 136 on page 168 shows the primary network node definition dialog. This dialog is reached by selecting the **NN/FP/DLUR...** menu option in Figure 135.

The APPN network node name for this NN consists of the network identifier (NETID) and the control point name (CP). The name of the NN is USIBMRA.NN061A. The network management focal point was defined as the VTAM network node, USIBMRA.RAK. The dependent LU server (DLUS) is also defined as the VTAM network node. This allows the 3746 NN to function as a dependent LU requester (DLUR) also.

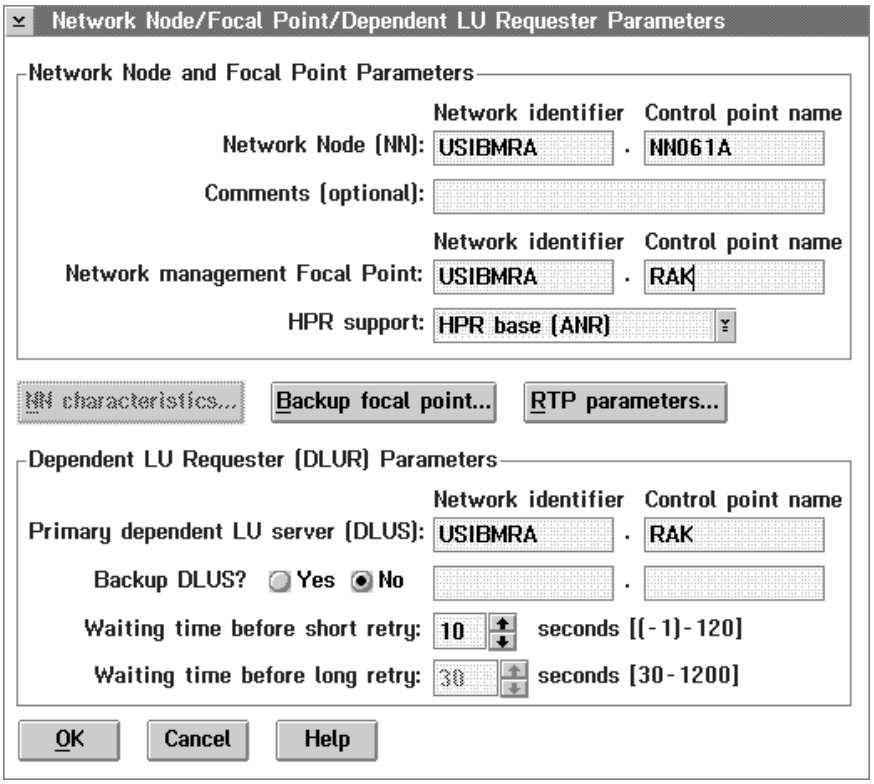

Figure 136. Primary NNP Definition

6.2.2 3746 Frame Relay Port 2336 APPN Definitions

Figure 137 shows the main CCM configuration screen.

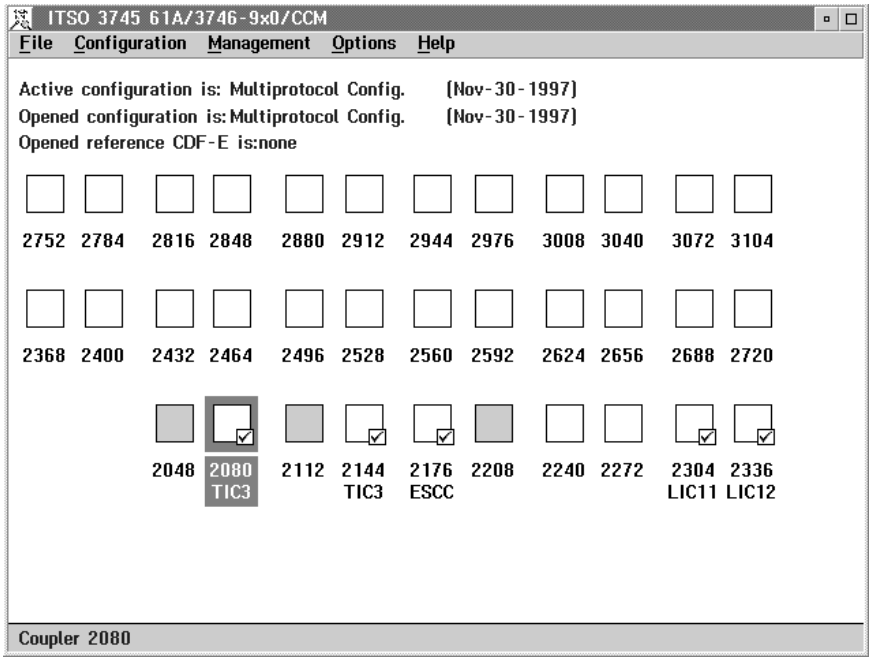

Figure 137. 3746 Ports Configuration

Figure 138 on page 169 shows the basic LIC12 port configuration for the 3746 frame relay port 2336. This dialog is reached by selecting the icon **2336/LIC12** in Figure 137. The frame relay port can be shared by the 3746 APPN and IP protocol stacks. We need to select the check box **APPN** to allow us to make APPN definitions on this port.

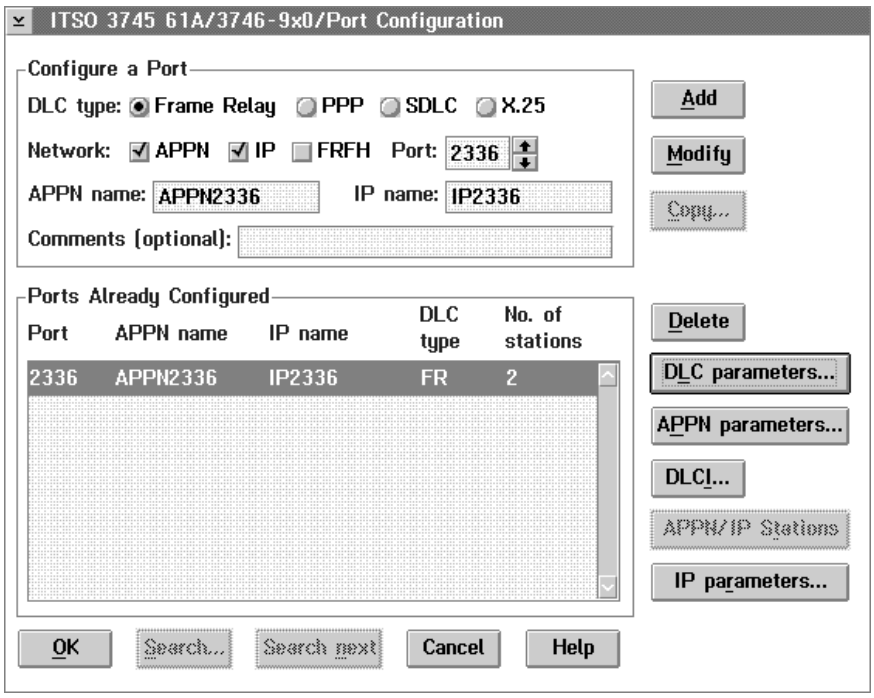

Figure 138. LIC12 Port Configuration

The data link control (DLC) parameters to be defined for the LIC12 frame relay port are shown in Figure 83 on page 129 and Figure 84 on page 129.

Figure 139 on page 170 shows APPN parameters for the LIC12. This dialog is reached by selecting the button **APPN parameters...** in Figure 138. The parameter values shown are the defaults.

The parameter Accept any incoming call defines whether incoming calls are accepted from other stations or whether we want to allow APPN connections to explicitly defined stations only. This is important for how we are making our APPN definitions on the other stations. We can explicitly define the MAC addresses of the other workstations here on the 3746, and disable the Accept any incoming call parameter. This means the 3746 must activate the connection to the workstation by calling out. No undefined workstation can get an APPN connection to this 3746 NN if we do this.

Or, we can leave all workstations undefined, and enable Accept any incoming call. This will allow any APPN station with connectivity to this port to establish an APPN connection to the 3746. This is the method we used in our test scenarios. Whether frame relay or token-ring connected, each workstation is configured with the MAC address of the 3746 and calls in when activated. When a call in is accepted, a dynamic APPN station is created in the 3746. These have names such as @010 to allow the operator to differentiate between predefined and dynamic link stations.

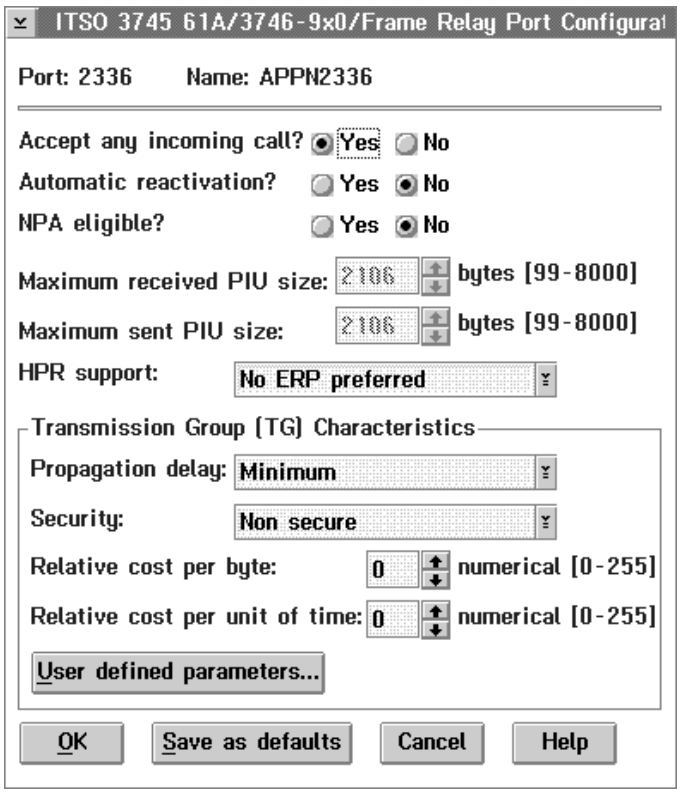

Figure 139. APPN Parameters for the Frame Relay Port

On port 2336, two DLCIs are defined. Figure 140 on page 171 shows the APPN parameters for DLCI 32. This dialog is reached by selecting the button **DLCI...** in Figure 138 on page 169. The DLCI number must be between 16 and 991 and unique on this port. The default is 32. We must define each DLCI that we expect to receive an APPN call in on.

The following definitions show how to pre-define an APPN link station. As previously discussed we are accepting any call in so this is not necessary for our test scenarios.

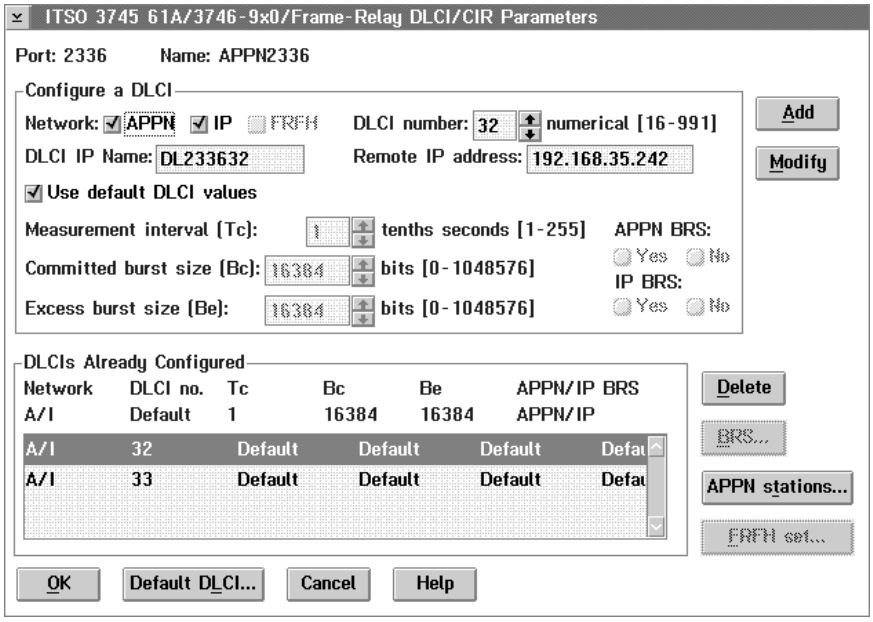

Figure 140. APPN Parameters for the Frame Relay DLCI 32

Figure 141 on page 172 shows the APPN station configuration dialog for frame relay DLCI 32. This dialog is reached by selecting the button **APPN stations** in Figure 140.

A single frame relay attached APPN station named ST233632 is configured. The frame format that is used for this station is defined as bridged frame format, and is used by boundary access node devices. The remote MAC address must be set to the MAC address of the station we will call out to.

If the frame format used for this station were defined as routed frame format (used by boundary network node devices), no remote MAC address can be defined. Frames that arrive at the remote BNN device are routed according to the DLCI used and remote SAP value specified.

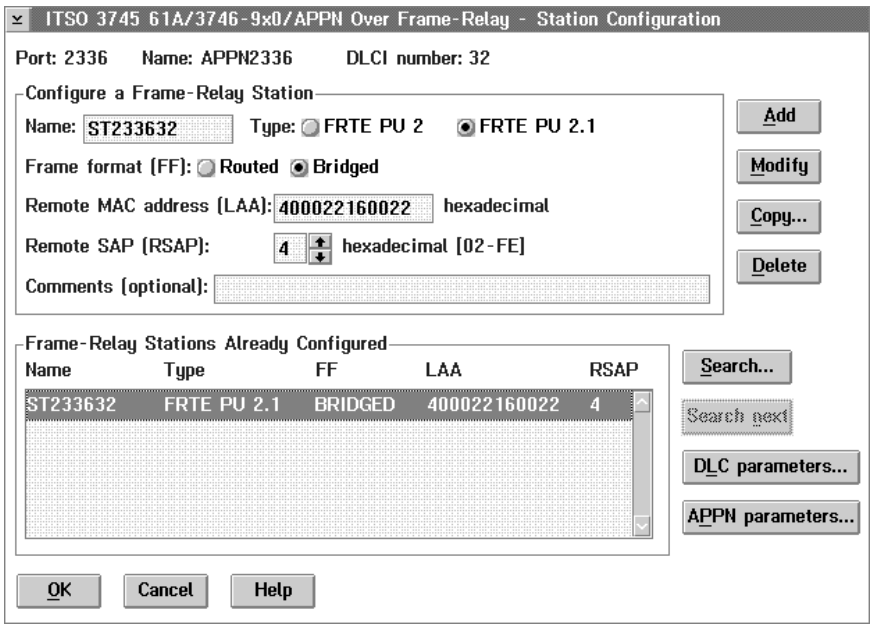

Figure 141. BAN APPN Station Configuration

Figure 142 shows APPN parameters for the pre-defined APPN station. This dialog is reached by selecting the button **APPN parameters...** in Figure 141. You can also specify multilink transmission group (MLTG), activation on demand (AOD), and dependent LU requester (DLUR) parameters here. A different DLUS (for this station) to the one defined as default (see Figure 139 on page 170) can also be specified here.

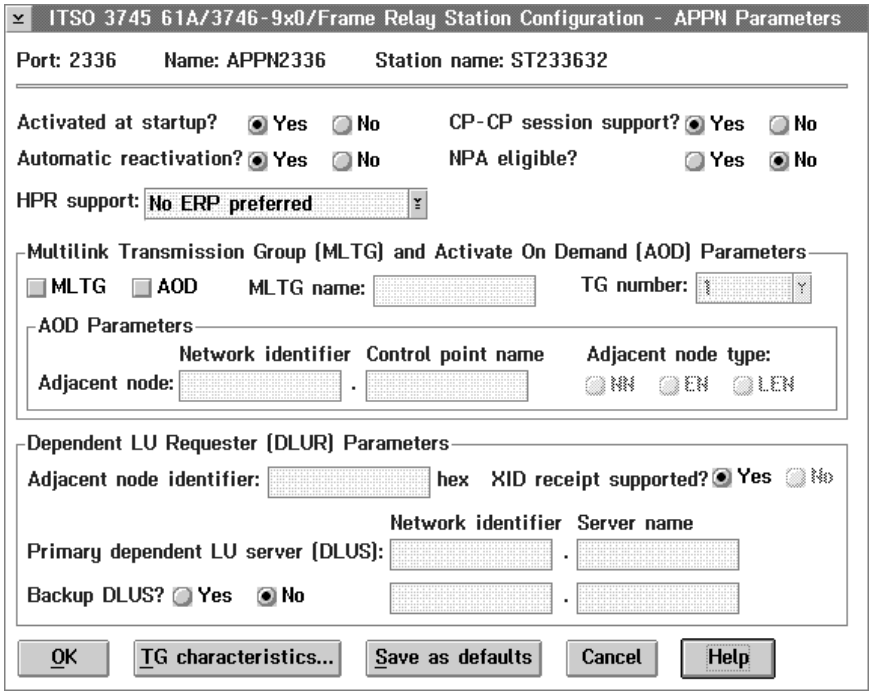

Figure 142. APPN Parameters for the Station

6.2.3 3746 ESCON Port 2176 APPN Definitions

ESCON port configuration and ESCON host link configuration dialogs are shown in Figure 143 and in Figure 144 on page 174. These dialogs are reached by selecting the icon **2176/ESCC** in Figure 137 on page 168 and by selecting the button **Host links...** in Figure 143.

In Figure 143, APPN must be selected, otherwise APPN definitions cannot be made in the dialogs that are used further on in the ESCON configuration process.

The ESCON port and host link must both be given APPN names.

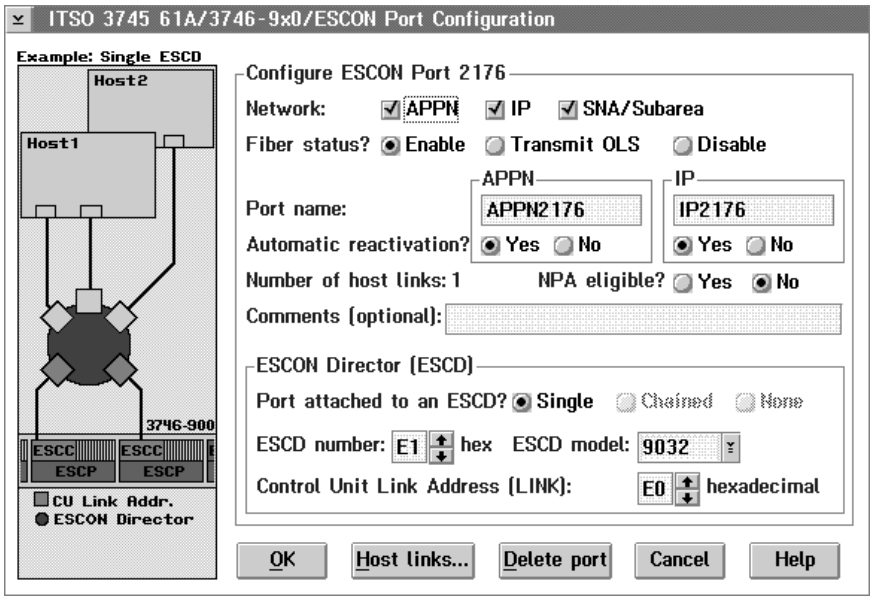

Figure 143. ESCON Port Configuration

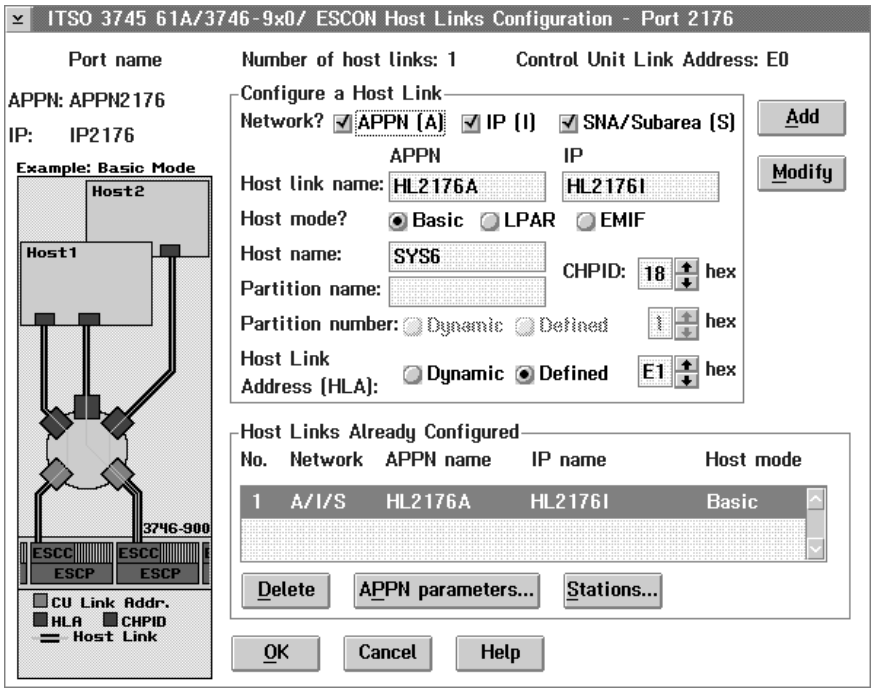

Figure 144. ESCON Host Link Configuration

Figure 145 shows the APPN parameters for the ESCON host link. This dialog is reached by selecting the button **APPN parameters...** in Figure 144. The parameters were left at the default settings.

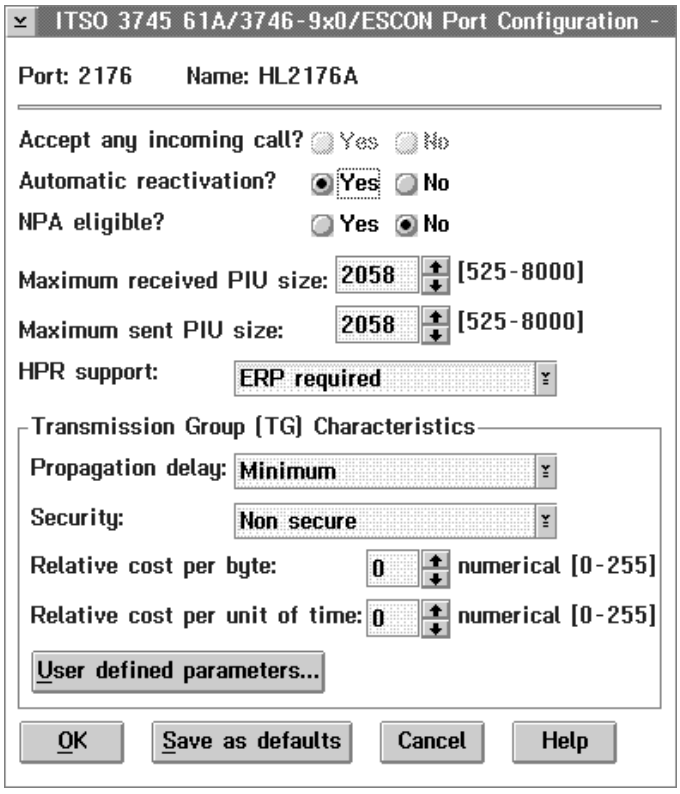

Figure 145. APPN Parameters for the ESCON Host Link

Figure 146 on page 175 shows the definitions for the ESCON station named ST92E that will be used by the 3746 APPN protocol stack. (See Figure 105 on page 145 for an overview of the stations used over ESCON.) The VTAM is defined as a Type 2.1 PU as it is an APPN network node. This dialog is reached by selecting the button **Stations ...** in Figure 144 on page 174.

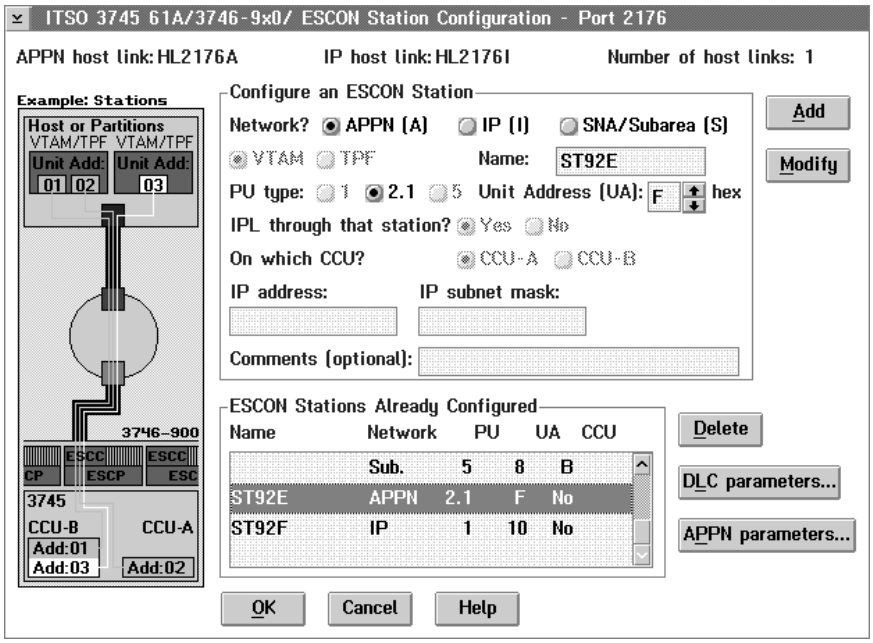

Figure 146. APPN Station Definition on the ESCON Host Link

6.2.4 3746 Token-Ring Port 2144 APPN Definitions

Figure 147 on page 176 shows the basic TIC3 token-ring port configuration. This dialog is reached by selecting the icon **2144/TIC3** in Figure 137 on page 168. The 3746 supports the sharing of a token-ring port by its APPN and IP protocol stacks, and the NCP.

The default value for the APPN local SAP is 8. NCP uses a SAP value of 4 for SNA subarea networking and this cannot be changed.

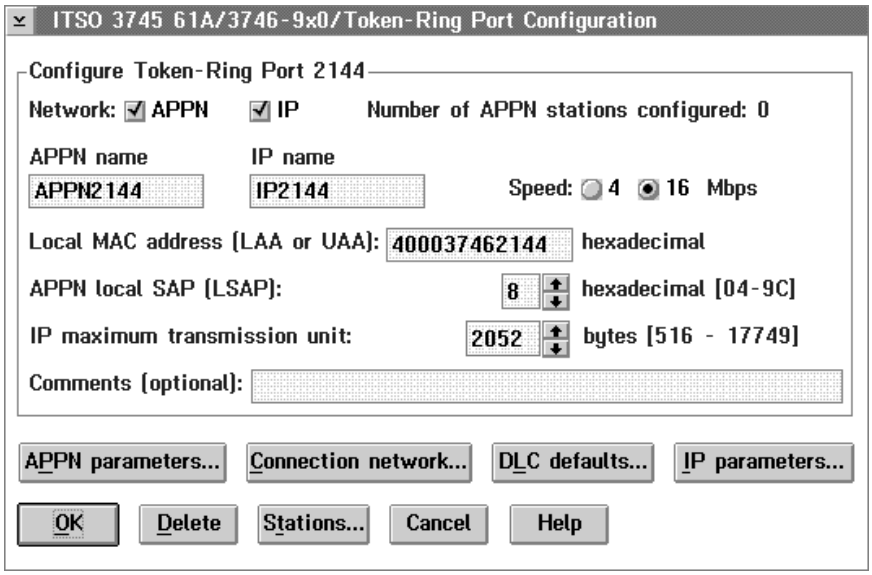

Figure 147. TIC3 Port Configuration

Figure 148 shows APPN parameters for the TIC3 token-ring port. This dialog is reached by selecting the button indicating **APPN parameters...** in Figure 147.

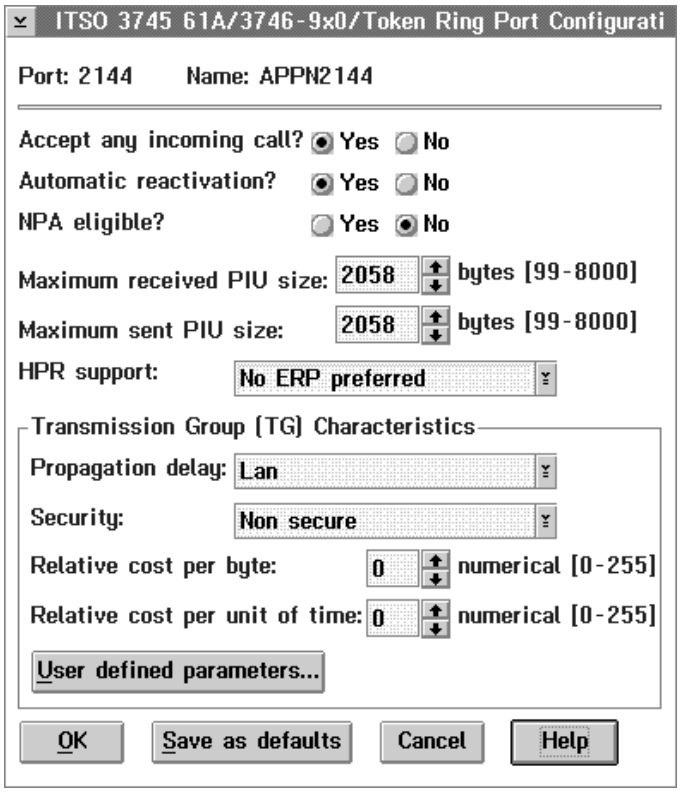

Figure 148. APPN Parameters for the Token-Ring Port

Figure 149 on page 177 shows a sample APPN station configuration for a token-ring network. This dialog is reached by selecting the button **Stations...** in Figure 147. As previously discussed the definition of a station is only necessary when the station should have parameters other than the default parameters, or when Accept any incoming call is not enabled, or when call out is required.

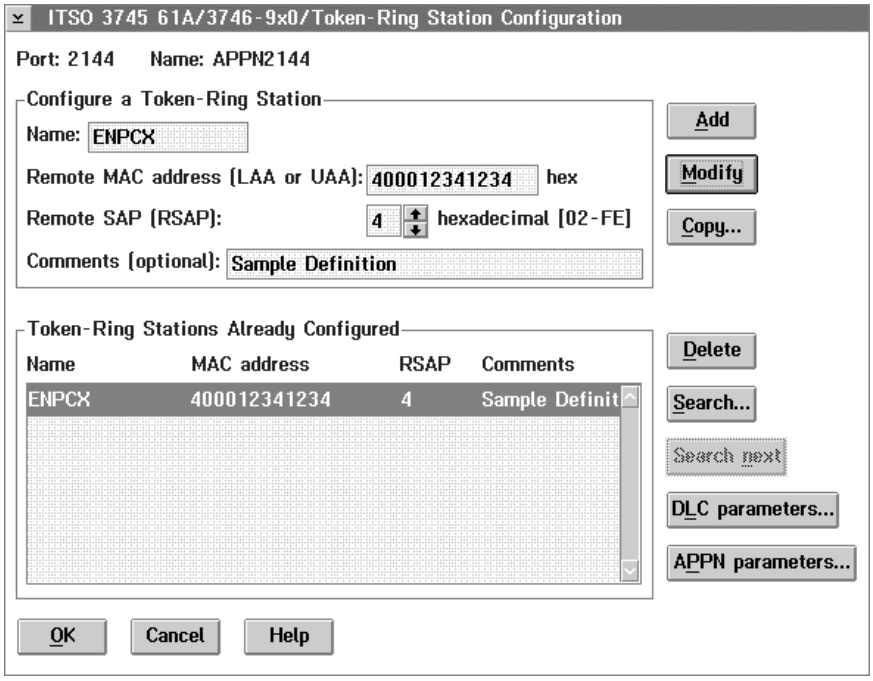

Figure 149. Sample APPN Station Configuration

6.2.5 Router 2216A BAN Definitions

In this scenario, we use the same hardware setup as in the previous scenario as far as the routers are concerned. This means that the TCP/IP configuration is the base for this scenario too. Therefore, only the additional configuration steps needed for BAN are explained in this section.

After starting the configuration tool, the previous configuration was loaded.

Refer to Figure 95 on page 135 for the frame relay interface definition window. We now need to configure frame relay BAN on the frame relay link that was previously defined. Click on the **BAN** tab to get to the BAN configuration page (see Figure 150 on page 178).

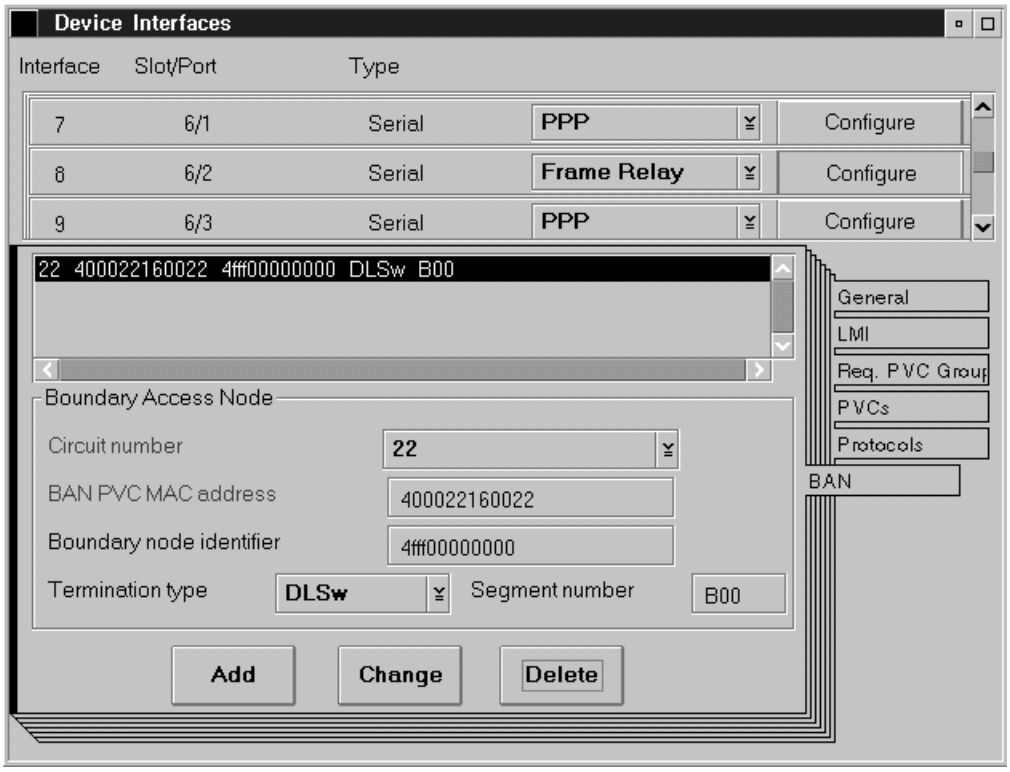

Figure 150. Device Interfaces: Frame Relay BAN

The definitions here include:

Circuit Number
This is the frame relay DLCI that is used to carry the BAN traffic.

BAN PVC MAC address
The endstation sends data to this address in the 2216. The data is then sent out on DLCI 22 to the boundary node identifier address shown here.

Boundary node identifier
This is the MAC address the BAN router sends data to across DLCI 22. The address must match the value defined in the 3746-900.

Termination type
The choices are Bridged (BAN1), or DLSw (BAN2).

Segment number
This is the segment number of the virtual LAN segment formed by the DLCI.

The next step is to configure DLSw. We must configure DLSw for our frame relay PVC (DLCI 22) as the APPN traffic from the 2210 router is transported by DLSw over the PPP link between the 2210 and 2216.

In the navigation screen, go down to the menu item **DLSw** and click on **General**.

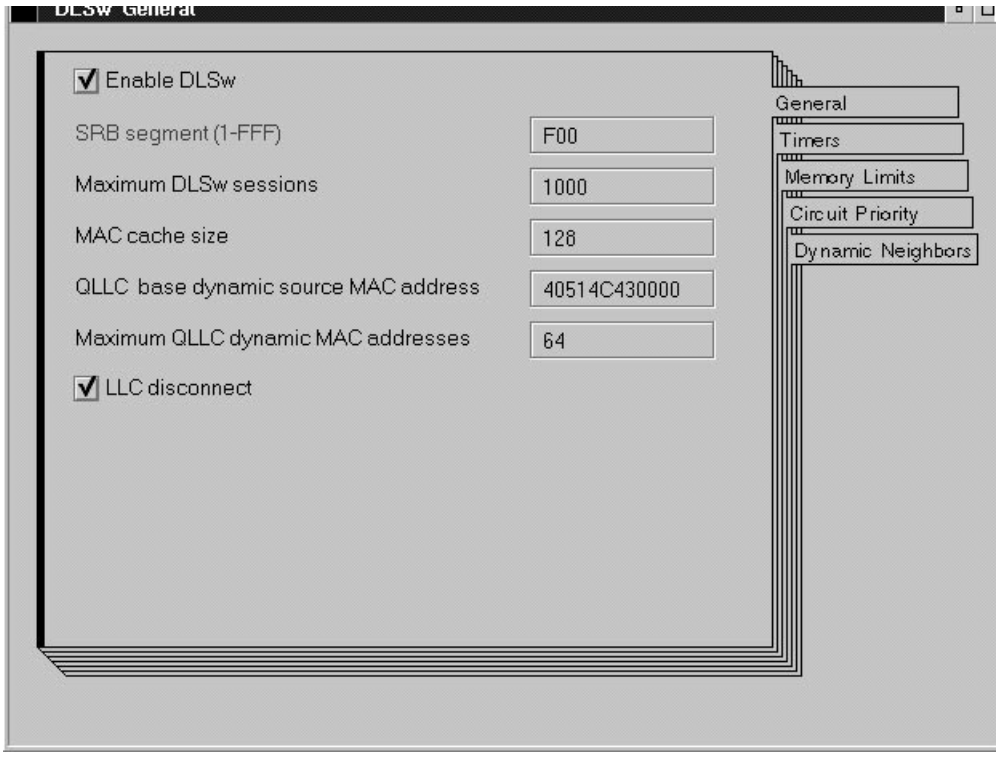

Figure 151. DLSw General

Figure 151 shows the DLSw notebook page. DLSw must be enabled here, and we must define the DLSw segment name. We used F00. Leave the other pages at the default. At this time, we are not using the option dynamic neighbors.

Next, select **DLSw TCP Connection** in the navigation menu.

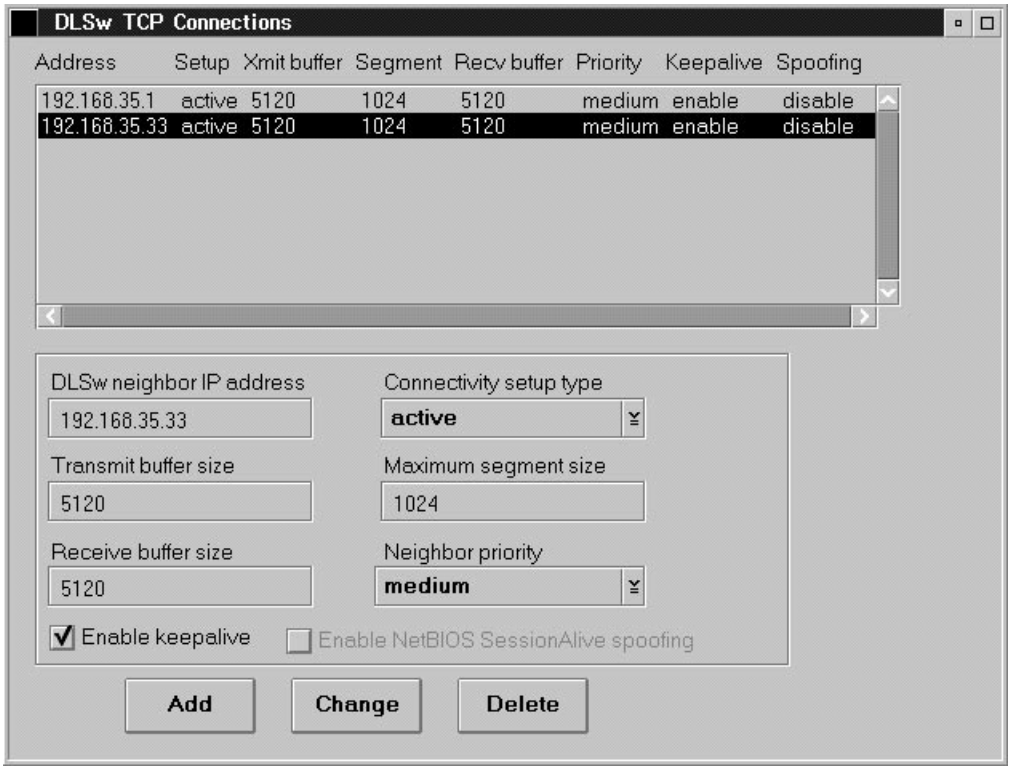

Figure 152. DLSw TCP Connection

As we are not using dynamic neighbors in DLSw, we need to define the DLSw neighbors here. The only true DLSw neighbor is the 2210B, so we added the internal IP address of the router 2210B (192.168.35.1).

In the case of BAN2, we need to define a TCP neighbor pointing to the 2216's own internal address too. The 2216 router's internal IP address is defined in the IP, General settings. It is an IP address that is not related to any particular interface. DLSw uses this local IP address to complete the DLSw connection.

Next, select **SAP LLC2 Customization** in the navigation menu (see Figure 153 on page 181). If you wish to customize any of the 802.2 LLC parameters for a specific SAP, this can be done from this notebook page.

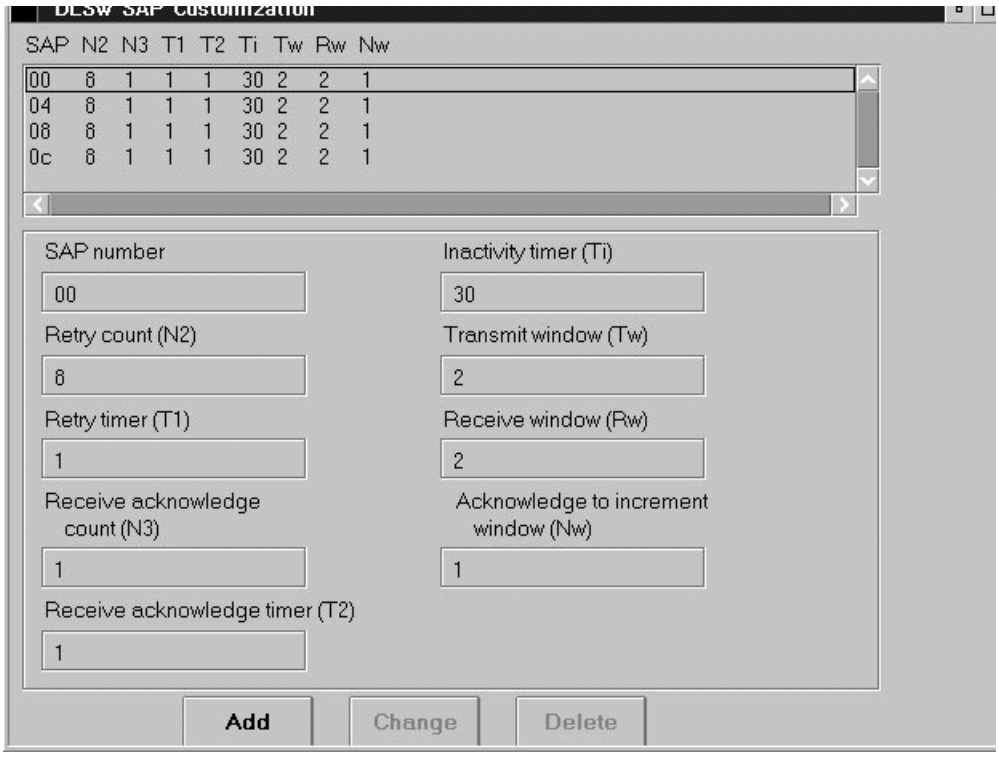

Figure 153. DLSw (LLC2) Customization

Next, we go down to DLSw Interfaces in the navigation menu. Add SAP 0, 4, 8, and C to the TR interface. This is done by selecting a SAP type of **SNA** (see Figure 154).

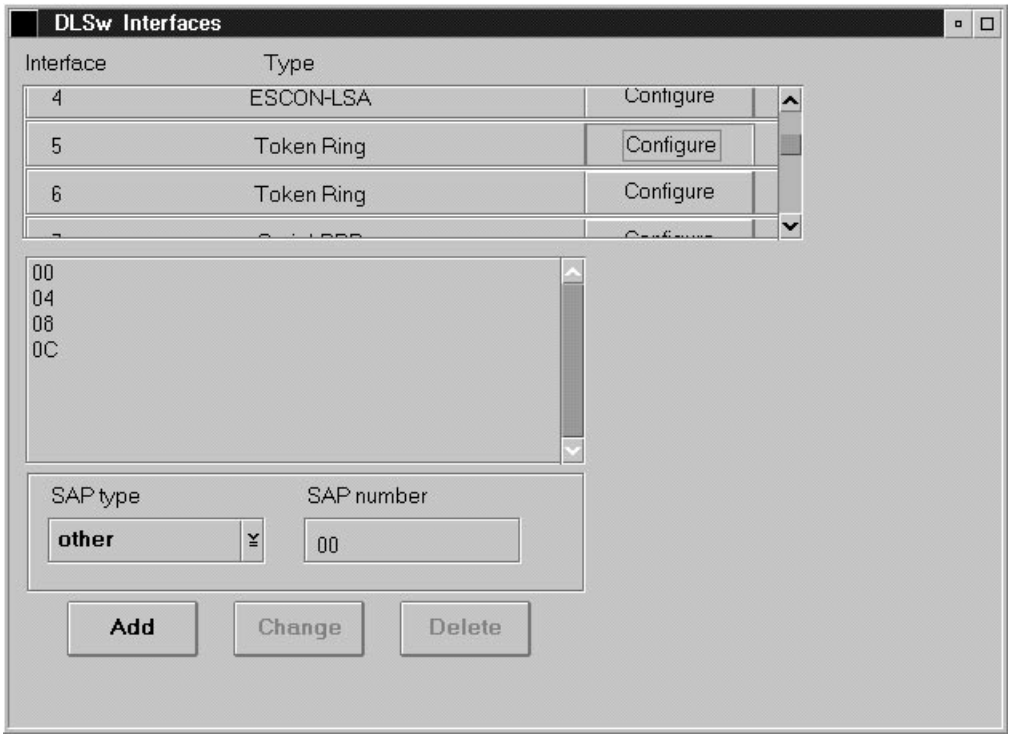

Figure 154. DLSw Interface TR Configure

At the frame relay interface, we need to add SAP 0, 4, 8, and C as well (see Figure 155 on page 182).

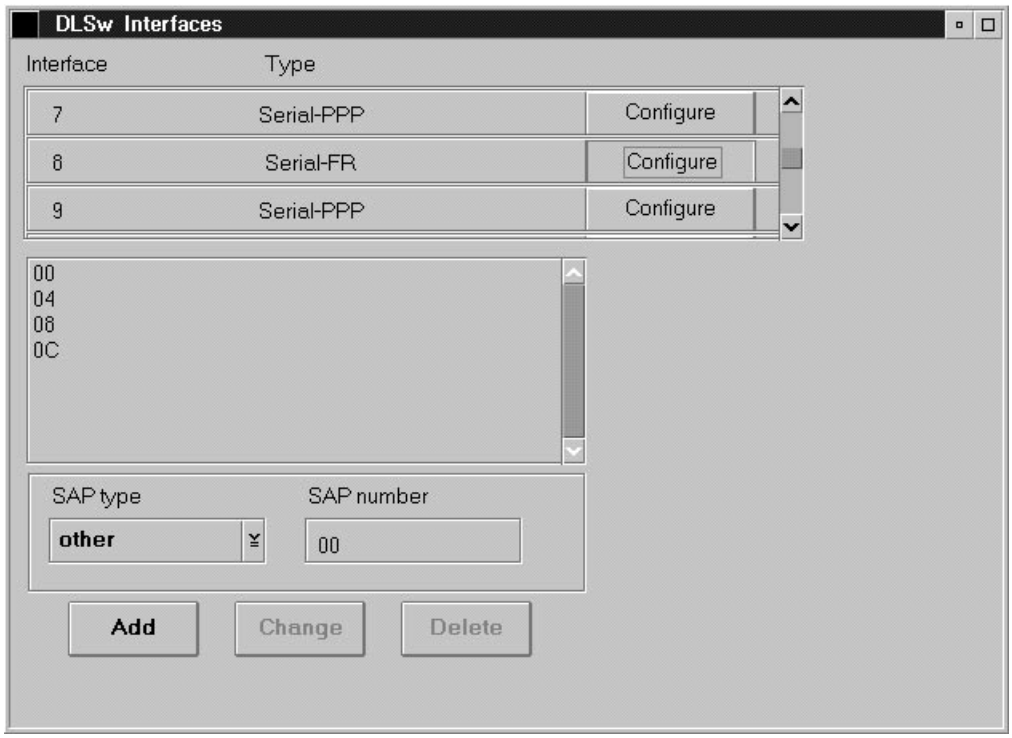

Figure 155. DLSw Interfaces Frame Relay

Next, we go down to **Bridging** in the navigation menu and click on **General**.

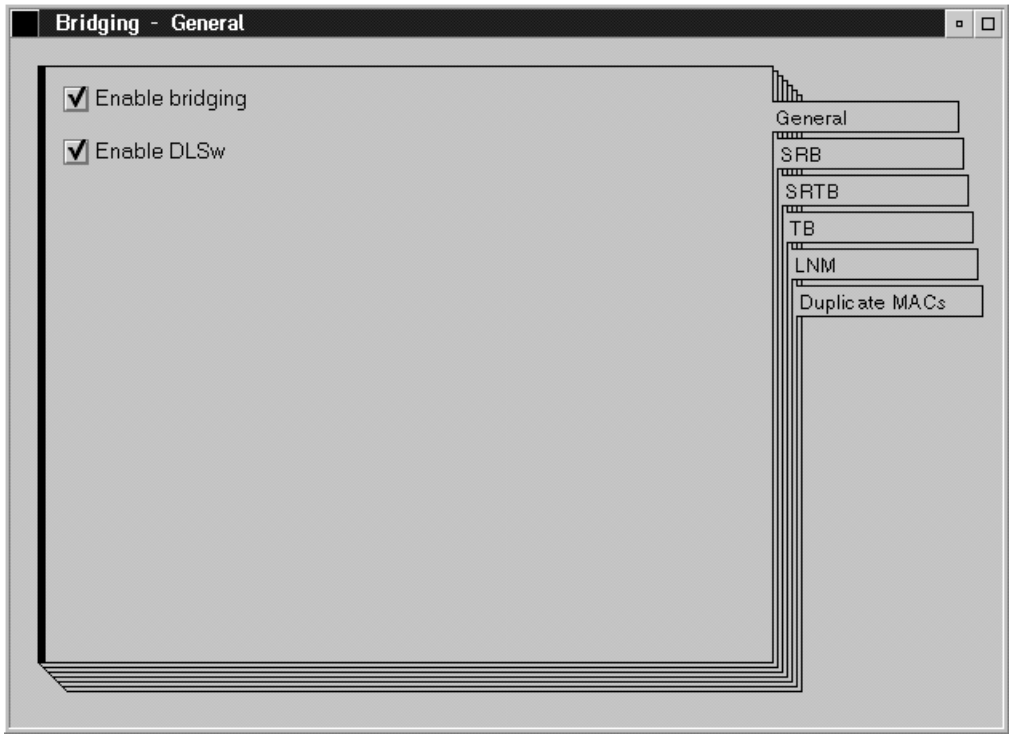

Figure 156. Bridging General

Bridging and DLSw must be enabled for BAN to work. Now select the tab **SRB**. Here we specify the bridge number and the internal virtual segment number. The router uses the internal virtual segment only if there are more than two ports defined for bridging. Leave the other fields at their default values.

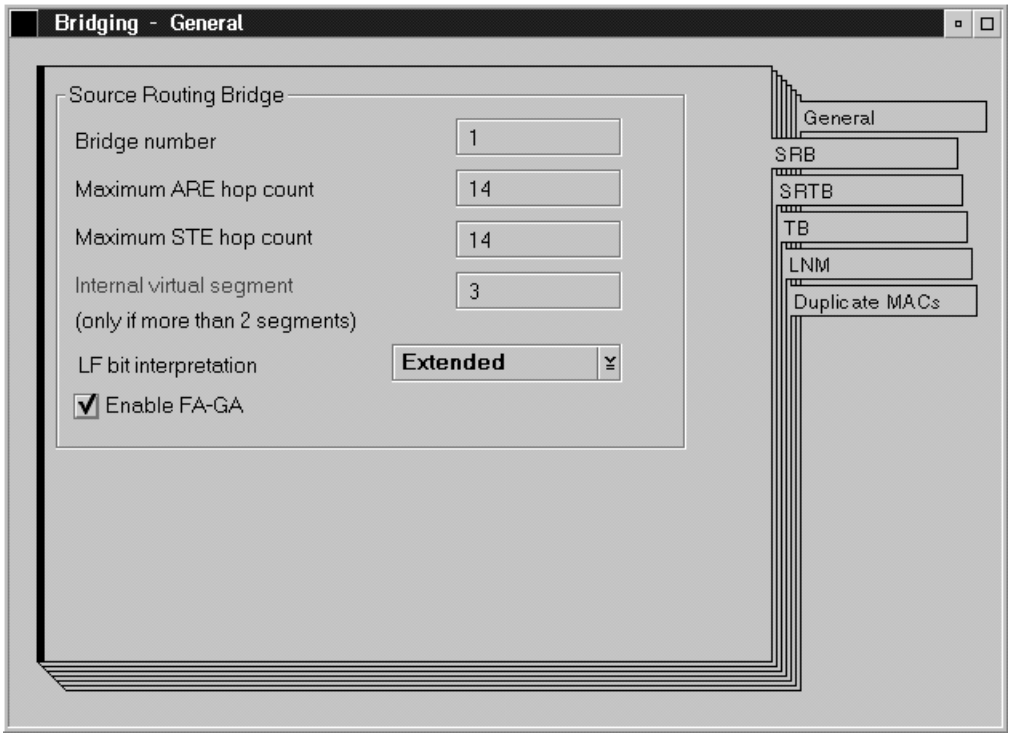

Figure 157. Bridging SRB

On the navigation menu select **Interfaces** under the menu item **Bridging**. Enable bridging on the token-ring interface and then select **Configure**.

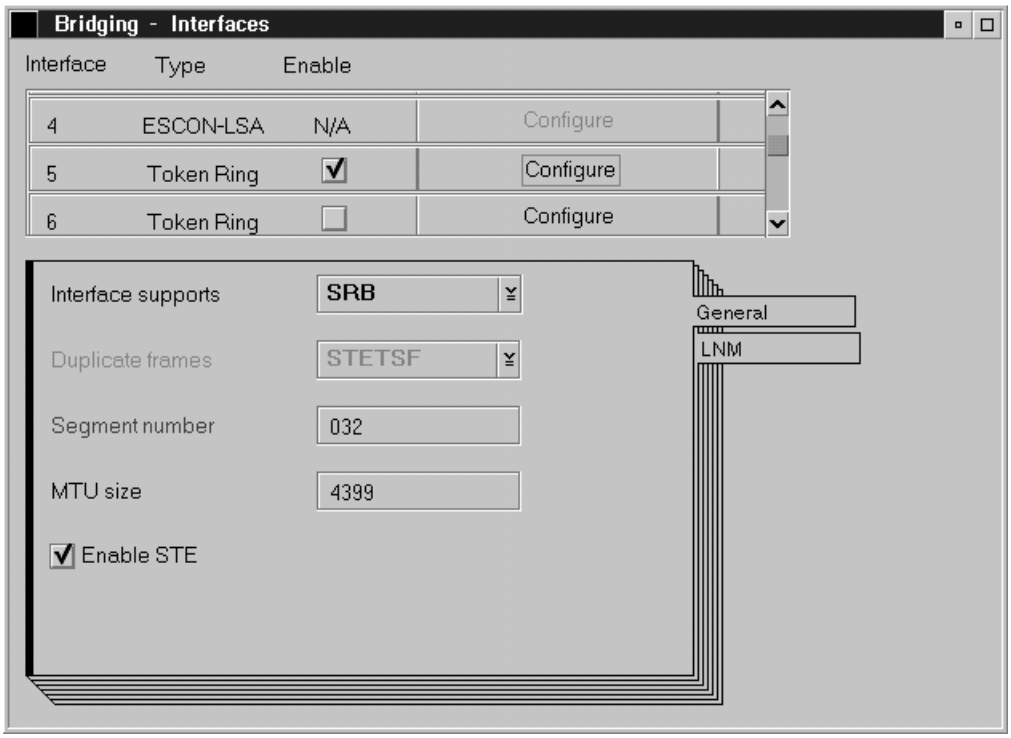

Figure 158. Bridging Interfaces Token-Ring

For the field Interface supports select **SRB**.

This concludes the 2216A configuration.

6.2.6 Router 2210B DLSw Definitions

In the frame relay BAN scenario we are using the same IP configuration as before. Therefore, only the additional steps needed to configure DLSw are described here.

After starting the configuration tool, the configuration used in the IP test scenario was loaded. At the navigation window, go down to the **DLSw** menu item and select **General**.

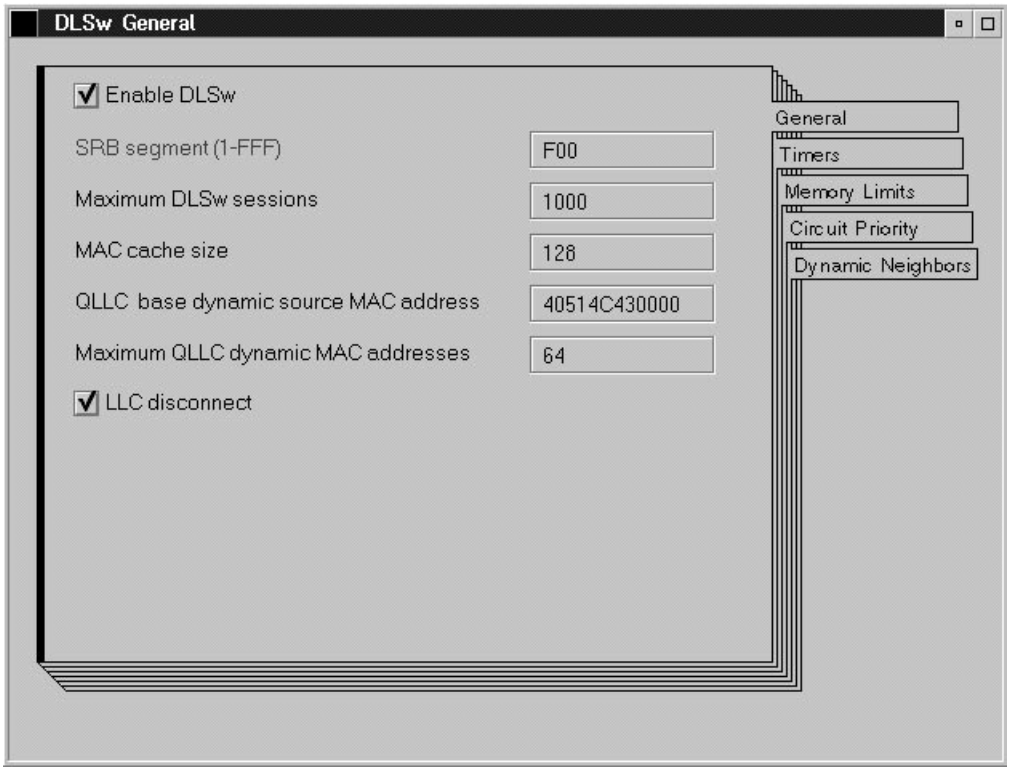

Figure 159. DLSw General

Only the DLSw segment must be defined here. Leave the other pages at their default. The DLSw segment number must match the segment number configured in the 2216A router. We used F00 (see Figure 151 on page 179). At this time, we are not using the option dynamic neighbors.

Next, select **DLSw TCP Connections** in the navigation menu.

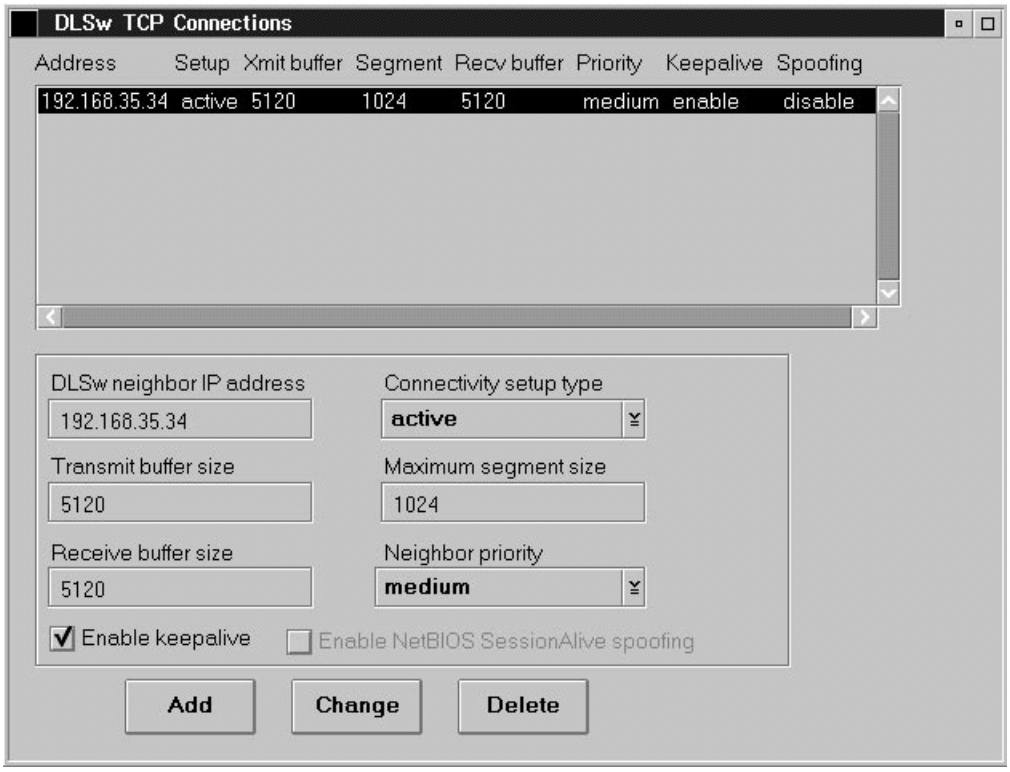

Figure 160. DLSw TCP Connection

The TCP connections notebook page is shown in Figure 160. Since we are not using dynamic neighbors, we must to define the DLSw neighbor or neighbors here. The only DLSw neighbor to the 2210 is the 2216A, so we added the 2216A's internal IP address (192.168.35.34).

Next, we go down to SAP LLC2 Connection in the navigation menu. Add SAP 0, 4, 8, and C to the LLC connection as shown in Figure 161 on page 187.

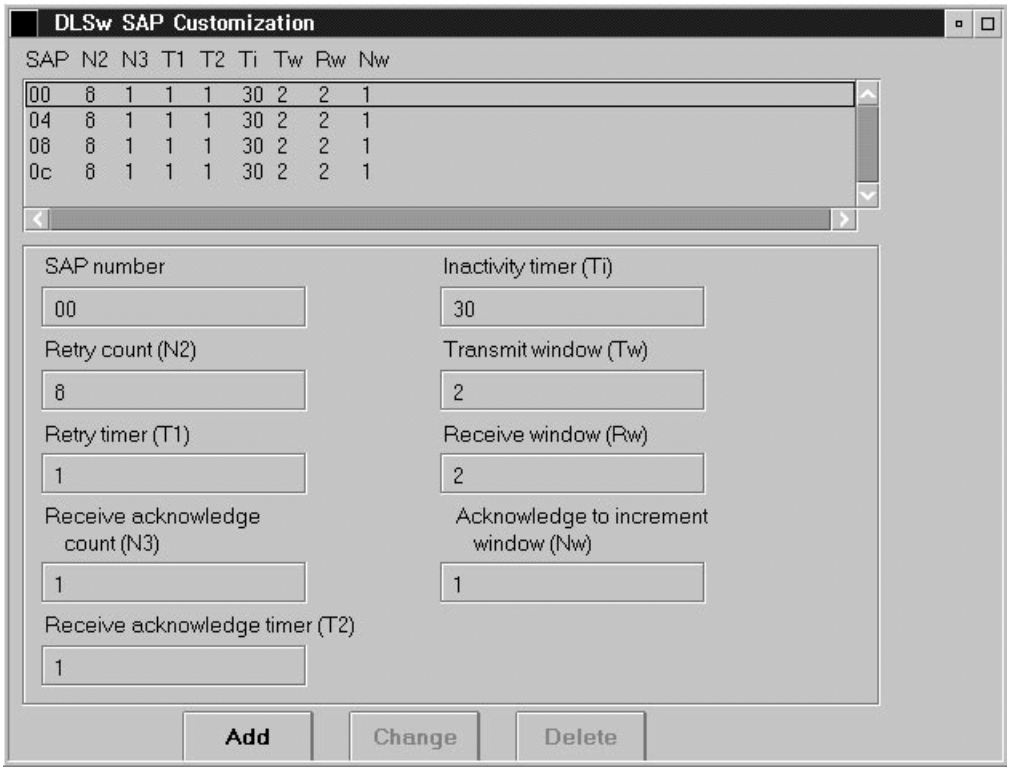

Figure 161. DLSw (LLC2) Customization

Next, we go down to DLSw Interfaces in the navigation menu. Click **Configure** beside the token-ring interface and add SAP 0, 4, 8, and C as shown in Figure 162.

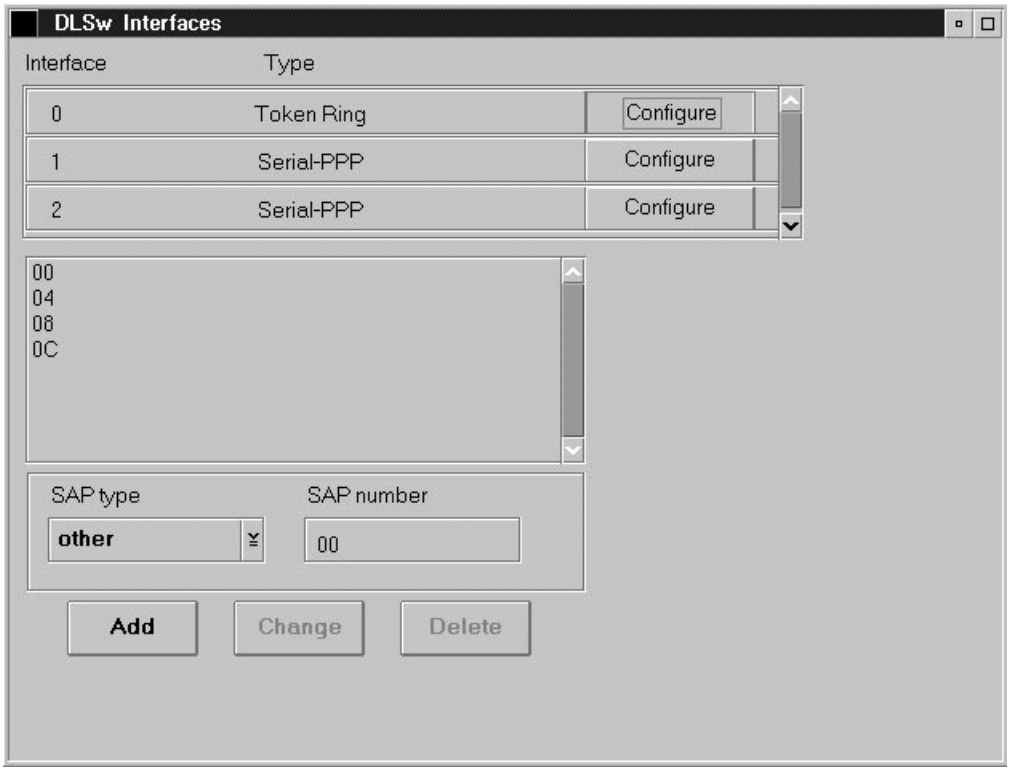

Figure 162. DLSw Interface TR Configure

Next, under Bridging in the navigation menu, select **General**. Bridging and DLSw must be enabled (see Figure 163 on page 188).

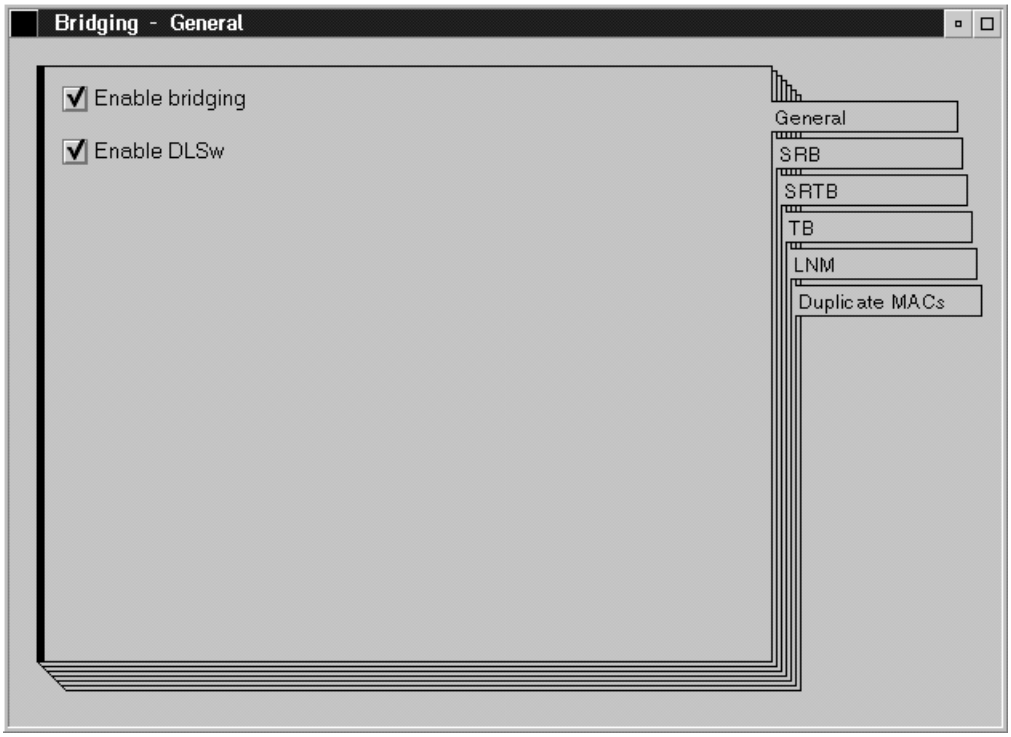

Figure 163. Bridging General

Now select the tab **SRB** (see Figure 164 on page 189). Here we specify the bridge number and the internal virtual segment number. The router uses the internal virtual segment only if there are more than two ports defined for bridging. Leave the other fields at their defaults.

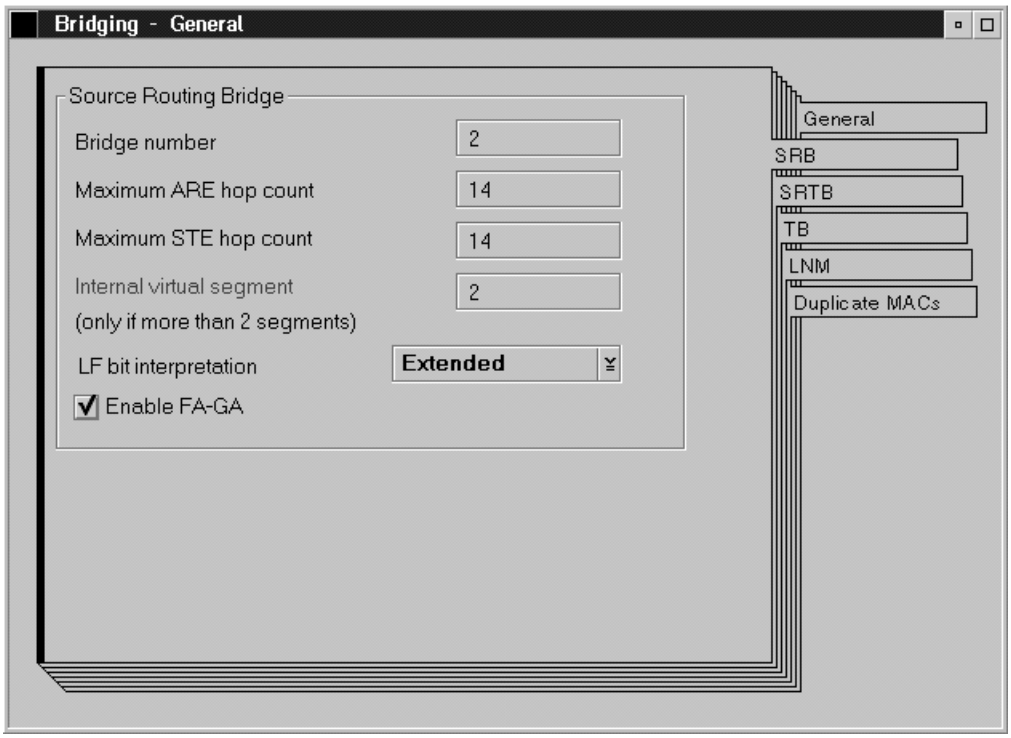

Figure 164. Bridging SRB

Next, under Bridging select **Interfaces**. Enable the token-ring interface for bridging and then click the **Configure** button (see Figure 165).

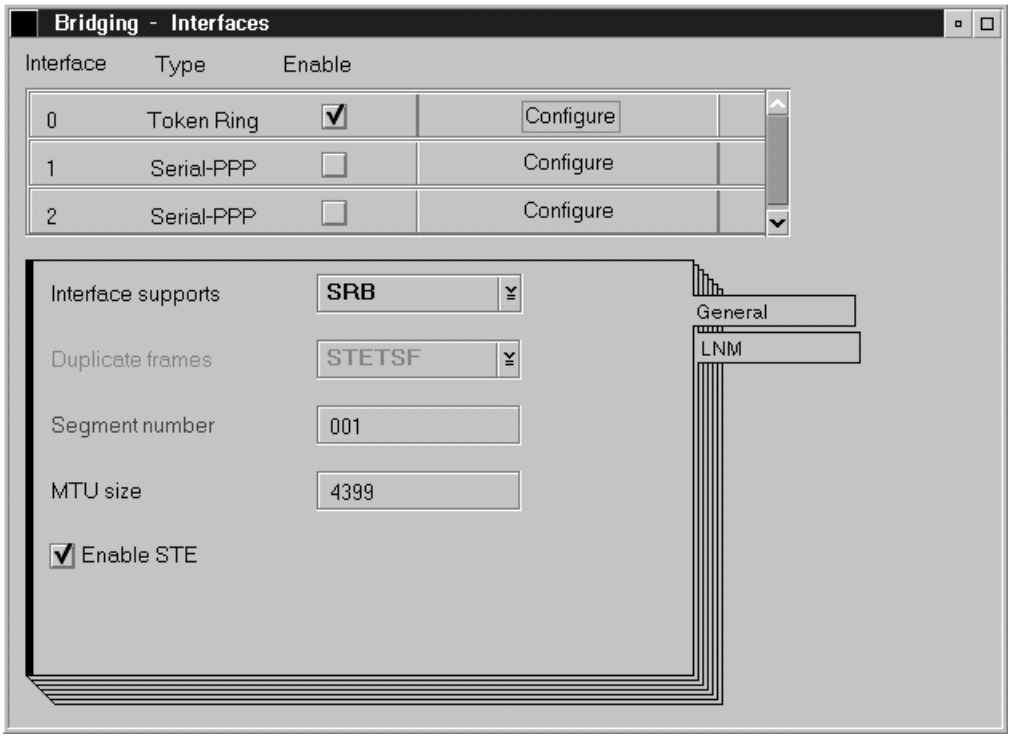

Figure 165. Bridging Interfaces Token-Ring

For the field Interface supports select **SRB**. Enter the segment number for the token-ring (001).

The next step is optional. We go back to the top of the navigation menu. Under the menu item **System**, select **General** (see Figure 166 on page 190).

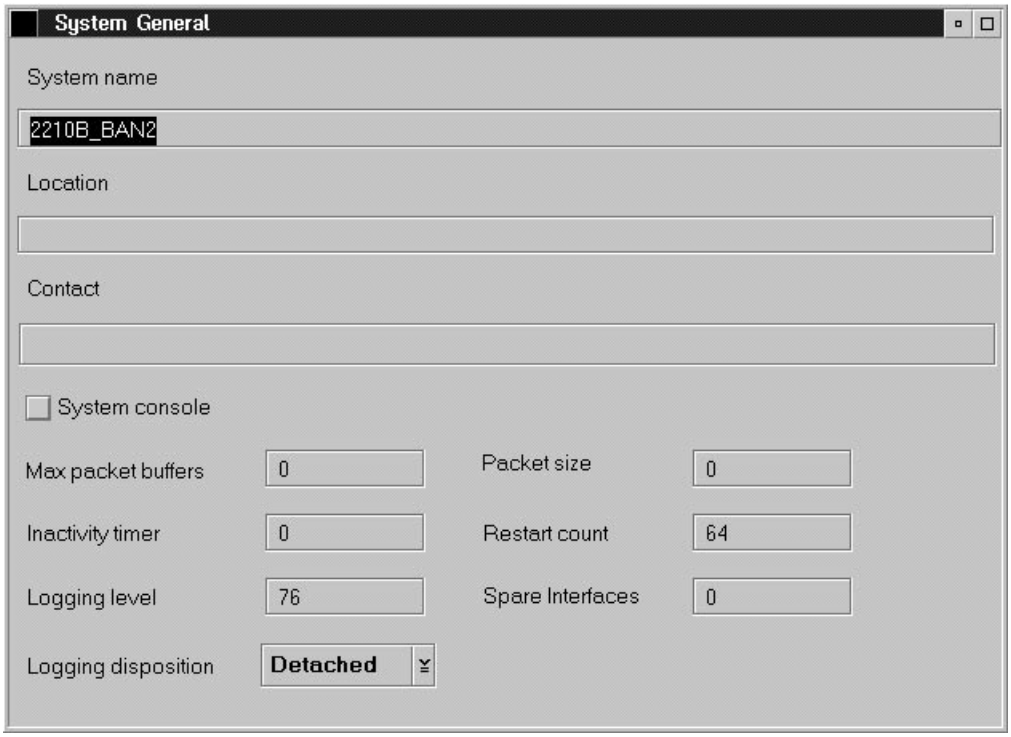

Figure 166. System General

On this page we can enter System Name, Location and Contact. The system name is the prompt string displayed by the router on a telnet or ASCII terminal connection.

That concludes the 2210B configuration.

6.3 Workstation Definitions

OS/2 Access Feature V5 is used for this scenario on the PS/2. The PS/2s are configured to work as APPN end nodes. There is one interface on each PS/2 that is attached to the token-ring network. The 3746 is the network node server for all PS/2s in this test scenario. As we are using a BAN connection between the 3746 and the router, all PS/2 end nodes are adjacent to the 3746 from the APPN point of view.

The OS/2 Access Feature provides the APPN configuration program. The following screens show the definition sequence on ENPC1.

Figure 167 on page 191 shows that we must select **APPC APIs over Token-Ring** as the communications type we define.

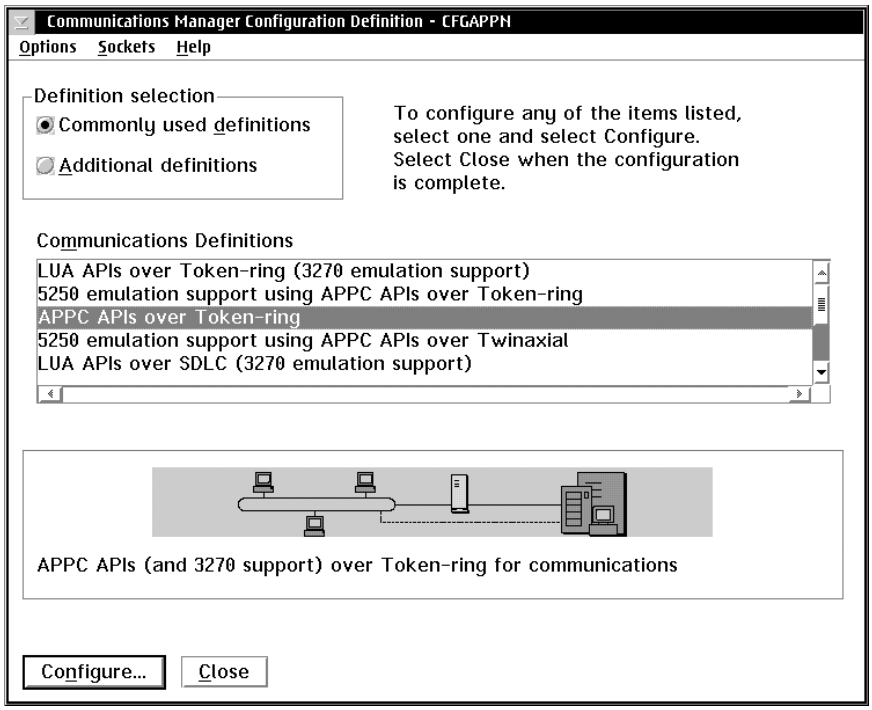

Figure 167. PS/2 Communication Configuration Type Definition

Figure 168 shows the primary APPN definitions. The workstation is an end node, named ENPC1. The APPN network is USIBMRA. The MAC address of the network node server is 400037462144. In this case, the PS/2 communicates directly via a token-ring network to the 3746, so the MAC address we defined is the MAC address of the 3746's TIC adapter port 2144.

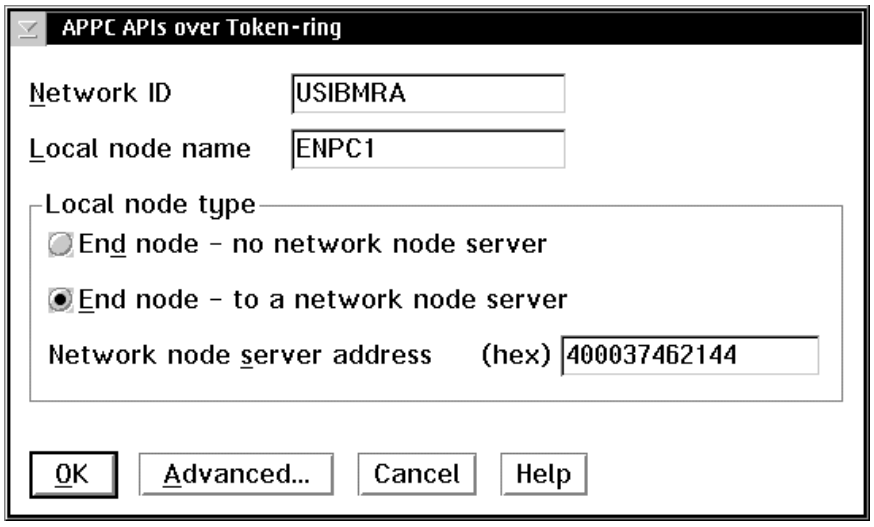

Figure 168. Primary APPN Definition

In the case of the other two workstations, the workstations communicate either to the 2216 directly through a token-ring network, or the workstation communicates to the 2210 (which connects to the 2216 using a PPP link controlled by DLSw) directly through a token-ring network. In both cases, the destination MAC address is the BAN PVC MAC address that was defined in Figure 150 on page 178. The MAC address used is 400022160022.

Figure 169 on page 192 shows several profiles that may be defined. Definitions in the optional profile SNA connections must be made if we wish to dial out.

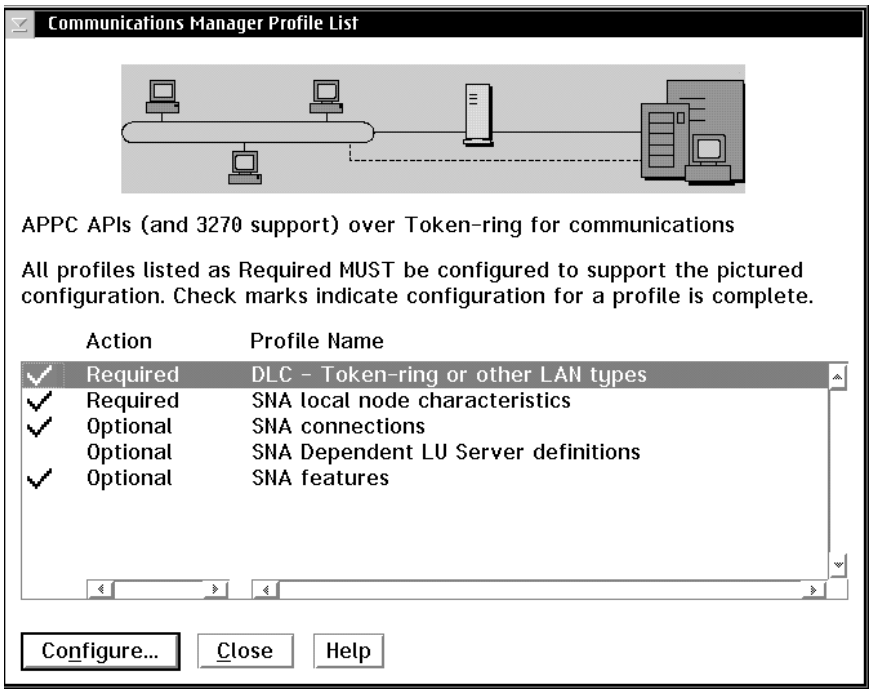

Figure 169. Communication Manager Profile List

To dial out we must pre-define a link to the adjacent node. We need the destination MAC address and the destination SAP we want to use. Figure 170 shows these parameters.

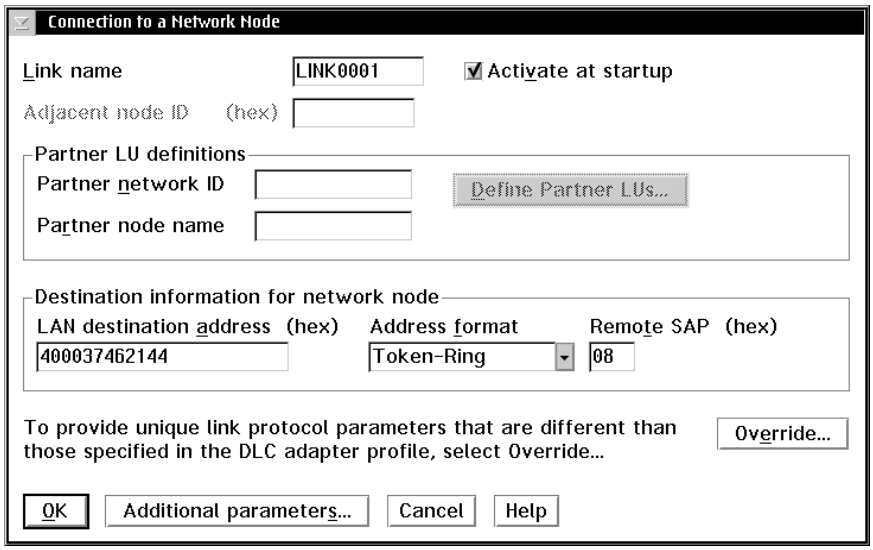

Figure 170. Destination Information for an Adjacent Node

6.4 Frame Relay BAN Details

In this frame relay BAN test scenario, we focus mostly on the 2216 BAN router. The 3746 knows nothing of BAN. It receives frames across a frame relay PVC from the 2216. The workstation connected to the 2210 knows nothing of BAN. The workstation connected to the 2216 via token-ring also knows nothing of the BAN function. What may change for these devices is the destination MAC address used on the workstations.

Each FR PVC used for BAN traffic can be seen as a virtual token-ring segment with virtual MAC addresses assigned on both ends. There can be, of course, multiple DLCIs carrying FR BAN traffic. In this case, there is a virtual LAN segment on each DLCI. The virtual MAC address at the 3746-900 side (BNI) is the same for all segments, whereas the router side can have different MAC addresses for each virtual segment or they can all be the same.

The 3746 allows the definition of a boundary node identifier (BNI) MAC address on a frame relay port. When using the bridged frame format, stations would normally use the BNI MAC address as the destination MAC address. When using the 2216 as a BAN router, the 2216 implements what is called a BAN PVC MAC address. This is an internal MAC address inside the 2216. Each BAN PVC originating at the 2216 has a BAN PVC MAC address defined. Frames with a destination MAC address equal to the BAN PVC MAC are forwarded onto the frame relay PVC. In addition, for each BAN PVC the BNI of the node at the other end of the PVC must be defined. See Figure 150 on page 178 for an example of these definitions. Frames sent to the BAN PVC MAC address on the 2216 that have the destination MAC address replaced by the BNI MAC address, are encapsulated according to RFC1490, and then forwarded onto the PVC. The 3746 receives frames with a destination MAC equal to its BNI MAC address.

When using frame relay BAN1, the BAN router consumes less resources, but the device at the other end of the frame relay PVC will have more processing to do as it must terminate all LLC connections.

In our test scenario, we are using DLSw between the routers to transport our SNA traffic. This is recommended as it ensures session integrity especially in cases of high line utilization. It would be possible to use pure bridging between the 2210 and 2216, but this is not recommended for SNA traffic. Whether DLSw or bridging is used, there are consequences for the BAN configuration of the 2216. The following is dictated:

- When DLSw is used between the 2210 and the 2216, stations off the 2210 can only use BAN2.

  DLSw by nature terminates the LLC of the SNA sessions. Therefore, it is not possible to have a single LLC session end-to-end. For this reason, DLSw traffic must used BAN2 as it provides the LLC session over the frame relay PVC.

- When pure bridging is defined between the two routers, all stations can use BAN1 or BAN2.

To determine whether frame relay BAN1 or BAN2 will be used on a PVC, the termination type of either bridged (BAN1) or DLSw (BAN2) must be selected, this can be seen in Figure 150 on page 178.

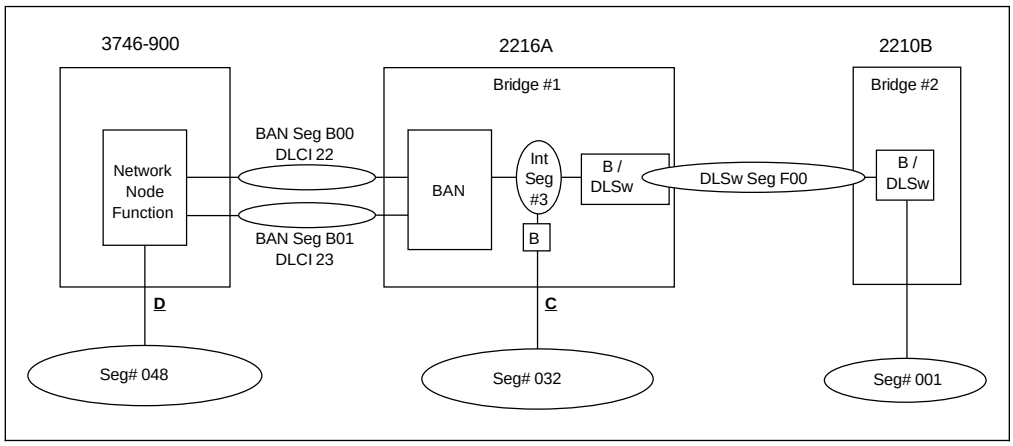

Figure 171. Frame Relay BAN Logical View

For further clarification of the LLC sessions, the following information was displayed at the 2216 and 2210 routers. From the 2216 TELNET command line interface, go to the protocol DLSw and enter the command `list llc session all`. The following information is displayed:

```
2216A_BAN2 DLSw>list llc ses all

    SAP  Int. Remote Addr    Local Address   State       RIF
1.  04    22  400022160022   400052005128    CONTACTED   081E F001 0031 B000  3
2.  04    22  400022160022   40007E108440    CONTACTED   081E F001 0031 B000  4
3.  08    5   40007E108440   400022160022    CONTACTED   089E 0321 0031 F000  2

2216A_BAN2 DLSw>
```

Figure 172. BAN Segment Display

In order to interpret the RIF field, please refer to Figure 171. Here is a summary of all segments:

- F00 is the DLSw segment (2210 and 2216).
- B00 is the BAN PVC segment (PVC 22).
- 003 is the 2216's internal segment.
- 032 is the 2216's token-ring segment.
- 001 is the 2210's token-ring segment.

The following is an explanation of Figure 172, and Figure 173 on page 195:

- **3** LLC session of the workstation ENPC3 from the DLSw function in the 2210 to the BAN PVC segment in the 2216.

  The LLC session from the workstation over the token-ring to the 2210 (**1**) is not visible at the 2216 as it is terminated by the 2210 DLSw function.
- **4** LLC session of the workstation ENPC1 from DLSw to the BAN PVC segment.
- **2** LLC session of the workstation ENPC1 from token-ring to DLSw segment.

In order to see the first LLC session from the workstation ENPC3, we need to issue the same command on the 2210 router.

```
2210B_BAN2 DLSw>list llc2 ses all
    SAP   Int. Remote Addr    Local Address   State        RIF
1.  04    0   40007E108440    400022160022    CONNECTED    069E 0012 F000 1
2210B_BAN2 DLSw>
```

Figure 173 shows the LLC sessions we previously discussed:

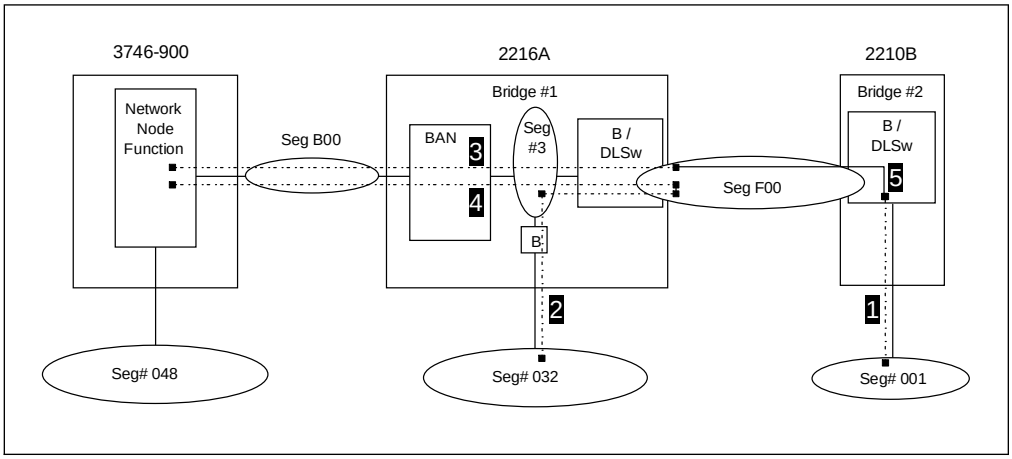

Figure 173. Frame Relay BAN LLC Sessions

On the 2216 we can also display the BAN segment by issuing the following commands:

```
Talk 5
Protocol ASRT
BAN
```

This returns the following data:

```
2216A_BAN2 BAN>list
bridge  BAN                 Boundary            bridged or
port    DLCI MAC Address    Node Identifier     DLSw terminated  Status
3       40:00:22:16:00:22   4F:FF:00:00:00:00   terminated       Up
2216A_BAN2 BAN>
```

6.5 Frame Relay BAN Details Displays

The following section shows various management displays of the frame relay BAN test network.

6.5.1 VTAM APPN Displays

The VTAM display for the local SNA major node is shown in Figure 174.

```
D NET,ID=NN061AXX,E

IST097I DISPLAY ACCEPTED
IST075I NAME = NN061AXX, TYPE = LCL SNA MAJ NODE
IST486I STATUS= ACTIV, DESIRED STATE= ACTIV
IST084I NETWORK RESOURCES:
IST089I CP90061A TYPE = PU_T2.1          , ACTIV--L-- ,CUA=092E
IST1500I STATE TRACE = OFF
IST314I END
```

Figure 174. VTAM Display of NN061AXX

The VTAM display for the PU in the local SNA major node is shown in Figure 175.

```
D NET,ID=CP90061A,E

IST097I DISPLAY ACCEPTED
IST075I NAME = CP90061A, TYPE = PU_T2.1
IST486I STATUS= ACTIV--L--, DESIRED STATE= ACTIV
IST1043I CP NAME = NN061A, CP NETID = USIBMRA, DYNAMIC LU = YES
IST1589I XNETALS = YES
IST1105I RESOURCE STATUS TGN CP-CP TG CHARACTERISTICS
IST1106I CP90061A AC/R    21 YES   988D000000000000000000014C00808080
IST1482I HPR = RTP - OVERRIDE = N/A - CONNECTION = YES
IST1510I LLERP = REQUIRED - RECEIVED = REQUIRED
IST136I LOCAL SNA MAJOR NODE = NN061AXX
IST077I SIO = 08518 CUA = 092E
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST355I LOGICAL UNITS:
IST080I ENPC2    ACT/S----Y NN061A   ACT/S----Y
IST314I END
```

Figure 175. VTAM Display of CP90061A

The VTAM display for the 3746-900 NNP network node is shown in Figure 176 on page 197.

```
D NET,ID=NN061A,E

IST097I DISPLAY ACCEPTED
IST075I NAME = USIBMRA.NN061A, TYPE = ADJACENT CP
IST486I STATUS= ACT/S----Y, DESIRED STATE= ACTIV
IST1402I SRTIMER =   120  SRCOUNT =    60
IST1447I REGISTRATION TYPE = NO
IST977I MDLTAB=***NA*** ASLTAB=***NA***
IST1333I ADJLIST = ***NA***
IST861I MODETAB=***NA*** USSTAB=***NA*** LOGTAB=***NA***
IST934I DLOGMOD=CPSVCMG USS LANGTAB=***NA***
IST597I CAPABILITY-PLU ENABLED  ,SLU ENABLED  ,SESSION LIMIT NONE
IST231I CDRSC MAJOR NODE = ISTCDRDY
IST1184I CPNAME = USIBMRA.NN061A - NETSRVR = ***NA***
IST1044I ALSLIST = ISTAPNPU
IST082I DEVTYPE = INDEPENDENT LU / CDRSC
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST171I ACTIVE SESSIONS = 0000000003, SESSION REQUESTS = 0000000000
IST206I SESSIONS:
IST1081I ADJACENT LINK STATION = CP90061A
IST634I NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I RAK      ACTIV/CP-S D2333BDF45F899C8 0DEE 0001       USIBMRA
IST635I RAK      ACTIV/SV-P F8D3D16428ED8710 0002 0002  0  0 USIBMRA
IST635I RAK      ACTIV/CP-P F8D3D16428ED8702 0001 0DF9  0  0 USIBMRA
IST924I -------------------------------------------------------------
IST075I NAME = USIBMRA.NN061A, TYPE = DIRECTORY ENTRY
IST1186I DIRECTORY ENTRY = DYNAMIC NN
IST1184I CPNAME = USIBMRA.NN061A - NETSRVR = ***NA***
IST1402I SRTIMER =   120  SRCOUNT =    60
IST314I END
```

Figure 176. VTAM Display of NN061A

The VTAM display for the PS/2 end node LU is shown in Figure 177 on page 198.

```
D NET,ID=ENPC2,E

IST097I DISPLAY ACCEPTED
IST075I NAME = USIBMRA.ENPC2, TYPE = CDRSC
IST486I STATUS= ACT/S----Y, DESIRED STATE= ACTIV
IST1402I SRTIMER =   120  SRCOUNT =    60
IST1447I REGISTRATION TYPE = NO
IST977I MDLTAB=***NA*** ASLTAB=***NA***
IST1333I ADJLIST = ***NA***
IST861I MODETAB=***NA*** USSTAB=***NA*** LOGTAB=***NA***
IST934I DLOGMOD=***NA*** USS LANGTAB=***NA***
IST597I CAPABILITY-PLU ENABLED  ,SLU ENABLED  ,SESSION LIMIT NONE
IST231I CDRSC MAJOR NODE = ISTCDRDY
IST479I CDRM NAME = RAK, VERIFY OWNER = NO
IST1184I CPNAME = USIBMRA.ENPC2 - NETSRVR = ***NA***
IST1044I ALSLIST = ISTAPNPU
IST082I DEVTYPE = INDEPENDENT LU / CDRSC
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST171I ACTIVE SESSIONS = 0000000002, SESSION REQUESTS = 0000000000
IST206I SESSIONS:
IST1081I ADJACENT LINK STATION = CP90061A
IST634I NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I RAK      ACTIV-S    EDBF2F5F2788494F 0004 0003  0  0 USIBMRA
IST635I RAK      ACTIV/SV-S EDBF2F5F2688494F 0001 0001  0  0 USIBMRA
IST314I END
```

Figure 177. VTAM Display of ENPC2

VTAM APING display samples are shown in Figure 178 and Figure 179 on page 199.

```
D NET,APING,ID=NN061A

IST097I DISPLAY ACCEPTED
IST1489I APING SESSION INFORMATION
IST1490I DLU=USIBMRA.NN061A SID=F8D3D16428ED8D79
IST933I LOGMODE=#INTER  , COS=*BLANK*
IST875I APPNCOS TOWARDS SLU = #INTER
IST1460I TGN  CPNAME             TG TYPE      HPR
IST1461I  21  USIBMRA.NN061A     APPN         ANR
IST314I END
IST1457I VTAM APING VERSION 2R33 (PARTNER TP VERSION 2R33)
IST1490I DLU=USIBMRA.NN061A SID=F8D3D16428ED8D79
IST1462I ECHO IS ON
IST1463I ALLOCATION DURATION: 126 MILLISECONDS
IST1464I PROGRAM STARTUP AND VERSION EXCHANGE: 367 MILLISECONDS
IST1465I         DURATION      DATA SENT   DATA RATE   DATA RATE
IST1466I      (MILLISECONDS)    (BYTES)  (KBYTE/SEC)  (MBIT/SEC)
IST1467I              124           200           1           0
IST1467I              114           200           1           0
IST1468I TOTALS:      238           400           1           0
IST1469I DURATION STATISTICS:
IST1470I MINIMUM = 114 AVERAGE = 119 MAXIMUM = 124
IST314I END
```

Figure 178. APING to NN061A

```
D NET,APING,ID=ENPC2

IST097I DISPLAY ACCEPTED
IST1457I VTAM APING VERSION 2R33 (PARTNER TP VERSION 2R43)
IST1490I DLU=USIBMRA.ENPC2 SID=EDBF2F5F2788494F
IST1462I ECHO IS ON
IST1463I ALLOCATION DURATION: 50 MILLISECONDS
IST1464I PROGRAM STARTUP AND VERSION EXCHANGE: 369 MILLISECONDS
IST1465I         DURATION      DATA SENT   DATA RATE   DATA RATE
IST1466I      (MILLISECONDS)    (BYTES)  (KBYTE/SEC)  (MBIT/SEC)
IST1467I               29           200           6           0
IST1467I               28           200           7           0
IST1468I TOTALS:       57           400           7           0
IST1469I DURATION STATISTICS:
IST1470I MINIMUM = 28 AVERAGE = 28 MAXIMUM = 29
IST314I END
```

Figure 179. APING to ENPC2

6.5.2 3746-900 CCM APPN Management Displays

Figure 180 shows the CCM APPN management menu options.

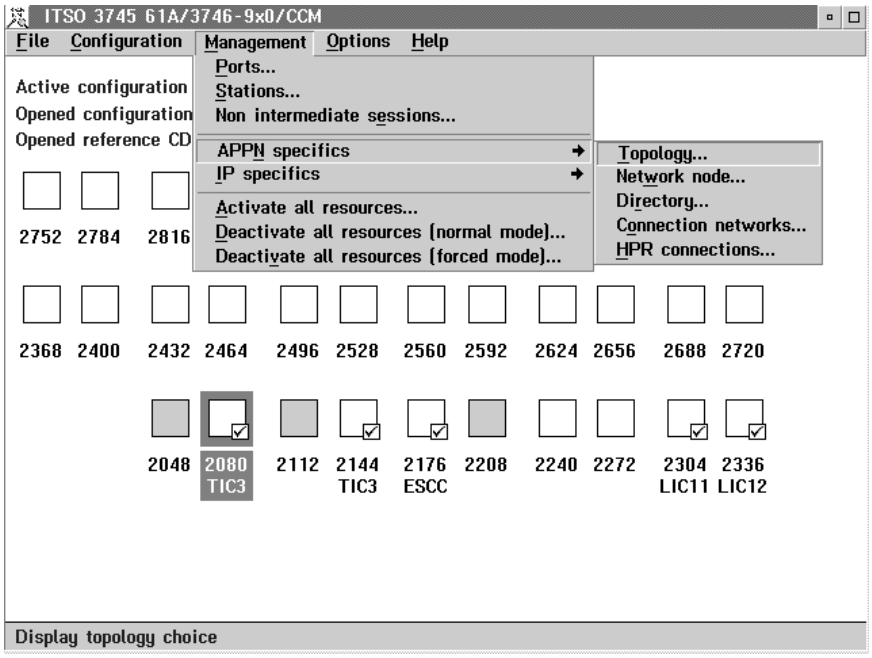

Figure 180. CCM Display Tags for APPN Protocol Management

Figure 181 on page 200 shows the ports management display. This display is reached by selecting the **Ports...** menu item in Figure 180. In this figure the current status for all ports is shown.

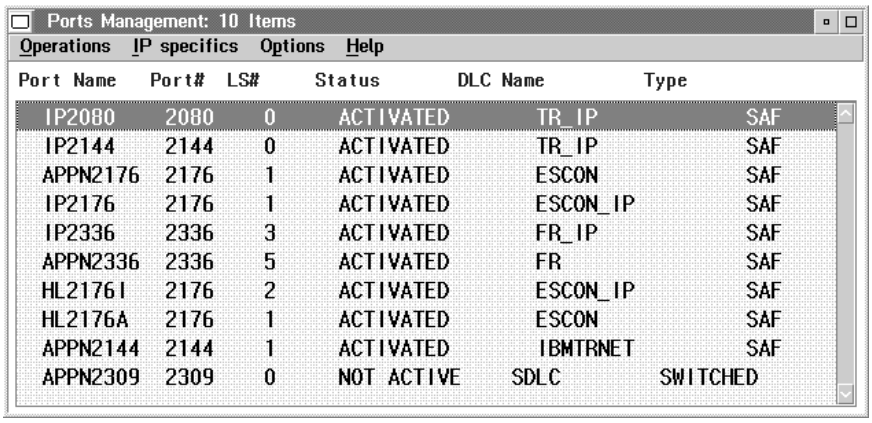

Figure 181. CCM Ports Management Display

Figure 182 shows the stations management display. This display is reached by selecting the **Stations...** menu item in Figure 180 on page 199. In this figure the current state for all stations is shown. Dynamically defined stations can be recognized by the leading @ in the station name. Double-clicking on an active station will bring up a display showing the APPN sessions to that station.

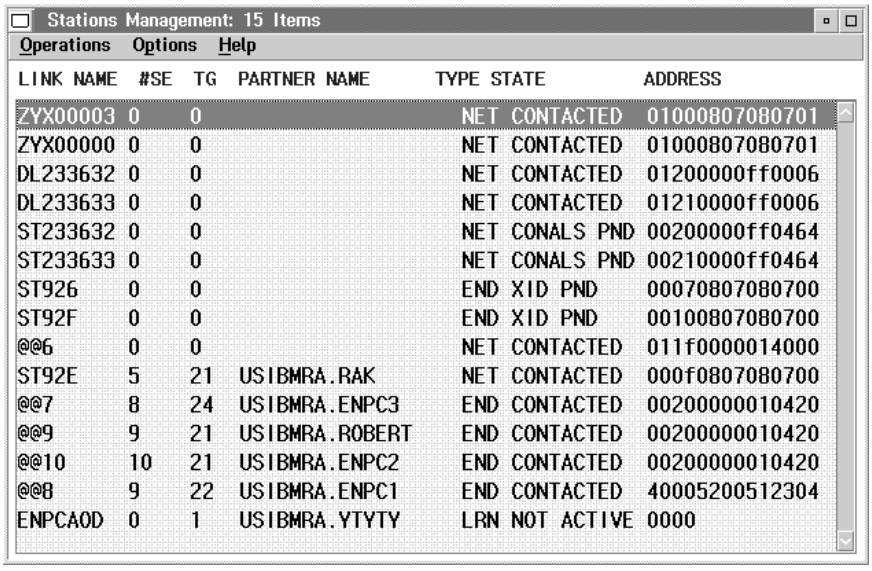

Figure 182. CCM Stations Management Display

Figure 183 on page 201 shows the APPN directory information display. This display is reached by selecting the **Directory...** menu item in Figure 180 on page 199. This display shows the directory of the 3746 NN. All registered LUs served by this NN are shown, along with the NETID.CPNAME of the APPN node where this LU is located. The adjacent network node USIBMRA.RAK (VTAM) is also shown.

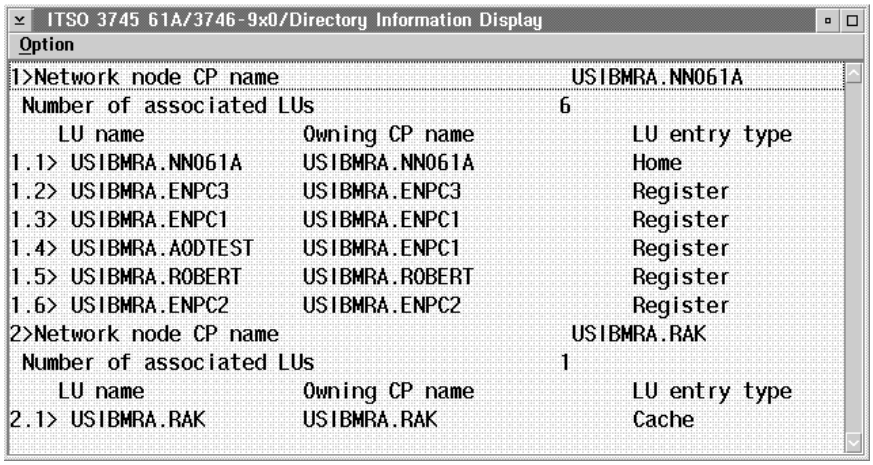

Figure 183. APPN Directory Information Display

6.5.3 Workstation Displays

The pmdsplay.exe program supplied with the OS/2 access feature can provide information about how a workstation is defined as well as current APPN status information. The following figures show some of the output from pmdsplay.exe, which shows the APPN definitions made on ENPC1.

Figure 184 shows global SNA information about the workstation, in particular the name USIBMRA.ENPC1, and that it is an end node.

```
************************************
*       SNA Global Information       *
************************************
Network name                              USIBMRA
Control point (CP) name                   ENPC1
Physical unit (PU) name                   ENPC1
Node ID (for XID)                         X'05D00000'
CP alias                                  ENPC1
Node type                                 End node
CP local address                          Not used (independent LU)
Workstation serial number                 -
Machine type                              0000
Machine model number                      X'000000'
Communications Server version             5.0
Branch extender support                   No
Search required                           No
```

Figure 184. SNA Global Information Display

Figure 185 on page 202 shows the APPN links that were defined on ENPC1. One APPN link was defined to destination MAC address 400037462144, SAP 08. This link is to the preferred network node server, which will be the 3746.

```
************************************
*    Link Definition Information     *
************************************
Number of links                          1
1>Link name                              LINK0001
  Adjacent node CP name
  Adjacent node type                     Network node
  DLC name                               IBMTRNET
  Adapter number                         0
  Destination DLC address                X'40003746214408'
  CP-CP session support                  Yes
  Preferred NN server                    Yes
  Auto-activate link                     Yes
  Transmission group number              0
  Limited resource                       No
  Limited resource timeout               30
  Inactivity timeout                     0
  Solicit SSCP session                   No
  Init self                              No
  BIND support                           Yes
  Link station role                      Negotiable
  Line type                              Switched
  Effective capacity                     16000000 bits per second
  Cost per connect time                  0
  Cost per byte                          0
  Propagation delay                      384.00 microseconds (LAN)
  User defined parameter 1               128
  User defined parameter 2               128
  User defined parameter 3               128
  Security                               Nonsecure
  Max activations attempts               0
  Physical unit (PU) name                ENPC1
  OCDE name
  Node id received                       X'00000000'
  Node id sent                           X'05D00000'
  Use CP as PU                           No
  Permanent connection                   No
  HPR support                            Yes
  Backup host link                       No
  Primary link name
  Branch extender uplink                 No
  HPR link error recovery type           No ERP preferred
```

Figure 185. Link Definition Information

Figure 186 shows the type 6.2 logical units defined on ENPC1. These are the LUs ENPC1 and AODTEST. As they are SSCP-independent logical units they have a local address of 0x00 (LOCADDR=0).

```
************************************
*     LU Definition Information     *
************************************
Number of logical units (LUs)            2
1>LU name                                ENPC1
  LU alias                               ENPC1
  LU type                                6.2
  LU local address                       X'00'
2>LU name                                AODTEST
  LU alias                               AODTEST
  LU type                                6.2
  LU local address                       X'00'
```

Figure 186. LU Definition Information

Pmdsplay.exe also provides information about the current status of APPN on the workstation. Figure 187 on page 203 shows the current directory information for ENPC1.

```
************************************
*       Directory Information       *
************************************
Total directory entries                    3
Network node entries                       1
1>Network node CP name                     USIBMRA.NN061A
  Number of associated LUs                 1
1.1>LU name                                USIBMRA.NN061A
    Owning CP name                         USIBMRA.NN061A
    LU entry type                          Home
Local and adjacent node entries            1
1>Owning CP name                           USIBMRA.ENPC1
  Number of associated LUs                 2
1.1>LU name                                USIBMRA.ENPC1
    LU entry type                          Home
1.2>LU name                                USIBMRA.AODTEST
    LU entry type                          Home
```

Figure 187. Directory Information

Figure 188 on page 204 shows information about the currently active links on ENPC1. We can see the link to our network node server (NNS) USIBMRA.NN061A (the 3746). The link was locally activated, which means that ENPC1 did the call out for this link.

```
************************************
*      Active Links Information      *
************************************
Number of active links                      1
1>Link name                                 LINK0001
  DLC name                                  IBMTRNET
  Adapter number                            0
  Destination DLC address                   X'40003746214408'
  Link activated                            Locally
  Link state                                Active
  Deactivating link                         No
  Limited resource                          No
  Limited resource timeout                  30
  Inactivity timeout                        0
  Active and activating sessions            9
  Max I-field size                          2058
  Adjacent node CP name                     USIBMRA.NN061A
  Adjacent node type                        Network node
  CP-CP session support                     Yes
  Connection type                           Peer
  Link station role                         Primary
  Line type                                 Switched
  Two-way simultaneous                      Yes
  Transmission group number                 22
  Effective capacity                        16000000 bits per second
  Cost per connect time                     0
  Cost per byte                             0
  Propagation delay                         384.00 microseconds (LAN)
  User defined parameter 1                  128
  User defined parameter 2                  128
  User defined parameter 3                  128
  Security                                  Nonsecure
  Physical unit (PU) name                   ENPC1
  HPR enabled                               Yes
  Local ANR label                           X'8002'
  Adjacent ANR label                        X'A320308E80'
  HPR SAP                                   X'08'
  Adjacent node HPR level                   ANR
  Primary link name
  Branch extender uplink                    No
  HPR link error recovery type              None
  Transmission group effective capacity     16000000 bits per second
  Transmission group cost per connect time  0
  Transmission group cost per byte          0
  Transmission group propagation delay      384.00 microseconds (LAN)
```

Figure 188. Active Links Information

Figure 189 on page 205 shows information about the partner logical units that ENPC1 currently knows.

```
************************************
* Partner LU Definition Information  *
************************************
Number of partner logical units             5
1>Partner LU name                           USIBMRA.NN061A
  Partner LU alias                          @I000000
  Partner LU uninterpreted name
  Maximum logical record send size          32767
  Conversation security                     No
  Parallel sessions                         Supported
  Default routing preference                Native first
2>Partner LU name                           USIBMRA.ENPC3
  Partner LU alias                          @I000001
  Partner LU uninterpreted name
  Maximum logical record send size          32767
  Conversation security                     No
  Parallel sessions                         Supported
  Default routing preference                Native first
3>Partner LU name                           USIBMRA.ROBERT
  Partner LU alias                          @I000002
  Partner LU uninterpreted name
  Maximum logical record send size          32767
  Conversation security                     No
  Parallel sessions                         Supported
  Default routing preference                Native first
4>Partner LU name                           USIBMRA.ENPC2
  Partner LU alias                          @I000003
  Partner LU uninterpreted name
  Maximum logical record send size          32767
  Conversation security                     No
  Parallel sessions                         Supported
  Default routing preference                Native first
5>Partner LU name                           USIBMRA.RAK
  Partner LU alias                          @I000004
  Partner LU uninterpreted name
  Maximum logical record send size          32767
  Conversation security                     No
  Parallel sessions                         Supported
  Default routing preference                Native first
```

Figure 189. Partner LU Information

Chapter 7. 2216 Network Node and DLUR

In this scenario, we convert the 2216 to an APPN network node. In addition to that it provides DLUR services for PU2.0 3270 host access. In the frame relay BAN scenario, as we did not define any connection networks on our LAN segments, sessions between the end nodes all passed through the network node server (3746). This means that sessions between stations on LAN segment 032, or between stations on LAN segment 001, or between those segments all traversed the WAN link between the 2216 and 3746.

Using connection networks is one solution to allow direct connectivity between end nodes on a shared transport medium. This would not solve the problem though for communication between the two LAN segments. Converting the 2216 to an APPN network node will add that processing load to the 2216, but will allow stations on the two LANs to communicate without the sessions being routed via the 3746.

For the NN-to-NN traffic between the 3746 and the 2216, we defined an additional FR PVC (DLCI 21 at the 2216 side and 31 at the 3746 side). This has been done just for clarity. All traffic (IP, BAN and APPN) could run over the same DLCI.

The 2216 will become the network node server (NNS) and DLUR for all workstations in segments 032 and 001. In addition, each W/S has 3270 access to the host using the DLUR service.

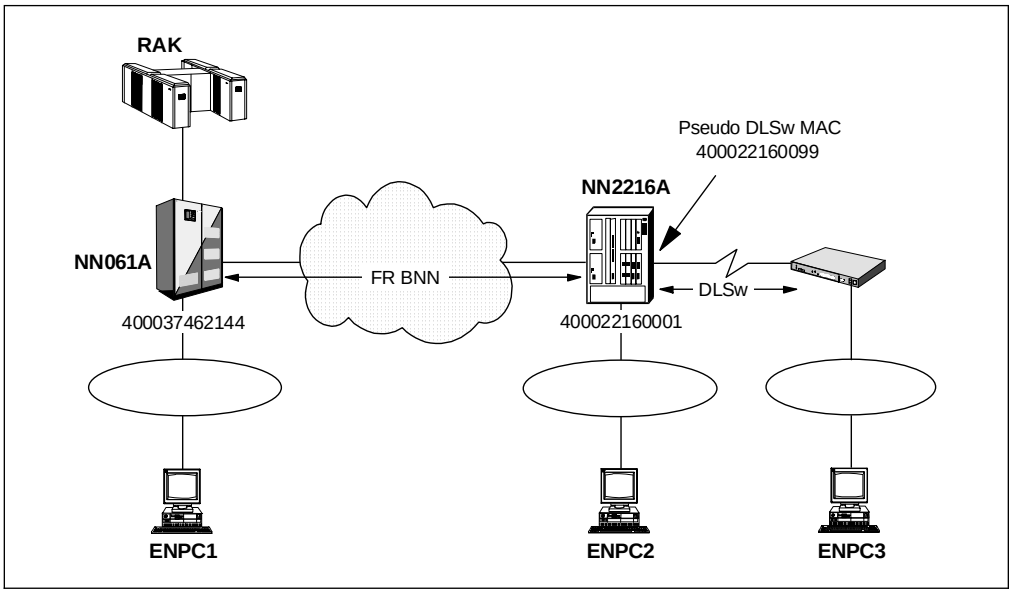

Figure 190. 2216 APPN NN and DLUR

The APPN support in the router introduces a so-called Pseudo DLSw interface. The purpose of this interface is to establish a data path from the DLSw segment to the APPN CP in the router. A MAC address needs to be configured for this interface. Endstations send their SNA/APPN data over a DLSw pipe to this MAC address. This MAC address is not related to a DLCI nor to a specific port. It is the address of an internal data path from the DLSw segment to the APPN function (see Figure 191 on page 208), it is the path between MAC address A and MAC address B.

This MAC address can be the same for all routers in the whole network, so all workstations can be configured with the same destination MAC address.

Figure 191 shows the logical view of the configuration.

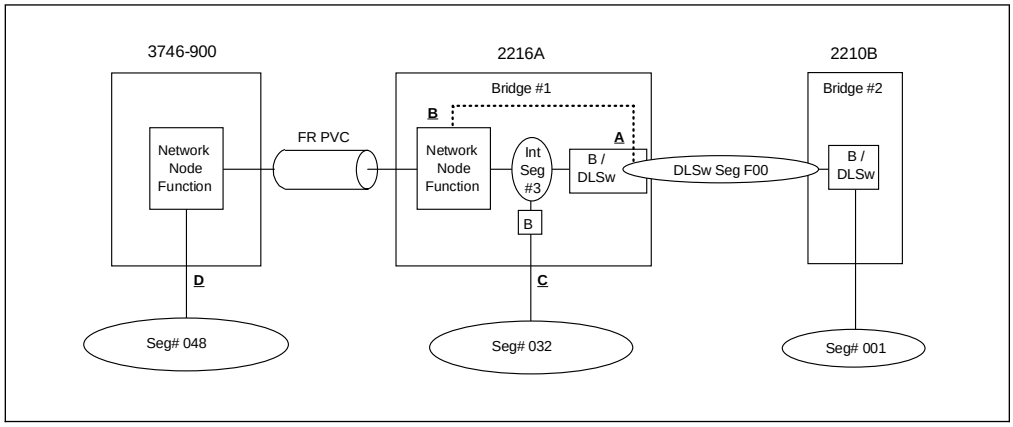

Figure 191. 2216 Network Node Logical View

The 2210, PPP, and LAN configuration in this scenario has not changed from previous scenarios. The 2210 router is serving a remote location and is connected to the 2216 over a PPP link. The SNA/APPN data are transported to the 2216A router using a DLSw connection.

For the wide area transport, we used a frame relay network. This transport network is protocol-independent and therefore is the same for all scenarios described in this book.

For a detailed description of the frame relay network please refer to Chapter 4, “Transport Network” on page 119. The IP addressing scheme is the same for all scenarios. For a detailed description of the IP addressing structure please refer to 5.1, “IP Subnets and Addresses” on page 143.

7.1 Definitions

The following section shows the definitions made for the 2216 network node test scenario.

7.1.1 VTAM Definitions

In this scenario, VTAM provides the dependent LU server (DLUS) function, and the 3746 and 2216 both provide dependent LU requester (DLUR) functions.

When VTAM is started as an APPN network node, the DLUS function is automatically active without the definition of specific parameters. Switched major node definition for dependent LUs and DLURs are required when the session will be started from the VTAM side. Because in our test scenario the session is initiated from the DLUR side, we do not need pre-defined dial out information. For the dependent LU definitions, the VTAM Configuration Services XID Exit Routine (ISTEXCCS) is used. This creates dynamic dependent LU definitions.

7.1.2 3746 APPN and Frame Relay Definitions

On the 3746 side we needed to create definitions for the new PVC that we are using to communicate with the 2216 for this scenario. At the 3746 side this PVC is using DLCI 31. This DLCI was defined on port 2336 in addition to the two DLCIs previously used.

On the frame relay port, three DLCIs are now defined. Figure 192 shows an overview of the DLCIs on port 2336.

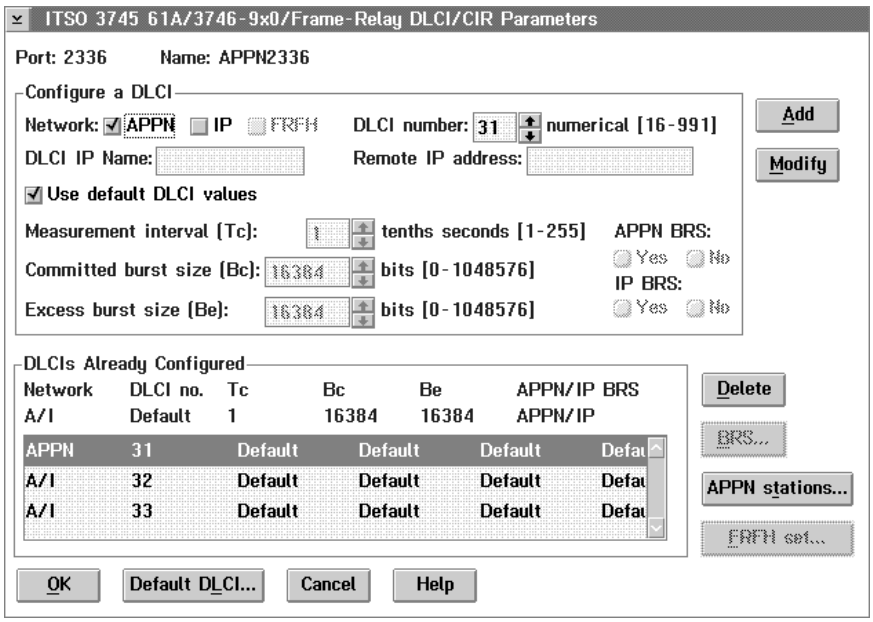

Figure 192. Port 2336 DLCIs Overview

Figure 193 on page 210 shows an APPN station configuration on the frame relay DLCI. This dialog is reached by selecting DLCI 31, then selecting the button **APPN stations...** in Figure 192. If the parameter Accept any incoming call in the APPN parameters for the frame relay port was not specified, then a station representing the 2216 must be predefined.

We defined a frame relay APPN station named ST233631 to allow us to call out to this station. The frame format that is used is routed. The routed frame format is used by the frame relay boundary network node (BNN) function. A remote MAC address is not needed for BNN, but a remote SAP address (RSAP) must be set. The RSAP must match the local SAP used for APPN in the 2216 (see Figure 200 on page 215).

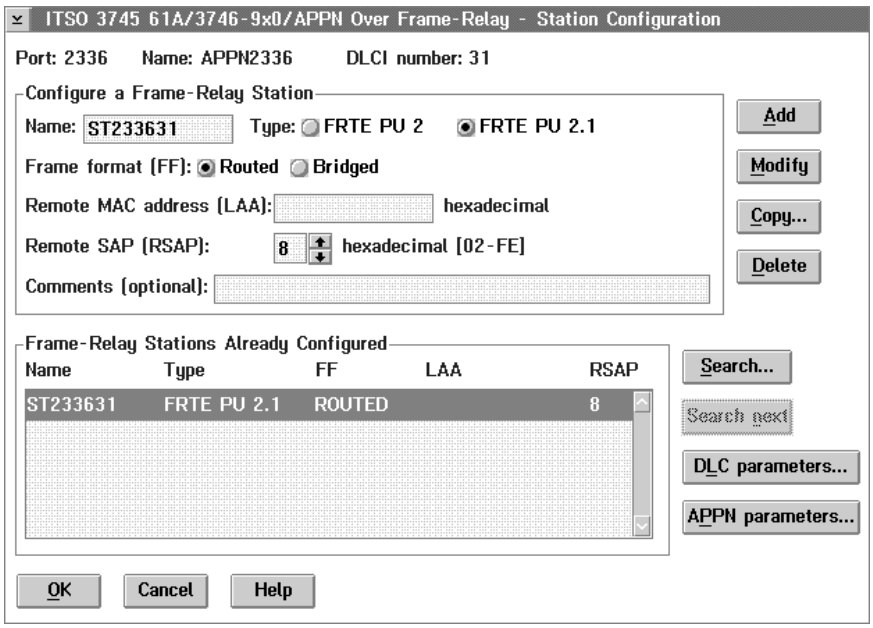

Figure 193. BNN APPN Station Configuration

Figure 194 shows the primary network node definitions and DLUR parameters for the 3746. This dialog is reached by selecting the **NN/FP/DLUR...** menu item from the **Configuration** menu option on the CCM main screen. The 3746 is configured as a DLUR, and its dependent LU server (DLUS) is defined as the VTAM network node USIBMRA.RAK.

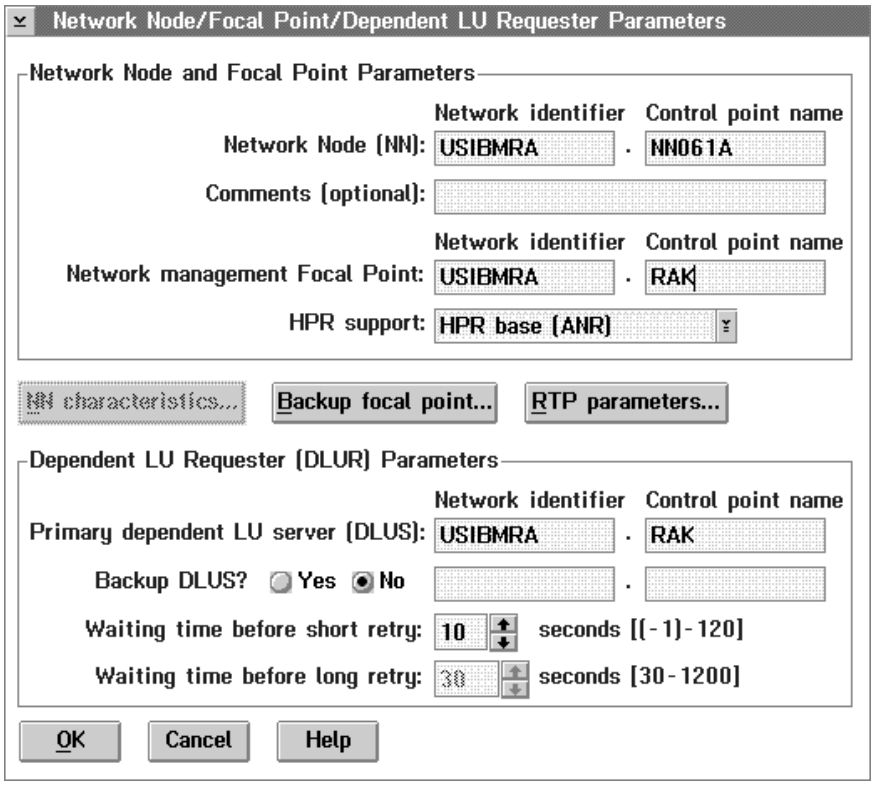

Figure 194. DLUR Parameters

7.1.3 Router 2216A Definitions

In this scenario we are use the same hardware setup as the previous APPN scenario. As far as the 2216 is concerned, the TCP/IP and BAN configuration has not changed. In this section we show the changes made to allow the 2216 to function as an APPN network node (NN2216A), and we add a single PVC for connection to the 3746. The PVC uses DLCI 21 at the 2216 end and DLCI 31 at the 3746 end.

The previous configuration was loaded into the configuration tool and used as a base for these definitions.

Refer to Figure 151 on page 179 for the frame relay interface definition window. Here we need to configure an additional DLCI. Click on the **PVC** tab to display to the notebook PVC page.

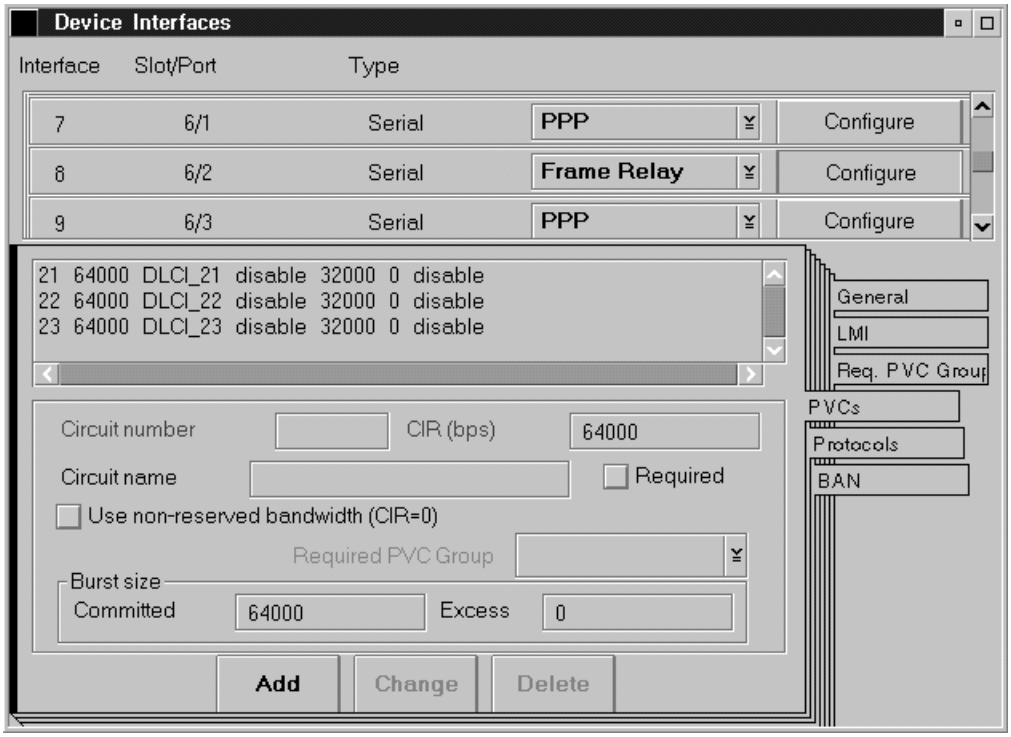

Figure 195. Device Interfaces: Frame Relay PVC

Here we added DLCI 21. Do not assign any protocol to this DLCI. We will assign this port to APPN later on. The DLCI assignment is now as follows:

- *DLCI 21*. No protocol assigned to this DLCI
- *DLCI 22*. FR BAN
- *DLCI 23*. TCP/IP

As mentioned previously, for clarity we have used a separate PVC for each protocol, all these protocols could be run on a single DLCI.

Now we need to define APPN. At the navigation window, go down to the menu item **APPN General**.

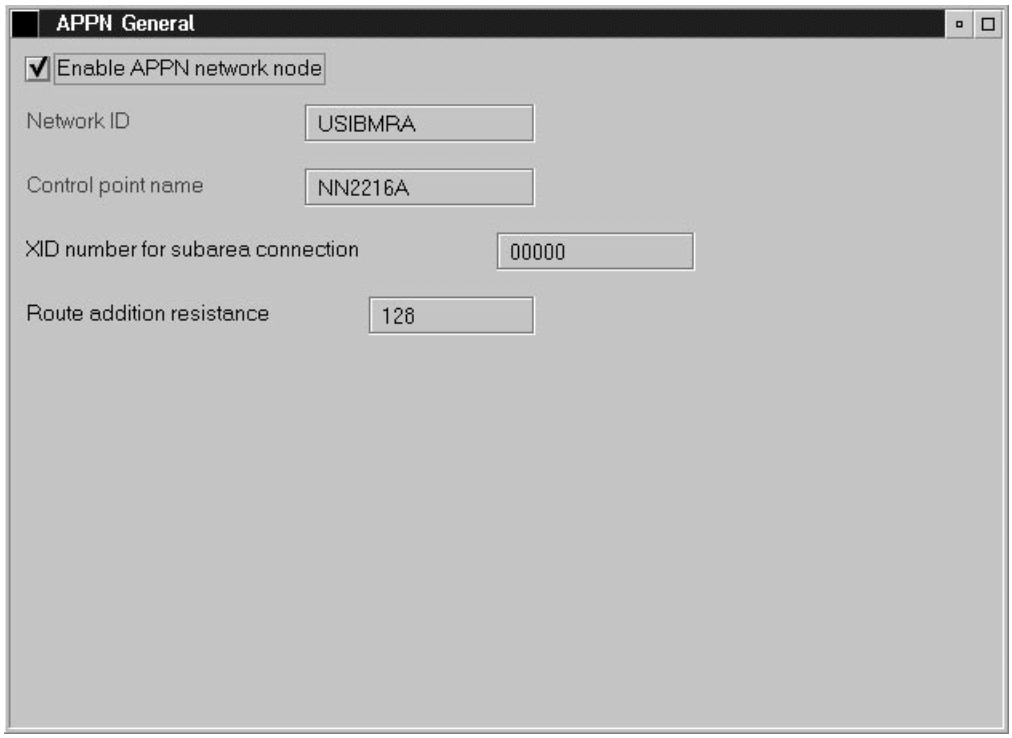

Figure 196. APPN General

Here we must enable APPN on the 2216, define the Network ID and define the Control point name.

Next, in the navigation window, go down to DLUR. The definition screen (Figure 197 on page 213) is self-explanatory. Enable DLUR and enter the DLUS name. A backup DLUS may also be defined here.

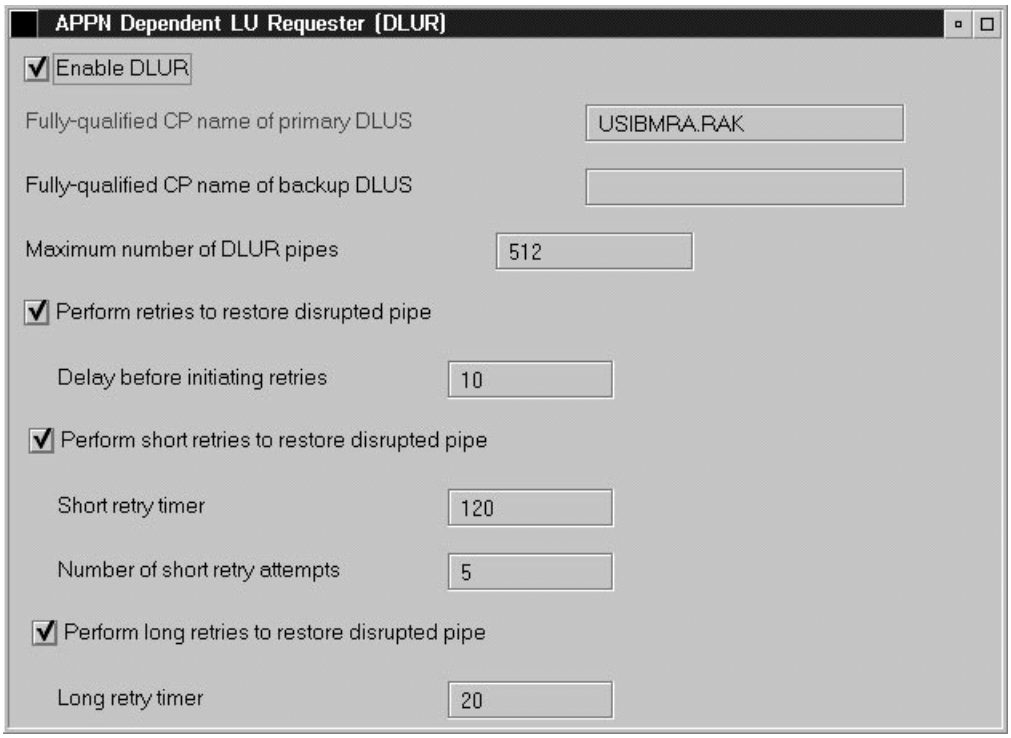

Figure 197. APPN DLUR

Next, in the navigation window, go down to **Interfaces**. Here we need to enable certain interfaces for APPN routing. The first interface to enable is the token-ring. Check the enable box and click on **Configure** (see Figure 198).

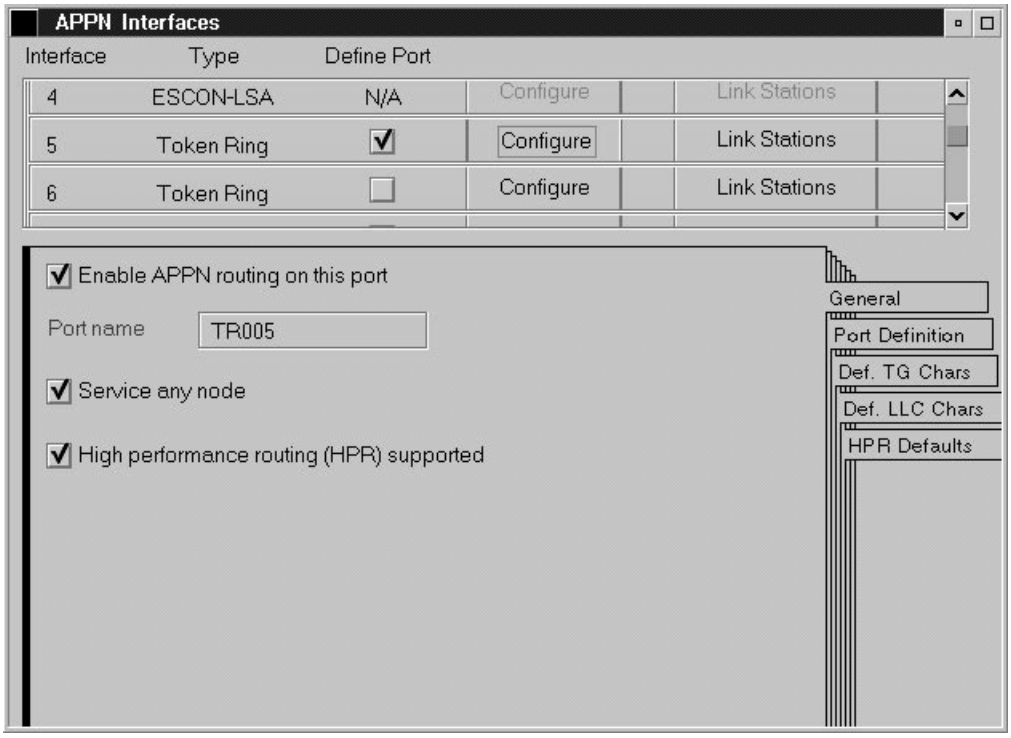

Figure 198. APPN Interfaces Token-Ring

The fields shown in the previous figure are:

Port name
Enter a port name or accept the default.

Service any node
When this parameter is enabled, the network node accepts any requests it receives from another node to establish a connection. When this parameter is disabled, only connection requests from nodes that are explicitly defined are accepted.

HPR supported
An HPR connection will be established when both link stations indicate HPR support during the XID exchange.

Now go down to the frame relay interface and click on **Configure** (see Figure 199).

The dialog is the same as for token-ring.

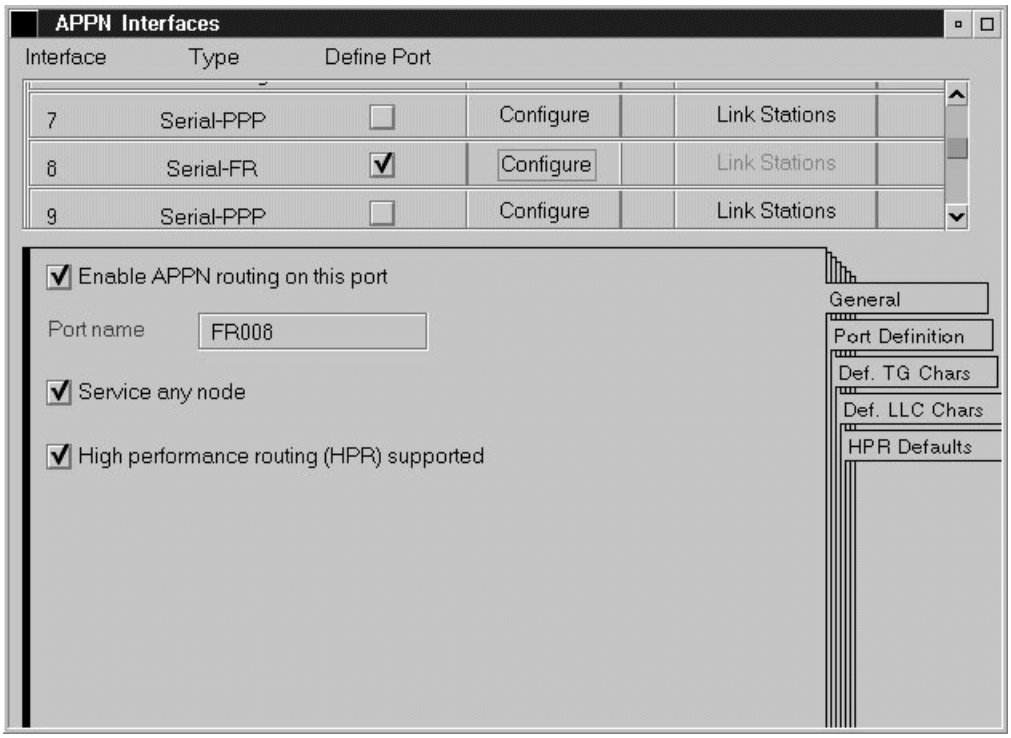

Figure 199. APPN Interfaces Frame Relay

Now go down to the Pseudo DLSw and click on **Configure** (see Figure 200 on page 215).

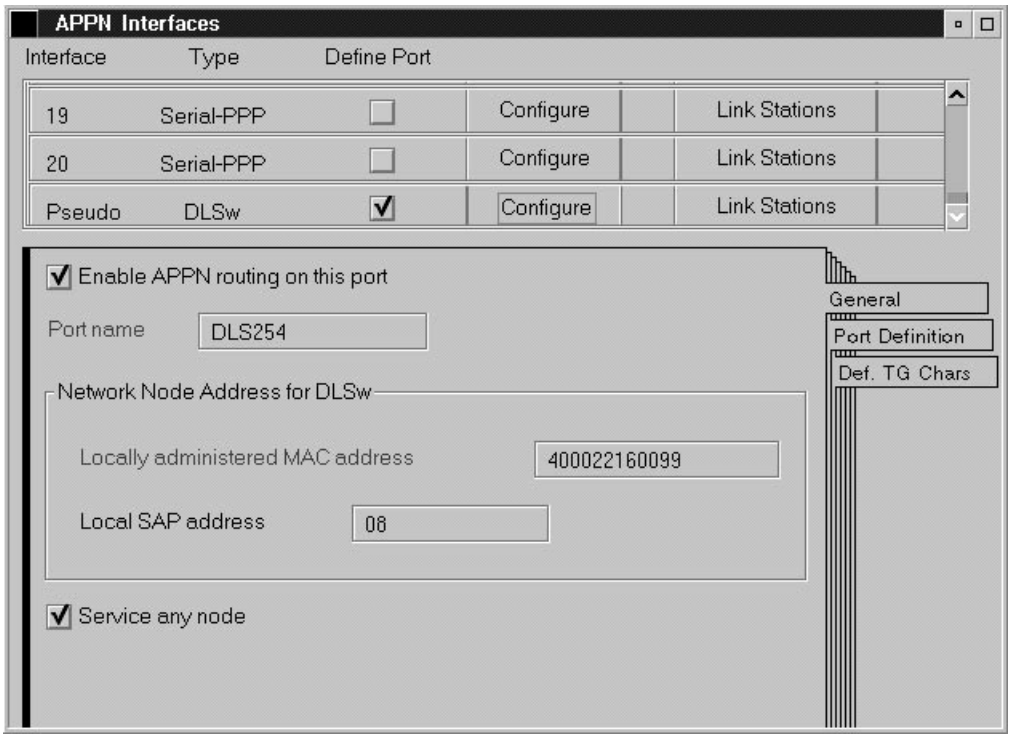

Figure 200. APPN Interfaces Pseudo DLSw

This is the APPN Pseudo DLSw Interface. This internal APPN interface provides a data path between the APPN and DLSw functions. We enable APPN routing on this port to allow DLSw-attached link stations to be connected to this network node. This pseudo interface also has a MAC address and a port name assigned to it. The MAC address, in this example 400022160099, must be unique. The remote workstations send their APPN data to this MAC address.

This MAC address serves a similar purpose as the BAN PVC MAC address used in the FR BAN. In fact, the remote workstation has the choice to define as the destination MAC address the:

- Pseudo DLSw address: The remote workstation connects to the 2216 APPN function.
- BAN PVC MAC address: The remote workstation connects to the 2216 BAN function and goes through the 2216 to the 3746 APPN function.

Note: 2216 APPN does not support locally attached link stations via DLSw. It only supports link stations that are attached remotely via DLSw.

In the navigation window, go down to FR PVC Stations (see Figure 201 on page 216). This displays a window with all defined FR PVCs. In this case, we want to run APPN traffic to the 3746 over PVC 21. So click on **Link Station** beside PVC 21.

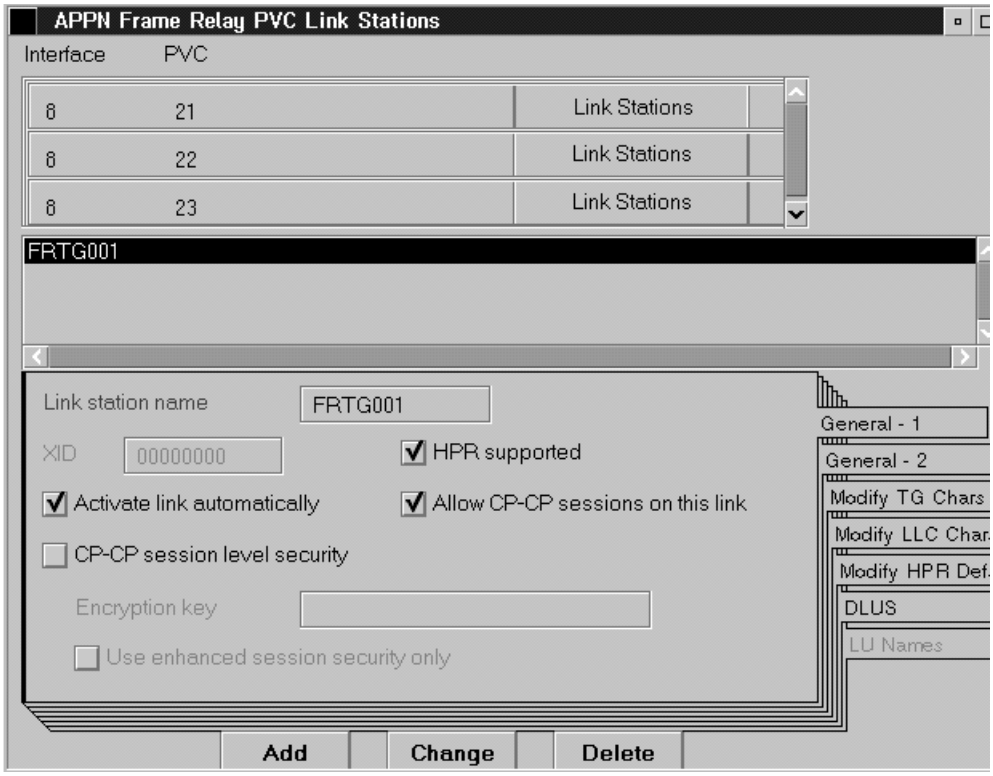

Figure 201. APPN FR PVC Link Stations, Page 1

Here we add a link station to allow the 2216 to call out to the 3746. The fields to define are:

Link station name
Enter a name or accept the default.

HPR supported
Allow HPR connections on this link.

Note that HPR connections will only be established when both link stations indicate HPR support during the XID exchange. This parameter may override the port definition.

Activate link automatically
If checked, the network node automatically activates the link to the adjacent node when started.

Allow CP-CP sessions on this link
This should be enabled as the adjacent node is a network node.

Now click **General - 2** to get the second page of the link station definition (see Figure 202 on page 217).

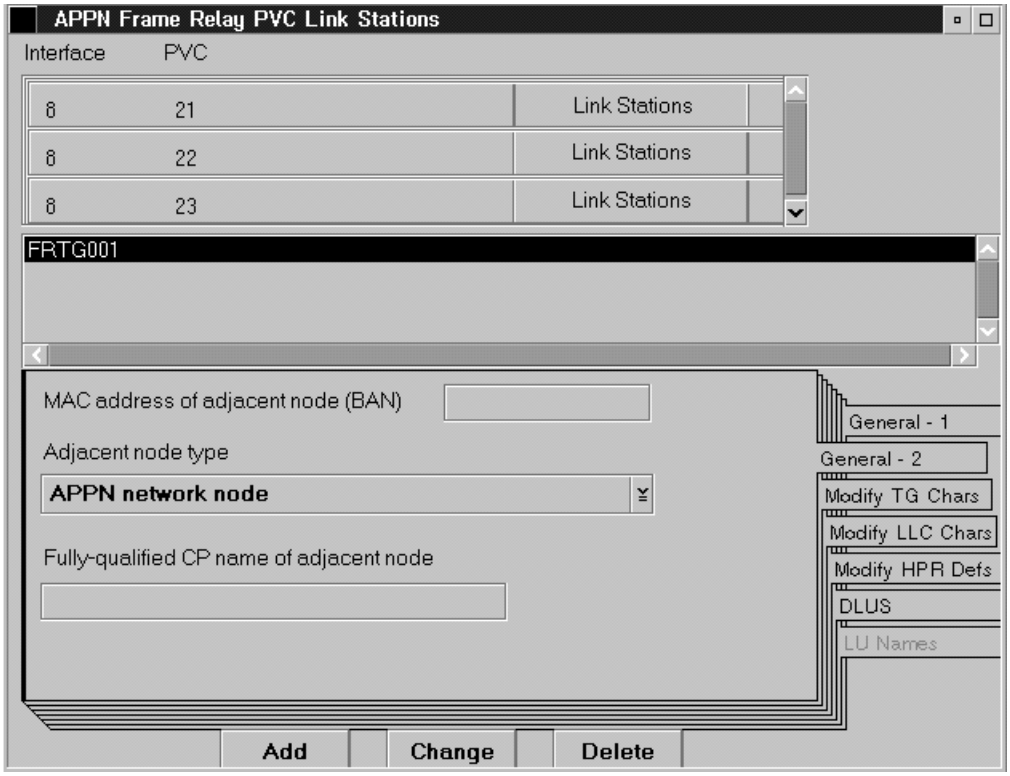

Figure 202. APPN FR PVC Link Stations, Page 2

For the Adjacent network node type select **APPN network node**.

That concludes the APPN and DLUR definitions for the 2216A.

7.1.4 Router 2210B Definitions

In this scenario, we used the same configuration as in the previous APPN scenario. The SNA/APPN traffic from the workstations connected to this router is still sent to the 2216 over the DLSw pipe, so for the 2210 no configuration changes are necessary.

7.1.5 Workstation Definitions

For this test scenario, we need to make changes in the workstations in LAN segments 032 and 001. The 2210 is still not running APPN, therefore for the workstations in LAN segment 001, we can now define either the 2216 or the 3746 as the network node server. The logical choice is to select the 2216. To do this we must change the network node server MAC address, and the destination MAC address on any predefined links.

On ENPC3, we did this on the following screen:

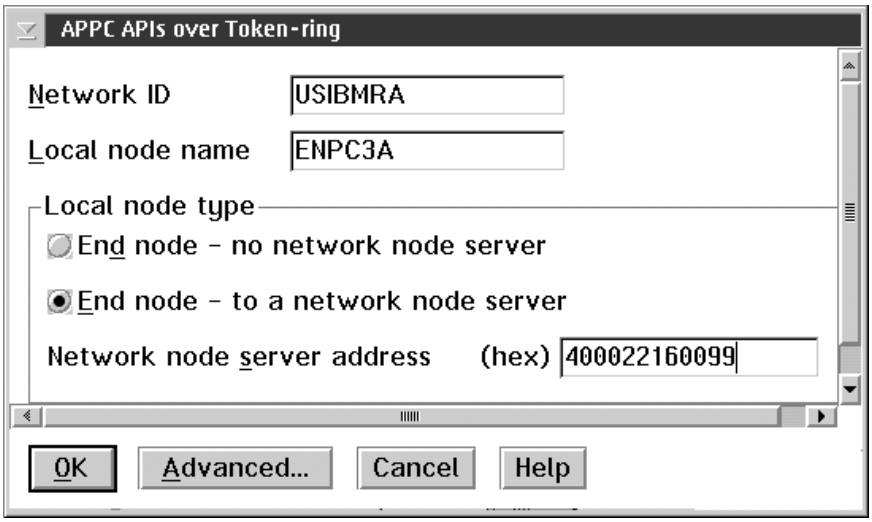

Figure 203. Primary APPN Definition

Figure 203 shows the workstation's primary APPN definition. As the 2210 uses DLSw to communicate with the 2216, we defined the 2216's pseudo DLSw MAC address. This was defined on Figure 193 on page 210.

For workstations in segment 032 (not connected via DLSw to the 2216), the MAC address of the 2216 network node server is 400022160001. This is the MAC address of the 2216 in segment 032.

As we left our BAN configuration definitions in place, workstations in segments 032 and 001 could still define the NNS MAC address as the BAN PVC MAC address of the 2216 (400022160022). This would cause the APPN traffic to be bridged through the 2216 onto the frame relay PVC. The traffic would bypass the 2216 APPN function and use the 3746 as NNS. The workstation would appear (from the APPN point of view) to be adjacent to the 3746.

For workstations in segment 048 (connected to the 3746), the MAC address of the network node server is 400037462144. This is the MAC address of the 3746 port 2144 TIC3 in segment 048.

7.2 Management Displays

The following section shows the management displays for the 2216 network node test scenario.

7.2.1 VTAM APPN Displays

The VTAM display for the 2216 network node is shown in Figure 204 on page 219.

```
D NET,ID=NN2216A,E
IST097I DISPLAY ACCEPTED
IST075I NAME = USIBMRA.NN2216A, TYPE = DIRECTORY ENTRY
IST1186I DIRECTORY ENTRY = DYNAMIC NN
IST1184I CPNAME = USIBMRA.NN2216A - NETSRVR = ***NA***
IST1402I SRTIMER =   120  SRCOUNT =    60
IST314I END
```

Figure 204. VTAM Display for NN2216A

The VTAM display for a PS/2 end node is shown in Figure 205. Note that the 2216 (NN2216A) is now the network node server (NETSRVR) for this end node.

```
D NET,ID=ENPC3A,E
IST097I DISPLAY ACCEPTED
IST075I NAME = USIBMRA.ENPC3A, TYPE = DIRECTORY ENTRY
IST1186I DIRECTORY ENTRY = DYNAMIC EN
IST1184I CPNAME = USIBMRA.ENPC3A - NETSRVR = USIBMRA.NN2216A
IST1402I SRTIMER =   120  SRCOUNT =    60
IST314I END
```

Figure 205. VTAM Display for ENPC3A

The VTAM display for the dependent LU requesters is shown in Figure 206.

```
D NET,DLURS
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = DLURS
IST1352I DLUR NAME          DLUS CONWINNER STATE  DLUS CONLOSER STATE
IST1353I USIBMRA.NN061A     ACTIVE                ACTIVE
IST1353I USIBMRA.NN2216A    ACTIVE                ACTIVE
IST314I END
```

Figure 206. VTAM Display for DLURs

The VTAM display for the 3746 DLUR and the two dynamically defined PUs served by it is shown in Figure 207 on page 220.

```
D NET,ID=NN061A,E
IST097I DISPLAY ACCEPTED
IST075I NAME = USIBMRA.NN061A, TYPE = ADJACENT CP
IST486I STATUS= ACT/S----Y, DESIRED STATE= ACTIV
IST1402I SRTIMER =   120  SRCOUNT =    60
IST1447I REGISTRATION TYPE = NO
IST977I MDLTAB=***NA*** ASLTAB=***NA***
IST1333I ADJLIST = ***NA***
IST861I MODETAB=***NA*** USSTAB=***NA*** LOGTAB=***NA***
IST934I DLOGMOD=CPSVCMG USS LANGTAB=***NA***
IST597I CAPABILITY-PLU ENABLED  ,SLU ENABLED  ,SESSION LIMIT NONE
IST231I CDRSC MAJOR NODE = ISTCDRDY
IST1184I CPNAME = USIBMRA.NN061A - NETSRVR = ***NA***
IST1044I ALSLIST = ISTAPNPU
IST082I DEVTYPE = INDEPENDENT LU / CDRSC
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST171I ACTIVE SESSIONS = 0000000005, SESSION REQUESTS = 0000000000
IST206I SESSIONS:
IST1081I ADJACENT LINK STATION = CP90061A
IST634I NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I RAK      ACTIV/SV-S D2333BDF4603D151 000B 0003  0  0 USIBMRA
IST635I RAK      ACTIV/DL-S D2333BDF4603D14E 002D 0000  0  0 USIBMRA
IST635I RAK      ACTIV/CP-S D2333BDF4603D14A 04DC 0001       USIBMRA
IST635I RAK      ACTIV/DL-P F8D3D1642B4FB9B0 0000 002E  0  0 USIBMRA
IST635I RAK      ACTIV/CP-P F8D3D1642B4FB9AF 0001 04D8  0  0 USIBMRA
IST1355I PHYSICAL UNITS SUPPORTED BY DLUR USIBMRA.NN061A
IST089I W05123   TYPE = PU_T2.1          , ACTIV---X-
IST089I W05124   TYPE = PU_T2.1          , ACTIV---X-
IST924I -------------------------------------------------------------
IST075I NAME = USIBMRA.NN061A, TYPE = DIRECTORY ENTRY
IST1186I DIRECTORY ENTRY = DYNAMIC NN
IST1184I CPNAME = USIBMRA.NN061A - NETSRVR = ***NA***
IST1402I SRTIMER =   120  SRCOUNT =    60
IST314I END
```

Figure 207. VTAM Display for NN061A

The VTAM display for the 2216 DLUR and a dynamically defined PU served by it is shown in Figure 208 on page 221.

```
D NET,ID=NN2216A,E
IST097I DISPLAY ACCEPTED
IST075I NAME = USIBMRA.NN2216A, TYPE = ADJACENT CP
IST486I STATUS= ACT/S----Y, DESIRED STATE= ACTIV
IST1402I SRTIMER =   120  SRCOUNT =    60
IST1447I REGISTRATION TYPE = NO
IST977I MDLTAB=***NA*** ASLTAB=***NA***
IST1333I ADJLIST = ***NA***
IST861I MODETAB=***NA*** USSTAB=***NA*** LOGTAB=***NA***
IST934I DLOGMOD=***NA*** USS LANGTAB=***NA***
IST597I CAPABILITY-PLU ENABLED  ,SLU ENABLED  ,SESSION LIMIT NONE
IST231I CDRSC MAJOR NODE = ISTCDRDY
IST479I CDRM NAME = RAK, VERIFY OWNER = NO
IST1184I CPNAME = USIBMRA.NN2216A - NETSRVR = ***NA***
IST1044I ALSLIST = ISTAPNPU
IST082I DEVTYPE = INDEPENDENT LU / CDRSC
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST171I ACTIVE SESSIONS = 0000000002, SESSION REQUESTS = 0000000000
IST206I SESSIONS:
IST1081I ADJACENT LINK STATION = CNR0000B
IST634I NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I RAK      ACTIV/DL-S D1C38C9F6F3D39C3 0017 0000  0  0 USIBMRA
IST635I RAK      ACTIV/DL-P F8D3D1642B4FC049 0000 0018  0  0 USIBMRA
IST1355I PHYSICAL UNITS SUPPORTED BY DLUR USIBMRA.NN2216A
IST089I W05125   TYPE = PU_T2.1          , ACTIV---X-
IST924I -------------------------------------------------------------
IST075I NAME = USIBMRA.NN2216A, TYPE = DIRECTORY ENTRY
IST1186I DIRECTORY ENTRY = DYNAMIC NN
IST1184I CPNAME = USIBMRA.NN2216A - NETSRVR = ***NA***
IST1402I SRTIMER =   120  SRCOUNT =    60
IST314I END
```

Figure 208. VTAM Display for NN2216A

The VTAM display for the dynamically defined PU (W05123) on ENPC1, served by NN061A, is shown in Figure 209 on page 222.

```
D NET,ID=W05123,E
IST097I DISPLAY ACCEPTED
IST075I NAME = W05123, TYPE = PU_T2.1
IST486I STATUS= ACTIV---X-, DESIRED STATE= ACTIV
IST1043I CP NAME = ENPC1, CP NETID = USIBMRA, DYNAMIC LU = YES
IST1589I XNETALS = YES
IST1354I DLUR NAME = NN061A             MAJNODE = ISTDSWMN
IST136I SWITCHED SNA MAJOR NODE = ISTDSWMN
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST355I LOGICAL UNITS:
IST080I W0512302 ACTIV---X- W0512303 ACTIV---X- W0512304 ACTIV---X-
IST080I W0512305 ACTIV---X- W0512306 ACTIV---X- W0512307 ACTIV---X-
IST080I W0512308 ACTIV---X- W0512309 ACTIV---X- W051230A ACTIV---X-
IST080I W051230B ACTIV---X- W051230C ACTIV---X- W051230D ACTIV---X-
IST080I W051230E ACTIV---X- W051230F ACTIV---X- W0512310 ACTIV---X-
IST080I W0512311 ACTIV---X- W0512312 ACTIV---X- W0512313 ACTIV---X-
IST080I W0512314 ACTIV---X- W0512315 ACTIV---X-
IST314I END
```

Figure 209. VTAM Display for W05123

7.2.2 3746 CCM Management Displays

Figure 210 shows the 3746 ports management display. On the APPN port APPN2336 we see that three link stations are active.

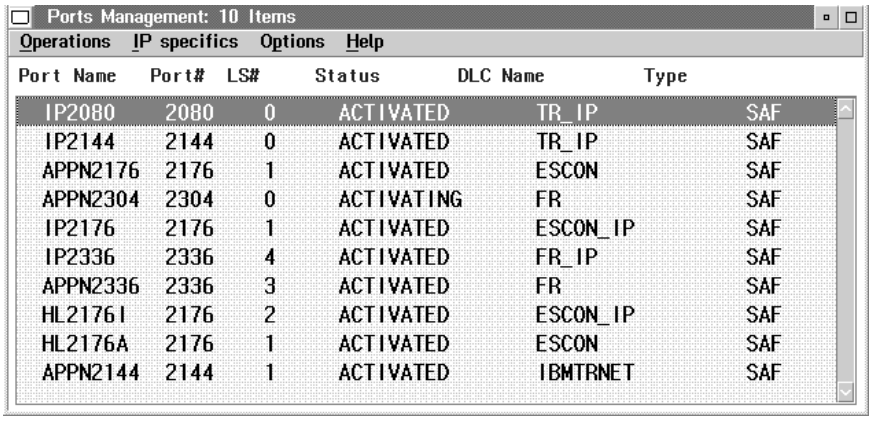

Figure 210. CCM Ports Management Display

Figure 211 on page 223 shows the 3746 stations management display. Here we see that station ST233631 is active, we are connected to USIBMRA.NN2216A over this link, and we are using our predefined link station definitions.

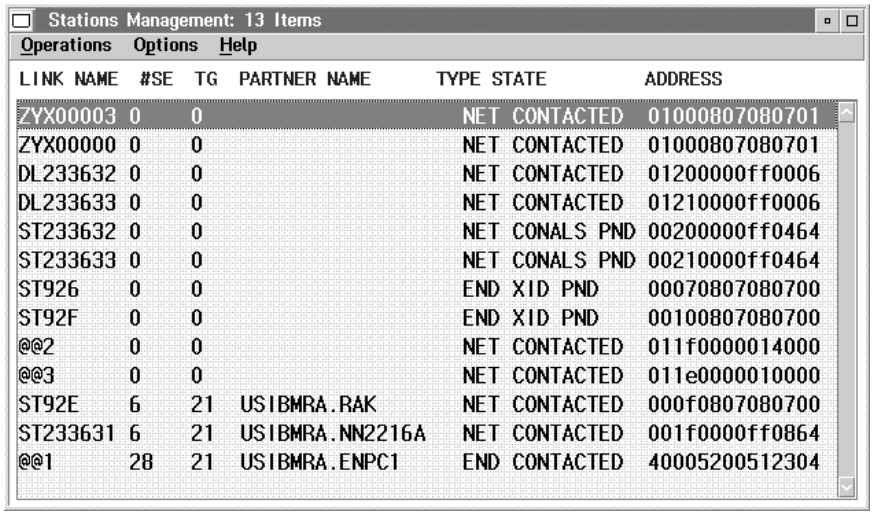

Figure 211. CCM Stations Management Display

Figure 212 shows the 3746 directory information display. The 2216 is now recognized as a network node (NN2216A) and we can see two end nodes (ENPC3A and ENPC2A) served by the 2216 network node.

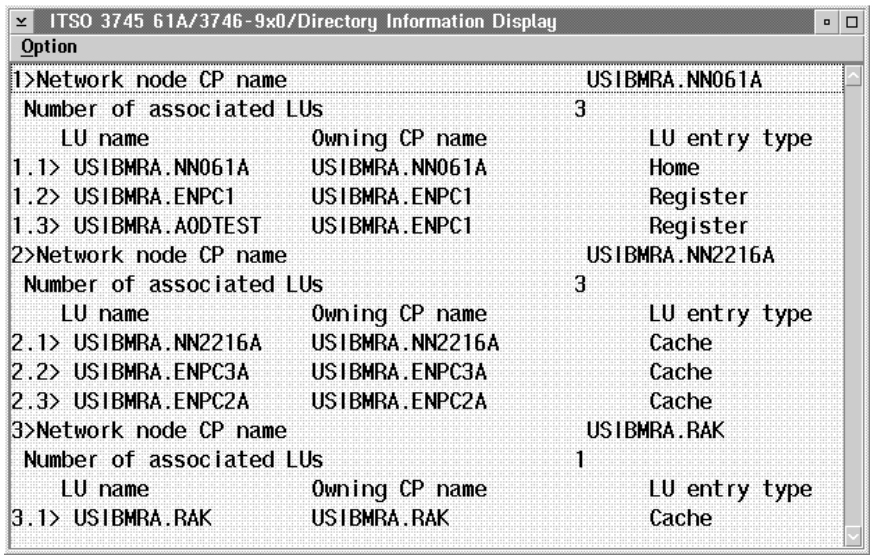

Figure 212. Directory Information Display

7.2.3 2216 APPN Displays

The following figures show APPN information for ENPC3.

```
2216A_BAN2NN APPN >list link
    Name   Port Name  Intf       Adj CP Name  Type       HPR      State
=============================================================================
 FRTG001       FR008    8     USIBMRA.NN061A    NN    ACTIVE     ACT_LS
     @@0      DLS254  254     USIBMRA.ENPC3A    EN  INACTIVE     ACT_LS
     @@2       TR005    5     USIBMRA.ENPC2A    EN  INACTIVE     ACT_LS
2216A_BAN2NN APPN >
```

Figure 213. APPN Links Display

At this point, we started the 3270 emulation at the workstation ENPC3A to get 3270 host access via DLUR in the 2216. Then we issued the same display command again.

```
2216A_BAN2NN APPN >list link
    Name   Port Name  Intf       Adj CP Name  Type       HPR      State
=============================================================================
 FRTG001       FR008    8     USIBMRA.NN061A    NN    ACTIVE     ACT_LS
     @@0      DLS254  254     USIBMRA.ENPC3A    EN  INACTIVE     ACT_LS
     @@2       TR005    5     USIBMRA.ENPC2A    EN  INACTIVE     ACT_LS
     @@3      DLS254  254    DLUR..@00000001    EN  INACTIVE     ACT_LS
2216A_BAN2NN APPN >
```

Figure 214. APPN Links Display

```
2216A_BAN2NN APPN >list cp
           CP Name      Type    Status  Connwiner ID    Conloser ID
=============================================================================
    USIBMRA.NN061A       NN     Active    348E882E        348E8830
    USIBMRA.ENPC3A       EN     Active    348E8867        348E8863
    USIBMRA.ENPC2A       EN     Active    348E8ED3        348E8ECF
2216A_BAN2NN APPN >
```

Figure 215. APPN Control Points Display

```
2216A_BAN2NN APPN >list isr
  Adjacent CP Name  TG Number   ISR Sessions
================================================
    USIBMRA.NN061A      21             2
    USIBMRA.ENPC3A      21             1
   DLUR..@00000001       0             0
    USIBMRA.ENPC2A      22             2
2216A_BAN2NN APPN >
```

Figure 216. APPN Intermediate Session Routing Sessions Display

```
2216A_BAN2NN APPN >list ses
     Origin CP Name          Primary LU        Secondary LU  Mode Name
=============================================================================
     USIBMRA.ENPC3A
      USIBMRA.ENPC1
     USIBMRA.ENPC3A
     USIBMRA.ENPC3A
     USIBMRA.ENPC3A
     USIBMRA.ENPC3A
2216A_BAN2NN APPN >
```

Figure 217. APPN Sessions Display

```
2216A_BAN2NN APPN >list rtp
RTP PARTNER TABLE:
 Remote Partner Name  Remote Boundary Name  TG Number
=====================================================
         USIBMRA.RAK           USIBMRA.RAK     -1
RTP CONNECTION TABLE:
  TCID             CP Name  ISR  APPC  Pathswitch    Alive    COS TPF   TG Number
=============================================================================
367453A0       USIBMRA.RAK    1    0    00000708  00000708     #INTER      0
367753C8       USIBMRA.RAK    0    2    00000708  00000708   SNASVCMG      0
2216A_BAN2NN APPN >
```

Figure 218. APPN RTP Connections Display

```
2216A_BAN2NN APPN >list port
Intf          Name       DLC Type        HPR          State
=============================================================
254         DLS254            DLS      FALSE       ACT_PORT
  5          TR005       IBMTRNET       TRUE       ACT_PORT
  8          FR008             FR       TRUE       ACT_PORT
2216A_BAN2NN APPN >
```

Figure 219. APPN Ports Display

7.2.4 Workstation APPN Displays

The workstation ENPC3 was configured to use the 2216 as its NNS. Once this configuration was activated, it was possible to connect to all other APPN nodes in our test network.

The following figures show APPN information for ENPC3.

```
************************************
*       SNA Global Information       *
************************************
Network name                                USIBMRA
Control point (CP) name                     ENPC3A
Physical unit (PU) name                     ENPC3A
Node ID (for XID)                           X'05D00000'
CP alias                                    ENPC3A
Node type                                   End node
CP local address                            Not used (independent LU)
Workstation serial number                   -
Machine type                                0000
Machine model number                        X'000000'
Communications Server version               5.0
Branch extender support                     No
Search required                             No
```

Figure 220. ENPC3 Global SNA Information

```
************************************
*    Link Definition Information     *
************************************
Number of links                          1
1>Link name                              LINK0001
  Adjacent node CP name
  Adjacent node type                     Network node
  DLC name                               IBMTRNET
  Adapter number                         0
  Destination DLC address                X'40002216009908'
  CP-CP session support                  Yes
  Preferred NN server                    No
  Auto-activate link                     Yes
  Transmission group number              0
  Limited resource                       No
  Limited resource timeout               30
  Inactivity timeout                     0
  Solicit SSCP session                   No
  Init self                              No
  BIND support                           Yes
  Link station role                      Negotiable
  Line type                              Switched
  Effective capacity                     16000000 bits per second
  Cost per connect time                  0
  Cost per byte                          0
  Propagation delay                      384.00 microseconds (LAN)
  User defined parameter 1               128
  User defined parameter 2               128
  User defined parameter 3               128
  Security                               Nonsecure
  Max activations attempts               0
  Physical unit (PU) name                ENPC3A
  OCDE name
  Node id received                       X'00000000'
  Node id sent                           X'05D00000'
  Use CP as PU                           No
  Permanent connection                   No
  HPR support                            Yes
  Backup host link                       No
  Primary link name
  Branch extender uplink                 No
  HPR link error recovery type           No ERP preferred
```

Figure 221. ENPC3 Link Definition Information

```
************************************
*     LU Definition Information      *
************************************
Number of logical units (LUs)               2
1>LU name                                   ENPC3A
  LU alias                                  ENPC3A
  LU type                                   6.2
  LU local address                          X'00'
2>LU name                                   AODTEST
  LU alias                                  AODTEST
  LU type                                   6.2
  LU local address                          X'00'
```

Figure 222. ENPC3 LU Definition Information

Three kinds of current status information are listed below. The first list, Directory Information, gives us a convenient way to know the current connection status.

```
************************************
*       Directory Information       *
************************************
Total directory entries                     3
Network node entries                        1
1>Network node CP name                      USIBMRA.NN2216A
  Number of associated LUs                  1
1.1>LU name                                 USIBMRA.NN2216A
    Owning CP name                          USIBMRA.NN2216A
    LU entry type                           Home
Local and adjacent node entries             1
1>Owning CP name                            USIBMRA.ENPC3A
  Number of associated LUs                  2
1.1>LU name                                 USIBMRA.ENPC3A
    LU entry type                           Home
1.2>LU name                                 USIBMRA.AODTEST
    LU entry type                           Home
```

Figure 223. ENPC3 Directory Information

```
************************************
*      Active Links Information      *
************************************
Number of active links                        1
1>Link name                                   LINK0001
  DLC name                                    IBMTRNET
  Adapter number                              0
  Destination DLC address                     X'40002216009908'
  Link activated                              Locally
  Link state                                  Active
  Deactivating link                           No
  Limited resource                            No
  Limited resource timeout                    30
  Inactivity timeout                          0
  Active and activating sessions              7
  Max I-field size                            2048
  Adjacent node CP name                       USIBMRA.NN2216A
  Adjacent node type                          Network node
  CP-CP session support                       Yes
  Connection type                             Peer
  Link station role                           Secondary
  Line type                                   Switched
  Two-way simultaneous                        Yes
  Transmission group number                   21
  Effective capacity                          16000000 bits per second
  Cost per connect time                       0
  Cost per byte                               0
  Propagation delay                           384.00 microseconds (LAN)
  User defined parameter 1                    128
  User defined parameter 2                    128
  User defined parameter 3                    128
  Security                                    Nonsecure
  Physical unit (PU) name                     ENPC3A
  HPR enabled                                 No
  Local ANR label                             X'8040'
  Adjacent node HPR level                     None
  Primary link name
  Branch extender uplink                      No
  Transmission group effective capacity       16000000 bits per second
  Transmission group cost per connect time    0
  Transmission group cost per byte            0
  Transmission group propagation delay        384.00 microseconds (LAN)
```

Figure 224. ENPC3 Active Links Information

```
************************************
* Partner LU Definition Information  *
************************************
Number of partner logical units             4
1>Partner LU name                           USIBMRA.NN2216A
  Partner LU alias                          @I000057
  Partner LU uninterpreted name
  Maximum logical record send size          32767
  Conversation security                     No
  Parallel sessions                         Supported
  Default routing preference                Native first
2>Partner LU name                           USIBMRA.NN061A
  Partner LU alias                          @I000058
  Partner LU uninterpreted name
  Maximum logical record send size          32767
  Conversation security                     No
  Parallel sessions                         Supported
  Default routing preference                Native first
3>Partner LU name                           USIBMRA.ENPC1
  Partner LU alias                          @I000059
  Partner LU uninterpreted name
  Maximum logical record send size          32767
  Conversation security                     No
  Parallel sessions                         Supported
  Default routing preference                Native first
4>Partner LU name                           USIBMRA.RAK
  Partner LU alias                          @I000064
  Partner LU uninterpreted name
  Maximum logical record send size          32767
  Conversation security                     No
  Parallel sessions                         Supported
  Default routing preference                Native first
```

Figure 225. ENPC3 Partner LU Information

Chapter 8. 3746 Frame Switching

This scenario is intended to show the use of the 3746 as a frame switching node. The frame switching function of the 3746 is called frame relay frame handler (FRFH).

FRFH takes the frames arriving on a DLCI and switches all those frames to another DLCI. The DLCIs may be on the same or differing adapters. The protocol contents of these frames are not examined by the 3746. This means even unsupported protocols (for example, IPX) when transported over frame relay can be switched by the 3746. The FRFH switching function takes place at a very low level, minimizing the load on the 3746.

Prior to ECA 170 (EC level D46130), the 3746 could switch frames between its adapters, or between 3746 and 3745 adapters, but the only method available to load the FRFH definitions was via NCP. This meant that the 3746-950 could not support FRFH. Figure 226 shows the FRFH configurations possible on the 3746. Any FRFH definitions that use 3745 adapters must be loaded from NCP. FRFH definitions that only use 3746 adapters may be loaded from NCP or CCM.

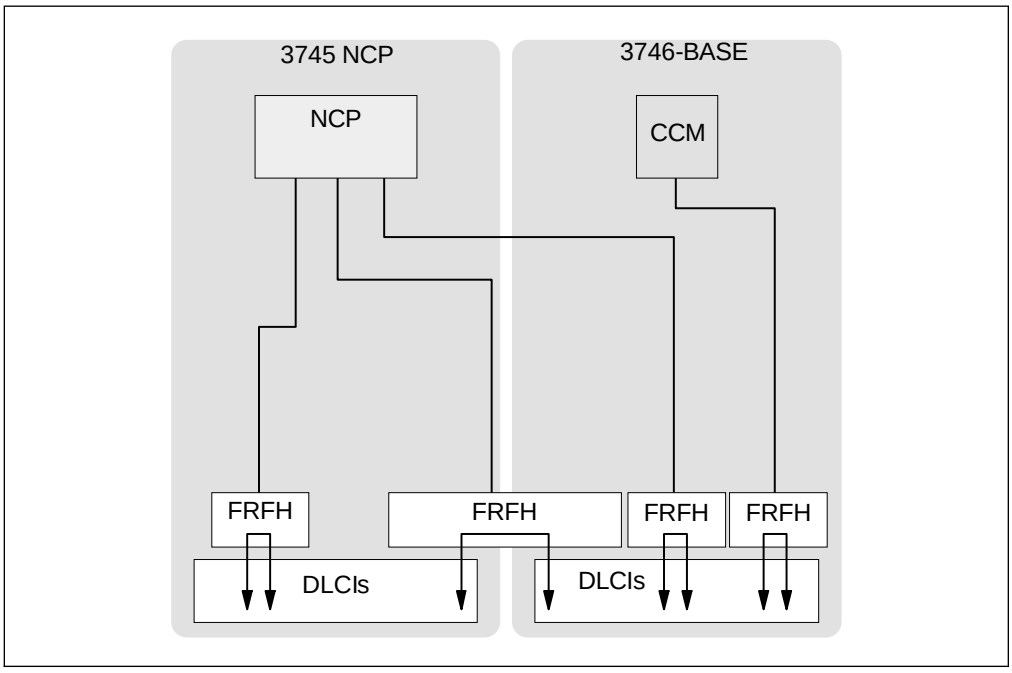

Figure 226. 3746 Frame Switching

In our test scenario, we connected the 3746 NN041A to NN061A via frame relay. This was in addition to our original connection between NN061A and NN2216A. All these nodes are still configured as APPN network nodes.

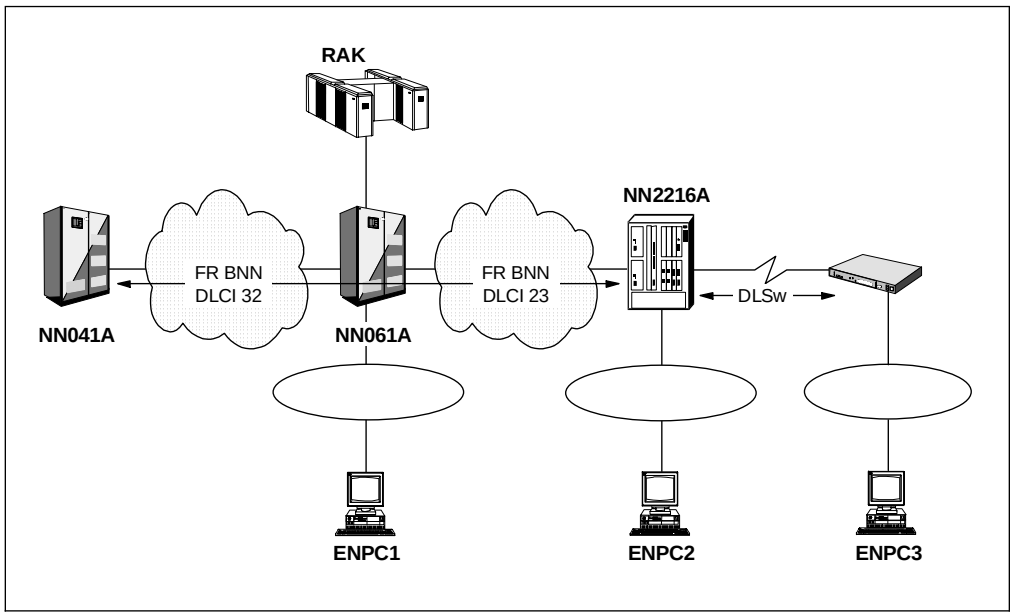

Figure 227. FRFH Test Scenario

Note: For these tests, the 2220 was no longer available. This caused us to change the FR DLCI used between NN061A and NN2216A. Originally we left NN061A on DLCI 33 and entered NN2216A on DLCI 23. The DLCI swapping was done by the 2220. For this test we are using DLCI 23 at both ends of the connection and our 3746 and 2216 are connected back-to-back. We also now use PORT 2316 in NN061A for the connection to NN2216A.

3746 NN061A was configured as an FRFH node for the DLCI 23 from NN2216A and DLCI 32 from NN041A. The effect of this from APPN's point of view is that NN041A and NN2216A now become adjacent, and CP-CP sessions are set up between these nodes. What was not expected, was that NN061A and NN2216A also set up CP-CP sessions, these sessions also ran over DLCI 23. Figure 227 shows the main port configuration screen. When FRFH is selected, the APPN check box is automatically selected too. In addition it is greyed-out and cannot be deselected. This means that APPN is automatically active on FRFH ports. This is because the APPN code must be loaded on the processor, even if only frame switching functions are needed. Hence, the APPN Parameters button is also still active. We needed to set Accept any incoming call to No on the APPN parameters definitions screen to avoid having CP-CP sessions between NN061A and NN216A.

8.1 3746 Frame Switching Definitions

FRFH definitions from CCM on the 3746 are divided into two phases:

1. FRFH DLCI definitions

   Here once the hardware port is defined, we specify which DLCI will be used for frame switching.

2. FRFH Set definitions

   Frame handler sets define the DLCI-DLCI switching of frames between the DLCIs which we previously defined.

The two stages are described in detail in the following sections.

8.1.1 3746 Frame Handler DLCI Definitions

Figure 228 shows the port configuration dialog for port 2304 on NN061A. The connection to NN2216A is via port 2304. The connection to NN041A is via PORT 2316. We configured both ports for FRFH functions only.

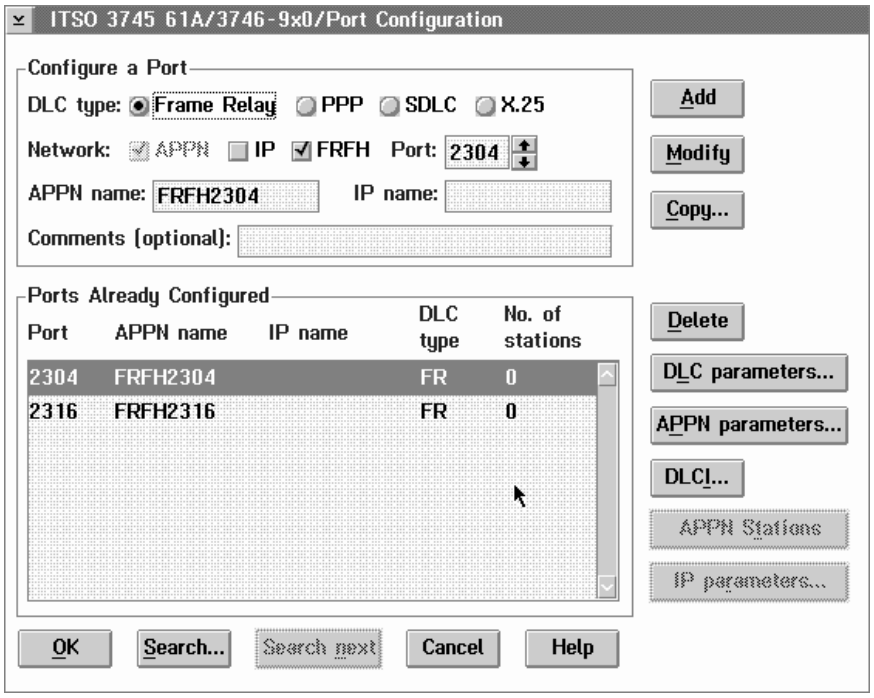

Figure 228. Port 2304 Definitions

Figure 229 on page 234 shows the DLCI configuration dialog for DLCI 32 on port 2304. We gave the DLCI the FRFH name FH230432. To return from this dialog select **OK**. In this case CCM will report an error informing you that the FRFH DLCI does not belong to an FRFH set. Select **Yes** to continue with configuration, we will define the FRFH set once all DLCIs are defined.

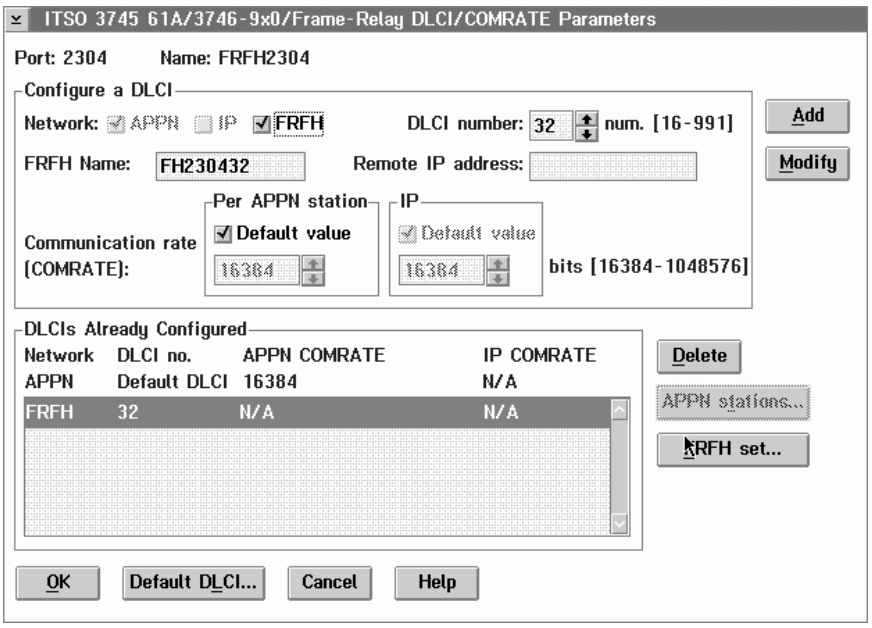

Figure 229. Port 2304 DLCI 32

Figure 230 shows the DLCI configuration dialog for DLCI 23 on PORT 3216. We gave the DLCI the FRFH name FH231623.

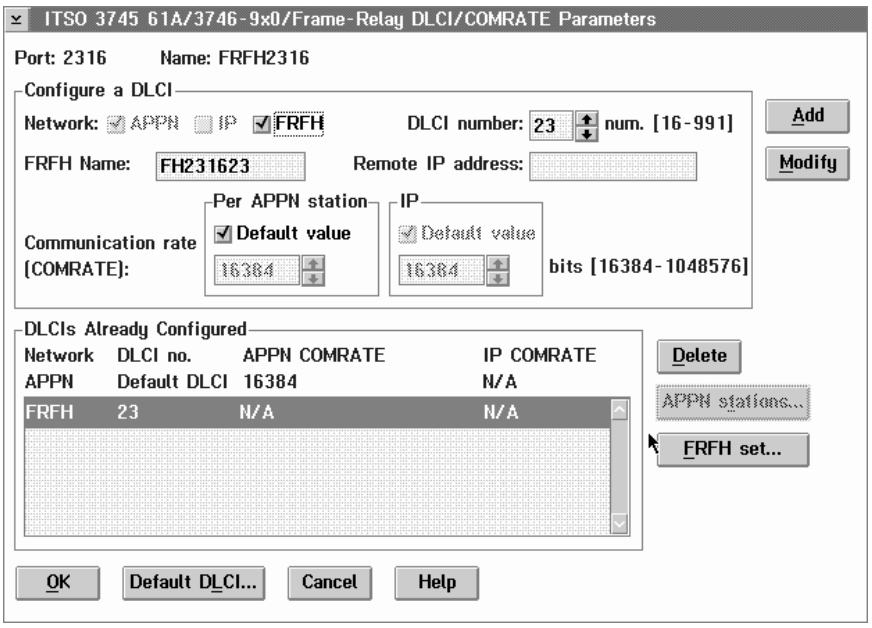

Figure 230. Port 2316 DLCI 23

8.1.2 3746 Frame Handler Set Definitions

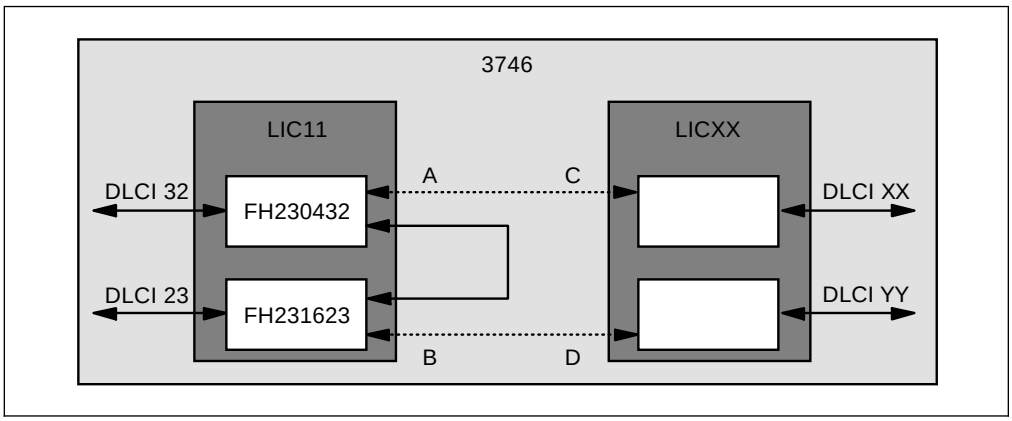

Figure 231. FRFH Set Partners

FRFH set definitions are connections made between FRFH partners. These partners are either *primary partners* or *substitute partners*. The primary switching path is between the partners A and B. For each primary partner you can also define a substitute partner. In the normal case A switches to B, but if problems occur with partner B, then the substitute partner C will be used instead. If partner A has a problem, then B will switch to D instead. This is shown in Figure 231, and Figure 232.

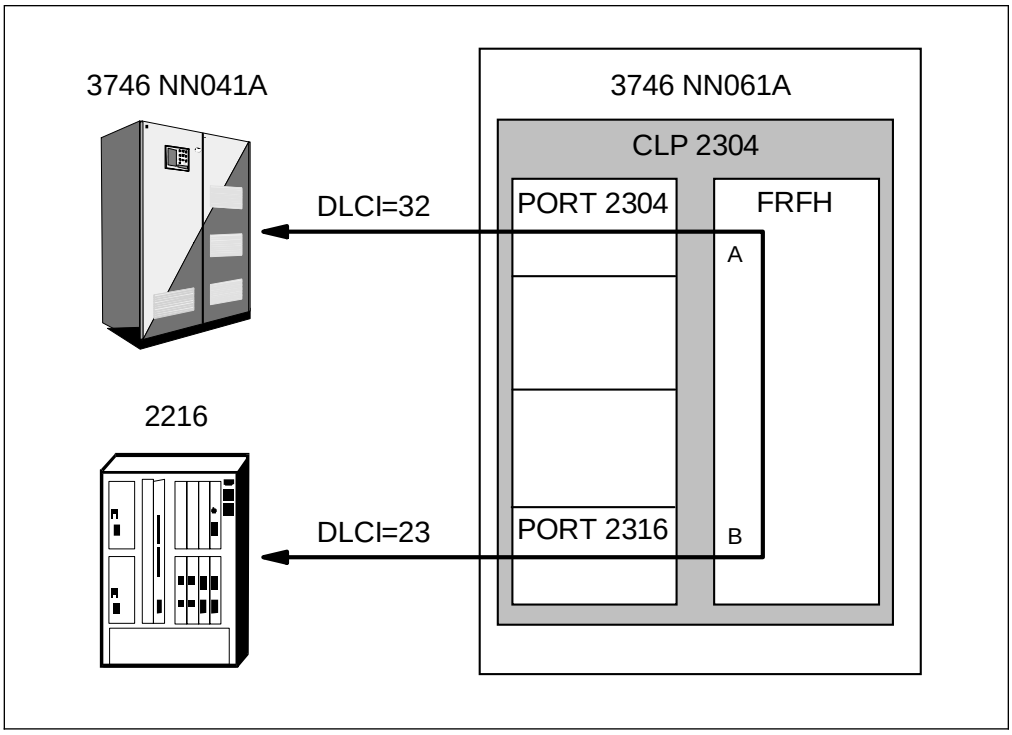

Figure 232. FRFH Configuration

The FRFH set definitions dialog can be reached in two ways:

1. From the main CCM configuration screen as shown in Figure 233 on page 236.
2. From the Frame-Relay DLCI/COMRATE parameters dialog as shown in Figure 229 on page 234 or Figure 230 on page 234.

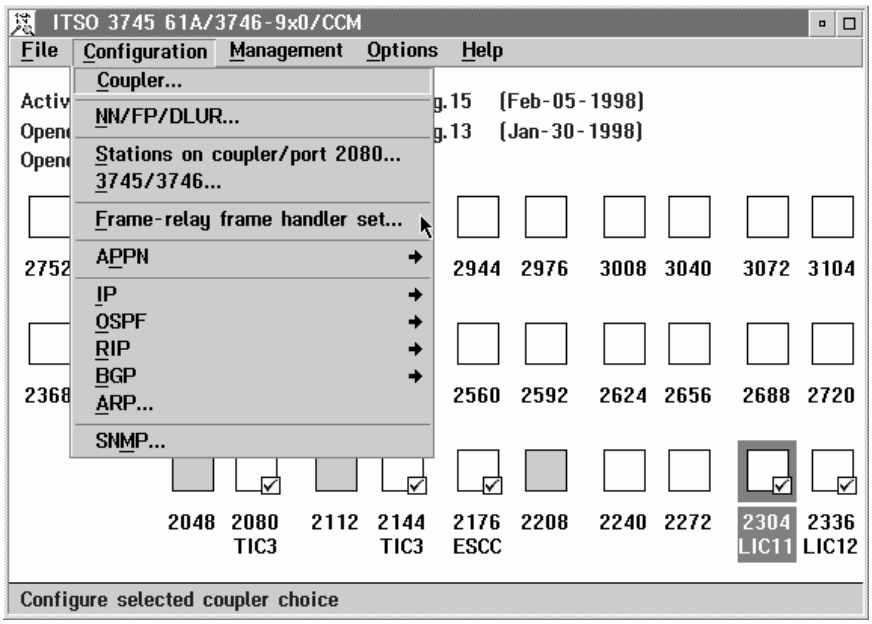

Figure 233. Frame Handler Dialog

Figure 234 shows the FRFH set definitions dialog. We gave our FRFH set the name FHSET_1, and we defined it such that the traffic from the primary partner A (port 2304 DLCI 32) will be switched to the primary partner B (port 2316 DLCI 23). In this case we did not specify any backup or substitute partners (C or D).

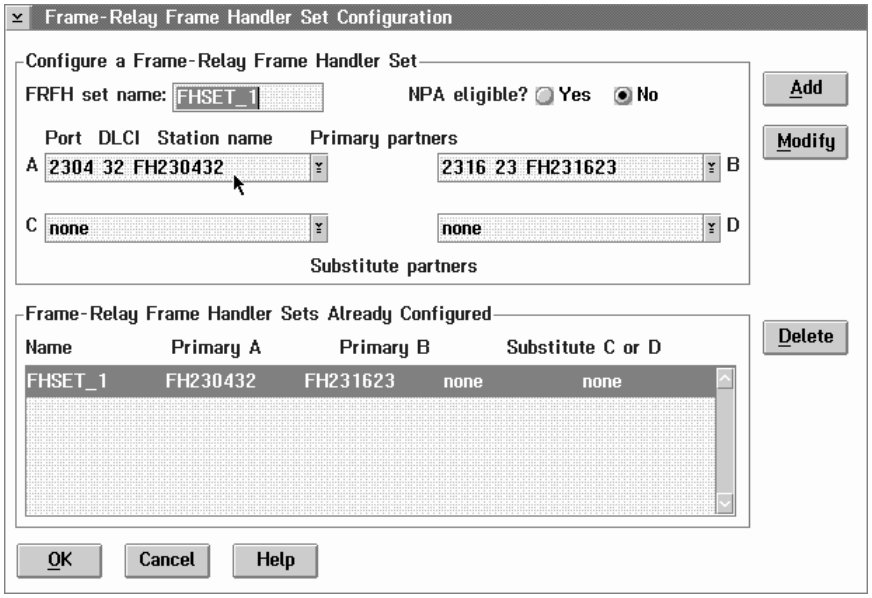

Figure 234. Frame Handler Set FHSET_1 Definitions

Figure 234 shows that when selecting one of the FRFH partners, a drop-down list appears which displays all previously defined FRFH DLCIs. Each partner can then be selected from this list.

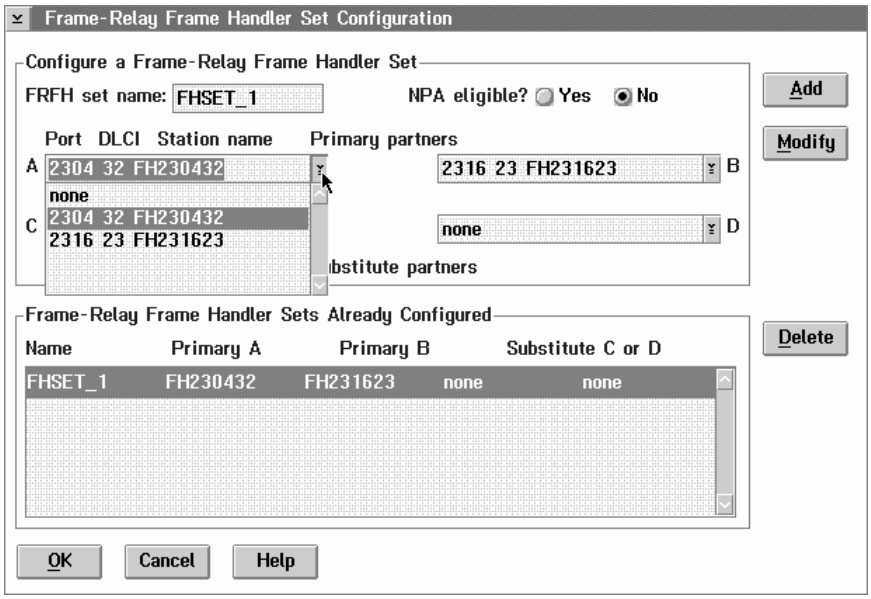

Figure 235. Frame Handler Set FHSET_1 Definitions

Appendix A. IEEE Logical Link Control 802.2

As described in 2.2, “Frame Relay, an International Standard” on page 48 all equipment attaching to a frame relay network must have implemented the layer two (DLC) DL-CORE functions. As frame relay does not guarantee message delivery and frames may be dropped, it is important that endstations implement additional DLC functions for frame acknowledgment and recovery from errors. Due to its limited congestion control frame relay networks assume also that endstations implement adequate functions for flow control.

ITU-T Q.922 specifies these additional functions and procedures within its DL-CONTROL. DL-CONTROL and DL-CORE together comprise LAPF. ITU-T Q.922 DL-CONTROL however, is optional and as the specifications were not finalized at the time product implementers started shipping their first products, most vendors use different functions and procedures.

Note: The additional DLC procedures provide end-to-end functions; it is therefore essential that endstations on either end of a virtual circuit use the same procedures.

It has been noted before that the delineation of the frame relay DLC functions is similar to the delineation of IEEE LAN standards between the Logical Link Control (LLC) and Media Access Control (MAC) sublayers. For the transport of SNA and NetBIOS traffic therefore, IBM is using IEEE 802.2 LLC, which largely resembles ITU Q.922 DL-CORE, for the additional frame relay DLC functions.

The remainder of this section explains some of the IEEE 802.2 LLC functions and procedures. For a more detailed description see the appropriate IEEE documents or refer to IBM token-ring documentation on this subject, for example, *Token-Ring Network: Architecture Reference*.

Figure 236 on page 240 depicts the format of the IEEE 802.2 LLC Protocol Data Units (PDUs). The PDUs are imbedded in RFC 1490 frames (with ITU-T Q.922 frame relay trailer and header) as shown in Figure 33 on page 71.

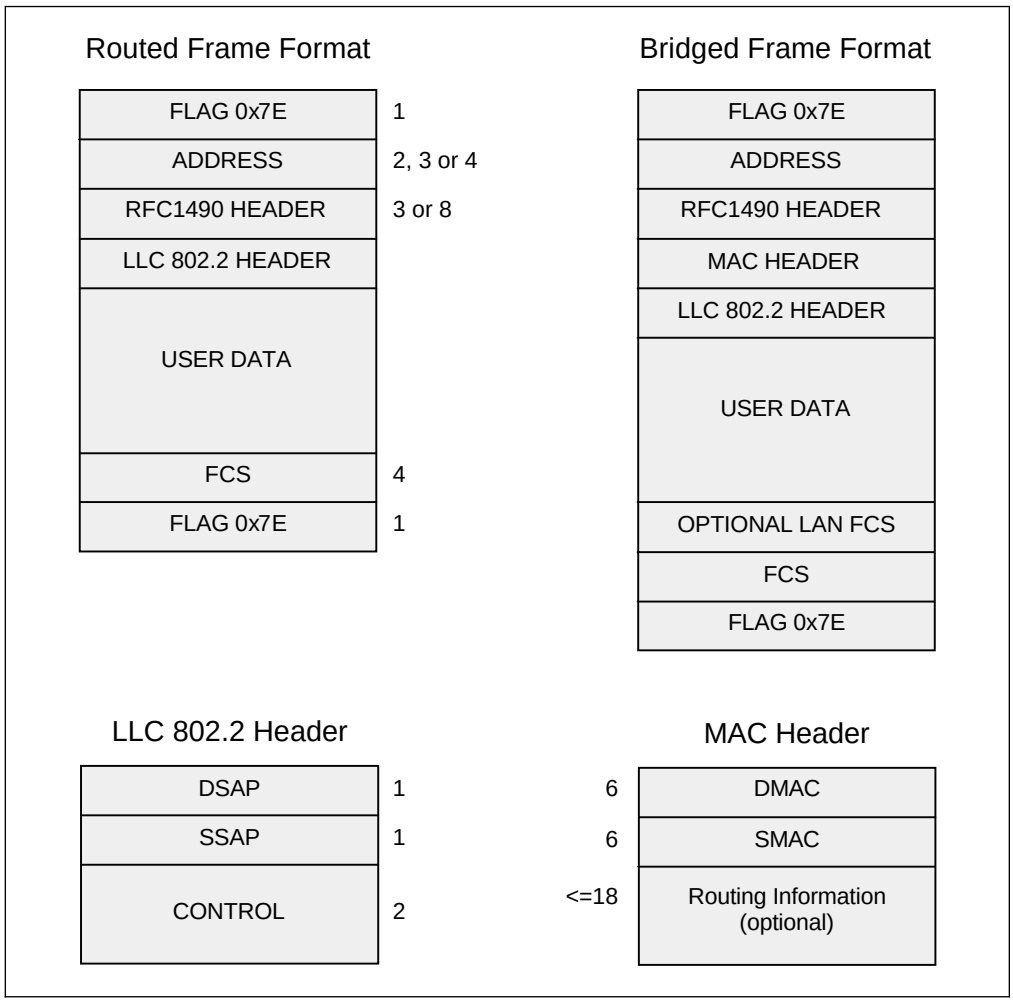

Figure 236. IEEE 802.2 LLC Protocol Data Unit Encapsulation

The PDUs contain a one-octet Destination Service Access Point (DSAP) and a one-octet Source Service Access Point (SSAP). The SAPs can be considered as the logical point at which the data link layer and the network layer connect. For SNA traffic SAPs of X'04', or multiples of this value, have been reserved.

The PDUs contain a Control field which is either one- or two-octet(s) long. For an overview of the PDUs identified by this control field see Table 20.

Table 20 (Page 1 of 2). IEEE 802.2 LLC PDUs and Control Field(s) Encoding

Format	Command	Response	Control Field(s) Encoding							
			7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
I_n	I	I	N_s							0
			N_r							P/F
S_n	RR	RR	0	0	0	0	0	0	0	1
			N_r							P/F
S_n	RNR	RNR	0	0	0	0	0	1	0	1
			N_r							P/F

Table 20 (Page 2 of 2). IEEE 802.2 LLC PDUs and Control Field(s) Encoding									
Format	**Command**	**Response**	**Control Field(s) Encoding**							
			7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
S_n	REJ	REJ	0	0	0	0	1	0	0	1
			N_r							P/F
S_u	SABME		0	1	1	P	1	1	1	1
S_u		DM	0	0	0	F	1	1	1	1
S_u	UI		0	0	0	P	0	0	1	1
S_u	DISC		0	1	0	P	0	0	1	1
I_u		UA	0	1	1	F	0	0	1	1
S_u		FRMR	1	0	0	F	0	1	1	1
S_u	XID	XID	1	0	1	P/F	1	1	1	1

PDUs are either command or response PDUs. PDUs are either supervisory (S) or information (I) format PDUs. S-format PDUs are used for layer two control messages; I-format PDUs allow the transport of higher layer data. PDUs are either numbered or unnumbered. Unnumbered PDUs contain a one-octet control field, numbered PDUs a two-octet control field. Numbered S-format PDUs are used for data acknowledgment and keep-alive messages, unnumbered S-format PDUs for connection establishment or termination. Numbered information (I_n) PDUs, or I-PDUs enable acknowledged data transport, unnumbered information (I_u) PDUs, or UI-PDUs, unacknowledged data transport. I-PDUs can only be sent after an IEEE 802.2 LLC connection has been established. UI-PDUs can be transmitted without a pre-established connection.

An I-PDU contains both a send sequence number (N_s) and a receive sequence number (N_r). The (numbered) S-PDUs contain a receive sequence number only. The send sequence number allows the receiving side to acknowledge the PDU. The receive sequence number is used to acknowledge received PDUs. Both sides maintain a send state variable (V_s) to denote the send sequence number of the next in sequence I-PDU to be sent and a received state variable (V_r) to denote the next in-sequence I-PDU to be received.

The IEEE 802.2 LLC procedures and functions use the following timers (Table 21) and counters (Table 22 on page 242).

Table 21. IEEE 802.2 LLC Timers	
Name	**Description**
T1	Reply timer
T_i	Inactivity timer
T2	Receiver acknowledgment timer

Table 22. IEEE 802.2 LLC Counters	
Name	**Description**
N1	Maximum length of information PDUs
N2	Maximum number of retransmissions
N3	Number of I-PDUs received before sending acknowledgment
N_w	Number of acknowledgments needed to increment the working window
T_w	Maximum transmit window size
R_w	Maximum receive window size

T1 is used to detect a failure to receive a required acknowledgment or response from a remote endstation. T1 is set when an I-PDU is sent, supposing T1 is not already running, or when a command PDU with the P bit set is transmitted. T1 is reset when it receives an acknowledgment of the I-PDU or a command PDU with the F bit set. When T1 expires an S-format command PDU with the P bit is set to solicit remote link status, or the U-format PDU that has not been responded to the first time is resent. This recovery is tried N2 times after which the connection is declared inoperative.

T_i is used to send keep-alive messages, which are used to detect an inoperative condition in either the remote link station or in the transmission medium. T_i is set when the last PDU sent has been acknowledged. When T_i expires the remote station is solicited using an S-format command PDU.

T2 is used in conjunction with counter N3. An endstation uses T2 to delay the sending of an acknowledgment for a received I-PDU. T2 is started when the PDU is received and reset when an acknowledgment in either an I-PDU or S-PDU is returned. When T2 expires an acknowledgement must be sent.

N1 indicates the maximum number of information bytes in an information, TEST or XID PDU. It includes the DSAP, SSAP and control field; see Figure 236 on page 240.

N2 defines the maximum number of times that a PDU is sent following the expiration of T1.

N3 is used in conjunction with T2 to reduce acknowledgment traffic by not immediately acknowledging received I-PDUs. The counter is decremented when a to-be-acknowledged I-PDU is received. When it reaches zero an acknowledgment sent. After sending an acknowledgment, which can also be due to expiration of T2, the counter is reset.

N_w is the number of I-PDUs sent and acknowledged before the working window (W_w) is incremented by one. For details see A.1.1, “Dynamic Window Algorithm” on page 243.

T_w is the maximum number of I-PDUs that an endstation may have outstanding waiting acknowledgment. For details see A.1.1, “Dynamic Window Algorithm” on page 243. R_w denotes the maximum of sequentially numbered I-PDUs that an endstation can receive from a remote link station.

Note: This value is sent in the XID information field during connection establishment. The XID receiver should set its T_w to a value less than or equal to the R_w of the XID sender to avoid overloading the XID sender.

A.1.1 Dynamic Window Algorithm

Frame relay networks assume that endstations reduce their traffic rate when the network is experiencing congestion. Endstations will either be explicitly informed about network congestion, when receiving frames with the congestion indicator (FECN, BECN) bits set, or implicitly, when the network is dropping frames.

To control their data flow certain IBM implementations (for example ACF/NCP) use the *dynamic window* algorithm as specified within IEEE 802.2 LLC. A send window indicates the number of I-PDUs an endstation is allowed to send before acknowledgment is required. A distinction is made between the working window (W_w) which is the actual send window being used, and a maximum send window (T_w) which indicates an upper boundary for W_w. W_w decreases when congestion is detected and increases, up to T_w, when the congestion is relieved.

When an endstation detects congestion through the loss of frames or (and here implementation may vary) the receipt of frames with BECN set, it resets its W_w to one. With a send window of one the link station must wait for an acknowledgment for each I-format PDU before the next I-format PDU can be sent. When N_w consecutive I-format PDUs are sent, and successfully acknowledged, W_w is increased by one. In this way W_w gradually increases until it reaches its maximum limit.

To prevent T2 and N3, which in essence try to delay acknowledgement, to interact with the dynamic window algorithm, the sending station will set the P bit to B'1' within any last I-format PDU within its working window. The remote station must then send a responding PDU with the F-bit set to B'1' as soon as possible.

Appendix B. Special Notices

This publication is intended to help system engineers, system planners, system programmers, and network administrators to understand the frame relay support offered by the 3746 Model 900 and the 3746 Nways Controller. The information in this publication is not intended as the specification of any programming interfaces that are provided by the 3746 Model 900 and the 3746 Nways Controller. See the PUBLICATIONS section of the IBM Programming Announcement for the 3746 Model 900 and the 3746 Nways Controller for more information about what publications are considered to be product documentation.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

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

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any pointers in this publication to external Web sites are provided for convenience only and do not in any manner serve as an endorsement of these Web sites.

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

ACF/VTAM
AIX
AS/400
DB2
IBM
MVS/ESA
Nways
OS/390
RS/6000
S/390
VTAM
IBM®

Advanced Peer-to-Peer Networking
APPN
CUA
ESCON
MVS
NetView
OS/2
PS/2
RXR/2
System/390
400

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc.

Java and HotJava are trademarks of Sun Microsystems, Incorporated.

Microsoft, Windows, Windows NT, and the Windows 95 logo are trademarks or registered trademarks of Microsoft Corporation.

PC Direct is a trademark of Ziff Communications Company and is used by IBM Corporation under license.

Pentium, MMX, ProShare, LANDesk, and ActionMedia are trademarks or registered trademarks of Intel Corporation in the U.S. and other countries.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Other company, product, and service names may be trademarks or service marks of others.

Appendix C. Related Publications

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

C.1 International Technical Support Organization Publications

For information on ordering these ITSO publications see “How to Get ITSO Redbooks” on page 249.

- *IBM 3746 Nways Controller Model 950 and IBM Model 900 - APPN Implementation Guide*, SG24-2536
- *IBM 3746 Nways Controller Model 950 and IBM Model 900 - IP Implementation Guide*, SG24-4845
- *IBM Frame Relay Guide*, GG24-4463
- *TCP/IP Tutorial and Technical Overview*, GG24-3376
- *The Basics of IP Network Design*, SG24-2580

C.2 Redbooks on CD-ROMs

Redbooks are also available on CD-ROMs. **Order a subscription** and receive updates 2-4 times a year at significant savings.

CD-ROM Title	Subscription Number	Collection Kit Number
System/390 Redbooks Collection	SBOF-7201	SK2T-2177
Networking and Systems Management Redbooks Collection	SBOF-7370	SK2T-6022
Transaction Processing and Data Management Redbook	SBOF-7240	SK2T-8038
Lotus Redbooks Collection	SBOF-6899	SK2T-8039
Tivoli Redbooks Collection	SBOF-6898	SK2T-8044
AS/400 Redbooks Collection	SBOF-7270	SK2T-2849
RS/6000 Redbooks Collection (HTML, BkMgr)	SBOF-7230	SK2T-8040
RS/6000 Redbooks Collection (PostScript)	SBOF-7205	SK2T-8041
RS/6000 Redbooks Collection (PDF Format)	SBOF-8700	SK2T-8043
Application Development Redbooks Collection	SBOF-7290	SK2T-8037

C.3 Other Publications

These publications are also relevant as further information sources.

- *IBM 3746 Nways Multiprotocol Controller Models 900 and 950 Overview*, GA33-0180
- *3745 Models A and 3746 Model 900 Migration and Planning Guide*, GA33-0183
- *IBM 3746 Nways Multiprotocol Controller Models 900 and 950 Migration and Planning Guide*, GA33-0349
- *IBM 3746 Nways Multiprotocol Controller Model 950 Safety Information*, GA33-0400
- *Input/Output Equipment Installation Manual Physical Planning,* GC22-7064
- *3745 All Models Advanced Operation Guide*, SA33-0097
- *3745 Connection and Integration Guide*, SA33-0129

- *IBM 3746 Expansion Unit Model 900 Console Setup Guide*, SA33-0158
- *3745 Master Index*, SA33-0172
- *IBM 3745 and IBM 3746 Alert Reference Guide*, SA33-0175
- *IBM 3746 Expansion Unit Model 900 Basic Operations Guide*, SA33-0177
- *IBM 3746 Nways Multiprotocol Controller Model 950 User's Guide*, SA33-0356
- *3745 Maintenance Information Procedures*, SY33-2054
- *3745 Service Functions*, SY33-2055
- *3745 Maintenance Information Reference*, SY33-2056
- *3745 Installation Guide*, SY33-2057
- *3745 Service Master Index*, SY33-2080
- *3746 Model 900 Installation Guide*, SY33-2088
- *IBM 3746 Nways Multiprotocol Controller Model 950 Service Guide*, SY33-2108

How to Get ITSO Redbooks

This section explains how both customers and IBM employees can find out about ITSO redbooks, CD-ROMs, workshops, and residencies. A form for ordering books and CD-ROMs is also provided.

This information was current at the time of publication, but is continually subject to change. The latest information may be found at `http://www.redbooks.ibm.com/`.

How IBM Employees Can Get ITSO Redbooks

Employees may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **PUBORDER** — to order hardcopies in United States
- **GOPHER link to the Internet** - type `GOPHER.WTSCPOK.ITSO.IBM.COM`
- **Tools disks**

  To get LIST3820s of redbooks, type one of the following commands:

  ```
  TOOLS SENDTO EHONE4 TOOLS2 REDPRINT GET SG24xxxx PACKAGE
  TOOLS SENDTO CANVM2 TOOLS REDPRINT GET SG24xxxx PACKAGE (Canadian users only)
  ```

  To get BookManager BOOKs of redbooks, type the following command:

  ```
  TOOLCAT REDBOOKS
  ```

  To get lists of redbooks, type the following command:

  ```
  TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET ITSOCAT TXT
  ```

  To register for information on workshops, residencies, and redbooks, type the following command:

  ```
  TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ITSOREGI 1998
  ```

  For a list of product area specialists in the ITSO: type the following command:

  ```
  TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ORGCARD PACKAGE
  ```

- **Redbooks Web Site on the World Wide Web**

  `http://w3.itso.ibm.com/redbooks/`

- **IBM Direct Publications Catalog on the World Wide Web**

  `http://www.elink.ibmlink.ibm.com/pbl/pbl`

  IBM employees may obtain LIST3820s of redbooks from this page.

- **REDBOOKS category on INEWS**
- **Online** — send orders to: USIB6FPL at IBMMAIL or DKIBMBSH at IBMMAIL
- **Internet Listserver**

  With an Internet e-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an e-mail note to `announce@webster.ibmlink.ibm.com` with the keyword `subscribe` in the body of the note (leave the subject line blank). A category form and detailed instructions will be sent to you.

Redpieces

For information so current it is still in the process of being written, look at "Redpieces" on the Redbooks Web Site (`http://www.redbooks.ibm.com/redpieces.html`). Redpieces are redbooks in progress; not all redbooks become redpieces, and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

How Customers Can Get ITSO Redbooks

Customers may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **Online Orders** — send orders to:

	IBMMAIL	**Internet**
In United States:	usib6fpl at ibmmail	usib6fpl@ibmmail.com
In Canada:	caibmbkz at ibmmail	lmannix@vnet.ibm.com
Outside North America:	dkibmbsh at ibmmail	bookshop@dk.ibm.com

- **Telephone orders**

United States (toll free)	1-800-879-2755
Canada (toll free)	1-800-IBM-4YOU
Outside North America	(long distance charges apply)
(+45) 4810-1320 - Danish	(+45) 4810-1020 - German
(+45) 4810-1420 - Dutch	(+45) 4810-1620 - Italian
(+45) 4810-1540 - English	(+45) 4810-1270 - Norwegian
(+45) 4810-1670 - Finnish	(+45) 4810-1120 - Spanish
(+45) 4810-1220 - French	(+45) 4810-1170 - Swedish

- **Mail Orders** — send orders to:

IBM Publications	IBM Publications	IBM Direct Services
Publications Customer Support	144-4th Avenue, S.W.	Sortemosevej 21
P.O. Box 29570	Calgary, Alberta T2P 3N5	DK-3450 Allerød
Raleigh, NC 27626-0570	Canada	Denmark
USA		

- **Fax** — send orders to:

United States (toll free)	1-800-445-9269
Canada	1-403-267-4455
Outside North America	(+45) 48 14 2207 (long distance charge)

- **1-800-IBM-4FAX (United States)** or **(+1)001-408-256-5422 (Outside USA)** — ask for:

  Index # 4421 Abstracts of new redbooks
  Index # 4422 IBM redbooks
  Index # 4420 Redbooks for last six months

- **Direct Services** - send note to `softwareshop@vnet.ibm.com`
- **On the World Wide Web**

Redbooks Web Site	http://www.redbooks.ibm.com/
IBM Direct Publications Catalog	http://www.elink.ibmlink.ibm.com/pbl/pbl

- **Internet Listserver**

  With an Internet e-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an e-mail note to `announce@webster.ibmlink.ibm.com` with the keyword `subscribe` in the body of the note (leave the subject line blank).

Redpieces

For information so current it is still in the process of being written, look at "Redpieces" on the Redbooks Web Site (`http://www.redbooks.ibm.com/redpieces.html`). Redpieces are redbooks in progress; not all redbooks become redpieces, and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

IBM Redbook Order Form

Please send me the following:

Title	Order Number	Quantity

First name Last name

Company

Address

City Postal code Country

Telephone number Telefax number VAT number

□ Invoice to customer number

□ Credit card number

Credit card expiration date Card issued to Signature

We accept American Express, Diners, Eurocard, Master Card, and Visa. Payment by credit card not available in all countries. Signature mandatory for credit card payment.

Index

Numerics

A

B

C

D

E

F

H

I

S

T

U

V

X

Z

ITSO Redbook Evaluation

3746, 2210, 2216, and 2220 Interconnectivity: Frame Relay and Related Functions
SG24-2146-00

Your feedback is very important to help us maintain the quality of ITSO redbooks. **Please complete this questionnaire and return it using one of the following methods:**

- Use the online evaluation form found at http://www.redbooks.ibm.com
- Fax this form to: USA International Access Code + 1 914 432 8264
- Send your comments in an Internet note to redbook@us.ibm.com

Please rate your overall satisfaction with this book using the scale:
(1 = very good, 2 = good, 3 = average, 4 = poor, 5 = very poor)

Overall Satisfaction _____

Please answer the following questions:

Was this redbook published in time for your needs? Yes____ No____

If no, please explain:

_____

_____

_____

_____

What other redbooks would you like to see published?

_____

_____

_____

Comments/Suggestions: **(THANK YOU FOR YOUR FEEDBACK!)**

_____

_____

_____

_____

_____

SG24-2146-00
Printed in the U.S.A.