Subarea to APPN Migration: HPR and DLUR Implementation

Jerzy Buczak, Daniela di Casoli, Gérard le Guen, Alnashir Sumar

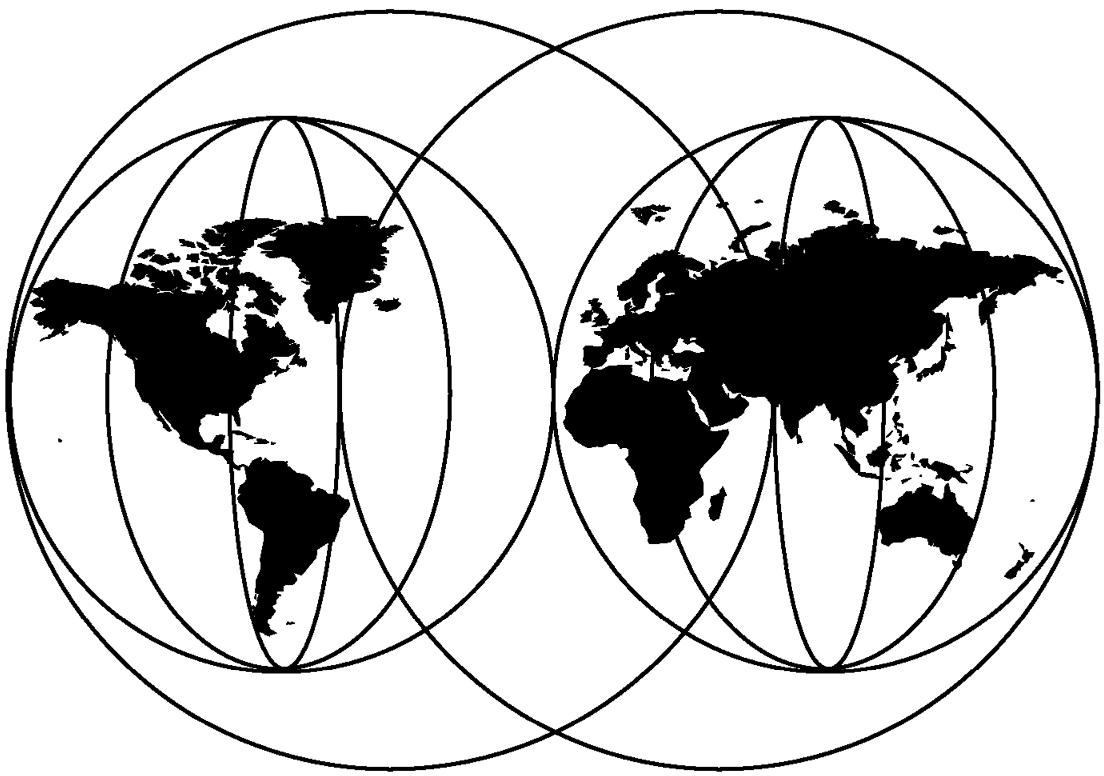

International Technical Support Organization

http://www.redbooks.ibm.com

SG24-5204-00

IBM

International Technical Support Organization

SG24-5204-00

Subarea to APPN Migration: HPR and DLUR Implementation

September 1998

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix E, “Special Notices” on page 323.

First Edition (September 1998)

This edition applies to the following products:

OS/390 Version 2, Release 5, Program Number 5647-A01
3746 Licensed Internal Code, Levels D46130 and F12380
Nways Multiprotocol Access Services Version 2, Release 2
Nways Multiprotocol Routing Services Version 2, Release 2
eNetwork Communications Server for OS/2, Version 5

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

Contents

Figures

Preface

This redbook is the second of two volumes containing detailed coverage of the migration of a subarea network to an APPN network. It focuses on the migration steps and requirements that can be used as guidelines by others to accomplish the migration in a simple and constructive manner. The first volume (SG24-4656-01) covers the implementation of the basic VTAM APPN functions including Border Node, VR-TG and network management. The second volume completes the coverage of a typical customer′s network using HPR, DLUR and APPN/HPR routers such as 3746, 2216 and 2210. In each case tested examples and definitions illustrate the theory.

This redbook will be helpful to anyone considering, planning, or supporting migrations of their existing subarea networks to APPN/HPR. Some knowledge of SNA subarea networks and familiarity with the functions, terms and data flows of APPN networks is assumed.

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.

Jerzy Buczak is an IT Consultant at the Systems Management and Networking ITSO Center, Raleigh. He writes extensively and teaches IBM classes worldwide on VTAM and APPN. Before joining the ITSO in 1996, Jerzy worked for Networking Systems in the UK. He has 17 years experience in SNA networking, network management, and a wide variety of product implementations. Jerzy holds an M.A. degree in mathematics from Cambridge University, England.

Daniela di Casoli is an IT Specialist with IBM Italy. She has eight years of experience in defect and non-defect networking support on the MVS platform.

Gérard le Guen is a certified IT/AP specialist with IBM France. He supports Enterprise Systems networking products, mainly in tuning and performance activities. Before taking up his current position, Gérard was the national specialist for OSI, and produced a redbook on the interoperability of open systems. Gérard has also been a customer engineer for 20 years, holding country-level positions in both hardware and software support.

Alnashir Sumar is a Networking IT Specialist with IBM Canada. Before joining IBM, Alnashir spent 10 years as a Senior Technical Analyst specializing in SNA networking, VTAM and CICS, serving the insurance industry.

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

Andrew Arrowood
APPN Development, Research Triangle Park

Roy Brabson
CS OS/390 Development, Research Triangle Park

Paul Braun
Systems Management and Networking ITSO Center, Raleigh

Jason Cornpropst
APPN Development, Research Triangle Park

Mac Devine
CS OS/390 Development, Research Triangle Park

Brian Dorling
Systems Management and Networking ITSO Center, Raleigh

Nancy Gates
Networking Systems Center, Gaithersburg, MD

Mike Haley
Systems Management and Networking ITSO Center, Raleigh

Johnathan Harter
CS OS/390 Development, Research Triangle Park

Lap Huynh
CS OS/390 Development, Research Triangle Park

Tim Kearby
Systems Management and Networking ITSO Center, Raleigh

Chris Mason
International Education Centre, La Hulpe, Belgium

John Parker
IBM United Kingdom

Sam Reynolds
CS OS/390 Development, Research Triangle Park

Juan Rodriguez
Systems Management and Networking ITSO Center, Raleigh

Carla Sadtler
Systems Management and Networking ITSO Center, Raleigh

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 337 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

Chapter 1. High-Performance Routing

In the first volume of this redbook (SG24-4656-1) we described the processes necessary to migrate a subarea SNA network to APPN. We covered:

- Conversion of subarea VTAM domains to APPN nodes
- Conversion of LEN connections to APPN connections
- Implementation of VTAM as an APPN border node
- Implementation of VR-TG to combine subarea and APPN protocols on the same connections
- How SSCP takeover is affected by APPN migration
- How VTAM combines APPN and subarea network architectures when performing network searches and route calculation
- Considerations when implementing APPN in a sysplex
- Network management in a combined subarea/APPN network

This volume has two major objectives that were not addressed by the first volume:

- To describe the implementation of high-performance routing (HPR). HPR is an extension to APPN which can improve performance and availability with minimal impact on network definition and administration. Although the conversion from base APPN to HPR is much more straightforward than that from subarea to APPN, we feel that HPR is sufficiently important to justify a separate redbook.
- To extend the scope of our tested scenarios outside the narrow realm of VTAM and NCP. In the first volume most of VTAM′s partner APPN nodes were PCs running Communications Manager or Communications Server, with dependent LUs directly attached to a subarea node. Here we describe:
  - The use of APPN/HPR routers such as the 3746, the 2216 and the 2210 to build a comprehensive SNA network.
  - The use of dependent LU requester/server functions to free the network from the restrictions imposed by the subarea architecture. Although an APPN network can support dependent LU sessions without this function, DLUR/S permits such sessions to take the optimum route *all the way* to the APPN node supporting the dependent LUs. By the same token, these sessions can enjoy the benefits of HPR all the way from the dependent LU requester node to the application.

Note

This pair of redbooks covers migration from subarea to APPN, and then to HPR, in two separate volumes. We must emphasize that this is purely to help the reader understand APPN and HPR, and is not intended as a recommendation to perform the migration in two phases. The difference between APPN and HPR is small compared to the difference between subarea and APPN. You may find it advantageous to let VTAM′s defaults take you all the way to HPR as you implement APPN.

1.1 Why HPR?

High-performance routing (HPR) is a set of enhancements for APPN whose main objectives are:

- **Improved APPN data routing**

  With faster, more reliable communication lines it is neither necessary nor desirable to perform error recovery, flow control and complicated routing functions at intermediate nodes in a network. HPR takes full advantage of modern technology to eliminate these functions, by ensuring that they are performed only at the endpoints of a session path. With HPR, intermediate nodes have only a minimal switching function to perform.

- **Improved APPN reliability**

  Customers have always wanted the network to recover from errors, and to find an alternate path without requiring the end user to take action. HPR switches session routes to bypass link and node failures if an acceptable alternate path is available. This occurs transparently to the sessions; in other words, the session is not disrupted.

- **Compatibility and easy migration**

  HPR implements many functions in exactly the same way as APPN does. The same topology, the same directory, the same search methods, and the same route calculations are used. Priority queueing, connection networks, DLUR/S, VR-TG and cross-network sessions are all supported by HPR. The same management data is carried on the same sessions in the same ways. HPR functions are invoked only as the session initiation request flows through nodes capable of performing those HPR functions.

  Migration is easy and low-cost because nodes can be upgraded from APPN to HPR individually without affecting the rest of the network. As the network is migrated, HPR will gradually be utilized on larger and larger tracts of it while maintaining the same appearance to the users.

- **Migration path to gigabit networks**

  The HPR architecture has similar design objectives to ATM, and is built on similar principles. As link speeds overtake node processing capabilities, it becomes necessary to remove as much processing as possible from intermediate routing nodes and divert it to endpoint nodes, so that the intermediate nodes have very little to do. HPR therefore provides a smooth migration path between the legacy subarea networks of yesterday and the integrated gigabit networks of tomorrow.

The intermediate nodes in an HPR network have only one function, to route data quickly and efficiently from one link to another. This technique is known as *automatic network routing (ANR)*.

All the rest of the HPR function is implemented in the endpoints of a connection. The major part of this function is called *rapid transport protocol (RTP)*. RTP provides a reliable end-to-end connection for sessions using HPR, and includes the following components:

- *End-to-end error recovery* removes the need for intermediate routing nodes to detect, check, acknowledge and retransmit packets in error. The only checking done at intermediate nodes is for data integrity (CRC checking), which is normally performed in the adapter and does not incur a processing overhead. Packets that fail the CRC check are simply discarded.

- *Adaptive rate-based (ARB)* flow control provides a congestion avoidance and control mechanism that removes from intermediate nodes the need to do adaptive pacing.
- The *non-disruptive path switch* function changes a session route without affecting the flow of data on the session.
- The *APPN/HPR boundary function* allows a network to include both base APPN and HPR portions, providing translation of the appropriate protocols at the boundaries.

1.2 Automatic Network Routing

Automatic network routing is a low-level routing mechanism that minimizes processing cycles and storage requirements for routing packets through intermediate nodes.

ANR is a source-routing protocol; the routing information (ANR label) for every node on the path is contained in the packet header. Furthermore, the ANR label represents the onward link for each node, *not* the session as with APPN ISR. There is *no* session awareness in a node performing ANR routing. All it has to do is inspect the first ANR label in the packet header, strip it off, and forward the packet to the correct outbound link.

This label stripping technique is much more efficient then the label swapping technique (ISR) used by base APPN nodes. ISR requires that the node inspects the LFSID in the incoming packet, uses it to look up a session table, swaps it to the outbound LFSID, and forwards the packet.

Traffic flowing on HPR connections uses network layer packets (NLPs) as opposed to FID-2 PIUs. Since HPR nodes must be able to support *both* NLPs and FID-2s on the same link, they are distinguished by the first two bits in the header. A FID-2 PIU starts with B′0010′, whereas an NLP starts with B′1100′ or B′1101′.

Figure 1 on page 4 illustrates the way ANR labels are used. 1.3, “Rapid Transport Protocol” on page 4 describes how they are assigned in the first place.

1. Node A sends an NLP to node B with ANR labels 21 / 33 / 65 / FF in the NLP header as shown.
2. Node B looks in the header for the first ANR label. This is 21, so node B removes it from the header and transfers the truncated NLP to the link it knows as 21.
3. Node C receives the NLP, removes the next ANR label (33), and sends the NLP on the link it knows as 33.
4. Node D receives the NLP and recognizes that the next ANR label (65) represents not a link but the endpoint of the RTP connection. Therefore, it passes the data in the NLP to the higher protocol layers for processing.
5. The response to this message takes a similar course through the network in the opposite direction.

In the example, the ANR labels are one byte in length but this is not necessarily so; different products implement different lengths of label. Since each node on an HPR path assigns the ANR labels which it is to interpret, there is no need for any other node to be aware of their length or meaning. Each node will find its own label at the start of any NLP it receives.

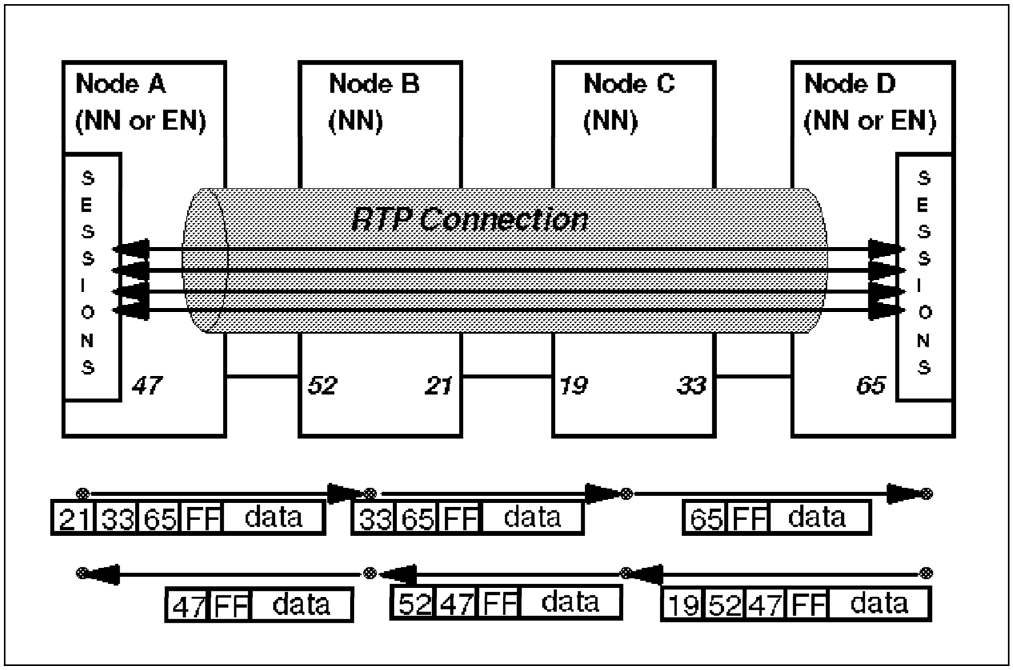

Figure 1. Use of ANR Labels

On a FID-2 connection, the session information held in each node contains the transmission priority. With HPR there is no session awareness in ANR nodes, therefore the transmission priority is carried in the NLP headers. Priority queueing is implemented on HPR links (or combined HPR/FID-2 links) in the same way as on base APPN (FID-2 only) links.

ANR is significantly faster than current APPN routing as seen in VTAM testing. The intermediate nodes have between 3 and 10 times less code to execute to route with ANR. This is offset to some extent by the increase in code at the RTP endpoints. Also, no intermediate node storage is required to maintain routing tables and no pre-committed buffers are necessary for each session.

1.3 Rapid Transport Protocol

Rapid transport protocol creates a logical connection between two HPR endpoints, such that ANR routing can be used for sessions between them. Although an RTP connection is conventionally represented as a *pipe*, this can be misleading. There is no connection information held at intermediate nodes; only the endpoints know the route. If the route should change, the intermediate nodes are not made aware of the fact. In some respects, an RTP connection can be compared to a TCP connection.

Nearly all the search and session setup mechanisms in HPR are exactly the same as in base APPN. The searches flow the same way and the session RSCV is calculated in exactly the same way, except that some extra HPR-related information flows at certain times. Note that HPR capability is not a node or TG characteristic for the purpose of APPN route calculation. The initial calculation of an APPN route takes no account of HPR, in order to preserve compatibility with base APPN. When a path switch is required, however, the opposite is true; an HPR-capable route *must* be chosen for the new path.

Session endpoints may reside in HPR nodes, or in APPN nodes. If only part of a session route is HPR-capable, there will be one or more APPN/HPR boundaries on the path. The nodes containing the boundary function act as RTP endpoints for the session. The RTP endpoints are identified by *network connection endpoints* (NCEs); an NCE identifier looks like an ANR label and is the last label in the routing part of the network layer header (NHDR). An NCE may represent an LU, a group of LUs, or an APPN/HPR boundary. If the node supports Control Flows over RTP (see 1.8, "HPR Control Flows over RTP Option" on page 12), the NCE may also represent a CP or the Route Setup function. A successful reply to a search will contain the NCE representing the destination LU, if CP(DLU) supports RTP.

The point where base APPN and HPR processing diverge comes as the BIND is about to flow across the network. In base APPN, each node on the session path examines the RSCV in the BIND and builds suitable session tables before sending the BIND to the next node. With HPR, this changes as soon as the BIND reaches an RTP-capable node; this node may be CP(PLU) itself, the session endpoint and the BIND origin.

When an RTP-capable node on the session path sees the BIND, it does the following:

- It examines the session RSCV, starting with its own entry, for a continuous sequence of ANR-capable nodes. Once it reaches the last ANR-capable node on the path it scans back until it finds an RTP-capable node. This sequence (between this node and the furthest RTP-capable node) is now a candidate for an RTP connection. It has an RTP-capable node at each end with any intermediate nodes able to perform ANR.
- If no such RTP-capable route is found, the BIND is forwarded into the network in the base APPN fashion after the session tables have been built.
- If an RTP-capable route is found, the node then checks whether there is already a suitable RTP connection that this new session can use. Such an RTP connection must have the same RSCV (for this part of the route) and the same APPN class of service as the new session. If the RTP partner is also the session endpoint (and thus the NCE to be used for the new session is known), then the NCEs must also match. If the RTP partner is an APPN/HPR boundary function, then the existing pipe must also be to such a function.
- If a suitable RTP connection already exists, the BIND is sent on that pipe as an NLP. The intermediate nodes on the RTP connection do not see the BIND.
- If no suitable RTP connection exists, a request called a *Route Setup* is sent to the RTP partner. Route Setup flows like a BIND, using the session RSCV (or that part of it between the RTP partners) to navigate through the network.

  The Route Setup flow is what establishes the ANR labels for an RTP pipe. As Route Setup passes through the network, each node converts the RSCV to ANR routing information; the RSCV tells it which TG to which node is next on the route, and the ANR node assigns a label to represent that TG. By the time Route Setup has reached the destination node, the ANR routing information is complete for this direction. Now the Route Setup Reply must flow in the opposite direction, picking up new ANR labels representing the same links in the reverse direction. When the Route Setup reply reaches the CP at the PLU end of the RTP connection, the complete set of ANR labels (including the NCEs) in each direction is known. Note that the NCE of an APPN/HPR boundary function is discovered at Route Setup time, not at the time of a successful search. In Figure 1 on page 4 the NCEs at the ends of the RTP connection have labels 65 and 47, and these labels appear at the end of the ANR routing information for NLPs flowing in the appropriate direction.

- From now on all flows are carried in NLPs, but more information is needed to complete the RTP connection setup. The RTP connection is identified at each end by a *transport connection identifier* (TCID); the NCEs alone are not enough to identify an RTP pipe since two RTP endpoints may have many RTP connections between them (different COSs, different routes). The two TCIDs are assigned in the first flows on a new RTP pipe (the *connection setup* and *connection identifier* segments), and thereafter are carried in the transport layer header (THDR).
- Once the new RTP connection has been set up, the BIND flows as an NLP to the RTP partner. Thus the entire session route is established by:
  - The BIND, on non-HPR parts of the route.
  - The Route Setup, on HPR parts of the route. The BIND flows as an NLP on those parts of the route where the Route Setup has established RTP connections.

Figure 2 shows the basics of APPN session establishment, while Figure 3 on page 7 shows the HPR version. Note that in this example, the RTP endpoints are also the session endpoints. The Locate flows and the BIND are always between the session endpoints, whereas the Route Setup and the transport connection flows are between the RTP endpoints.

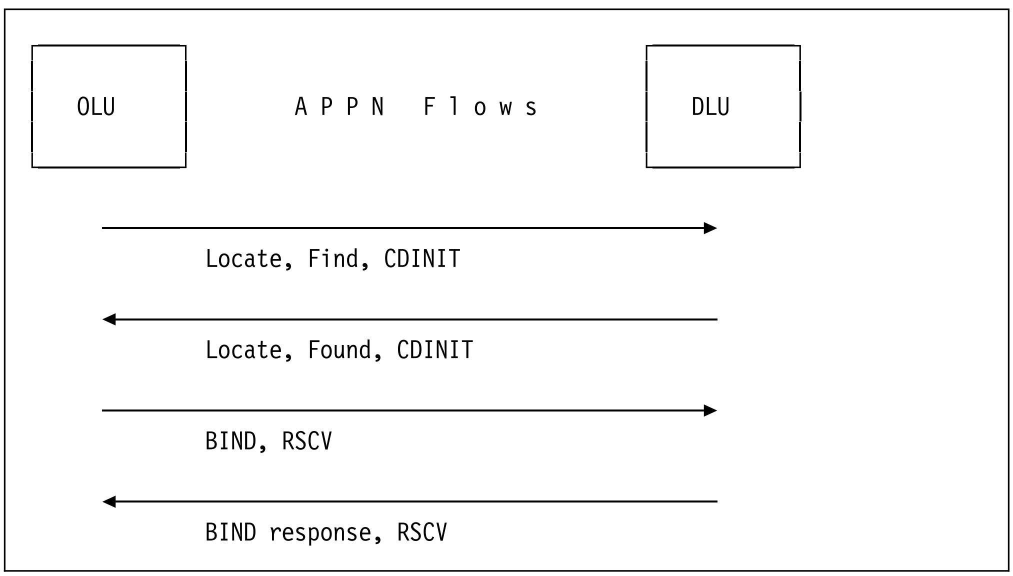

Figure 2. APPN Session Setup

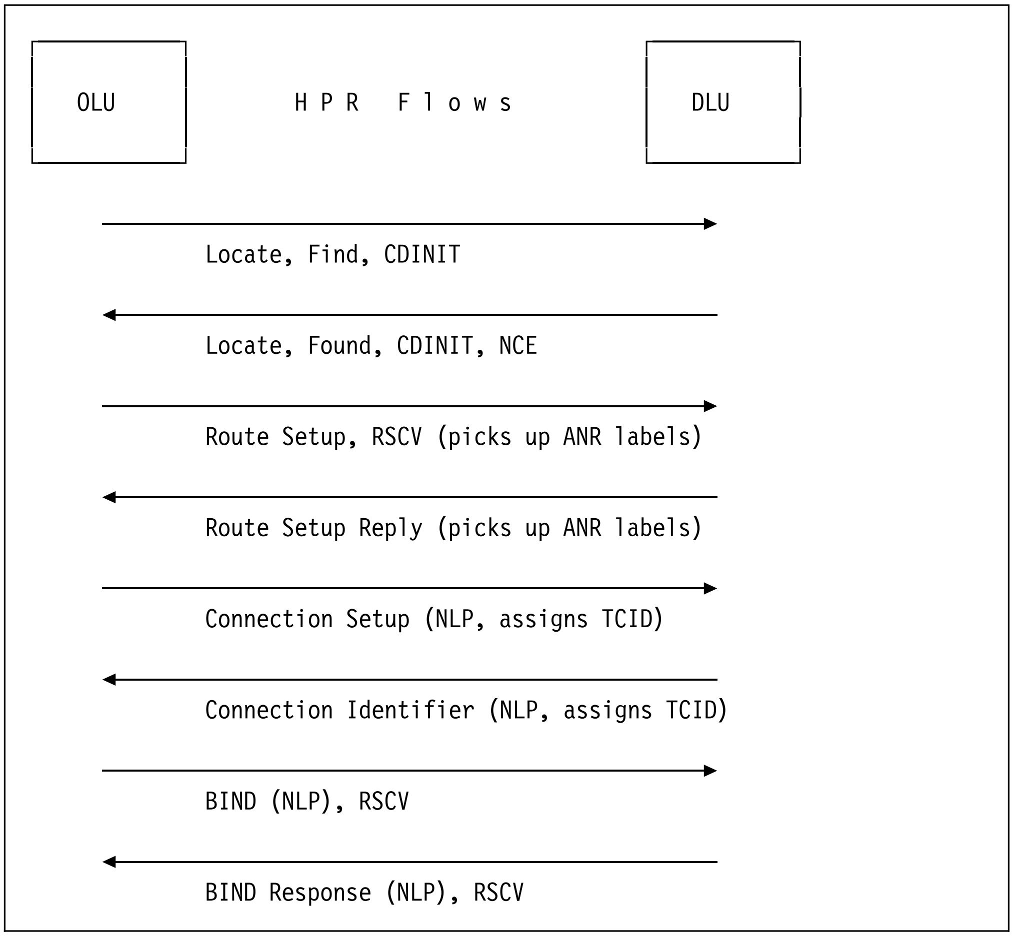

Figure 3. HPR Session Setup

The TCIDs are sufficient to identify an RTP pipe, but not enough to identify a session since many sessions may use the same pipe. Since there is no LFSID in HPR, sessions are distinguished by a *session address*. This address is assigned by the RTP endpoints as the BIND and BIND response flow, and is carried in a new FID-5 header within the NLP. The session address is unique in each direction.

Thus an NLP will contain the network layer header (with ANR labels), a transport layer header (with a TCID) and, if there is data, a FID-5 header (with a session address). Figure 4 illustrates.

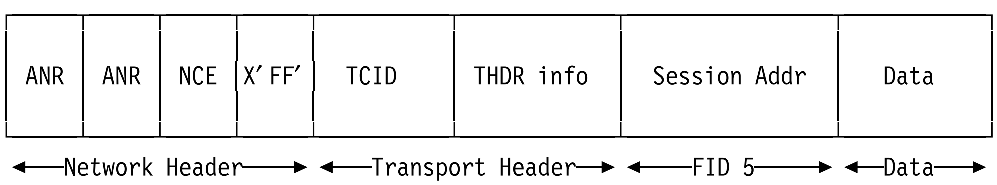

Figure 4. Format of Network Layer Packet

As with base SNA, data can be segmented to fit into NLPs and reassembled at the far end of the connection. Segmenting and reassembly is done by the RTP endpoints, never at the intermediate nodes. However, NLP headers (NHDR, THDR, FID5) cannot be segmented. As the reader will appreciate these headers can be quite lengthy; therefore, HPR demands that each link on an RTP connection can support a minimum of 768 bytes in an NLP. The maximum packet size for each link on an RTP connection is discovered on the Route Setup flows, so that both RTP partners know when a packet must be segmented.

Because an RTP connection may traverse multilink TGs as well as connectionless transport networks, it is possible for packets to arrive at their destination in the wrong order. RTP endpoints, therefore, must have the ability to resequence packets that have arrived out of order.

1.4 End-to-End Error Recovery

In the past, error recovery on each link (hop) of a network route has been necessary because of relatively high link error rates. Improvements in link quality have made it feasible and desirable to provide end-to-end error recovery in place of error recovery at each hop. HPR provides this capability by:

- Utilizing existing data link controls in such a way that they bypass link-level error recovery. For example, traffic over a LAN may be sent as unnumbered information (UI) frames.

  Note that some DLCs do not permit bypassing error recovery (channel and X.25 connections, for example). Also, individual product implementations may not be compatible with each other on a particular DLC.

- Using the rapid transport protocol (RTP) to do end-to-end error recovery on pipes. RTP retransmits only those packets that have failed to reach the receiver (selective retransmission).

  The HPR data stream is treated by the RTP endpoints as a continuous byte stream, rather like TCP. Each NLP that carries data contains the byte sequence number of the first data byte in the NLP. By counting the lengths of the data portions and comparing them to the byte counts the receiver can determine if data packets are missing. If data loss is detected, the receiver sends a *status segment* in a transport header to inform the sender of where, and how long, the missing portion(s) are. The sender can then re-transmit only those portions of data that have been lost.

Note that the sender must buffer all packets sent until an acknowledgment is received from the endpoint. The RTP acknowledgment is requested by the RTP endpoints at the RTP level, not the session level.

1.5 ARB Flow and Congestion Control

The APPN hop-by-hop window-based congestion control algorithm (adaptive session-level pacing) is inappropriate for high-speed networks, and in any case is not usable on an RTP connection where the intermediate nodes have no session awareness. Adaptive pacing is still used in HPR for end-to-end flow control, but the whole RTP connection is treated as a single hop. To control throughput on a high-speed RTP connection, HPR has developed an algorithm called adaptive rate-based (ARB) flow/congestion control. This regulates the flow of data over a pipe by adaptively changing the rate at which the sender is allowed to send data into the network, based on the rate at which the receiver and the network are able to process the data. This algorithm allows for high link utilization and prevents congestion in the network or at the receiver. Figure 5 on page 9 shows the difference between APPN and HPR flow control.

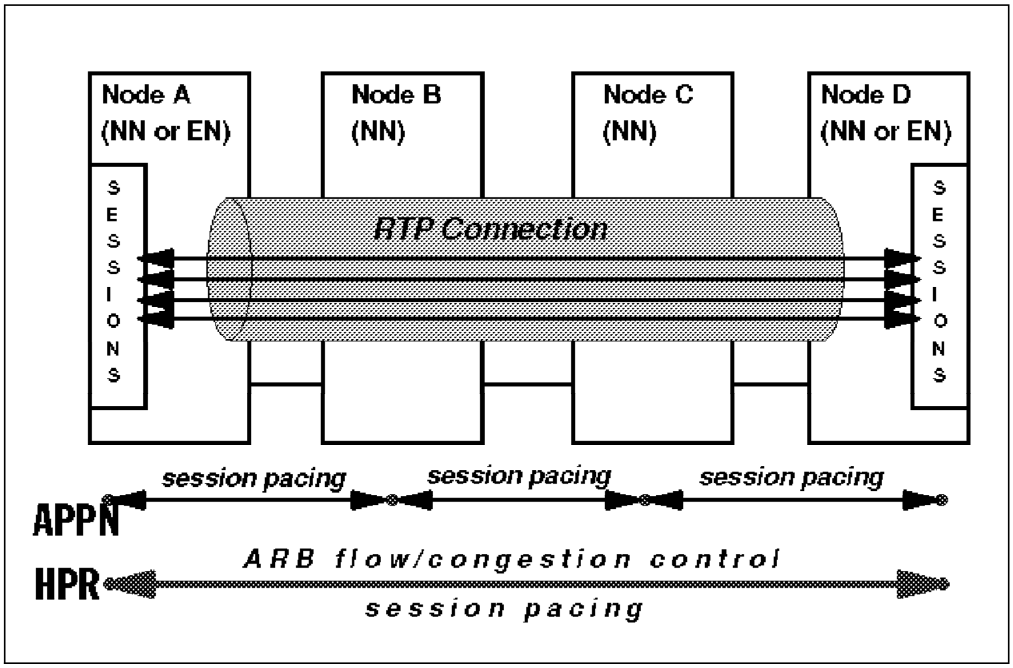

Figure 5. Flow Control in APPN and HPR

The ARB algorithm is self-adjusting in that it will balance itself between the traffic requirements and the load on the network. However, it needs to be initialized with reasonable values at the start of the RTP connection. ARB finds out the lowest capacity link on the HPR route, and uses 10% of that value as a starting point for the sending rate. This is felt to be high enough to ensure that capacity is attained quickly, but low enough not to cause congestion due to the initial extra load on the network.

The Route Setup and its reply are used to determine the lowest capacity link on the connection; each node on the path checks the APPN CAPACITY characteristic of the next link on the RSCV, and substitutes it into the Route Setup if it is lower than the value it finds there. It is important to get the CAPACITY values correct; if ARB is initialized with too low a value (for example, 8K instead of 32M), it can take a very long time to reach capacity and the performance gains from HPR may not be realized. Prior to the introduction of HPR, CAPACITY was used only for the purpose of route calculation.

Most other congestion control schemes merely react once congestion has occurred. ARB is designed to detect the onset of congestion, thereby ensuring that congestion leading to the discarding of data packets does not occur.

For a more detailed description of the way ARB flow control works, please see Appendix A, “Adaptive Rate-Based Flow and Congestion Control” on page 231.

1.6 Non-Disruptive Path Switch

If an intermediate link or node on the path of an RTP connection fails, a new path for the RTP connection is selected that best fits the same class of service as the failed connection. The connection is then moved to the new path and data traffic is resumed (including recovery of lost packets) without notifying the higher layer protocols. This means that SNA sessions using this connection are not interrupted and session data flows will see a short delay, but no timeout or loss of data.

Path switch can also preserve a session across a failure even if there is no alternate route available. If the network connection (typically a shared facility such as frame relay or token-ring) recovers in a timely manner, the new path may take the same route as the old one.

A path switch may be initiated when a link failure is detected, or when a timeout procedure fails to elicit a response. A path switch will normally involve the same flows as when the connection was initially set up: Locate searches and Route Setup (but not, of course, the BIND).

The path switch need not be initiated by the node that requested the session (or the RTP connection) in the first place; therefore, the new path may be calculated by a different node from the one that calculated the original route.

At initial session setup, the route calculated by NNS(PLU) will take no account of HPR, even if CP(PLU) is HPR-capable. This is to preserve compatibility between APPN and HPR. However, at path switch time there is no choice; an RTP-capable route must be selected. Therefore, each NN that is capable of ANR must also be capable of calculating HPR-only routes. New bits in the TDUs (CV45 and CV46) carry HPR-related information, which is understood and stored only by HPR-capable nodes.

The HPR architecture allows RTP-capable nodes to define themselves as *stationary* or *mobile* when exchanging connection setup information. If both are stationary or both are mobile, either can initiate the path switch under any circumstances that demand it. If one is stationary and one is mobile, only the mobile partner can initiate the path switch after a connection failure. This is done to ensure that a static node does not waste resources trying to contact a partner that has moved physically to obtain an alternative connection. It is also used for multinode persistent sessions in a sysplex, where one partner can move to a new node after a failure.

1.7 Combined HPR/APPN Network

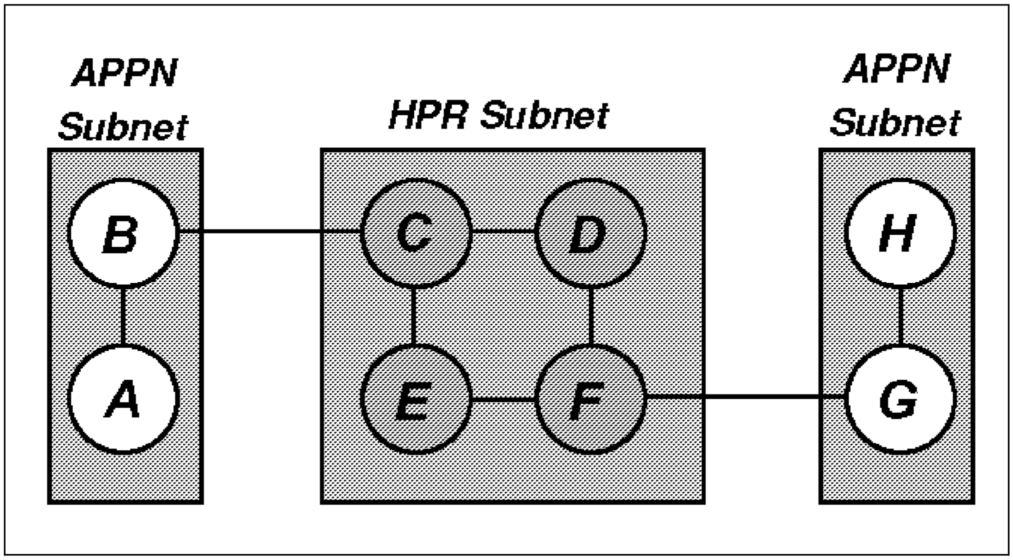

Figure 6. Combined APPN and HPR Network

Figure 6 shows a network where nodes A, B, G and H have no HPR capability, whereas nodes C, D, E and F understand HPR. All nodes are network nodes and may have CP-CP sessions between them. The CP-CP sessions are used for the usual APPN functions such as the exchange of topology information.

When all the nodes are active and all the desired CP-CP sessions have been established, each NN has a representation of the whole network in its topology database. Those nodes that do not understand HPR see no difference in capability between themselves and the other NNs. Those nodes that do have HPR capability are aware of the HPR capability in all other nodes. The APPN protocols relating to topology exchange ensure that the bits in the TDUs indicating HPR support are ignored, yet propagated, by the nodes that do not understand them.

As described above, it is a requirement of HPR that in order to establish an RTP connection, there must be at least two nodes in the session path that are capable of RTP, and that all intermediate nodes be capable of at least ANR if not RTP.

HPR and path switching can lead to some interesting session paths. For example, suppose C and D were the only RTP-capable nodes in Figure 6, and E and F were ANR-capable. A session from A to H might take the route ABCEFGH if A decided that was the best route. In that case there would be no HPR within the session since there are not two RTP nodes on the path. However, if the route chosen was ABCDFGH, then an HPR portion would be established between C and D. Now, if the link CD failed, a path switch would be attempted by C or D. The new route chosen for the RTP pipe would be CEFD, since that is the only valid alternative. Thus the session would take the route ABCEFDFGH. This route traverses the same link twice: once as base APPN and once as part of an RTP pipe. Only a path switch could cause such a route, since no network node would ever calculate it at session setup.

1.8 HPR Control Flows over RTP Option

Nodes that support ANR and RTP must be able to send and receive both FID-2 PIUs and NLPs on the same link, provided of course that the link itself can accommodate both types of packet. With just the ANR and RTP options of APPN implemented, CP-CP sessions and Route Setup messages flow as FID-2 packets and will be mixed with session NLPs.

A node that implements RTP can, if it wishes, also implement the Control Flows over RTP APPN option. This addition to the architecture permits both CP-CP sessions and Route Setup messages to flow over RTP connections.

Nodes supporting the Control Flows option use RTP connections (if both adjacent nodes support it) for CP-CP sessions. These connections (one or two depending on timing) will be dedicated to the CP-CP sessions because the CPSVCMG class of service is used only for these sessions. As with FID-2 CP-CP sessions there is no need for Locate flows, nor is there any need for Route Setup since the nodes have determined all the relevant information at the time the XIDs flow after link establishment.

When a link connecting two nodes that both support this option is needed for an LU-LU session, a long-lived RTP connection is established between the adjacent nodes solely for the purpose of forwarding Route Setup messages. This connection is called *long-lived* because it stays up as long as the link stays up, even though no sessions ever use it. The long-lived RTP connection is never path-switched, even though the CP-CP sessions may be switched. If a link breaks, its long-lived RTP connection will be deactivated, and may be re-established if the link is recovered.

Aside from providing additional resilience, the Control Flows option is required if the connection supports HPR only and does not permit FID-2 traffic.

1.9 Multilink Transmission Groups

A multilink transmission group (MLTG) comprises a number of physical links that appear as a single connection to the network topology. The potential benefits of such an arrangement are improved performance (through load sharing) and greater availability (since a physical failure need not affect TG operation).

Subarea networking has implemented MLTGs between NCPs for many years, but base APPN does not support them. One reason for this is the requirement for resequencing packets, which imposes extra processing overhead on the network. Packets transmitted on an MLTG may not arrive at their destination in the same sequence as they were sent, because the packet sizes, link speeds and link utilizations will vary. Therefore, a node on the path somewhere must buffer them and reorder them into their original sequence.

In subarea networking, the NCPs at either end of the MLTG perform the resequencing function. In a modern high-speed network this is not possible because the overhead on the intermediate nodes would be excessive. Moreover, nodes implementing RTP already have the capability of buffering quite large amounts of data because acknowledgements are end-to-end and can be infrequent. Therefore, an APPN MLTG must support HPR to the exclusion of FID-2 packets. In other words, ISR is not supported. This in turn implies that

CP-CP sessions and Route Setup messages must flow over HPR, and therefore the MLTG partners must support the Control Flows option.

As with subarea networking, the presence of an MLTG can be determined by predefined TG numbers. The partner nodes must declare MLTG support as each link in the group is activated. They must then ensure, by selecting one of the predefined TG numbers (1-20), that the same TG number is negotiated for all the links in the group.

Alternatively, if MLTG support is defined yet no TG numbers are predefined, such links will comprise an MLTG with the default number 240.

1.10 Summary of HPR Implementation Options

The APPN architecture defines the components of HPR in terms of APPN Option Sets, which describe what options a particular node implementation has for HPR support, and the prerequisites for each one. If a product wishes to implement a given feature of HPR, it should implement *all* the features within the same APPN option set as well as all the features in all the prerequisite option sets. Figure 7 shows a summary of the HPR options.

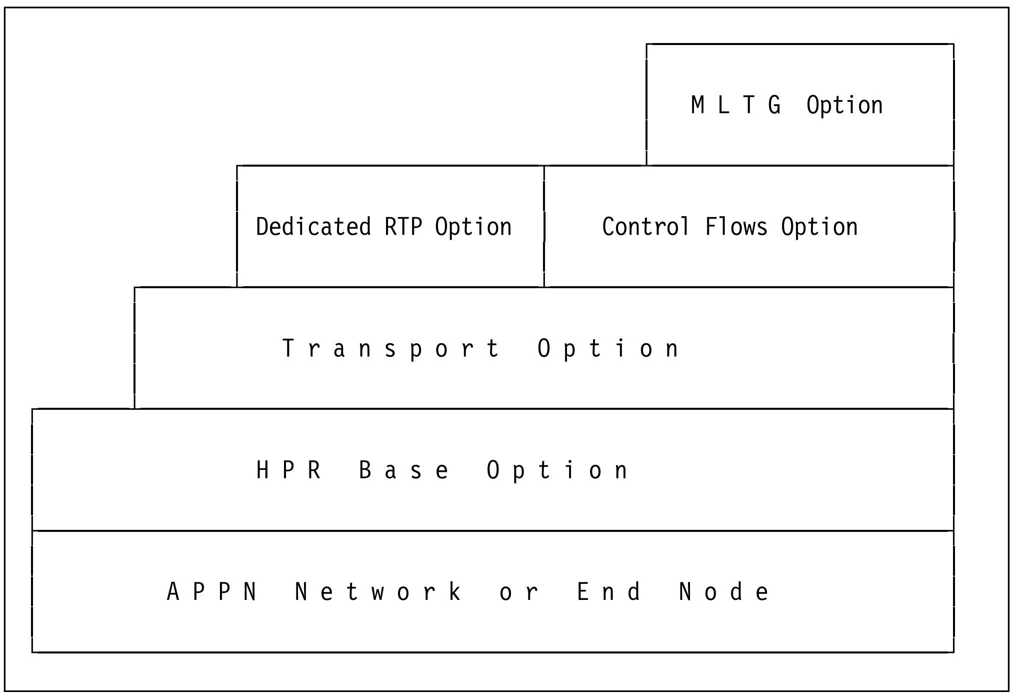

Figure 7. HPR Option Sets

Clearly any node implementing HPR must be an APPN end node or network node.

The HPR Base (Option Set 1400) includes the following functions:

- ANR Routing (NNs only)
- Ability to use FID-2 packets for CP-CP sessions
- Ability to share links between FID-2 packets and NLPs
- HPR support indication in topology updates
- Ability to send and receive Route Setup
- Support for at least 768-byte packets on HPR-capable links

- Ability to exchange HPR capability via XID-3
- Calculation of HPR-only routes (NNs only)
- Ability to use links without link-level error recovery

The HPR Transport option (Option Set 1401) includes the following functions:

- RTP, including end-to-end error recovery and ARB flow/congestion control
- APPN/HPR boundary function
- Nondisruptive path switch
- Ability to return NCE on successful search reply

The HPR Base is a prerequisite for the Transport option.

The Control Flows option (option set 1402) comprises the ability to set up CP-CP sessions over RTP connections, and to send Route Setup messages over RTP connections. RTP (the Transport option) is a prerequisite for Control Flows. Control Flows is required where a connection is unable to support packets other than NLPs; examples include ATM, Enterprise Extender, MLTGs and some implementations of MPC+.

The Dedicated RTP Connections option (Option Set 1403) allows a node to establish an RTP connection that is only used for one session. This is useful for applications that are able to specify their own quality of service. It is intended to be used with APPN Option Set 2005, which permits an ATM SVC to be dedicated to a single RTP pipe. Thus an application can have an ATM connection all to itself. RTP capability is a prerequisite for this option set.

HPR MLTG (Option Set 1404) allows multiple links to be included in a single TG. Because MLTG cannot accommodate base APPN packets, Control Flows is a prerequisite for MLTG.

Chapter 2. HPR Implementations

In this chapter we present an overview of the HPR implementations on current IBM APPN platforms. For those platforms we actually tested, we show detailed descriptions of configuration and operation in the appropriate chapters. It is important to note that the overview given here refers to the *current* releases of these products. Previous releases that implemented HPR may not have implemented as many options or configurations as the current releases.

HPR is supported by the following current products:

- eNetwork Communications Server for OS/390 Release 5
- Communications Server/2 Version 5
- Communications Server/NT Version 6
- Communications Server/AIX Version 5
- OS/400 Version 4 Release 2
- ACF/NCP Version 7 Release 6 (in conjunction with VTAM or CS OS/390)
- Personal Communications Version 4 Release 2 (for OS/2, Windows NT and Windows 95/98)
- 3746-9X0 Network Node Processor, microcode Releases D46130J and F12380; also the 3746 Multiaccess Enclosure
- 2216 Multiprotocol Access Services, Version 3 Release 1
- 2210 Multiprotocol Routing Services, Version 3 Release 1

Some older products, particularly the 3174 with LIC C6 and the 6611, also implement some HPR functions.

Whether HPR is actually used on a connection depends on the initial negotiation that takes place when two adjacent nodes make contact. The HPR capabilities of each node are exchanged during the XID-3 flows, using CV61 control vectors that carry the following information:

- Whether the HPR base is supported (the presence of CV61 indicates this)
- Whether RTP and the HPR Transport option is supported
- Whether the Control Flows over RTP option is supported
- Whether link level error recovery is required

The actual capabilities of the connection become the lowest common denominator of the support levels exchanged. In particular, for any kind of HPR traffic to flow, the link level error recovery capabilities must be compatible. Their possible values are:

- Not allowed
- Not preferred
- Required

If one node says “not allowed” and the other says “required”, then the connection becomes base APPN and HPR is suppressed. This is not normally a problem because most nodes now default to “not preferred” unless the DLC forces a particular value (required for channels or X.25, not allowed for native ATM). Some older versions of HPR products either insisted on link-level error recovery or prohibited it. This was a particular concern with early releases of the 3746-9X0 (required on a token-ring) and Communications Manager/2 (prohibited). Today most nodes permit the default value of “not preferred” to be overridden. If you do override these values, you need to make sure the partners have not specified conflicting requirements.

2.1 HPR Support in CS OS/390 R5 and NCP

HPR support in eNetwork Communications Server for OS/390 Release 5 is quite comprehensive, especially given the wide variety of connection types that a product capable of both APPN and subarea networking can have. The actual level of HPR support in VTAM can be configured as follows:

- By means of the HPR start option (RTP, ANR or NONE).
- By means of the HPR keyword (RTP, ANR or NONE) on individual link station definitions. By default this keyword takes the value of the second parameter of the HPR start option, and cannot be used to promote a TG above the value assigned to the node as a whole.

If VTAM is configured as an APPN end node or network node, then:

- An RTP connection can be established over *any* APPN TG, including VR-TG.
- ANR can be performed (if this VTAM is an NN) between any two APPN TGs, including VR-TG.
- HPR is not supported over LEN or IC-TG (subarea) connections since these are not APPN in the first place.
- Control Flows over RTP is supported over any APPN TG *except*:
  - A LAN connection through an XCA (3172, OSA or 2216 LSA)
  - A VR-TG
  - A connection to an NCP (this VTAM does not care but the NCP′s owner will prevent Control Flows)
- VTAM does not support directly attached APPN MLTGs. Although MPC and MPC+ act rather like MLTGs, the function is done at the DLC layer and the MPC/MPC+ connection appears as a single APPN link.

As described in *Subarea to APPN Migration: VTAM and APPN Implementation*, a composite network node comprising a VTAM and its owned NCPs appears as a single APPN network node. With this in mind, if VTAM is configured as an APPN network node, the the following additional considerations apply:

- An RTP connection can be established over any APPN TG, including VR-TG, through an NCP.
- An RTP connection must be terminated in the VTAM host, never an NCP. Thus a session involving an NCP-attached dependent LU has two choices:
  - Use base APPN on the first portion of a route, until an RTP-capable node is encountered.
  - Use subarea to its owning VTAM, then HPR. The VTAM start option HPRNCPBF controls whether the HPR part of such a session is allowed to pass through the NCP from which it started, thus doubling back on the path.
- Control Flows over RTP is not supported over *any* NCP-attached TG.
- ANR routing can be performed between *any* two APPN TGs, including VR-TGs.
- An NCP can act as an ANR routing node without needing to pass data to its owning VTAM, whether using a FID-2 connection (BF-TG) or a VR-TG.
- An APPN/HPR MLTG cannot be directly connected to an NCP.

To clarify these rules, we illustrate them with an example in Figure 8 on page 17.

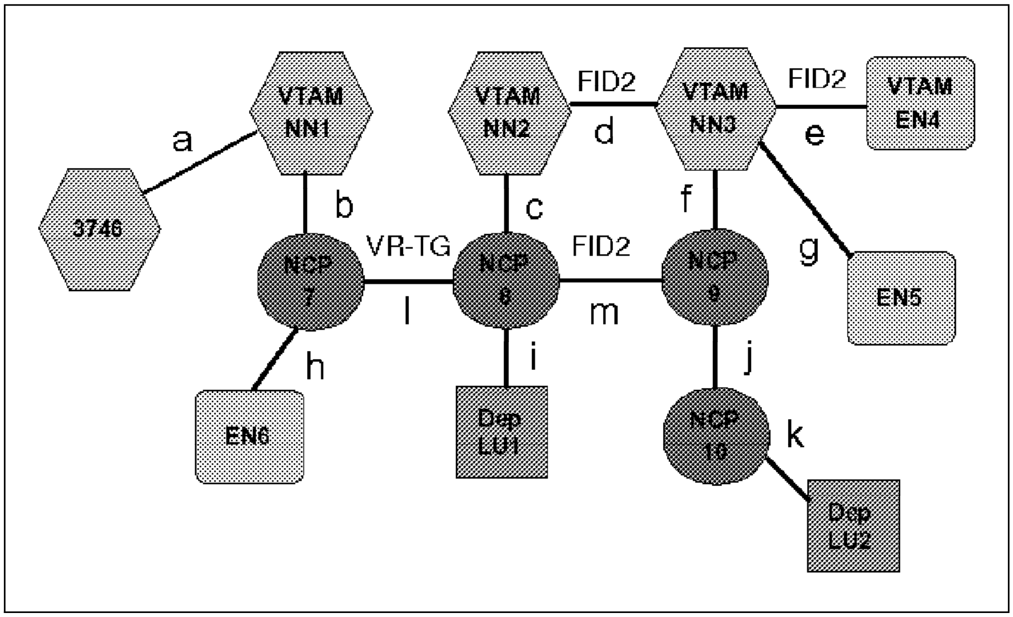

Figure 8. VTAM/NCP HPR Configurations

In the figure:

- b, c, f and j are FID-4 subarea connections that carry the NCP ownership sessions from the three VTAM NNs.
- l is a FID-4 connection on which a VR-TG has been defined between NN1 and NN2.
- a, d, e, g, h and m are FID-2 (APPN BF-TG) connections between VTAM nodes and other APPN nodes.
- i and k are peripheral subarea connections to type 2 nodes with dependent LUs.

Assuming that all nodes and all TGs are configured to support RTP, the following observations can be made:

- Control Flows over RTP is supported over connections a, d, e and g as long as they are not XCA connections. All the other APPN links are either VR-TG or NCP-attached.
- A session between EN6 and EN4 can run over an RTP connection using the route h-l-m-f-e. NCP 7, NCP 8 and NCP 9 all act as ANR routers, as does NN3. NCP 7 and NCP 8 translate HPR protocols between APPN BF-TGs and subarea explicit routes.
- A session between LU2 and NN2 over the connections k-j-m-c will not use HPR. The first RTP-capable node on this path is NN2, although to the APPN network NN3 is RTP-capable. Similarly, a session between LU2 and NN1 using k-j-m-l-b will not use HPR.
- A session between LU2 and NN2 over the connections k-j-f-d *will* use HPR on the d link alone. Both NN3 and NN2 are RTP-capable. A session between LU2 and NN1 using k-j-f-d-c-l-b will also use HPR between NN3 and NN1.
- If the start option HPRNCPBF is set to YES on NN3, the session between LU2 and NN2 can be established using the path k-j-f-f-m-c. The HPR portion of this route is f-m-c, which retraces the subarea (internal to the CNN) TG f. If

HPRNCPBF is NO, such a route will not be tolerated by NN3 except as a result of a path switch from the k-j-f-d path.

As stated above, VTAM and NCP support HPR over VR-TG connections. The Route Setup flows are the same as for BF-TG connections, and the VTAMs or NCPs at the APPN/subarea boundary assign ANR labels in the usual way. However, the NLPs do not flow on a virtual route (despite the name virtual route-based transmission group). A VR is used for the Route Setup flow itself, but if it is not required, it is deactivated after the Route Setup reply has been sent. Subsequently, NLPs flow on an ER between the VR-TG endpoints. Once again, this is done to minimize the amount of processing an ANR router has to do. There is no VR flow control over the VR-TG portion of an RTP connection.

A VTAM interchange node is designed to translate session protocols between subarea and APPN networks. Moreover, the APPN architecture requires that an RTP node is able to translate protocols between base APPN and HPR. However, an ICN cannot always do both together for the same session. Because of limitations in the subarea and APPN architectures, an ICN cannot translate directly between subarea (IC-TG) protocols and HPR protocols *unless* the HPR part of the session route can be guaranteed to be one hop only. The only way to guarantee this is if the adjacent HPR-capable node is an end node. VTAM checks the adjacent node′s role, and uses ISR (base APPN) protocols for the session concerned if the adjacent node is a network node or a border node managing a peripheral border (which pretends to be an end node). This works well unless:

- The adjacent node does not support base APPN flows over this type of TG. Today the 2216 and the 3746 MAE do not support base APPN over an MPC+ connection.
- The TG itself does not support base APPN flows. Native ATM connections using OSA and enterprise extender connections require HPR for all flows.

It is important to realize that this restriction applies only to subarea IC-TGs adjacent to an HPR-capable connection. It does *not* apply if the subarea part of the route is internal to a composite network node, or if it is represented by a VR-TG. Thus, in Figure 8 on page 17, the session between LU2 and NN2 using k-j-f-d does not fall foul of the restriction. If, however, LU2 was on another VTAM node and the connection k was a FID-4 link without VR-TG, then the session could not use HPR on link d.

2.1.1 High-Performance Data Transfer

High-performance data transfer was introduced in VTAM V4R4 for two purposes:

1. To improve the efficiency of high-speed links, especially for bulk data transfer
2. To provide a common DLC that both the SNA and the TCP/IP components of CS OS/390 could share.

HPDT achieves its objectives by two means:

- A function called Communication Storage Manager, running in its own address space, manages data transfer for HPDT connections. Certain applications (including VTAM and TCP/IP) make use of CSM to reduce the number of data moves required to send information across an HPDT link.

- Improved "seldom-ending" channel programs, which run continually while there is data to send, and allow multiple subchannels to be used concurrently at optimum utilization.

Multipath channel protocols utilizing HPDT are known as MPC+. CS OS/390 presently supports HPDT over three types of connection: multipath channel, XCF and ATM (which uses MPC+ protocols to communicate with the OSA). HPDT *requires* HPR. Both partners on an HPDT connection must support RTP and the Control Flows over RTP option, so that CP-CP sessions and Route Setup messages flow as NLPs, as well as LU-LU session traffic.

Because HPDT requires HPR, there is an issue relating to the translation between IC-TG and RTP mentioned above. If the HPDT connection is to another VTAM over an MPC or XCF connection, VTAM will use ISR routing for sessions which need to traverse the IC-TG/APPN boundary, even though all the other traffic uses HPR. However, this is not possible in two cases:

- An ATM connection cannot accept ISR traffic, because it does not allow link-level error recovery which ISR requires. The same consideration will apply to an enterprise extender connection if and when CS OS/390 supports such a connection.
- A 2216 configured for MPC+ (or a 3746 MAE) cannot send or receive ISR traffic over such a connection.

There is a second issue relating to the fact that HPDT is independent of both VTAM and TCP/IP. If two VTAM nodes at the requisite level (V4R4 or above) are connected over an HPDT-capable link, this link will be established as HPDT by default even before the VTAMs have exchanged capabilities. Thus, an HPDT connection between VTAMs where one VTAM has been configured *not* to support RTP will be usable only by TCP/IP. To run SNA over an HPDT connection, both VTAMs must be capable of RTP (which also makes them capable of Control Flows).

HPDT can be turned off on an MPC link by coding the MPCLEVEL=NOHPDT keywork on the TRLE definition. It cannot be turned off for an XCF or an ATM connection.

2.1.2 How VTAM Defines HPR Connections

As described in *Subarea to APPN Migration: VTAM and APPN Implementation*, an APPN-capable VTAM always has two "faces": a subarea side and an APPN side. Each side must be able to see the other in its own terms. Also, an RTP connection behaves like a DLC, and the (logically adjacent) RTP partner node has the appearance of a link station.

Therefore, VTAM treats an RTP connection as a dynamically defined PU, since PUs represent link stations. This RTP PU, by default, receives a name of the form CNR followed by a unique VTAM-chosen number. Sessions that traverse a TG on an RTP connection are regarded by VTAM as being associated with the RTP PU, not the PU (link station) associated with the physical TG. Thus, a display of the physical PU may well show no LUs using it, whereas the display of the RTP PU may show many.

The RTP PUs are stored by VTAM in the major node ISTRTPMN. Three distinct types of RTP connections are identified:

- LULU, for LU-LU sessions

- CPCP, for CP-CP sessions on TGs that support Control Flows
- RSTP, for long-lived RTP pipes on TGs that support Control Flows

Examples of displays of HPR resources are shown throughout this book.

By default, an APPN-capable VTAM supports HPR on all eligible connections unless the installation defines otherwise. The definitions available for the user to modify and tune HPR operation are summarized below, and are described more fully in Chapter 6, "High-Performance Routing on VTAM" on page 49 and Chapter 7, "HPR between CNN Nodes" on page 73. Although the default values are well-chosen, you should review them all because the defaults cannot cover all eventualities.

2.1.2.1 Start Options

Node-wide HPR capability can be defined using four VTAM start options.

HPR This specifies the level of HPR support provided by VTAM. The default is HPR=(RTP,RTP). The first operand defines the node capability and the second defines the default capability for link stations.

HPRPST This is the maximum time that VTAM waits for a path switch to be completed before terminating an RTP connection. The default is HPRPST=(8M,4M,2M,1M) for the four transmission priorities (the highest priority having the shortest timer).

PSRETRY This specifies the frequency at which VTAM attempts a path switch automatically, in order to find a better route for an RTP connection. The default is PSRETRY=(0,0,0,0) meaning no automatic path switch attempts.

HPRNCPBF This parameter, whose default is NO, determines whether VTAM will allow sessions to traverse the same NCP twice. This can only occur if one pass through the NCP is on an RTP connection and the other is not. This option allows the installation to maximize the HPR portion of a session path at the cost of extending the total session path.

2.1.2.2 Link Station Parameters

Note that the term *link station* may apply to the CDRM definitions as well as the PU definitions if VR-TG is used. These parameters apply to NCP-attached links as well as VTAM-attached links.

HPR This overrides the HPR start option (second operand) for individual links. A link cannot be upgraded to a level of HPR support above that of the VTAM node as a whole.

LLERP This specifies the level of link error recovery to be used on the connection. The default depends on the connection type. HPR does not require link-level error recovery, but some types of link (channel and X.25) and some partner nodes do.

2.1.2.3 NCP Parameters

These parameters are coded on the BUILD statements. Most of them are necessary because of the presence of internal connections in a composite network node that are on an HPR path but not in the APPN topology database.

HPR
This determines whether this NCP is capable of HPR; in other words whether it can send and receive NLPs on a BF-TG. An NCP in the middle of a VR-TG connection is not aware of HPR and need not be at a software level that supports HPR. The default is HPR=YES.

HPRATT
This specifies the average transmission time for 150 bytes across the CNN, and is used to initialize the ARB flow control algorithm. The default is 12 milliseconds. Coding this parameter is recommended only if the transit time is less than 200 milliseconds.

HPRMLC
This defines the capacity of the slowest internal link in the CNN, and is also used for ARB flow control. The default is 9 kbps, which may well result in RTP connections taking a very long time to reach optimum throughput.

HPRMPS
This specifies the maximum packet size that can be sent across the CNN. It is only required if there are pre-HPR NCPs (V7R2 or before) in your CNN, since HPR-capable NCPs can determine the value for themselves. The default is zero, meaning that NCP is to work out the value for itself.

If your subarea network has been designed to transport PIUs of a certain size (typically 4125 bytes for an RU size of 4096), then this is the appropriate value to set for HPRMPS if you have pre-HPR NCPs.

HPRQLIM
This defines the maximum amount of HPR data that can be queued to a BF-TG at any one time. The default is zero, meaning no limit.

Please refer to *NCP, EP and SSP Resource Definition Reference* for a complete description of these parameters. Note that that manual refers to a "composite ANR node", but the parameters apply equally to a composite RTP node where VTAM is the endpoint of the RTP connections passing through the NCP.

2.2 HPR Support in 3746-9X0

The 3746 network node processor is (as its name implies) configurable as an APPN NN only. Any level of HPR support can be configured at the node level: none, ANR, RTP or Control Flows. Individual links can be defined as HPR-capable or not. If the 3746 is configured to support Control Flows, then APPN MLTGs can be defined on the connections.

HPR is available over the following connection types:

- Token-ring.
- SDLC.
- Frame relay.
- X.25.

- ESCON channel. The ESCON channel processor uses the CDLC protocol and not MPC or MPC+. Thus, the VTAM definition required is a local SNA major node.

The 3746-9X0 can be upgraded by the attachment of the Multiaccess Enclosure (MAE). The MAE has exactly the same functions as the 2216, being based on 2216 technology. A 3746 with both NNP and MAE installed, therefore, looks the same to the APPN network as an NNP plus a 2216: two separate but connected network nodes. Note that this is not true for IP; the integration between the MAE and the 3746 IP functions is rather different from the APPN integration.

The MAE can usually be expected to support the same software levels as the 2216 within a month or two. At the time of writing a new release (V3R1) of the Multiprotocol Access Services software has just become available on the 2216, and will shortly be available on the MAE. This release adds support for session services extensions and the extended border node function.

A 3746 can be connected to a VTAM host, via the MAE, using MPC+. This requires RTP support in VTAM, and will prevent SNA sessions from crossing between the MAE and a subarea network unless VR-TG is available within the subarea network. 2.1.1, “High-Performance Data Transfer” on page 18 describes the reasons for this.

2.3 HPR Support in 2216 and 2210 Routers

The 2216 running Nways Multiprotocol Access Services cannot be configured as an APPN end node, but apart from that its APPN/HPR capabilities are among the most extensive of any IBM product. It can be a network node, a network node with branch extender support, or even (from summer 1998) an extended border node. If you enable APPN on a 2216 you automatically have support for ANR, RTP and Control Flows. HPR support is configured individually for each port and each predefined connection.

The 2216 supports HPR over a wide variety of connections, since it must be able to communicate with other multiprotocol routers as well as SNA workstations and servers. The connection types over which HPR is available are:

- Token-ring, Ethernet (DIX V2 or IEEE 802.3) and FDDI
- 100 Mbps token-ring and Ethernet
- Frame relay BNN and BAN, over serial or ISDN link
- PPP, over serial or ISDN link
- ATM, which is available in two flavors:
  - Forum-compliant LAN emulation.
  - Native. Native ATM support is an HPR-only connection and requires Control Flows over RTP in both partner nodes. Currently VTAM, the 2216, 2210 and MAE are the only products that can communicate using native SNA over ATM.

    Because the 2216 is always a network node, if a session traverses the native ATM link between the 2216 and VTAM, it cannot then continue into the subarea network (unless VR-TG is used in the subarea network). This is because of the architectural restriction described in 2.1.1, “High-Performance Data Transfer” on page 18.

- Parallel or ESCON channels, which can be configured in two ways:
  - LSA mode, which looks like a 3172 (XCA) connection to VTAM.
  - MPC+ mode. This corresponds to VTAM′s HPDT (see 2.1.1, "High-Performance Data Transfer" on page 18). As with ATM, this is an HPR-only connection and cannot carry ISR packets. Once again, a session arriving in VTAM across such a link cannot continue into the subarea network. VTAM itself can use ISR over an MPC+ connection but the 2216 cannot, thus a session request that needs to use such a route will fail.
- Enterprise extender. Please see Chapter 5, "Enterprise Extender" on page 45 for a description of this function.

SDLC, X.25 and DLSw connections do not support HPR.

The 2210 software, Multiprotocol Routing Services, is essentially the same code as 2216 MAS. However, the 2210 is a smaller and less powerful machine, and cannot concurrently support as many networking protocols and options as the 2216. Presently the only major functional difference between the 2210 and 2216 is that the 2210 has no channel attachments (ESCON or parallel).

2.4 HPR Support in Communications Server/2 and Communications Server/NT

CS/2 can be configured as an APPN end node, or a network node with or without branch extender support. In each case HPR capability is included, and need not be configured specifically. CS/2 supports ANR, RTP, Control Flows over RTP and APPN MLTGs over any type of connection. On each individual connection you can define whether HPR is to be available. If HPR has been defined on a connection, then CS/2 will act as an ANR or an RTP node depending on the node role and the negotiation with adjacent nodes. In addition, each connection can be defined as part of an MLTG. MLTGs can be assigned specific TG numbers (in the 1-20 range) or can be allowed to default to TG 240.

CS/2 supports HPR on the following types of connection:

- SDLC
- Frame relay
- Token-ring
- Ethernet
- ISDN

CS/2 can also use HPR over generalized DLC (GDLC) connections. GDLC is a developer API for support of intelligent adapters that provide the link protocol within the adapter.

CS/NT support for HPR is almost identical to that in CS/2, but there are four main differences:

- CS/NT supports Enterprise Extender.
- CS/NT supports HPR over X.25 connections.
- CS/NT does not support APPN MLTGs.
- CS/NT does not support HPR over frame relay through the Wide Area Connector (WAC card), although it does so using non-IBM cards.

2.5 HPR Support in Communications Server/AIX

CS/AIX Version 5 can be configured as an APPN network node, end node or LEN node. HPR is not available with LEN, but a CS/AIX machine configured for APPN automatically has support for both RTP and ANR. There is no distinction between ANR-level and RTP-level support on individual connections as there is with VTAM. When defining the connections, you simply specify whether HPR is, or is not, to be available on that link. If HPR support has been permitted on a link, then CS/AIX will use ANR or RTP according to the role it is performing on the session path.

Communications Server/AIX does not support Control Flows over RTP, therefore it cannot establish APPN MLTGs.

The connection types supported by CS/AIX are as follows:

- Parallel channel. This uses the CDLC protocol which requires a local SNA definition on the host (VTAM) side, similar to a FID-2 NCP attachment.
- ESCON channel. This can be configured in two ways:
  - CDLC protocol, as for the parallel channel support.
  - MPC protocol. CS/AIX communicates with VTAM using a multipath channel connection, thus the VTAM definitions required are a TRLE and a local SNA PU. CS/AIX does *not* support HPDT, but VTAM will ascertain this at connection time and does not require HPDT to be turned off in the definitions.
- Token-ring, Ethernet and FDDI connections. ATM is also supported but only as LAN emulation which appears as the relevant type of LAN to VTAM.
- SDLC and X.25 connections.

2.6 HPR Support in Personal Communications/3270

Personal Communications/3270 has had a varied history in terms of the communications protocols it supports. At times it has contained its own SNA stack; at others it has relied on companion products. The same PComm version has not always been consistent across the platforms on which it ran. However, if you need HPR, then the story is quite consistent:

1. HPR requires APPN.
2. The APPN protocol is provided for PComm either by the Access Feature or the full Communications Server product. These are available on OS/2, Windows 95 (Access Feature only), and Windows NT.
3. If you need HPR all the way from the PComm workstation, you need DLUR support (as described in Chapter 3, "Dependent LU Requester/Server" on page 27). A workstation APPN node whose dependent LUs do not use DLUR must use peripheral subarea attachment for those LUs.

The Access Feature is a subset of the Communications Server product on the appropriate platform, and is integrated with PComm in the most recent releases. The major difference (in APPN terms) between the Access Feature and the full CS product is that the Access Feature provides end node support only; the HPR capabilities are similar. Thus:

- If you are using DLUR, the HPR capability of PComm is that of the appropriate Communications Server product, in other words RTP all the way from the desktop to the RTP partner.
- If you are not using DLUR, HPR is not available directly to the PComm node, but may of course be present elsewhere on the session path.

2.7 HPR Support in OS/400

The AS/400 can be configured as an APPN end node or network node. OS/400 is always capable of ANR support, but RTP can be disabled or enabled on a node-wide basis. The Control Flows option, and therefore MLTGs, are not supported.

HPR support (and therefore ANR or RTP depending on the overall node capabilities) can be individually specified on each APPN connection. As with CS/AIX, it is not possible to configure some links as supporting ANR only and others supporting RTP. The connections capable of HPR are:

- SDLC
- X.25
- ISDN
- Token-ring
- Ethernet
- ATM (LAN Emulation)
- Frame relay
- FDDI and SDDI
- Wireless

Chapter 3. Dependent LU Requester/Server

It is important to remember that dependent LU sessions are supported across an APPN network. Provided that the session boundary function nodes (and their network node servers, if the boundary functions are on end nodes) implement session services extensions, then dependent LU sessions can traverse APPN FID-2 connections, and do not require subarea or VR-TG paths. What the dependent LU requester/server (DLUR/S) function does is to *enhance* the way that dependent LU sessions operate in an APPN network.

3.1 How DLUR/S Works

All dependent LUs, and the PUs that support them, require sessions to their owning SSCP. These sessions carry various control and management requests such as INIT-SELF, NOTIFY, NMVT, and USS messages. They always take the form of SSCP-PU and SSCP-LU sessions which, prior to DLUR/S, flow entirely within a single VTAM′s subarea domain. This means that a PU serving dependent LUs must *always* be directly connected either to its owning VTAM, or to an NCP owned by that VTAM. Cross-domain or cross-border ownership of dependent LUs is out of the question.

The other restriction affecting dependent LUs is that routing in a subarea network is always done at the subarea level. In other words, any session involving a dependent LU *must* pass through the same adjacent subarea node as the SSCP-LU session, even if the dependent LU happens to reside in an APPN node.

DLUR/S removes both these restrictions, providing the following functions:

- The session between each dependent LU (or PU) and its SSCP is now encapsulated within an LU 6.2 *pipe*. This pipe consists of a pair of sessions between the CPs in the DLUR and DLUS nodes, using the mode name CPSVRMGR and the APPN class of service SNASVCMG. The DLUR/S pipe can carry a number of SSCP-PU and SSCP-LU sessions, and need not be between adjacent CPs. The pipe can be carried on an HPR connection, and can cross APPN network boundaries.
- LU-LU session routing is now performed wholly by the APPN function, and does not require the subarea boundary function at an adjacent subarea node. In fact, the DLUR node itself provides the boundary function.

  When a PLU requests a search for a dependent LU, it will normally receive a positive response from the DLUS, not the DLUR. The response will indicate the DLUS as being the network node server for the dependent LU, even though it may not be the network node server for the DLUR node. The owning CP name, however, will be correct (the DLUR). The route will then be calculated directly to the DLUR, normally by the network node server of the primary LU (which is never the dependent LU supported by the DLUR). In some cases (where the DLUR supports cross-network DLUR/S control sessions) the DLUR itself may respond to a search, in which case the CP name and NNS name given are the correct ones.

Because the DLUS presents itself as the network node server for the dependent LUs, it must always be a network node.

DLUR/S support requires *no* changes to existing applications or dependent terminals. Figure 9 on page 28 and Figure 10 on page 29 show the differences between the traditional SSCP operation and the way DLUR/S operates.

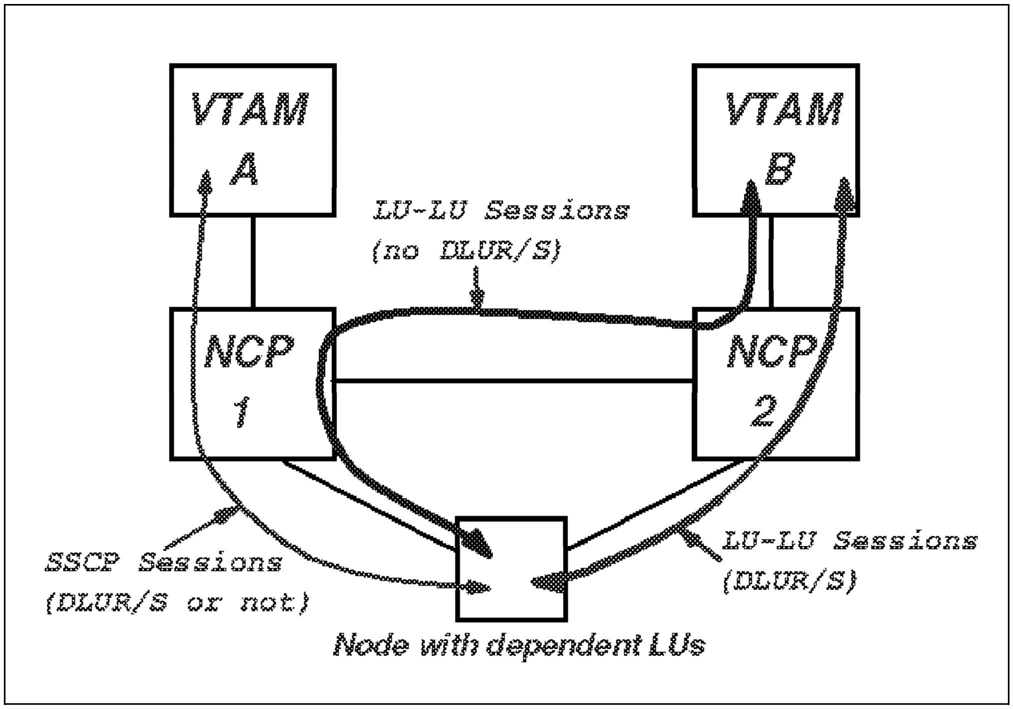

Figure 9. SSCP and DLUR Operation - Routing

Figure 9 shows a node with dependent LUs, connected to two VTAM APPN nodes via their NCPs. Without DLUR/S, the SSCP-PU and SSCP-LU sessions go directly to the VTAM (VTAMA) that owns the PU, while sessions between dependent LUs and applications on VTAMB must flow via NCP1 and NCP2 since NCP1 provides the boundary function.

With DLUR/S, the SSCP sessions still flow the same way, except they are now encapsulated in the DLUR/S pipe. However, the PU is now an APPN node and *itself* provides the boundary function, so that LU-LU sessions to VTAMB flow directly via NCP2.

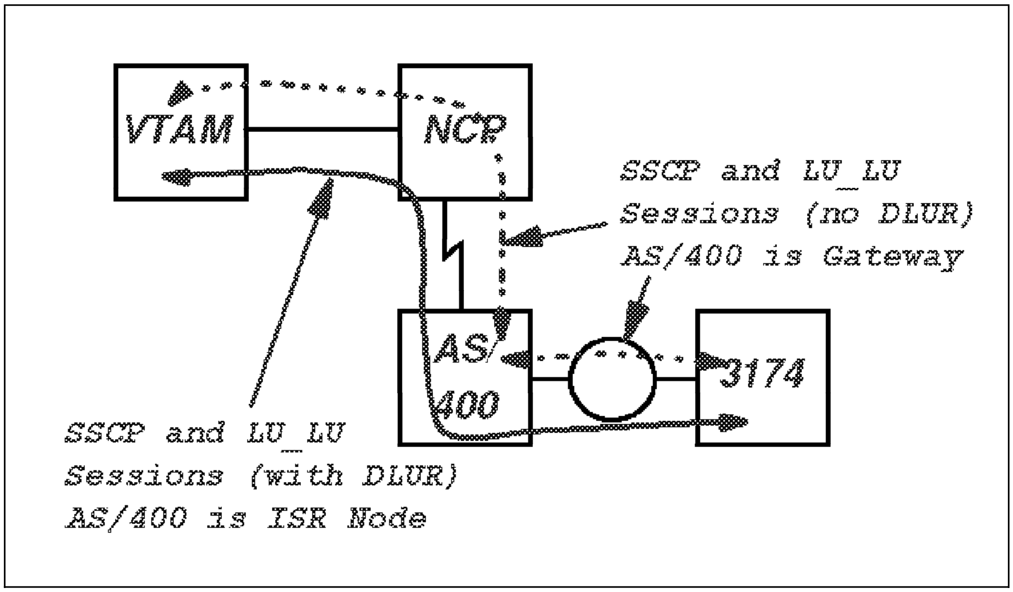

Figure 10. SSCP and DLUR Operation - Resource Utilization

Figure 10 shows the case where a 3174 is connected to an AS/400 in a location remote from the VTAM central site. Without DLUR/S, the SSCP function is performed by the AS/400, which acts as a gateway and itself has SSCP-PU and SSCP-LU sessions with the owning VTAM. Passthrough provides a mapping between the two sets of sessions. The LU-LU sessions are handled in the same way: one session between the dependent LU and the AS/400 passthrough application, the other session between the AS/400 and the VTAM application.

With DLUR/S, all these sessions flow directly to VTAM from the 3174. The AS/400 is now just an APPN network node on the session path; even the LU 2 sessions are routed by the AS/400 since they are carried in APPN FID-2 packets. If the DLUR node was RTP-capable (the 3174 is not), then the AS/400 could perform ANR instead of ISR.

The downstream PU gateway function in CS/2, CS/AIX or CS/NT operates in a manner similar to the AS/400 passthrough function shown here, and the use of DLUR would transform the role of such a node to that of an NN in the same way.

Even if your dependent LUs cannot be upgraded to a software level that supports DLUR, the DLUR/S architecture will allow you to gain some of the benefits. This involves connecting the dependent LU nodes to a node that provides the DLUR function on behalf of *external* type 2 nodes, instead of internal LUs. Such a DLUR node can be located anywhere within the APPN network, thus the dependent LU sessions can be APPN/HPR *almost* all the way to the actual nodes containing the dependent LUs. Figure 11 on page 30 illustrates this DLUR Passthrough function.

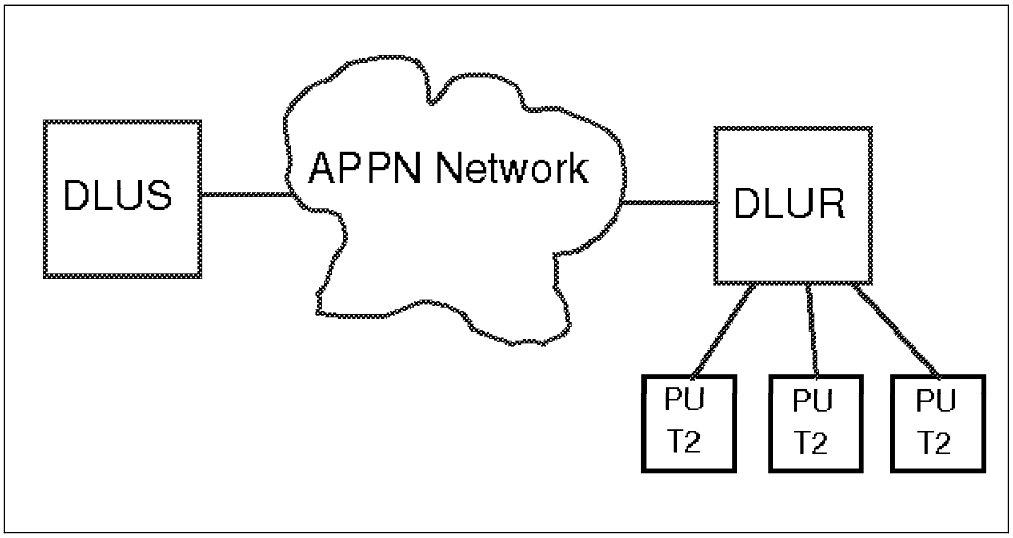

Figure 11. DLUR with External Type 2.0 Nodes

There is a major difference between DLUR passthrough and the SNA gateway configuration, in which the gateway node happens to use DLUR for its dependent LUs. In the gateway configuration, the dependent PUs and LUs are defined internally to the DLUR node, and therefore the SNA sessions (SSCP-PU, SSCP-LU and LU-LU) are split within the DLUR. For example, a session between TSO and a DLUR LU is terminated in the DLUR node and a second session goes from the DLUR node to the external dependent LU.

In the DLUR passthrough configuration all the sessions are preserved intact through the DLUR node. The SSCP-PU and SSCP-LU sessions are passed in and out of the DLUR/S pipe, so that the SNA identity of the dependent resources is preserved for the DLUS and any management product that is running with it. The LU-LU sessions, while still being required to traverse the DLUR node, are simply converted from type 2.0 FID-2 packets to APPN FID-2 (or NLPs, if HPR is used). This is similar to what an NCP would do under the same circumstances. The NCP would forward the packets as either FID-4 or APPN FID-2 packets depending on the session path; it could not convert them directly to NLPs because it cannot act as an RTP endpoint.

The examples shown later in this book include both internal LU and external LU configurations of DLUR nodes.

3.2 DLUR/S Sessions and Connections

With DLUR/S support, all the information that was carried on the SSCP-PU and SSCP-LU sessions now flows on the two CPSVRMGR sessions between the server and requester. Because a DLUR may be responsible for multiple PUs (either internal PUs or downstream external PUs using DLUR passthrough), it may have multiple session pairs with a number of DLUS nodes. Each DLUS may handle many DLUR PUs, complete with their associated LUs, on behalf of many DLUR nodes.

The DLUR/S sessions are established between the two control points in the server and the requester, but they are distinct from any CP-CP sessions that may be present if the DLUS and DLUR happen to be adjacent. The format of the RUs flowing on the DLUR/S sessions is similar to that on the CP-CP sessions, but they use different transaction programs. The transaction program used on the DLUR/S pipe is X‘22F0F0F6’ meaning “receive encapsulated FID-2 PIU”. The GDS variable is X‘1500’, and contains the complete PIU that would have been sent on a real SSCP session. This includes the transmission header, request header and RU. The encapsulated DLUR/S connection resembles a type 2.0 connection closely, except that XID-3 frames may be exchanged between the DLUS and nodes external to the DLUR. The FID-2 header on a DLUR/S connection is formatted according to type 2.0 rules, carrying the local address that identifies the PU or LU in question.

The DLUR/S function follows *switched* procedures, and thus the dependent resources are defined to VTAM (if at all) in switched major nodes. This is quite independent of whether the DLUR node, or any external downstream PUs, are connected to the network using leased or switched connections. The use of switched procedures in DLUR/S reflects the dynamic nature of APPN in general and DLUR/S in particular.

Connections other than DLUR/S that use switched procedures to contact VTAM (for example, token-ring connections or real switched SDLC lines) rely upon the exchange of XID frames after contact to provide identification and operational parameters. This identification allows VTAM to select suitable definitions to represent the LUs and PUs. The identification information can be either the node ID (IDBLK and IDNUM) or the CP name of the remote resource. When contact is made via an NCP, the XID image is sent to the owning VTAM on a REQCONT request.

In the case of a DLUR/S connection, the XID image is carried on a REQACTPU request encapsulated in the DLUR/S pipe. With DLUR passthrough, the XID image received from the external node is passed unchanged to VTAM. Thus the same choices exist for VTAM to identify the remote node and to create suitable definitions. With PUs internal to the DLUR, or where the external nodes do not send XIDs (because they are type 2.0 nodes connected to the DLUR by a leased link), an XID-0 is created by the DLUR node which contains IDBLK and IDNUM values customized at the DLUR. Such PUs, therefore, can only be identified by their node IDs and not by CP name.

In general, we recommend that the IDBLK/IDNUM method is used for all DLUR/S connections. There is always a possibility that a node contacting VTAM as a DLUR-supported PU is also an APPN or LEN node adjacent to VTAM. In such a case the same CP name might be called upon to resolve two quite different connections.

It is important to remember that the LU-LU sessions are *not* encapsulated; they flow natively on whichever link is chosen for them by APPN route selection.

Figure 12 on page 32 shows the relationship of PUs, LUs and control sessions in a DLUR/S environment.

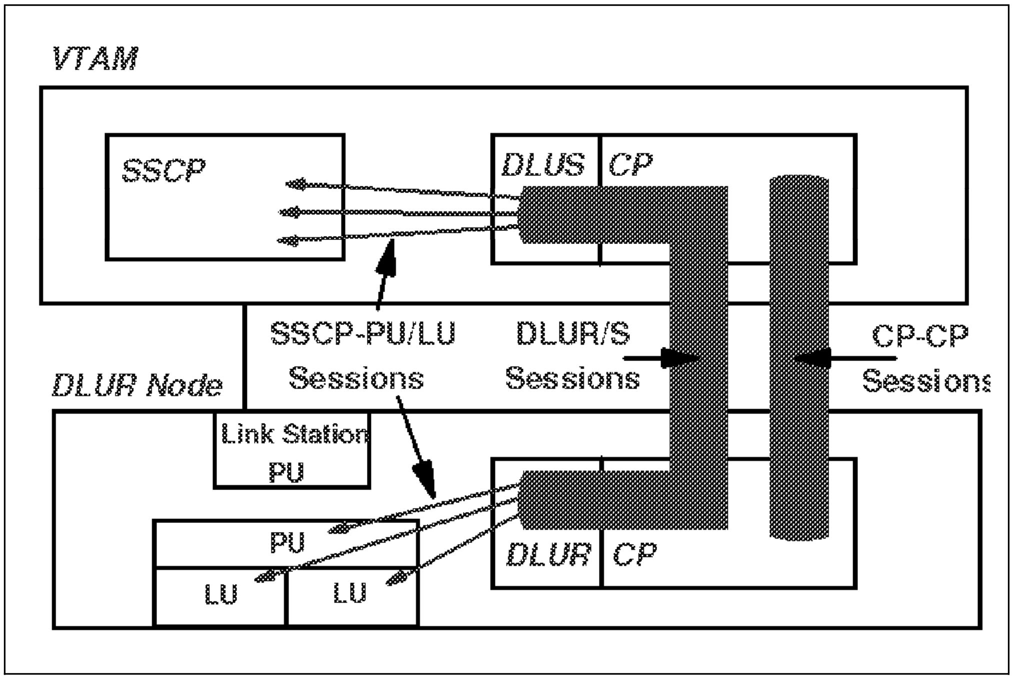

Figure 12. DLUR/S Network Resources and Sessions

The DLUR/S sessions and the SSCP sessions they encapsulate are always present in a DLUR/S environment. The adjacent link station PU is only known to VTAM if the DLUR node is directly connected to VTAM′s domain. The CP-CP sessions are present only if the node is directly connected, and if the two nodes agree to establish them.

3.3 DLUR/S Design Considerations

The DLUR/S function greatly improves the flexibility available to the network designer, by offering new options for both routing and connectivity. However, there are a number of points that must be considered:

- The DLUR/S pipe *must* be established over an APPN network with no subarea hops. Therefore, if your DLUS is in the subarea network you must define a VR-TG to ensure an APPN search path between DLUR and DLUS.

  LEN connections are not permitted for a DLUR/S pipe either, with one exception: the connection between partners implementing AnyNet SNA over TCP/IP is a LEN connection, yet supports DLUR/S. No other LEN connection is supported.

- The primary LU′s CP, and its network node server (if applicable), must support session services extensions. Today only VTAM and the 2216/2210 support session services extensions; thus functions such as the AS/400 Primary LU support cannot be used with DLUR LUs unless the AS/400 is in the subarea network and therefore uses a VTAM ICN as its APPN PLU node.

- The DLUS, DLUR and the PLU in session with the dependent LU may all reside in different networks. However, the following conditions must be met:

  – The DLUS and DLUR must be connected via an APPN border, not an SNI gateway.

- If the DLUS and the PLU are in different APPN subnetworks, both sides require an extended border node. In other words, the boundary between their networks must be an extended border.
- If the DLUS and DLUR are in different APPN subnetworks, an extended border node is required in the DLUS′s subnetwork. The border may be peripheral but an EBN must manage it from the DLUS side.
- If the search path between the PLU and the DLUR or DLUS goes via an interchange node, an extended border node is required in the same APPN subnetwork as the interchange node.

• If the DLUR is an end node, it must supply its TG vectors whenever a dependent LU session is requested. In base APPN, these TG vectors are usually registered to its network node server so that the end node need not transmit them on every session request. With DLUR/S, the DLUS acts as the NN server of the dependent LUs (even if the DLUR is a network node), although it may not be the NN server of the DLUR. Thus the DLUR registers its TG vectors with the DLUS. If the DLUR and DLUS are in the same subnetwork, the DLUS responds to searches for the dependent LUs without troubling the DLUR, and can thus provide the needed information to the node calculating the route at minimal overhead.

  If, however, the DLUR EN and the DLUS are in different subnetworks, the TGs from the DLUR are required in the DLUR′s subnetwork and cannot be supplied by the DLUS. In this case, the DLUR node has two choices:

  - It allows itself to be searched for dependent LUs on APPN broadcasts, so that its real network node server (in its own subnetwork) can provide the TGs. This can place significant searching overhead on the network.

  - It registers its TGs with its *network node server* rather than its DLUS. This requires additional intelligence on the part of the network node server, and is defined in APPN option 1116. Not all NNs support this, but VTAM introduced it in APAR OW26193.

• If a DLUR node is adjacent to its VTAM DLUS node, VTAM′s PU representations may cause confusion. In this case there are two distinct kinds of PU:

  1. The link station connection to the DLUR APPN node is normally identified (unless dynamically defined) by the CPNAME keyword on a PU statement.

  2. The PU supported by DLUR is normally identified (unless dynamically defined using a VTAM exit) by the IDBLK and IDNUM keywords on a separate PU statement.

  If the physical connection between VTAM and the DLUR node uses switched protocols, both these PUs will need to be defined in switched major nodes. The same PU statement will not cater for both.

• A DLUR may contain multiple type 2 PUs, and each may be served by a different DLUS. Therefore, several DLUR/S pipes may be set up from a single DLUR. In addition, a DLUR may allow the definition of an alternate DLUS in case the primary DLUS connection fails. The switch to the backup DLUS is accomplished simply by establishing a new pair of DLUR/S LU 6.2 sessions, and should not disrupt the dependent LU sessions. This works exactly the same way as SSCP takeover in a true subarea environment. The difference is that the boundary function is in the DLUR instead of the NCP, and the DLUR has sufficient intelligence to find the backup owner for itself; no VARY commands are required from the backup DLUS.

3.4 DLUS Implementation in VTAM

At present VTAM is the only APPN node which supports the dependent LU server function. A DLUS node, by definition, must act as a subarea SSCP and thus the potential candidates for DLUS are few in number:

- Nodes that act as gateways for dependent LU sessions, such as the Communications Server workstation products and the 3174. In these cases the sessions must be passed on to a real SSCP anyway, so it makes more sense to implement DLUR rather than DLUS on these nodes.
- Nodes that really do provide primary LU applications, such as the AS/400 or some non-IBM nodes. To our knowledge none of these presently offer the DLUS function.

All that is necessary for a VTAM to provide DLUS function is that it is configured as an APPN network node. There is no definition anywhere that turns the DLUS function on or off. Definitions may be required, however, to allow the DLUS to contact a DLUR and to provide session services for the dependent LUs on that DLUR. Either the DLUR or the DLUS may initiate the DLUR/S connection, and once that connection is established then either the DLUR or the DLUS may initiate the dependent LU/PU activation process.

If the DLUR initiates the DLUR/S connection, it sets up its CONWINNER LU 6.2 session in the normal APPN fashion, with the aid of its network node server if appropriate. VTAM responds by setting up its own CONWINNER session and the DLUR/S pipe is in place.

If VTAM is to initiate the DLUR/S connection, the process is similar, but this time VTAM needs a definition to enable it to locate the DLUR. In fact, VTAM treats the operation like a switched dial-out operation driven by PU activation. Thus the definition required is that of a PU in a switched major node. The PU defined thus is one of the PUs served by the DLUR, but the definition differs from normal dial-out practice in two ways:

- The DLURNAME on the PATH statement identifies the CP name of the DLUR node. This is sufficient information for VTAM to identify the DLUR node, since APPN searching will be used to locate it.
- The actual PU is identified using the DLCADDR keywords on the PATH statement. The node ID or the CP name are coded here. This information is transmitted to the DLUR node, on the DLUR/S pipe, when the PU is to be contacted.

Once the DLUR/S pipe is active, the DLUR may send a REQACTPU on behalf of any of its served PUs. This is treated by VTAM exactly like a switched type 2 connection and the same definition considerations apply:

- A switched major node may be defined in the normal way, specifying the type 2 DLUR PU and its dependent LUs.
- The configuration services exit ISTEXCCS can be used to define the resources dynamically, upon being presented with the IDBLK/IDNUM of the type 2 PU.
- If the DLUR node supports it, the dependent LUs can be defined using the dynamic dependent LU definition exit ISTEXCSD.

If VTAM is to activate a DLUR PU without a REQACTPU from the DLUR, it needs a switched major node definition as described above, namely one with DLURNAME to identify the DLUR and DLCADDR to identify the PU. Whether or not the DLUR/S sessions already exist, the activation of a switched PU with DLURNAME coded will cause VTAM to send ACTPU to the appropriate DLUR for the PU in question.

DLUR resources are treated by VTAM just like any other type 2.0 resources, and their displays have almost no indication that they are on a DLUR node. There is just one message in the PU type 2 display that tells you what the DLUR CP name is.

3.4.1 DLUR Takeover and Giveback

The ownership of DLUR resources can be transferred between VTAMs just as the ownership of NCP-attached resources can. Ownership of a DLUR PU can be relinquished by means of a VARY INACT,TYPE=GIVEBACK command, and LU-LU sessions will continue as long as ANS=CONT is in effect for the PU. To release all the PUs served by a DLUR, the operator can issue VARY INACT,TYPE=GIVEBACK against the DLUR name (the CDRSC representing the DLUR). This will terminate all the SSCP-LU and SSCP-PU sessions, as well as the DLUR/S pipe. Once a PU has been released by one VTAM, another can activate it as described above.

Many DLUR implementations will attempt to re-establish their DLUR/S pipe automatically when it is broken. To prevent this from happening when you want another DLUS to take over the DLUR PUs, the command format VARY INACT,TYPE=GIVEBACK,FINAL=YES can be issued against the DLUR. If this is done, VTAM will mark the DLUR as not available and will reject session requests for the DLUR/S pipe from that DLUR. A subsequent VARY ACT command against the DLUR will clear this condition.

The ability to release all the PUs on a DLUR with one command, together with the FINAL enhancement, was introduced by APAR OW25386.

3.5 DLUR Implementations

In this section we present an overview of the DLUR implementations on current IBM APPN platforms. For those platforms we actually tested, we show detailed descriptions of configuration and operation in the appropriate chapters. It is important to note that the overview given here refers to the *current* releases of these products. Previous releases that implemented DLUR may not have implemented as many options or configurations as the current releases. Only VTAM supports the DLUS function; an overview of that support is given in 3.4, “DLUS Implementation in VTAM” on page 34.

DLUR is supported by the following current products:

- Communications Server/2 Version 5
- Communications Server/NT Version 6
- Communications Server/AIX Version 5
- OS/400 Version 4 Release 2
- Personal Communications Version 4 Release 2 (for OS/2, Windows NT and Windows 95)
- 3746-9X0 Network Node Processor, microcode levels D46130J and F12380; also the 3746 Multiaccess Enclosure

- 2216 Multiprotocol Access Services Version 3 Release 1
- 2210 Multiprotocol Routing Services Version 3 Release 1

Some older products, particularly the 3174 with LIC C6, the 2217 and the 6611, also implement DLUR.

3.6 DLUR Support in 3746-9X0

The DLUR support in the 3746 NNP is mainly intended for the attachment of external nodes (DLUR passthrough). These external nodes can be connected via any SNA link that the 3746 supports: token-ring, frame relay, SDLC or X.25. In addition, the 3746 NNP uses DLUR for the dependent LU session required by NetView Performance Monitor for transmission of performance statistics to the host.

The attached external nodes can be type 2.0 or 2.1. Any node (such as SDLC-attached type 2.0) that does not support the exchange of XID information can have the relevant information (node ID) predefined for it in the 3746. Different PUs can be served by different DLU servers, and primary and backup DLU servers can be defined for any PU.

The 3746 MAE provides exactly the same DLUR functions as the 2216. These are identical to those of the NNP, with the addition of internal DLUR PUs for the TN3270E server function.

3.7 DLUR Support in 2216 and 2210 Routers

When the 2216 is configured as an APPN network node, DLUR support is an additional option that must also be configured. DLUR is available to downstream nodes via passthrough, thus the SNA identity and sessions of the dependent PUs and LUs are preserved across the network. DLUR is also available to the internal dependent LUs that represent the clients of the TN3270E function.

The 2216 provides DLUR functions for external type 2 devices attached to it by any supported method except PPP. This includes Ethernet, token-ring, frame relay, SDLC, X.25, ATM, and FDDI. The upstream connection to the APPN network can be any 2216-supported link including PPP, enterprise extender and ESCON or parallel channel.

Each dependent PU (internal or external) can be assigned to a different DLU server, and a backup DLU server can also be defined. Internal PUs can have a node ID specified, but external ones will normally use their own. If the external node does not provide a node ID (if it is a an SDLC-attached leased type 2 node), then one can be defined for it.

The 2210 DLUR support is the same as that provided by the 2216. Remember though, that the 2210 has less capacity than the 2216 and may not be able to support the same range of options concurrently.

3.8 DLUR Support in Communications Server/2 and Communications Server/NT

CS/2 comes with DLUR support as standard, but only as internal DLUR (not DLUR passthrough). Externally connected workstations use the SNA gateway function; their resources are mapped to internal CS/2 LUs instead of having their session setup flows passed directly to the DLU server.

DLUR support is available to LUA sessions (as used by PComm), TN3270 clients and SNA gateway clients.

CS/2 treats DLUR as being one more type of logical link. Multiple internal PUs can be defined, each with a different DLU server (or pair of primary/alternate DLU servers). LU definitions are then assigned to the PUs at the user′s discretion. DLUR PUs can be assigned a user-specified node ID to identify them to the DLUS.

CS/2 supports cross-border DLUR flows. It can also act as a DLUR on behalf of downstream type 2.0 nodes when it is acting as a branch extender, but it still uses gateway rather than passthrough functions for this.

The DLUR support in CS/NT is identical to that in CS/2, with the exception that CS/NT provides DLUR passthrough. This preserves the SNA identity and the SSCP control sessions all the way to the downstream node. There are no internal PUs and LUs defined at which the sessions are split. Thus management products such as NetView can see the correct session configuration.

3.9 DLUR Support in Communications Server/AIX

Communications Server/AIX has DLUR function configured as standard, and it cannot be turned off. The DLUR implementation on AIX supports local LUs only (internal DLUR). These internal LUs are used for downstream SNA PUs (when CS/AIX acts as an SNA gateway) and internal SNA PUs (when CS/AIX acts as a 3270 emulator for attached ASCII and Telnet devices). Multiple DLUR PUs can be configured to give flexibility in the assignment of PUs to DLU servers.

CS/AIX supports SSCP (DLUS) takeover, cross-border DLUR/S connections, and the use of a backup DLU server should the primary one fail.

3.10 DLUR Support in Personal Communications/3270

As with HPR, the DLUR functions available to PComm are the same as those available with Communications Server or the Access Feature on the relevant platform. PComm simply connects to the LU services provided by the SNA product, and does not know if the LU(s) in question utilize DLUR or traditional peripheral connection.

3.11 DLUR Support in OS/400

OS/400 comes with DLUR support as standard. It allows internal PUs to use DLUR, but does not provide DLUR passthrough. Thus nodes (such as 3174s) connected remotely to the AS/400 requiring DLUR will use the AS/400 gateway function. This means that DLUR PUs and LUs are defined internally to the AS/400, with user-selected node ID values, instead of having their own identities passed through to the DLUS. The attached remote nodes are then mapped to the internal PU/LU definitions.

OS/400 provides DLUR support for all its internal LUs except for dependent APPC LUs. This support therefore includes 3270 devices, 5250 devices (via 3270 emulation), RJE, and other AS/400-unique applications.

Configuration of DLUR is very flexible. A DLUR connection is treated as just another type of logical link, and a PU (controller) can be assigned to a DLUR link. The PU is then identified by node ID and individual LUs can be assigned to it. Primary and backup DLU servers are permitted, and each PU could be served by a different DLUS.

OS/400 supports cross-network DLUR/S operation. It will modify the session setup flows accordingly if the DLUS or the PLU is found to be in a different subnetwork.

Chapter 4. Branch Extender

The border node functions described in *Subarea to APPN Migration: VTAM and APPN Implementation* are very extensive in the two major areas for which they were designed: joining together two distinct APPN networks, and subnetting a single network to reduce topology and search traffic. However, there is one particular scenario where neither the peripheral border node nor the extended border node function fits the requirements closely. This is the case where an organization has many remote branch locations, each with a number of APPN nodes. In this case:

- The gateway between each branch and the backbone network must be a network node (or multiple network nodes for availability), else the branch workstations will be unable to communicate across the backbone.
- If the network is not subnetted, then each branch NN will have to maintain the topology database, and each will exchange broadcast search traffic and topology updates with the backbone and with each other.
- If extended or peripheral borders are used to isolate the branches, then the branch gateways are no longer in the backbone topology database. There are three possibilities:

  1. Extended borders throughout, in other words an EBN in each branch and one or more EBNs in the backbone
  2. Peripheral borders managed by the branch gateways, which appear as ENs to the backbone
  3. Peripheral borders managed by the backbone EBNs, which appear as ENs to the branches

  All three have drawbacks:

  – In all three cases, resources in the branches cannot be registered to a CDS in the backbone. A search for an unknown resource, whether from a branch or from the backbone, could be sent (at worst) to every branch.

  – In option 3, session routes may not take the best path. In particular, a session between two branches could traverse the backbone more than once, even if there is a direct route between the branches. This is because the branch gateway will search the backbone EBN (which looks like a served EN) before it searches other NNs in its own subnetwork. Similarly, it is possible for a session between two nodes in the *same* branch to pass through the backbone twice.

  – Options 1 and 3 do not permit a connection network to be used across the backbone. If the wide area backbone uses switched protocols, then a connection network can save much definition work. A connection network cannot cross an APPN border.

  – In option 2, using PBNs rather than EBNs for the branch gateways will greatly restrict the function. Sessions can cross only one border managed by a PBN, and such a border cannot support the same network ID on each side.

Therefore, the requirement is to subnet the network in such a way that:

- The branch gateway nodes do *not* appear as NNs to the backbone.
- Sessions between branches can take the direct route if one exists.

- Branch resources can be registered to a CDS in the backbone.
- The relatively complex and costly EBN function need not be installed in each branch.
- A connection network can be used across the backbone if required.

4.1 Branch Extender Operation

The new branch extender (or branch network node) function enhances the functionality of a network node. The branch extender is based on a concept close to that of the peripheral border node. Whereas the PBN is designed to interconnect two different networks (with different network names), the branch extender is designed to connect branch networks to the APPN backbone.

This allows the network to be organized in a departmental fashion to match the shape of the owning business.

Figure 13 illustrates the configuration of a typical branch extender.

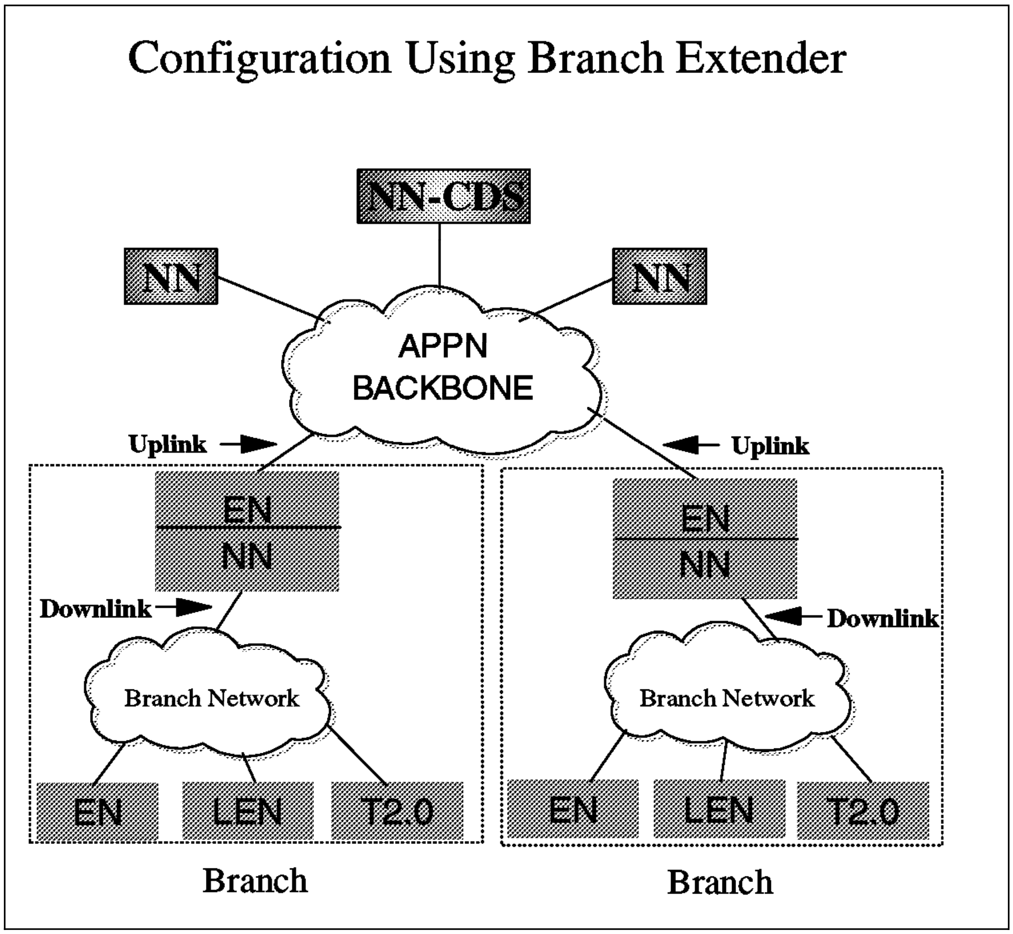

Figure 13. Branch Access Configuration

Like a PBN, a branch extender implements both the network node and the end node functionalities. It appears as a network node to the APPN nodes downstream of it, but presents an end node image to the APPN nodes upstream of it. Since each BX is seen by the backbone as an end node, *none* of the topology in the branches is maintained in the backbone topology database. This can save a large amount of topology update traffic, particularly if the branch configurations are volatile.

In fact, the backbone topology database includes unidirectional TGs between NNs and BXs adjacent to them. These TGs are marked *quiesced* and are used only for management purposes (to allow easy identification of BXs).

Since the BX looks like an end node upstream, it must have a network node server in the backbone network. The BX presents itself to its NN server as the owner of all resources downstream of itself. It uses a process called *resource hierarchy modification*, whereby it modifies search and session setup requests that cross its border. Downstream CPs are downgraded to LUs as far as the backbone is concerned. What is actually an LU on an EN served by the BX is presented upstream as an LU on the BX. Additional control vectors are added to the Locate requests to tell the upstream network of the *true* hierarchy, but this information is used only for management purposes and is not required for route calculation or session setup. Thus the nodes in the backbone network need not be aware of the branch extender′s presence or function.

4.1.1 Resource Registration

One of the major differences between the BX and PBN functions is the ability to register resources *across* the BX border. The BX can, of course, register its own LUs with its backbone NN server and/or with a backbone CDS. End nodes within the branch, similarly, register their own resources with the BX (which is their NN server), with or without a request to forward the registration to a CDS. The BX takes the registration request, modifies the hierarchy appropriately, and sends it into the backbone. Thus LUs registered at the BX are also registered at the BX′s backbone NN server. LUs registered at the BX with a CDS forwarding request are registered at the backbone NNS, together with the indication that CDS registration is required.

If a BX registers its downstream resources with its upstream NN server, it can request not to be searched on domain broadcasts, thus eliminating the need for searches to be propagated into the branch. However, this may not always be possible if the BX acts as a DLUR. Unless the BX supports NNS registration of DLUR resources, and the NNS implements APPN option set 1116 (see 3.3, "DLUR/S Design Considerations" on page 32), registration of DLUR resources cannot take place and the BX must allow itself (and thus, potentially, the branch network) to be searched on a domain broadcast.

The registration function, which never happens across extended or peripheral borders, can save a lot of search traffic. If (as is usual) every branch has the same network ID, a border node in the backbone would have to forward a search for an unknown resource to *every* branch in turn. The only way to eliminate this requirement with a border node would be to code suitable adjacent cluster tables and a DSME exit, relying on naming conventions to permit the exit to select the correct subnetwork. A BX removes the need for DSME, and provides better support where the naming conventions do not allow easy location of a resource.

4.1.2 BX Features and Restrictions

As can be gathered from the preceding discussion, the BX design assumes that APPN nodes downstream are all end nodes (or LEN nodes, whose capability is a subset of the EN capability). Network nodes are not permitted in the downstream network, nor are border nodes, which may look like ENs downstream of the BX. However, cascaded BXs *are* permitted.

There may be multiple BXs in a branch, with multiple connections to the backbone, for increased availability. However, each branch EN can only maintain CP-CP sessions and resource registration with one such BX at a time. The BX architecture means that all the sessions from the EN into the backbone will always pass through the same BX. Load balancing at the node level does not work, although ENs within the branch can be distributed between the available BXs as servers. If a BX connection fails, then the branch EN is free to choose another BX as its NN server, and (if HPR is used) its sessions can be switched nondisruptively to the new path.

Direct connections between branches are permitted, but these must be on the upstream side. These appear to the backbone as EN-to-EN connections, and are perfectly acceptable for session routes. A session from an EN in one branch to an EN in another will appear to the backbone as being between adjacent ENs (actually BXs), and should take the direct path provided the TG characteristics and COS tables have been properly set up.

Similarly, a direct connection between BXs within the same branch must be on the upstream side. This means that a session between ENs in the same branch but on *different* BXs must pass through the upstream link. This link, of course, can be physically within the branch so the backbone need not be troubled with the session.

A connection network can be defined on the upstream side, or the downstream side, or both, of a BX. However, if there are multiple BXs in a branch then their domains cannot share the same connection network. Because the BX-to-BX connection is upstream, a BX is not aware of TGs in another BX′s domain and the route must be calculated by an upstream (backbone) NN. Thus a session between two ENs in different BX domains will not take the direct path even if they are on the same connection network.

A BX can choose only one NN in the backbone as its NN server, and maintain CP-CP sessions with it, at any one time. This is not true of a border node managing a peripheral border, which can appear as a served EN to multiple NNs in the backbone. This restriction is imposed to prevent search looping; an EBN has additional logic to prevent this.

4.1.3 Branch Extender and DLUR

The ENs or BXs downstream of the backbone-connected BXN cannot be DLUR nodes. The cross-border DLUR/S function is complex and requires EBN-level logic to implement. If DLUR is used in a branch, the backbone-connected BX itself must supply the function. This means that the DLUR benefits (APPN all the way) cannot be realized if the nodes in the branch network are themselves capable of it. Some customers have decided to implement DLUR in the workstations (not the gateway) so BX may clash with the existing network design.

Note

The restriction on downstream DLUR nodes is the subject of a proposal to the APPN Implementers′ Workshop at the time of writing. It may well be removed shortly.

A BX cannot support native (non-DLUR-served) dependent LUs downstream. This means that VTAM cannot be downstream of a BX, even if the VTAM is an end node and the BX supports session services extensions.

4.1.4 Branch Extender and HPR

A BX can act as an RTP node upstream or downstream, as well as performing ANR function across the border. As with a border node, if the BX acts as an ANR router, then the RTP endpoints have to be aware of the *complete* session path and the BX cannot conceal either the upstream or the downstream part. This is accomplished using a control vector CV 4685 to transmit the "hidden" part of the route on the RSCV. CV 4685 is not new to BX; it is the same as is used in cross-border HPR connections.

Since a session RSCV has a maximum length of 255 bytes, the number of hops on a session path is restricted by this limit. In base APPN, a border isolates the route segments so that each subnetwork sees only a subset of the complete route. With HPR, however, the RTP endpoints must be aware of the entire route so that the Route Setup message flows from end to end. An HPR-capable border node, if faced with an excessive RSCV, can split the route into back-to-back RTP connections to alleviate the problem. A branch extender does not have this logic and so, in extreme cases, an HPR route across a BX may be restricted to as few as four hops.

4.2 Branch Extender Implementations

The branch extender function is presently supported by the following products:

- Communications Server/2, Version 5
- Communications Server/NT, Version 6
- 2216 Multiprotocol Access Services, Version 2 and Version 3 (and therefore the 3746 MAE with corresponding software levels)
- 2210 Multiprotocol Routing Services, Version 2 and Version 3

CS/2 or CS/NT must be configured as a network node before BX functions are available, and BX support must then be configured for the node. The same physical port can be used for both upstream and downstream connections, but it must be enabled for upstream support. Ports are so enabled by default on CS/NT, but must be manually configured on CS/2. Individual logical links will then allow themselves to be configured as branch extender uplinks or not. CS/2 and CS/NT support the DLUR, HPR and connection network options in conjunction with BX as described above.

The 2210, 2216 and MAE are always network nodes if they are APPN-capable. BX support can be configured at the node, adapter (physical port) or station (logical connection) level. There are no restrictions on sharing adapters; any combination of upstream and downstream connections can be configured on any set of ports.

Chapter 5. Enterprise Extender

One of the major changes in corporate networks in recent years has been the growth in Internet-based communication, whose explosive expansion has resulted in huge new business opportunities. At the same time this technology has made its way into internal business applications, since the same workstation with the same software can be used to access both internal and external application sites. The result has been a massive increase in the need for TCP/IP communication.

At the same time, many large organizations have retained and upgraded their SNA networks, because of the requirement for consistent and manageable service levels for critical applications. TCP/IP, because of its underlying connectionless Internet Protocol, is inherently unstable and unmanageable. Its former major advantage over SNA, the ability to reroute automatically around failure points, is now present in HPR. The challenge facing most customers today is how best to integrate the SNA and TCP/IP worlds while preserving the advantages of both.

There are many possible technical solutions to this challenge, most of which are beyond the scope of this book. However, no document describing HPR can be regarded as complete without at least an overview of the enterprise extender technology, which is one of these solutions.

Enterprise extender is aimed at those customers who have decided to implement an IP routing backbone network, yet require SNA-like consistency and predictability for critical applications. The objectives of enterprise extender are:

- To provide SNA connectivity over an IP backbone
- To support current SNA-exploiting functions such as those available in a Parallel Sysplex (generic resources and multinode persistent sessions)
- To allow SNA sessions to be prioritized at the user′s discretion, both among themselves and against TCP and UDP traffic on the same backbone
- To provide better levels of service (response time and throughput) than are available from previous SNA-over-IP technologies such as data link switching, bridging, and AnyNet
- To operate with minimal (or, ideally, none) changes to the typical IP backbone network
- To be configurable with minimum definition

Such a design requires that the IP network somehow takes account of the SNA class of service. Also, in order to compensate for the IP network′s inherent instability, it requires an end-to-end protocol that can tolerate lost packets and network outages with minimum effort. Fortunately, the former requirement can be met on many IP backbones thanks to the presence of a priority scheme in most routers. The latter requirement is met by HPR, which also happens to be a prerequisite for the full exploitation of a sysplex. The enterprise extender technology, therefore, is based on running HPR over UDP/IP.

5.1 Enterprise Extender Description

The enterprise extender architecture, called upon to carry HPR over an IP backbone, must treat the IP network as another type of DLC over which NLPs flow. It has three choices as to the method of transporting NLPs:

- A TCP connection would impose unacceptable overhead, and provides no additional function that HPR over UDP does not already have. RTP implements end-to-end packet loss detection and resequencing. The only function that HPR requires from its DLC at the link level is the detection of corrupted packets, which UDP can also provide. An ANR node should not be burdened with the administration of a TCP connection, and the re-establishment of such a connection when a path switch is required.
- Raw IP datagrams incur no such overhead. However, they do not provide any means of identifying the process at the destination node which is to handle the data. It is not a good idea to restrict an IP host to enterprise extender traffic *only*.
- UDP also uses datagrams, but has a built-in port number that can identify the process to which a packet is directed. Moreover, many routers can be configured to prioritize the IP traffic based on the UDP port number. UDP can also detect corrupted packets. Being connectionless it cannot detect lost packets or packets out of sequence, but RTP looks after all that.

Figure 14 illustrates the concept of an enterprise extender connection.

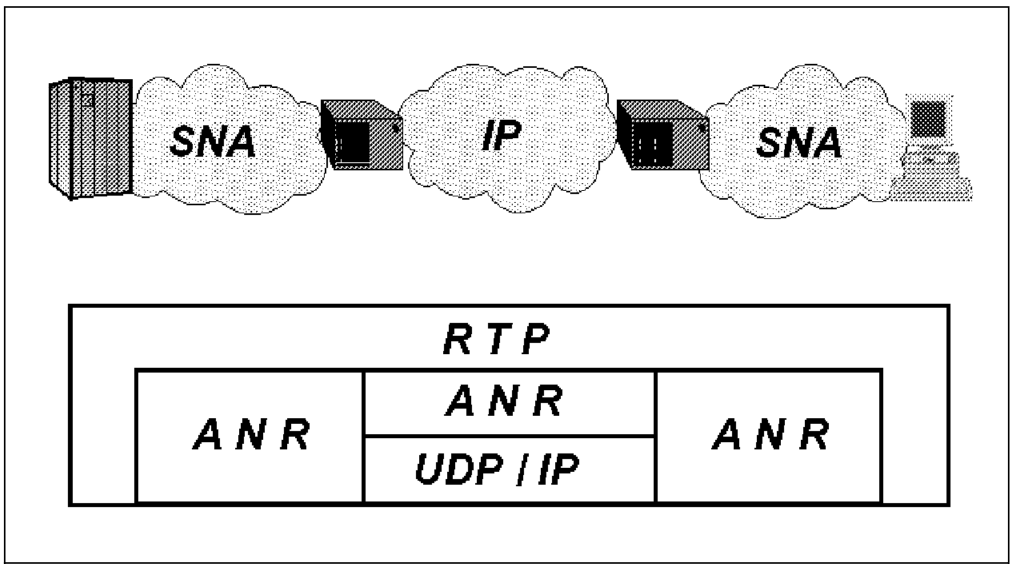

Figure 14. Enterprise Extender Operation

The enterprise extender function is very similar in concept to the way native SNA over ATM is implemented:

- The underlying transport network appears as an APPN TG, but uses logical data link control (LDLC) to exchange XIDs and NLPs. LDLC is a subset of LLC2 that eliminates much of the error handling and acknowledging that RTP makes unnecessary at link level. LDLC, used also for native SNA over ATM, includes only the XID, TEST, DISC, DM and UI frame types.
- The UDP port number identifies the destination of the datagram as being the partner IP host′s ANR routing function. Five UDP ports have been registered with the IANA for this purpose. One of these default ports is mapped to each

of the APPN transmission priority values, with the fifth being used for XID exchange. ANR labels are mapped to the partner's IP address.

- LDLC permits no link-level error recovery, which means that only HPR NLPs can be transported once the XID flows are completed. Therefore, both partner nodes must support control flows over RTP. Also, the restriction described on Page 18 applies; an enterprise extender connection cannot be adjacent to an IC-TG unless the other end of the connection is an end node.
- A connection network can be defined on the IP network, which uses logical enterprise extender links. Defining all such logical links between each pair of a large number of IP addresses would be an unpleasant task, just as it would be on a LAN.
- The SNA transmission priority is mapped to the UDP port number, which is why five UDP ports have been registered for enterprise extender use. The main reason for this is that many IP routers can be configured to prioritize traffic based on the port number. However, the enterprise extender architecture permits the use of the precedence bits in the IP header for the same purpose. These bits are reserved in the TCP/IP architecture for exactly this usage, but in practice few routers take account of them. Among those few are the 2216, 2210 and MAE, which set both the precedence bits and the UDP port number.

An enterprise extender connection can use ARB flow control, as described in 1.5, "ARB Flow and Congestion Control" on page 8, to ensure the smooth passage of SNA traffic across an RTP connection. However, the standard ARB algorithm is designed to promote fairness between various types of SNA traffic; IP traffic tends to take as much resource as it can and does not take account of any rival traffic. Therefore, a new version of ARB has been developed to allow enterprise extender traffic to compete fairly with IP traffic on the same IP network. This is called *responsive mode ARB*, and is an option for nodes that implement enterprise extender.

Because of its design, the enterprise extender architecture is extremely flexible. It can be used in all networks from the smallest to the largest, and provides the customer with a wide choice of where the SNA/IP boundary is placed. It enables remote branches or workstations to be connected to the SNA backbone using the Internet, with no application changes required, while maintaining SNA connectivity from end to end. Using DLUR/S, dependent LU sessions can be carried on an enterprise extender connection as easily as any others.

5.2 Enterprise Extender Implementations

At the time this redbook was written, there were five products which implemented, or planned to implement, enterprise extender:

- The 2216 with Multiprotocol Access Services Version 3, Release 1
- The 3746 MAE with the equivalent level of software (planned for autumn 1998)
- The 2210 with Multiprotocol Routing Services Version 3, Release 1
- Communications Server for Windows NT, Version 6
- Communications Server for OS/390. At the time of writing the schedule for this was not known.

The 221X family supports the responsive mode ARB option, but CS/NT does not.

Chapter 6. High-Performance Routing on VTAM

In the testing scenarios described in the latter chapters of this redbook, we start by implementing HPR on VTAM alone, and gradually extend it to the furthest reaches of the network via NCP, 3746, 2210 and 2210 all the way to the workstation. At the same time we extend DLUR support outwards from the 3746 to the workstation. Our objective in this first test is to get HPR working among MVS VTAM nodes.

6.1 VTAM As an APPN/HPR Node

The level of HPR support available on a VTAM node depends on the APPN role of the node itself, the VTAM release level and the capabilities of partner nodes. If desired, HPR support can be switched off or downgraded (from RTP to ANR) on a node-wide basis. Within a node providing HPR support, individual connections can be configured by using the appropriate keywords on the link definitions.

HPR can be utilized on all connections supported by a VTAM or its owned NCPs, including subarea connections if VR-TG is defined. HPR allows you to migrate existing NCP connections to APPN without incurring the additional overhead associated with converting FID-4 links to FID-2 links; NCP handles the former more efficiently than the latter. If the connections are converted directly to HPR, then NCP is not aware of the sessions passing across those connections, with a corresponding saving in processing power and storage.

Of course, the other main advantage of HPR is that a failure in a connection can be bypassed (without disrupting the sessions that were using that connection) if there is an alternative path. This applies equally to VTAM-attached, NCP-attached and VR-TG connections.

For a stand-alone VTAM node, RTP and ANR support has been available on BF-TGs (FID-2 connections) since V4R3. V4R4 introduced VR-TG support and Control Flows over RTP. The Control Flows option is not available, however, over XCA LAN connections even with CS OS/390 Release 5. HPR over an XCA-connected LAN was first made available in APAR OW26732.

6.1.1 Implementation Overview

There are three levels of HPR support for a VTAM node. If VTAM is defined as an APPN node, the default is RTP support (with the Control Flows option); if not, there is no HPR support available. The three levels are specified using the HPR start option:

- Rapid transport protocol (HPR=RTP)

  An RTP connection can exist between a VTAM node with APPN capability (NODETYPE is coded) and another RTP-capable APPN node.

- Automatic network routing (HPR=ANR)

  A VTAM network node (NODETYPE=NN) can also provide ANR-level support as an intermediate node on an HPR route. You can use the HPR=ANR start option to prevent the NN from being an RTP endpoint. For a stand-alone VTAM this has little practical application, but you may wish to use this in a CMC environment (see Chapter 7, “HPR between CNN Nodes” on page 73)

where you want NCPs to perform ANR routing but wish to minimize the load on the owning VTAM.

- No HPR support (HPR=NONE)

  HPR support can be disabled for a particular VTAM NN or EN using the HPR=NONE start option. If HPR is disabled, that VTAM cannot be an endpoint or an intermediate node for an RTP connection. A VTAM subarea node (NODETYPE not coded) cannot provide HPR support, and the HPR start option is ignored.

If the HPR start option is coded as above, all connections owned by the VTAM node by default take on the HPR capability of the node itself. Thus a VTAM that is capable of RTP can act as an RTP endpoint on all attached links. However, you may want to downgrade the capability of particular links so that (for instance) HPR is not permitted at all on some connections, or perhaps HPR traffic arriving on an NCP link can only be rerouted out of the node instead of terminating in VTAM. This downgrading can be done in two ways:

- The second operand on the HPR start option assigns the default HPR capability for all APPN connections. Thus HPR=(RTP,ANR) makes VTAM itself capable of RTP, but allows it to perform only ANR routing for traffic carried on all the attached links. Clearly this setup is the same as HPR=ANR unless at least one link is upgraded to RTP support. The permitted combinations are (RTP,ANR), (RTP,NONE) and (ANR,NONE).
- The HPR keyword on the link definition statement (PU for BF-TGs and CDRM for VR-TGs) can be used to upgrade a connection to VTAM′s level of HPR support, or downgraded to no HPR support. The allowable values are HPR=YES or NO meaning upgrade or downgrade. The greatest flexibility comes if the start option is (RTP,ANR). Then each APPN link can be upgraded to RTP (HPR=YES), downgraded to base APPN (HPR=NO) or left as ANR (no HPR coded).

There are three other definitions that you can use to influence the way VTAM implements HPR:

- The HPRPST start option (the path switch timer) controls how long VTAM will wait for a path switch to complete before giving up and terminating the connection. HPRPST comes into effect as soon as VTAM has detected a failure on an RTP connection.

  VTAM does not always initiate a path switch when it detects an RTP failure. When using multinode persistent sessions in a sysplex, the VTAM in the sysplex declares itself as *mobile* when establishing the RTP connection, whereas VTAM in all other circumstances declares itself as *stationary*. The architecture demands that if one partner is stationary and one is mobile, only the mobile partner can initiate a path switch. Thus the RTP partner *outside* the sysplex must simply wait for the MNPS partner to switch the path.

  There is one HPRPST value for each APPN transmission priority; the default is HPRPST=(8M,4M,2M,1M) for low, medium high, and network priorities respectively. It is recommended that the network priority value is much less than the others. This is because if CP-CP sessions are running on an RTP connection, it is better to terminate them quickly than to wait a long time for the path switch to fail. Those CP-CP sessions will be restarted after termination, and they may be needed to switch the LU-LU sessions that are waiting on the other HPRPST values.

There is an enhancement to the HPR architecture that permits a node to terminate CP-CP sessions immediately, instead of waiting for a path switch timer, if there is no alternative Control Flows-capable link available after a failure. This option is implemented on the 2210/2216, and planned for a future release of CS OS/390.

If HPRPST is too long (for any priority), you might tie up resources unnecessarily and leave an LU in an unusable state if there is no prospect of an alternative route. If HPRPST is too short, you might cause sessions to fail when an alternative is available but the network is busy and the flows take a long time to complete.

- The PSRETRY start option controls how often VTAM will attempt an automatic path switch for each RTP connection. When an RTP pipe is set up, VTAM sets the PSRETRY timer going. When the timer expires VTAM initiates a path switch for the pipe by requesting a new HPR-only route for it. If the newly calculated route differs from the old one, then VTAM completes the path switch by exchanging Route Setup messages and other connection information; if the new route is the same, then VTAM does nothing more except to restart the timer.

  CS OS/390 Development is working on a refinement to the PSRETRY function, which will permit the installation to specify alternative criteria for switching the path when the PSRETRY timer expires.

  The intention of PSRETRY is to allow RTP connections, which have moved due to a connection failure, to return to their original paths when the failure has been rectified. The alternative is for the operator (human or automated) to issue the MODIFY RTP command at intervals. However, if the network has many parallel TGs of equal weight then PSRETRY could result in excessive path switching due to VTAM′s habit of allocating session paths in turn between such TGs.

  PSRETRY has a value for each APPN transmission priority; the default is zero for each, meaning that no automatic switching is to be performed.

- The LLERP keyword is used on link station definitions to specify the link-level error recovery capabilities of a connection. Remember that the partner′s error recovery requirements must be compatible if an HPR connection is to be established. For most VTAM connections there is no choice; channel connections, X.25 and VR-TGs always have link-level error recovery but ATM connections forbid it. For switched LAN connections through an XCA (3172 or OSA), the default is LLERP=NOTPREF, which means it will not be used unless the partner demands it.

6.2 Example of HPR Implementation

For the first example, we have four MVS systems all running VTAM V4R4 as interchange nodes (ICNs) connected together as shown in Figure 15 on page 52. They are ICNs because we wish to demonstrate the compatibility between VR-TGs and BF-TGs when running HPR. Thus the VTAMs need both subarea and APPN functions.

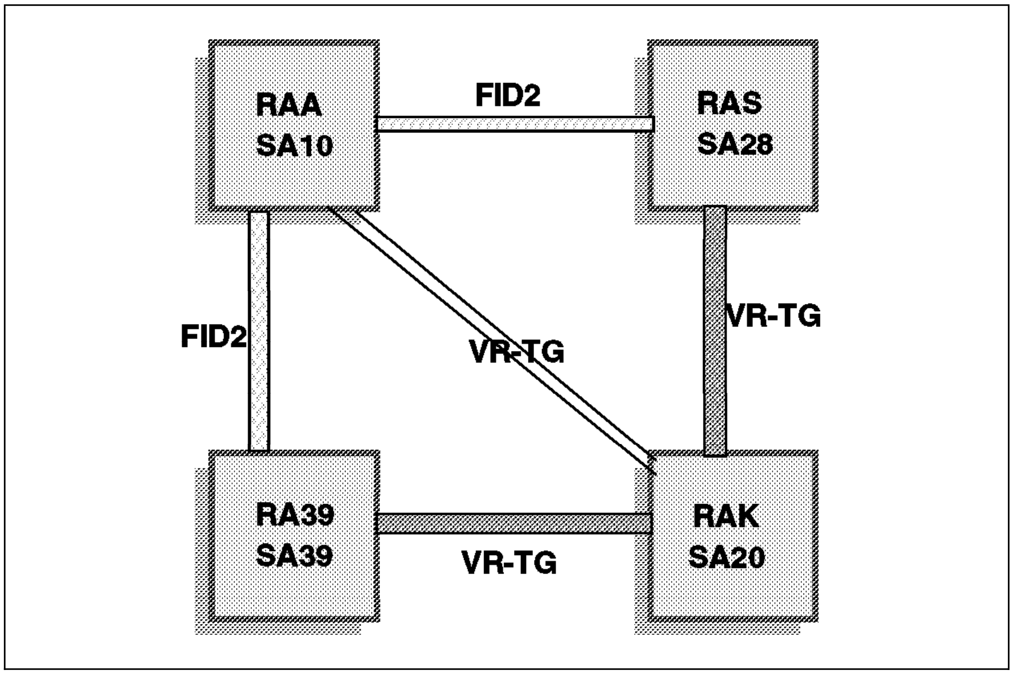

Figure 15. VTAM NN Network Configuration

RAK is connected to the other systems via virtual route-based transmission groups (VR-TGs). The connection between RAA and RAS is a FID-2 BF-TG as is the one between RAA and RA39.

Once VTAM is running as an APPN node, there are very few changes you need to make to the start options to enable HPR functions. We let them all default so the values in use were:

- HPR=RTP, so full RTP support is available on all connections. Before VTAM V4R4 this option was not valid on ICNs. In our case, if the ICN is in the middle of a session path it can also act as an ANR router. HPR=RTP can be overridden by the HPR keyword in the following definitions: local SNA, NCP, switched major node, and CDRM. Before VTAM V4R4 the HPR keyword was not valid on the CDRM statement.
- HPRPST=(8M,4M,2M,1M), so the time allowed to switch an RTP path is lower for the higher-priority connections. Before VTAM V4R4 there were only three values because the network priority was not supported.
- PSRETRY=(0,0,0,0) so no automatic path switch is to be attempted.
- Link-level error recovery will be done only where required by the partner node or by the DLC itself.

6.3 HPR and Path Switch

From local resources on RAK (in fact they were NetView/Access pseudo-terminals) we established sessions with TSO and NetView running on RAA. Before deactivating any links we displayed all the active RTP connections (pipes) on RAA by issuing the command `D NET,ID=ISTRTPMN,E` as shown in Figure 16 on page 53.

```
D NET,ID=ISTRTPMN,E
IST097I  DISPLAY  ACCEPTED
IST075I  NAME = ISTRTPMN         , TYPE = RTP MAJOR NODE
IST486I  STATUS= ACTIV     , DESIRED STATE= ACTIV
IST1486I    RTP NAME   STATE        DESTINATION CP  MNPS  TYPE
IST1487I    CNR00070   CONNECTED    USIBMRA.RAS     NO    RSTP 2
IST1487I    CNR0006F   CONNECTED    USIBMRA.RAS     NO    CPCP 3
IST1487I    CNR0006E   CONNECTED    USIBMRA.RAS     NO    CPCP 3
IST1487I    CNR0004A   CONNECTED    USIBMRA.RAK     NO    LULU 1
IST1487I    CNR00048   CONNECTED    USIBMRA.RAK     NO    LULU 1
IST1487I    CNR00046   CONNECTED    USIBMRA.RAK     NO    LULU 1
IST1487I    CNR0001F   CONNECTED    USIBMRA.RA39    NO    LULU
IST1487I    CNR0001E   CONNECTED    USIBMRA.RA39    NO    RSTP 2
IST1487I    CNR00019   CONNECTED    USIBMRA.RA39    NO    CPCP 3
IST314I  END
```

Figure 16. Display of Active RTP Connections on RAA

ISTRTPMN is a major node created by VTAM, to which an entry is added for each RTP connection established with any other node.

1 You can see that we have three RTP pipes between RAA and RAK, all labelled LULU. These are for LU-LU sessions that we have established with TSO. There are several possible reasons why the sessions are spread over three connections. Most likely they take different paths through the network. Possibly they have different APPN COS values, or some of them are the results of path switches.

2 These pipes, labelled RSTP, were created for Route Setup flows. They are the *long-lived* pipes described in 1.8, "HPR Control Flows over RTP Option" on page 12, and their presence proves that the partner VTAMs are using Control Flows over RTP on these links. RSTP pipes are only ever established between adjacent nodes, and are only ever used for Route Setup flows. They are never path-switched. There is no RSTP pipe between RAA and RAK because Control Flows are not supported over a VR-TG.

3 Because VTAM V4R4 supports Control Flows over RTP, the same pairs of nodes have a CPCP RTP pipe as have an RSTP pipe. The CPCP pipe is only used for CP-CP sessions, and *will* be path-switched if a suitable alternative route is available. Again, there is no CPCP pipe between RAA and RAK because their only connection is a VR-TG.

The CPCP and RSTP pipes are set up at different times and under different circumstances. The CPCP pipe is established immediately after the XIDs have been exchanged, provided both nodes request CP-CP sessions and the architecture permits their establishment at this time. There may be one, or at most two, CP-CP pipes between any two nodes. If there are two, it means that each partner has independently established the RTP connection at the same time as the other, even though the route and the APPN COS are the same. Usually on an EN to NN connection there is only one pipe, because the NN does not set up its CONWINNER session until the EN has done so; thus the NN will already have a suitable RTP pipe to re-use.

The RSTP pipe is established over *each* eligible connection between two nodes, when an LU-LU session requires that link. Therefore, there may be as many RSTP pipes as there are Control Flows capable links between two nodes.

The corresponding display on RAK is shown in Figure 17 on page 54. Note that the CNRnnnnn connection names do not match, since they are PU names and have only local significance. Since RAK′s connections are all VR-TG, there are no RSTP or CPCP connections to be seen.

```
D NET,ID=ISTRTPMN,E
IST097I DISPLAY ACCEPTED
IST075I NAME = ISTRTPMN, TYPE = RTP MAJOR NODE
IST486I STATUS= ACTIV, DESIRED STATE= ACTIV
IST1486I   RTP NAME   STATE            DESTINATION CP     MNPS  TYPE
IST1487I   CNR0000F   CONNECTED        USIBMRA.RAA        NO    LULU
IST1487I   CNR0000D   CONNECTED        USIBMRA.RAS        NO    LULU
IST1487I   CNR0000B   CONNECTED        USIBMRA.RAA        NO    LULU
IST1487I   CNR00009   CONNECTED        USIBMRA.RA39       NO    LULU
IST1487I   CNR00006   CONNECTED        USIBMRA.RAS        NO    LULU
IST1487I   CNR00005   CONNECTED        USIBMRA.RAA        NO    LULU
IST314I END
```

Figure 17. Display of ISTRTPMN on RAK

To see how many sessions and which LUs are mapped on to these RTPs, as well as the physical path they are using, a further display is needed. See Figure 18 for a display of one particular RTP pipe on RAA.

```
DISPLAY NET,ID=CNR0004A,SCOPE=ALL
IST097I  DISPLAY  ACCEPTED
IST075I  NAME = CNR0004A         , TYPE = PU_T2.1  4
IST1392I  DISCNTIM = 00010 DEFINED AT PU   FOR DISCONNECT
IST486I  STATUS= ACTIV--LX-, DESIRED STATE= ACTIV
IST1043I  CP NAME = RAK     , CP NETID = USIBMRA , DYNAMIC LU = YES
IST1589I  XNETALS = YES
IST933I  LOGMODE=***NA***, COS=#CONNECT  10
IST1476I  TCID X′1B6DB739000000AF′ - REMOTE TCID X′1EC14EC8000001EF′  5
IST1481I  DESTINATION CP USIBMRA.RAK       - NCE X′D000000000000000′  5
IST1587I  ORIGIN NCE X′D000000000000000′  5
IST1477I  ALLOWED DATA FLOW RATE =     18 KBITS/SEC  6
IST1516I  INITIAL DATA FLOW RATE =   6400  BITS/SEC  6
IST1511I  MAXIMUM NETWORK LAYER PACKET SIZE =   4046 BYTES  7
IST1478I  NUMBER OF UNACKNOWLEDGED BUFFERS =     0
IST1479I  RTP CONNECTION STATE = CONNECTED            - MNPS = NO
IST1480I  RTP END TO END ROUTE - PHYSICAL PATH
IST1460I  TGN  CPNAME            TG TYPE      HPR
IST1461I  255  USIBMRA.RAK       VRTG         RTP  8
IST231I  RTP      MAJOR NODE = ISTRTPMN
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I  STATE TRACE = OFF
IST355I  LOGICAL UNITS:  9
IST080I  RAKTX022 ACT/S----Y RAKTX017 ACT/S----Y RAKTX058 ACT/S----Y
IST080I  RAKTX046 ACT/S----Y RAKTX034 ACT/S----Y
IST314I  END
```

Figure 18. Display of Active Sessions Mapped to Pipe CNR0004A on RAA

Note the following:

4 VTAM treats an RTP connection as another kind of APPN link station (PU type 2.1).

5 The RTP connection is uniquely identified by the NCE and the TCID on each side. VTAM uses the same NCE value for all LU-LU sessions unless special processing is required such as for MNPS.

6 The initial and current ARB flow rate values are associated with each RTP pipe.

7 The maximum NLP size that can flow on this connection has been determined from the maximum packet sizes on each link, which in turn has been determined by XID exchange.

8 The route for this RTP pipe is directly over the VR-TG to the RTP partner. There is no way of displaying the subarea route that is used within the VR-TG.

9 The LU-LU sessions that are carried on this pipe are linked to the RTP connection, not to the physical connection. In base APPN you will find the LU session partners listed under the real PU, not the RTP PU.

10 This pipe is for the #CONNECT APPN class of service.

Now we display another (different) RTP connection on the RAK side (see Figure 19).

```
D NET,ID=CNR00005,E
IST097I DISPLAY ACCEPTED
IST075I NAME = CNR00005, TYPE = PU_T2.1 456
IST1392I DISCNTIM = 00010 DEFINED AT PU FOR DISCONNECT
IST486I STATUS= ACTIV--LX-, DESIRED STATE= ACTIV
IST1043I CP NAME = RAA, CP NETID = USIBMRA, DYNAMIC LU = YES
IST1589I XNETALS = YES
IST933I LOGMODE=***NA***, COS=#CONNECT
IST1476I TCID X'1EC14EBE000001F5' - REMOTE TCID X'1B6DB7350000008E'
IST1481I DESTINATION CP USIBMRA.RAA - NCE X'D000000000000000'
IST1587I ORIGIN NCE X'D000000000000000'
IST1477I ALLOWED DATA FLOW RATE = 14 KBITS/SEC
IST1516I INITIAL DATA FLOW RATE = 6400 BITS/SEC
IST1511I MAXIMUM NETWORK LAYER PACKET SIZE = 4046 BYTES
IST1478I NUMBER OF UNACKNOWLEDGED BUFFERS = 0
IST1479I RTP CONNECTION STATE = CONNECTED - MNPS = NO
IST1480I RTP END TO END ROUTE - PHYSICAL PATH
IST1460I TGN  CPNAME             TG TYPE      HPR
IST1461I 255  USIBMRA.RAA        VRTG         RTP
IST231I RTP MAJOR NODE = ISTRTPMN
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST355I LOGICAL UNITS:
IST080I RAAAN010 ACT/S----Y RAAAT02  ACT/S----Y RAAAN009 ACT/S----Y
IST080I RAAT420  ACT/S----Y
IST314I END
```

Figure 19. Display of Active Sessions Mapped to Pipe CNR00005 on RAK

You can tell this is not the same pipe as previously displayed, because the TCIDs are different. The NCEs, the route and the APPN class of service are the same. Both CNR00005 and CNR0004A are using the VR-TG between RAA and RAK, but which ER is actually used for each pipe cannot usually be determined from VTAM displays.

At this point we broke the direct connection between RAK and RAA by deactivating the CTC link on the RAA side. The messages we received on RAA are shown in Figure 20 on page 56.

```
17:53:21   V NET,INACT,ID=RAACTCA
17:53:21   IST097I  VARY     ACCEPTED
17:53:21   IST1494I  PATH SWITCH STARTED   FOR RTP CNR0004A  11
17:53:21 2 IST1133I  RAAPC13           IS NOW INACTIVE, TYPE = LINK STATION
17:53:22   IST526I  ROUTE FAILED FROM         10 TO         20 - DSA ...
17:53:22   IST526I  ROUTE FAILED FROM         10 TO         20 - DSA ...
17:53:22   IST526I  ROUTE FAILED FROM         10 TO         20 - DSA ...
17:53:22 2 IST1133I  RAACTCA           IS NOW INACTIVE, TYPE = CA MAJOR NODE
17:53:22   IST526I  ROUTE FAILED FROM         10 TO         20 - DSA ...
17:53:22   IST1097I  CP-CP SESSION WITH USIBMRA.RAK       TERMINATED
17:53:22   IST1280I  SESSION TYPE = CONWINNER - SENSE = 08420001
17:53:22   IST314I  END
17:53:22   IST1196I  APPN CONNECTION FOR USIBMRA.RAK       INACTIVE - TGN = 255
17:53:22   IST819I  CDRM RAK      COMMUNICATION LOST - RECOVERY IN PROGRESS
17:53:22   IST1097I  CP-CP SESSION WITH USIBMRA.RAK       TERMINATED
17:53:22   IST1280I  SESSION TYPE = CONLOSER  - SENSE = 08420001
17:53:22   IST314I  END
17:53:22   IST663I  INIT OTHER REQUEST               FAILED  , SENSE=80140001
17:53:22   IST664I  REAL  OLU=USIBMRA.RAA         REAL  DLU=USIBMRA.RAK
17:53:22   IST889I  SID = F7FF61644FEBEA1A
17:53:22   IST314I  END
17:53:22   IST1110I  ACTIVATION OF CP-CP SESSION WITH USIBMRA.RAK FAILED
17:53:22   IST1280I  SESSION TYPE = CONWINNER - SENSE = 80140001
17:53:22   IST1002I  RCPRI=0004 RCSEC=0000
17:53:22   IST314I  END
17:53:23   IST1494I  PATH SWITCH COMPLETED FOR RTP CNR0004A  11
17:53:23   IST1480I  RTP END TO END ROUTE - PHYSICAL PATH
17:53:23   IST1460I  TGN  CPNAME             TG TYPE      HPR
17:53:23   IST1461I   21  USIBMRA.RAS        APPN         RTP
17:53:23   IST1461I  255  USIBMRA.RAK        VRTG         RTP
17:53:23   IST314I  END
```

Figure 20. Path Switch for CNR0004A. The IST526I messages have been truncated to allow the time stamps to appear.

We only saw messages related to the CNR0004A pipe on RAA **11**. The other two LULU pipes (CNR00046 and CNR00048) were also recovered, but we must look to RAK for the related messages. The pipes are called CNR00005 and CNR0000B on RAK. See Figure 21 on page 57 for RAK′s log, where **12** and **13** show the path switch messages.

```
17:53:22   IST1097I  CP-CP SESSION WITH USIBMRA.RAA       TERMINATED
17:53:22   IST1280I  SESSION TYPE = CONLOSER  - SENSE = 08420001
17:53:22   IST314I  END
17:53:22   IST526I  ROUTE FAILED FROM        10 TO        20 - DSA ...
17:53:22   IST526I  ROUTE FAILED FROM        10 TO        20 - DSA ...
17:53:22   IST526I  ROUTE FAILED FROM        10 TO        20 - DSA ...
17:53:22   IST526I  ROUTE FAILED FROM        10 TO        20 - DSA ...
17:53:22   IST1097I  CP-CP SESSION WITH USIBMRA.RAA       TERMINATED
17:53:22   IST1280I  SESSION TYPE = CONWINNER - SENSE = 08420001
17:53:22   IST314I  END
17:53:22   IST1196I  APPN CONNECTION FOR USIBMRA.RAA       INACTIVE - TGN = 255
17:53:22   IST819I  CDRM RAA      COMMUNICATION LOST - RECOVERY IN PROGRESS
17:53:22   IST259I  INOP RECEIVED FOR RAKPC13  CODE = 01
17:53:22 4 IST619I  ID = RAKPC13  FAILED - RECOVERY IN PROGRESS
17:53:22   IST619I  ID = RAKPC13  FAILED - RECOVERY IN PROGRESS
17:53:22   IST521I  GBIND QUEUED FOR COS ISTVTCOS FROM RAK TO RAA
17:53:22   IST528I  VIRTUAL ROUTE NUMBER    0   1
17:53:22   IST523I  REASON = NO ROUTES OPERATIVE
17:53:22   IST1494I  PATH SWITCH STARTED   FOR RTP CNR00005  12
17:53:22   IST663I  INIT OTHER REQUEST              FAILED  , SENSE=80140001
17:53:22   IST664I  REAL  OLU=USIBMRA.RAK         REAL  DLU=USIBMRA.RAA
17:53:22   IST889I  SID = F8D3D164311A5DBD
17:53:22   IST314I  END
17:53:23   IST1110I  ACTIVATION OF CP-CP SESSION WITH USIBMRA.RAA FAILED
17:53:23   IST1280I  SESSION TYPE = CONWINNER - SENSE = 80140001
17:53:23   IST1002I  RCPRI=0004 RCSEC=0000
17:53:23   IST314I  END
17:53:23   IST1494I  PATH SWITCH COMPLETED FOR RTP CNR00005  12
17:53:23   IST1480I  RTP END TO END ROUTE - PHYSICAL PATH
17:53:23   IST1460I  TGN  CPNAME             TG TYPE      HPR
17:53:23   IST1461I  255  USIBMRA.RA39       VRTG         RTP
17:53:23   IST1461I   21  USIBMRA.RAA        APPN         RTP
17:53:23   IST314I  END
17:55:29   IST1494I  PATH SWITCH STARTED   FOR RTP CNR0000B  13
17:55:29   IST1494I  PATH SWITCH COMPLETED FOR RTP CNR0000B
17:55:29   IST1480I  RTP END TO END ROUTE - PHYSICAL PATH
17:55:29   IST1460I  TGN  CPNAME             TG TYPE      HPR
17:55:29   IST1461I  255  USIBMRA.RA39       VRTG         RTP
17:55:29   IST1461I   21  USIBMRA.RAA        APPN         RTP
17:55:29   IST314I  END
```

Figure 21. Path Switch for CNR00005. Again the IST526I messages have been truncated.

The path switch is started by the first RTP endpoint that detects the failure. The criteria used to detect the failure include loss of the local link (the first link on the pipe), or a timeout. It is possible for both endpoints to start the path switch process at the same time because both detect the problem at the same time. In this case the two nodes may even calculate (or have calculated on their behalf) different alternative routes for the new RTP path. If this happens, the path (RSCV) calculated on behalf of the active partner (the one that set up the pipe originally) will be chosen. For a deeper understanding of the mechanism that triggers a path switch, please refer to *Inside APPN - The Essential Guide to the Next-Generation SNA*, SG24-3669-03, or to the *APPN/HPR Architecture Reference*, SV40-1018-02.

When VTAM initiates the path switch, you will see the IST1494I message sequence in the log (see **13** in Figure 21 for example). This is not so if VTAM did not initiate the path switch. Thus we have to look at both logs to understand the full picture.

In Figure 21 you can see at **13** that the path switch messages relating to CNR0000B were received two minutes after the failure, whereas CNR00005 **12** was switched within one second. This may happen when no data is actually flowing across the RTP connection, so the nodes rely on a periodic *keep alive* message that is sent on the pipe. Indeed, CNR0000B was carrying just one control session with mode and COS SNASVCMG.

Let us now take a look at the pipe CNR0004A on RAA. All the sessions are still active and it is now a two-hop path passing through RAS **14** (see Figure 22 and Figure 23 on page 59).

```
DISPLAY NET,ID=CNR0004A,SCOPE=ALL
IST097I  DISPLAY  ACCEPTED
IST075I  NAME = CNR0004A         , TYPE = PU_T2.1
IST1392I  DISCNTIM = 00010 DEFINED AT PU   FOR DISCONNECT
IST486I  STATUS= ACTIV--LX-, DESIRED STATE= ACTIV
IST1043I  CP NAME = RAK     , CP NETID = USIBMRA , DYNAMIC LU
IST1589I  XNETALS = YES
IST933I  LOGMODE=***NA***, COS=#CONNECT
IST1476I  TCID X'1B6DB739000000AF' - REMOTE TCID X'1EC14EC8000001EF'
IST1481I  DESTINATION CP USIBMRA.RAK       - NCE X'D000000000000000'
IST1587I  ORIGIN NCE X'D000000000000000'
IST1477I  ALLOWED DATA FLOW RATE =   6000  BITS/SEC
IST1516I  INITIAL DATA FLOW RATE =   6400  BITS/SEC
IST1511I  MAXIMUM NETWORK LAYER PACKET SIZE =   4046 BYTES
IST1478I  NUMBER OF UNACKNOWLEDGED BUFFERS =     0
IST1479I  RTP CONNECTION STATE = CONNECTED            - MNPS = NO
IST1480I  RTP END TO END ROUTE - PHYSICAL PATH
IST1460I  TGN  CPNAME             TG TYPE      HPR
IST1461I   21  USIBMRA.RAS        APPN         RTP   14
IST1461I  255  USIBMRA.RAK        VRTG         RTP
IST231I  RTP      MAJOR NODE = ISTRTPMN
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I  STATE TRACE = OFF
IST355I  LOGICAL UNITS:
IST080I  RAKTX022 ACT/S----Y RAKTX017 ACT/S----Y RAKTX058 ACT/S----Y
IST080I  RAKTX046 ACT/S----Y RAKTX034 ACT/S----Y
IST314I  END
```

Figure 22. CNR0004A Display on RAA

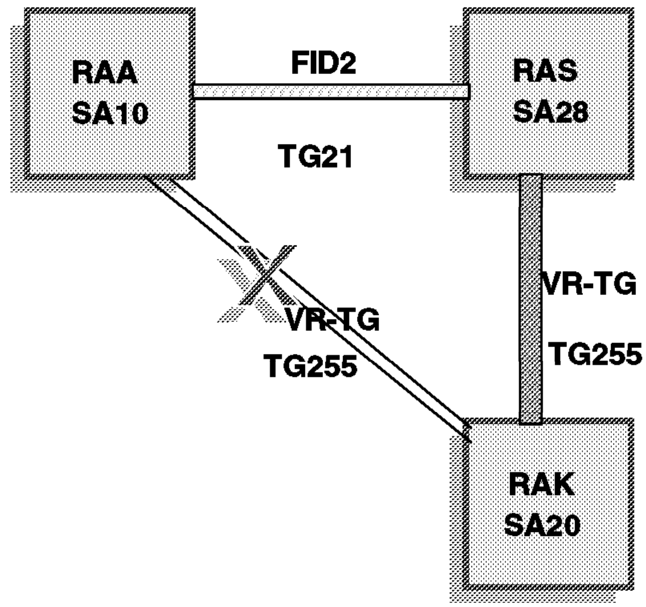

Figure 23. RTP Connection Path after Path Switch

Similarly, the RTP connection named CNR00005 on RAK is now a two-hop connection passing through RA39, as demonstrated at **15** in Figure 24.

```
D NET,ID=CNR00005,E
IST097I DISPLAY ACCEPTED
DISPLAY NET,ID=CNR00005,E
IST075I NAME = CNR00005, TYPE = PU_T2.1
IST1392I DISCNTIM = 00010 DEFINED AT PU FOR DISCONNECT
IST486I STATUS= ACTIV--LX-, DESIRED STATE= ACTIV
IST1043I CP NAME = RAA, CP NETID = USIBMRA, DYNAMIC LU = YES
IST1589I XNETALS = YES
IST933I LOGMODE=***NA***, COS=#CONNECT
IST1476I TCID X'1EC14EBE000001F5' - REMOTE TCID X'1B6DB7350000008E'
IST1481I DESTINATION CP USIBMRA.RAA - NCE X'D000000000000000'
IST1587I ORIGIN NCE X'D000000000000000'
IST1477I ALLOWED DATA FLOW RATE = 6000 BITS/SEC
IST1516I INITIAL DATA FLOW RATE = 6400 BITS/SEC
IST1511I MAXIMUM NETWORK LAYER PACKET SIZE = 4046 BYTES
IST1478I NUMBER OF UNACKNOWLEDGED BUFFERS = 1
IST1479I RTP CONNECTION STATE = CONNECTED - MNPS = NO
IST1480I RTP END TO END ROUTE - PHYSICAL PATH
IST1460I TGN  CPNAME             TG TYPE      HPR
IST1461I 255  USIBMRA.RA39       VRTG         RTP   15
IST1461I  21  USIBMRA.RAA        APPN         RTP
IST231I RTP MAJOR NODE = ISTRTPMN
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST355I LOGICAL UNITS:
IST080I RAAAN010 ACT/S----Y RAAAT02  ACT/S----Y RAAAN009 ACT/S-----Y
IST080I RAAT420  ACT/S----Y
IST314I END
```

Figure 24. CNR00005 Display on RAK

The RTP connections CNR00005 and CNR0004A have been recovered using different alternative paths. As the two routes have the same weight for the APPN COS being used, they could have been recovered on the same path.

The TSO users continued to work during our test without noticing any failure or delay. The HPR path switch was very quick as you can see from the log time stamps.

This test highlights some of the new facilities offered by VTAM V4R4 and later releases, which include:

- HPR over VR-TG
- The ability of an interchange node to act as an RTP endpoint
- Control Flows over RTP

6.4 Forcing a Path Switch

VTAM allows the operator to initiate a path switch, whereupon VTAM (or its NNS) will recalculate the connection route and switch to it if it differs from the existing route. To perform this function you can issue the following operator command:

```
F netproc,ID=CNRnnnnn,RTP
```

To verify this function we first reactivated the subarea connection (and therefore the VR-TG) between RAA and RAK, to ensure that the optimum route was available again. You can see in Figure 25 that the SSCP session, the VR-TG and the CP-CP sessions between RAA and RAK start as soon as the subarea link is restored **16**. As soon as the SSCP-SSCP session has been re-established the two nodes are adjacent again (in APPN terms) so they set up the APPN connection.

We then issued the MODIFY command to switch the path, with the result seen at **17**. The RTP connection has been switched back to the original, direct path.

```
V NET,ACT,ID=RAACTCA
IST097I  VARY     ACCEPTED
IST1132I  RAACTCA           IS ACTIVE, TYPE = CA MAJOR NODE
IST464I  LINK STATION RAAPC13  HAS CONTACTED RAK      SA         20
IST1132I  RAAPC13           IS ACTIVE, TYPE = LINK STATION
IST1086I  APPN CONNECTION FOR USIBMRA.RAK       IS ACTIVE - TGN = 21
IST1132I  USIBMRA.RAK       IS ACTIVE, TYPE = CDRM
IST1096I  CP-CP SESSIONS WITH USIBMRA.RAK       ACTIVATED   16
*
F NETA0,RTP,ID=CNR0004A
IST097I  MODIFY   ACCEPTED
IST1494I  PATH SWITCH COMPLETED FOR RTP CNR0004A   17
IST1480I  RTP END TO END ROUTE - PHYSICAL PATH
IST1460I  TGN  CPNAME             TG TYPE      HPR
IST1461I  255  USIBMRA.RAK        VRTG         RTP
IST314I  END
```

Figure 25. RAA Log during Forced Path Switch

An enhancement implemented with VTAM V4R4 and above is PSRETRY. This was introduced by APAR OW25288, and controls the interval between automatic attempts to switch an RTP pipe to a better path. This feature is controlled by VTAM start options. It was initially turned off (the default is no automatic switching) in our test, so we turned it on using the MODIFY VTAMOPTS command.

The following displays show what happened when we used the PSRETRY option on RAA. In Figure 26 we enabled the automatic path switch feature **18**, specifying that RTP pipes of network priority should be switched every 60 seconds, while the others should be switched every 90 seconds. We then displayed an RTP connection, CNR00001, between RAA and RAK. This connection used the path via RAS, then a VR-TG between RAS and RAK, because the direct route had been deactivated.

```
F NETA0,VTAMOPTS,PSRETRY=(90,90,90,60)  18
*
IST1189I  PSRETRY  = LOW           90S    PSRETRY  = MEDIUM        90S
IST1189I  PSRETRY  = HIGH          90S    PSRETRY  = NETWRK        60S
*
D NET,ID=CNR00001,E
IST097I  DISPLAY  ACCEPTED
IST075I  NAME = CNR00001         , TYPE = PU_T2.1
IST1392I  DISCNTIM = 00010 DEFINED AT PU   FOR DISCONNECT
IST486I  STATUS= ACTIV--LX-, DESIRED STATE= ACTIV
IST1043I  CP NAME = RAK     , CP NETID = USIBMRA , DYNAMIC LU = YES
IST1589I  XNETALS = YES
IST933I  LOGMODE=***NA***, COS=#CONNECT
IST1476I  TCID X'2538DF2500000090' - REMOTE TCID X'1EC14ECF000001F9'
IST1481I  DESTINATION CP USIBMRA.RAK       - NCE X'D000000000000000'
IST1587I  ORIGIN NCE X'D000000000000000'
IST1477I  ALLOWED DATA FLOW RATE =     10 KBITS/SEC
IST1516I  INITIAL DATA FLOW RATE =   6400  BITS/SEC
IST1511I  MAXIMUM NETWORK LAYER PACKET SIZE =   4046 BYTES
IST1478I  NUMBER OF UNACKNOWLEDGED BUFFERS =     1
IST1479I  RTP CONNECTION STATE = CONNECTED            - MNPS = NO
IST1480I  RTP END TO END ROUTE - PHYSICAL PATH
IST1460I  TGN  CPNAME             TG TYPE      HPR
IST1461I   21  USIBMRA.RAS        APPN         RTP
IST1461I  255  USIBMRA.RAK        VRTG         RTP
IST231I  RTP      MAJOR NODE = ISTRTPMN
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I  STATE TRACE = OFF
IST355I  LOGICAL UNITS:
IST080I  RAKTX024 ACT/S----Y RAKTX016 ACT/S----Y RAKTX021 ACT/S----Y
IST314I  END
```

Figure 26. Display of CNR00001 before Path Switch Takes Place

We then activated the direct VR-TG connection between RAA and RAK, as shown in Figure 27 on page 62, and waited. The messages relating to APPN connection and CP-CP session establishment are not shown in the display.

```
V NET,ACT,ID=RAACTCA
IST097I  VARY     ACCEPTED
IST1132I  RAACTCA           IS ACTIVE, TYPE = CA MAJOR NODE
IST464I  LINK STATION RAAPC13  HAS CONTACTED RAK      SA
IST1132I  RAAPC13           IS ACTIVE, TYPE = LINK STATION
*
IST1494I PATH SWITCH COMPLETED FOR RTP CNR00001   19
IST1480I RTP END TO END ROUTE - PHYSICAL PATH
IST1460I TGN  CPNAME             TG TYPE      HPR
IST1461I 255  USIBMRA.RAK        VRTG         RTP
IST314I END
```

Figure 27. PSRETRY Forced Path Switch for CNR00001

Exactly 90 seconds after our MODIFY VTAMOPTS command, the route was switched **19**. The PSRETRY timer runs individually for each RTP connection, to ensure that the path switch attempts are not all tried at the same time. When we initialized the PSRETRY timer a path switch was attempted immediately for CNR00001, but did not work because there was no alternative path. After 90 seconds the switch was attempted again, and this time completed successfully because the new path was available.

Note: Specifying too short an interval might have an adverse impact on network performance, especially if there are a large number of RTP connections.

6.5 Path Switch over VR-TG

Next, we performed a similar test but this time we tried to switch a path *to* a VR-TG instead of *from* a VR-TG. In fact, VTAM will also switch a path *within* a VR-TG; if a subarea route fails, it is possible for the RTP connection to be switched to another subarea route within the same VR-TG, always without terminating the sessions.

If you refer to Figure 23 on page 59, this time we have established LU-LU sessions between RAA and RAS over the MPC connection between them. The alternate path available is that comprising the two VR-TGs from RAA to RAK and from RAK to RAS.

While all the connections were active, we displayed the RTP pipe major node from RAA, as shown in Figure 28. This shows two LU-LU session pipes, and the expected CP-CP and long-lived pipes because the MPC connection supports Control Flows.

```
IST1487I    CNR00078   CONNECTED        USIBMRA.RAS       NO    LULU
IST1487I    CNR00076   CONNECTED        USIBMRA.RAS       NO    LULU
IST1487I    CNR00075   CONNECTED        USIBMRA.RAS       NO    RSTP
IST1487I    CNR00074   CONNECTED        USIBMRA.RAS       NO    CPCP
```

Figure 28. ISTRTPMN before MPC Failure

Next, we deactivated the MPC connection. The results are shown in Figure 29 on page 63.

```
V NET,INACT,ID=RAAAHHC
IST097I  VARY     ACCEPTED
IST1196I  APPN CONNECTION FOR USIBMRA.RAS       INACTIVE - TGN =  21
IST1494I  PATH SWITCH STARTED   FOR RTP CNR00078  1
IST1494I  PATH SWITCH STARTED   FOR RTP CNR00076  1
IST1133I  RAAHRAS           IS NOW INACTIVE, TYPE = PU_T2
IST1494I  PATH SWITCH STARTED   FOR RTP CNR00074  1
IST1133I  RAAAHHC           IS NOW INACTIVE, TYPE = LCL SNA MAJ NODE
IST1488I  INACTIVATION FOR RTP CNR00075 AS PASSIVE PARTNER COMPLETED  2
IST1416I  ID = CNR00075          FAILED - RECOVERY IN PROGRESS
IST1136I  VARY INACT CNR00075          SCHEDULED - UNRECOVERABLE ERROR
IST1133I  CNR00075          IS NOW INACTIVE, TYPE = PU_T2.1
IST871I  RESOURCE CNR00075 DELETED
IST1494I  PATH SWITCH FAILED    FOR RTP CNR00076   5
IST1495I  NO ALTERNATE ROUTE AVAILABLE
IST1494I  PATH SWITCH FAILED    FOR RTP CNR00074   3
IST1495I  NO ALTERNATE ROUTE AVAILABLE
IST1097I  CP-CP SESSION WITH USIBMRA.RAS       TERMINATED  4
IST1280I  SESSION TYPE = CONLOSER  - SENSE = 80020000
IST314I  END
IST1488I  INACTIVATION FOR RTP CNR00076 AS ACTIVE  PARTNER COMPLETED
IST1416I  ID = CNR00076          FAILED - RECOVERY IN PROGRESS
IST1136I  VARY INACT CNR00076          SCHEDULED - UNRECOVERABLE ERROR
IST1097I  CP-CP SESSION WITH USIBMRA.RAS       TERMINATED   4
IST1280I  SESSION TYPE = CONWINNER - SENSE = 80050000
IST314I  END
IST1488I  INACTIVATION FOR RTP CNR00074 AS ACTIVE  PARTNER COMPLETED
IST1416I  ID = CNR00074          FAILED - RECOVERY IN PROGRESS
IST1136I  VARY INACT CNR00074          SCHEDULED - UNRECOVERABLE ERROR
IST1133I  CNR00076          IS NOW INACTIVE, TYPE = PU_T2.1
IST1110I  ACTIVATION OF CP-CP SESSION WITH USIBMRA.RAS       FAILED  12
IST1280I  SESSION TYPE = CONWINNER - SENSE = 80140001
IST1002I  RCPRI=0004 RCSEC=0000
IST314I  END
IST871I  RESOURCE CNR00076 DELETED
IST1133I  CNR00074          IS NOW INACTIVE, TYPE = PU_T2.1
IST871I  RESOURCE CNR00074 DELETED
IST663I  INIT OTHER REQUEST               FAILED  , SENSE=80140001
IST664I  REAL  OLU=USIBMRA.RAA         REAL  DLU=USIBMRA.RAS
IST889I  SID = F7FF61644FEBFFF3
IST314I  END
IST1494I  PATH SWITCH FAILED    FOR RTP CNR00078   5
IST1495I  NO ALTERNATE ROUTE AVAILABLE
IST314I  END
IST1488I  INACTIVATION FOR RTP CNR00078 AS PASSIVE PARTNER COMPLETED
IST1416I  ID = CNR00078          FAILED - RECOVERY IN PROGRESS
IST1136I  VARY INACT CNR00078          SCHEDULED - UNRECOVERABLE ERROR
IST1133I  CNR00078          IS NOW INACTIVE, TYPE = PU_T2.1
IST871I  RESOURCE CNR00078 DELETED
```

Figure 29. Path Switch after VR-TG Failure

The following messages were expected:

1 The path switch was initiated for the CP-CP and LU-LU session pipes.

2 The long-lived pipe was deactivated immediately, since such pipes are never switched. There is no need to switch them.

3 The path switch failed for the CP-CP session pipe. Control Flows are not supported over a VR-TG, so there is no eligible alternate route for the CPCP connection.

4 Because the pipe could not be switched, the CP-CP sessions were terminated and needed restarting over another connection. In this case there was no alternative one-hop connection.

However, the rest of the display was not expected. The LU-LU session pipes were *not* switched **5**, even though there seemed to be a valid alternate APPN route available. Although we were careful to set up the APPN environment for our tests, we had forgotten the old rules of subarea networking.

A VR-TG is an APPN connection between two subarea-capable VTAMs, which:

- Uses the SSCP-SSCP session to establish the connection.
- Uses subarea VRs and ERs between the VTAMs′ domains to carry APPN traffic.

We displayed the active subarea routes (Figure 30) and active CDRMs (Figure 31) from RAA to see what was going on.

```
D NET,ROUTE,DESTSA=28
IST097I  DISPLAY  ACCEPTED
IST535I  ROUTE DISPLAY   4 FROM SA         10 TO SA         28
IST808I  ORIGIN PU = ISTPUSA0 DEST PU = ***NA*** NETID = USIBMRA
IST536I  VR  TP    STATUS   ER       ADJSUB  TGN  STATUS  CUR MIN MAX
IST537I   0   0    INACT     0            8    1  INOP
IST537I   0   1    INACT     0            8    1  INOP
IST537I   0   2    INACT     0            8    1  INOP
IST537I   1   0    INACT     1            5    1  INOP
IST537I   1   1    INACT     1            5    1  INOP
IST537I   1   2    INACT     1            5    1  INOP
IST537I   2   0    INACT     2            5    1  INOP
IST537I   2   1    INACT     2            5    1  INOP
IST537I   2   2    INACT     2            5    1  INOP
IST537I                      3           28    1  INOP
IST537I   3   0    INACT     4           28    1  INOP
IST537I   3   1    INACT     4           28    1  INOP
```

Figure 30. Routes between RAA and RAS

There were no virtual routes defined between RAA and RAS through RAK (subarea 20).

```
D NET,CDRMS
IST097I  DISPLAY  ACCEPTED
IST350I  DISPLAY TYPE = CDRMS
IST089I  RAACDRM  TYPE = CDRM SEGMENT     , ACTIV
IST1546I  CDRM     STATUS    SUBAREA ELEMENT NETID    SSCPID
IST1547I  RAA       ACTIV         10      1  USIBMRA      99
IST1547I  RAK       ACTIV         20      1  USIBMRA      20
IST1454I           2 RESOURCE(S)  DISPLAYED
```

Figure 31. Active CDRMs from RAA

Not surprisingly, there was no SSCP-SSCP session between RAA and RAS either. That would have required a virtual route.

In fact we had no definitions for the missing routes and CDRMs, so we corrected the problem by defining path tables on all three nodes (RAA, RAS, and RAK) and adding an extra CDRM definition to both RAA and RAS. Figure 32 on page 65 illustrates the new CDRM major node on RAA. In fact VRTG=YES and HPR=YES are the defaults, but VRTG=YES is not always desirable as we discuss in 6.6, "HPR with VR-TG Considerations" on page 67. Here, indeed, VRTG=YES was unnecessary.

```
*******************************************************************
*                                                                 *
* CDRM MAJORNODE FOR RAA                                          *
*                                                                 *
*******************************************************************
        VBUILD TYPE=CDRM
       NETWORK NETID=USIBMRA
RAA      CDRM  SUBAREA=10,CDRDYN=YES,CDRSC=OPT
RAK      CDRM  SUBAREA=20,CDRDYN=YES,CDRSC=OPT,TGP=CHANNEL
RAS      CDRM  SUBAREA=28,CDRDYN=YES,CDRSC=OPT,TGP=CHANNEL
*              VRTG=YES,HPR=YES
RA39     CDRM  SUBAREA=39,CDRDYN=YES,CDRSC=OPT,TGP=CHANNEL
```

Figure 32. CDRM Major Node on RAA

We now activated the new path tables. A display of the path table from RAA includes the entries seen in Figure 33.

```
D NET,PATHTAB
IST097I  DISPLAY  ACCEPTED
IST350I  DISPLAY TYPE = PATH TABLE CONTENTS
IST516I  DESTSUB   ADJSUB TGN  ER  ER STATUS   VR(S)
IST517I       28       20   1   0   INACT          0
IST517I       28       20   1   1   INACT          1
```

Figure 33. Path Table from RAA

Now we had some VRs but no SSCP-SSCP session. We activated our new CDRMs to see the display in Figure 34.

```
VARY NET,ACT,ID=RAACDRM1
IST097I  VARY     ACCEPTED
IST1132I  RAACDRM1          IS ACTIVE, TYPE = CDRM SEGMENT
IST1086I  APPN CONNECTION FOR USIBMRA.RAS       IS ACTIVE - TGN = 255   6
IST1132I  USIBMRA.RAS       IS ACTIVE, TYPE = CDRM
IST1096I  CP-CP SESSIONS WITH USIBMRA.RAS       ACTIVATED   7
```

Figure 34. CDRM and VR-TG Activation

As soon as the SSCP-SSCP session was established, both the VR-TG **6** and the CP-CP sessions **7** were set up. The APPN connections were now as shown in Figure 35 on page 66.

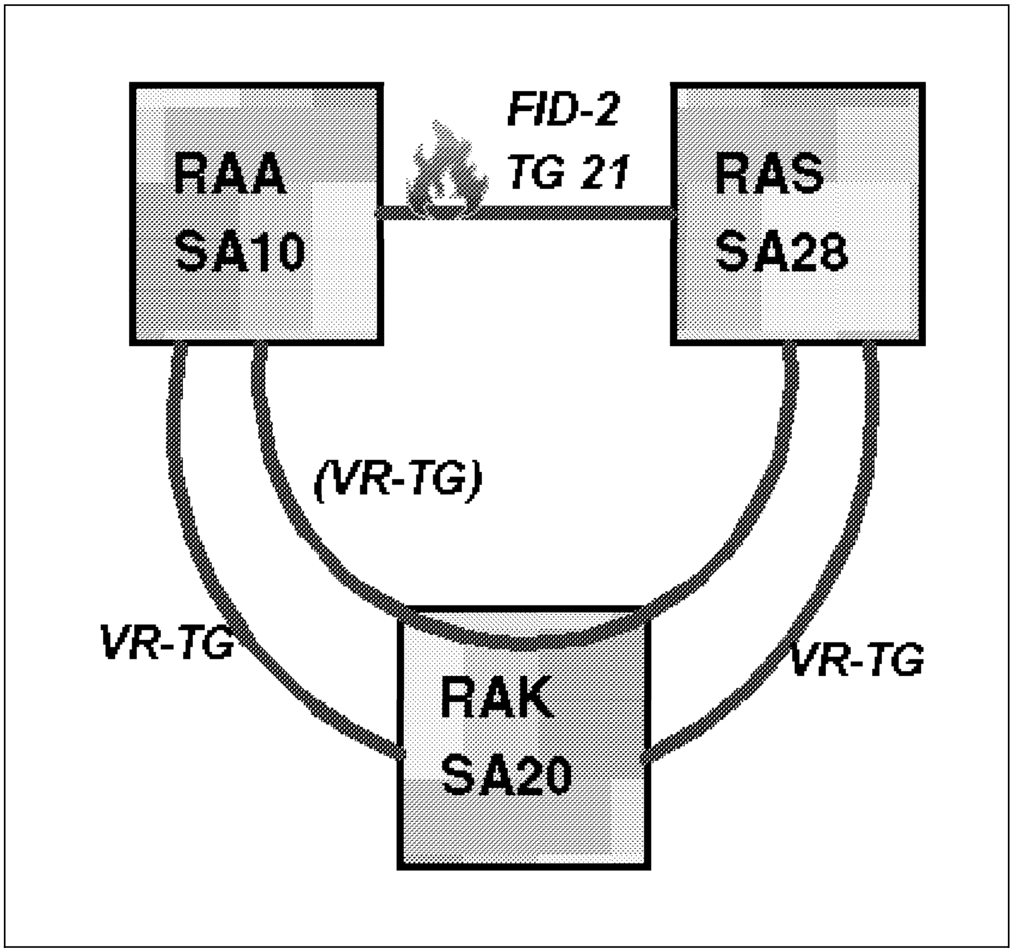

Figure 35. A VR-TG with an Intermediate Node

After the corrections were made, we restored the failed MPC connection, set up the LU-LU sessions again, and restarted our test. Again, we started by displaying the RTP pipes from RAA as shown in Figure 36.

```
D NET,ID=ISTRTPMN,E
IST097I  DISPLAY  ACCEPTED
IST075I  NAME = ISTRTPMN         , TYPE = RTP MAJOR NODE
IST486I  STATUS= ACTIV     , DESIRED STATE= ACTIV
IST1486I    RTP NAME   STATE            DESTINATION CP     MNPS  TYPE
IST1487I    CNR0007A   CONNECTED        USIBMRA.RAS        NO    LULU
IST1487I    CNR00079   CONNECTED        USIBMRA.RAS        NO    RSTP
```

Figure 36. RTP Connections from RAA

We see again the LU-LU session pipe, and the long-lived pipe which was used to set up the LU-LU session pipe. Although the Route Setup messages flow over an RTP connection (RSTP pipe), the CP-CP sessions themselves are on the VR-TG and therefore have no CPCP pipe displayed.

Next, we inactivated the MPC connection again and this time we saw the messages in Figure 37 on page 67.

```
V NET,INACT,ID=RAAAHHC
IST097I  VARY     ACCEPTED
IST1196I  APPN CONNECTION FOR USIBMRA.RAS       INACTIVE - TGN =  21   8
IST1494I  PATH SWITCH STARTED   FOR RTP CNR0007A  9
IST1488I  INACTIVATION FOR RTP CNR00079 AS PASSIVE PARTNER COMPLETED
IST1133I  RAAHRAS           IS NOW INACTIVE, TYPE = PU_T2
IST1416I  ID = CNR00079          FAILED - RECOVERY IN PROGRESS 10
IST1133I  RAAAHHC           IS NOW INACTIVE, TYPE = LCL SNA MAJ NODE
IST1136I  VARY INACT CNR00079          SCHEDULED - UNRECOVERABLE ERROR
IST1133I  CNR00079          IS NOW INACTIVE, TYPE = PU_T2.1
IST871I  RESOURCE CNR00079 DELETED
IST1494I  PATH SWITCH COMPLETED FOR RTP CNR0007A 11
IST1480I  RTP END TO END ROUTE - PHYSICAL PATH
IST1460I  TGN  CPNAME            TG TYPE      HPR
IST1461I  255  USIBMRA.RAS        VRTG         RTP
IST314I  END
```

Figure 37. Path Switch to VR-TG Success

Now everything is as expected:

8 The APPN connection is broken.

9 The path switch commences for the LU-LU pipe.

10 The long-lived pipe is broken and does not recover.

11 The path switch completes successfully, and the new route for the LU-LU sessions takes the VR-TG.

A display of the RTP major node (Figure 38) shows that the LU-LU RTP pipe CNR0007A is now the only active RTP connection between RAA and RAS.

```
D NET,ID=ISTRTPMN,E
IST097I  DISPLAY  ACCEPTED
IST075I  NAME = ISTRTPMN         , TYPE = RTP MAJOR NODE
IST486I  STATUS= ACTIV     , DESIRED STATE= ACTIV
RTP NAME   STATE             DESTINATION CP     MNPS  TYPE
CNR0007A   CONNECTED         USIBMRA.RAS        NO    LULU
```

Figure 38. RTP Major Node after Path Switch

6.6 HPR with VR-TG Considerations

The first volume of this book discussed some issues and recommendations relating to the use of VR-TG in your APPN network. The use of HPR with VR-TG introduces some new considerations and emphasizes some of the old ones.

6.6.1 Path Switch Timer

When an RTP connection is switched from one path to another, a new route calculation is required in most cases. This calculation is performed by the network node server of the path switch initiator, and normally requires a chain of CP-CP sessions between the RTP partners. Therefore, if the failure that causes the path switch also breaks the CP-CP session path between the partners, a new CP-CP path must be established before the RTP path switch can take place.

This is not usually a problem because APPN nodes will try to restore CP-CP sessions over an alternate route as soon as they fail.

However, if you look at the scenario described in 6.5, “Path Switch over VR-TG” on page 62, you will see that this does not happen. In Figure 29 on page 63, the LU-LU session pipe is terminated (the path switch timer expires) for CNR00076 **5** *before* RAA attempts to set up new CP-CP sessions **12**. RAA cannot make this attempt before the existing CP-CP sessions are terminated. Those sessions are currently on a CP-CP RTP connection, and will not be terminated until their path switch timer expires **3**. Therefore, you must make sure that CP-CP sessions that traverse an RTP connection either have an alternative Control Flows capable connection available, or will be terminated quickly.

Note

Ensure that the path switch timer for network priority is much lower than those for the other three priorities. This will give failed CP-CP sessions a chance to recover before LU-LU session pipes time out.

Note, however, that some products will terminate CP-CP sessions immediately if there is no Control Flows-capable alternative connection available. This removes the need to tune the path switch timer for CP-CP sessions.

6.6.2 Obey the Subarea Rules

The test we conducted in 6.5, “Path Switch over VR-TG” on page 62 showed that while VR-TG is APPN, it is very much subarea networking as well. Therefore, the rules of subarea networking must be strictly observed. This is especially true with HPR when you may find routes being chosen at path switch time that you did not expect. RAA and RAS are in the same subarea network and therefore must be in session with each other. The fact that they are not is masked by the fact that they have CP-CP sessions with each other and can use APPN flows to establish sessions. Once the APPN FID-2 connection is broken they have nothing.

This will not usually be a problem if you are converting a working subarea network to APPN, but is quite likely to occur if you add a subarea-capable sysplex to an existing subarea network.

Note

In a subarea network, every VTAM must have an SSCP-SSCP session with every other VTAM. If this is not so, session setup will fail.

Also, every pair of subarea nodes between which an LU-LU session will be established must have a suitable VR defined between them. If this is not so, session setup will fail.

6.6.3 Do Not Apply Subarea Rules to APPN

The reader will observe that when we defined RAS as a CDRM to RAA, we allowed VRTG=YES (and VRTGCPCP=YES) to default. This is not necessary. The subarea rules are observed if a VR and an SSCP-SSCP session exist between RAA and RAS. The APPN rules are observed if there is a chain of APPN connections and a chain of CP-CP sessions between RAA and RAS. Neither the connections nor the CP-CP sessions need to be meshed in the way that SSCP sessions must be meshed, and that applies to VR-TG as well as to FID-2 links. If there had been no VR-TG (and therefore no CP-CP sessions over it) between RAA and RAS, then the following would have happened when the MPC link was broken:

1. The path switch timer starts for the LU-LU session pipe CNR00076 (for example).
2. RAA (as the network node detecting the failure) calculates a new route for CNR00076. It can use the CP-CP session chain RAA - RAK - RAS to reach RAS if it needs to. In this case it does not need to do so, because its RTP partner is in the APPN topology database and its location is known.
3. The new route is the optimum route available, RAA - (TG255) - RAK (TG255) - RAS. There is no direct route because there is no direct VR-TG.
4. Because the session route contains two VR-TGs in succession, RAA combines them into a single VR-TG to make the route RAA - (VR-TG) - RAS. It does this secure in the knowledge that the subarea rules are being observed, and that two VRs in succession are not permitted.
5. Now RAA, knowing that TG255 needs to be converted into a subarea route, chooses a VR list from the subarea session COS and selects a VR between itself and RAS.
6. RAA activates the VR, sends a Route Setup message across it, waits for the route setup reply, and deactivates the VR.
7. From now on the NLPs for the pipe flow on the ER (not the VR) between RAA and RAS.

Because of step 4, there is no need to define a VR-TG between every possible pair of nodes in a subarea network. This feature is called RSCV pruning.

Note

Do not define a VR-TG on every subarea connection in your network. It is even less desirable to establish CP-CP sessions between every possible pair of nodes. APPN requires a contiguous chain of CP-CP sessions, not meshed sessions.

Please see Chapter 3 in *eNetwork Communications Server for OS/390 Network Implementation Guide* for detailed recommendations on where you should define VR-TGs and CP-CP sessions in your subarea network.

6.6.4 Get the TG Characteristics Right

The observant reader will have noticed in Figure 18 on page 54 that the initial rate **6** for the ARB flow control algorithm for the VR-TG between RAA and RAK was 6400 bits per second. Since ARB starts operating on a value of 10% of the defined capacity of the link, this indicates that the characteristics of the VR-TG connection include a CAPACITY value of only 64 kbps. Since this is rather less than the true value, it may result in the VR-TG connection being rejected at route calculation time when it is, in fact, better than the alternative. In addition, it has the result that the ARB flow control algorithm could take a very long time to reach maximum capacity, thereby affecting throughput.

VTAM is not aware of the actual speed of a VR-TG, since a VR-TG could constitute a number of routes of various speeds. It assumes a CAPACITY of 8 kbps for such a VR-TG unless you tell it otherwise. Therefore, it is important that you make sure that VTAM is aware of the true capacity. This can be done by coding CAPACITY= on *each* CDRM definition for a VR-TG connection. However, a better way is to use TG profiles as we illustrate in Figure 32 on page 65. When we corrected the problem of the missing CDRM definitions, we added the keyword TGP=CHANNEL to both CDRM definitions. This points to an entry named CHANNEL in the TG profiles member of VTAMLST. An extract from the TG profiles we were using (IBMTGPS) is shown in Figure 39.

```
CHANNEL  TGP   COSTTIME=0,COSTBYTE=0,SECURITY=SECURE,
               PDELAY=NEGLIGIB,CAPACITY=36M
```

Figure 39. TG Profile for Channel

When we display the topology (Figure 40 is an extract) we can see that CAPACITY=35M **13** for the VR-TG connection.

```
IST097I  DISPLAY  ACCEPTED
IST350I  DISPLAY TYPE = TOPOLOGY
IST1299I  TRANSMISSION GROUPS ORIGINATING AT CP USIBMRA.RAA
IST1357I                                             CPCP
IST1300I  DESTINATION CP    TGN      STATUS   TGTYPE      VALUE WEIGHT
IST1301I  USIBMRA.RAS       255      OPER     INTERM VRTG YES   *NA*
IST1579I                    ------------------------------------------
IST1163I                    RSN               HPR         TIME LEFT
IST1164I                    2                 YES         15
IST1579I                    ------------------------------------------
IST1302I                    CAPACITY PDELAY   COSTTIME    COSTBYTE
IST1303I               13   35M      NEGLIGIB 0           0
IST1579I                    ------------------------------------------
IST1304I                    SECURITY UPARM1   UPARM2      UPARM3
IST1305I                    SECURE   128      128         128
```

Figure 40. Topology Display of VR-TG

Why does the display show a value of 35M when the definition is 36M? This is because the CAPACITY value is stored internally (as are all the TG and node characteristics) as a single byte. Thus precision may be lost, and numbers that are close together will be represented by the same value. The encoding of the CAPACITY value is in a form similar to floating point encoding, and has the effect that high values are less granular than low values. The actual coded value can be seen in the display of an APPN link station (message IST1106I). The CAPACITY is the second byte of the string of hex under TG characteristics. A table of coded CAPACITY values is given in the “TG Profiles” section of *eNetwork Communications Server for OS/390 Resource Definition*. The one to watch for is X′2D′ which means 8 kbps and has probably been defaulted in the absence of better knowledge.

Note

Make sure your TG characteristics are correct. This is particularly important with VR-TG (and with NCP-attached lines) because VTAM has no way of knowing the truth. Code TGP= on both sides of a connection, since TGs are bidirectional.

Chapter 7. HPR between CNN Nodes

This chapter covers HPR implementation on combined VTAM/NCP nodes. This combination is known as a composite network node (CNN), and appears to the APPN network as a single network node. The levels of VTAM and NCP that you need to support HPR in a CNN are as follows:

- To use HPR over an NCP-attached BF-TG you must have, at a minimum, NCP Version 7 Release 3 (and therefore a 3745 controller). Link types supported are channel, token-ring, SDLC and frame relay.
- To use HPR over an X.25 connection requires NCP V7R5 or above, and a 3746-900.
- To use HPR over a subnetwork boundary where an NCP forms one or both sides of the boundary requires NCP V7R5 or above.
- To use HPR over a VR-TG that passes *through* an NCP requires no HPR or even APPN awareness in the NCP.
- The VTAM that owns an NCP with an HPR-capable BF-TG must be V4R3 or above.
- If an NCP with an HPR-capable BF-TG is to send HPR traffic to (or through) its owner, that owner must be V4R4 or above.
- If the VTAM owner of *any* NCP is to be an RTP endpoint, it must be V4R4 or above.
- If an HPR VR-TG passes *through* an NCP, its owner need not be aware of APPN, let alone HPR.

7.1 HPR Definitions in NCP

This section describes the HPR parameters that can be coded in the NCP. None of these is required, since an NCP of the correct release level will be HPR-capable by default. However, some of these parameters affect the ARB algorithm across and within a CNN, and therefore the performance of the network. They apply only to NCPs that have HPR-capable BF-TGs, since an NCP in the middle of an HPR VR-TG is not aware of HPR.

We use the term *composite ANR node* to describe an HPR-capable CNN. Because the NCP is not capable of being an RTP endpoint, each NCP can act only as an ANR node. An NCP can route ANR traffic between a BF-TG and any of:

- Its VTAM owner, via an explicit route
- Another HPR-capable NCP or VTAM, via an ER
- Another HPR-capable BF-TG, directly

The ARB parameters required for an RTP connection are normally obtained from the XID exchanges on the TGs between nodes, and made known to the RTP endpoints via the Route Setup and connection setup flows. If the HPR connection is actually a VR-TG, this is still true, since the SSCP-SSCP session is used to convey the XID parameters. However, if the HPR connection is simply an ER within a CNN there is no XID exchange and the ARB parameters must be coded. They are coded within the NCP source, on the BUILD statement. Please refer to Appendix A, "Adaptive Rate-Based Flow and Congestion Control" on page 231 for details of the ARB algorithms.

Another reason for additional NCP coding relates to the fact that there is no subarea flow control across an HPR VR-TG. NLPs across the subarea network flow on ERs, whereas subarea flow control is done using VR pacing. Therefore, additional parameters are provided in the NCP to address this issue.

7.1.1 Defining the ARB Flow Control Parameters for a CNN

You define these values using the BUILD statement. The characteristics of the composite network node that NCP can provide values for are:

- Whether HPR is to be supported on BF-TGs attached to this NCP (HPR=)
- The maximum packet size (HPRMPS=)
- The accumulated transmission time (HPRATT=)
- The minimum link capacity (HPRMLC=)

HPR on the BUILD definition statement enables NCP to be part of a composite ANR node. The default is HPR=YES.

HPRMPS defines the largest packet size that can be sent across the CNN without being segmented on any of the subarea links along the path. This value must be at least 768 bytes for HPR to work at all. If all the NCPs in the CNN are V7R3 or above, they are capable of determining the value for themselves so code HPRMPS=0 to let them do this. If some of the NCPs are at an earlier level, the later NCPs will not be able to work out the maximum size for those links owned by the back-level NCPs. There is a table in *NCP, SSP and EP Resource Definition Reference* to enable you to work out the correct value. If the HPRMPS value is too large, then extra overhead may be incurred by unnecessary segmenting and reassembly.

HPRATT is used to define the average time in microseconds that it takes to route 1200 bits across the composite ANR node′s subarea network. NCP estimates this time for the subarea network if the value is greater than or equal to 200,000 microseconds. If the value is less than 200,000 microseconds, code a value on the HPRATT keyword. The value you provide is used when transmission begins, being sent in the ARB segment at connection setup time. After transmission has begun the actual value, which the RTP endpoints track, is used.

There is a table in *NCP, SSP and EP Resource Definition Reference* to help you work out the correct values for HPRATT in your network.

HPRMLC specifies the capacity in kilobits per second of the slowest subarea transmission group in the composite ANR node′s subarea network that can carry APPN HPR data. Again, the *NCP, SSP and EP Resource Definition Reference* provides useful assistance in calculating the correct values. Note that the default is 9 kbps, which means that the initial data flow rate allowed by ARB will be just 900 bits per second. You should always code a value here corresponding to your actual network, otherwise performance may suffer.

You do not need to define a maximum packet size, accumulated transmission time, or a minimum link capacity for links to adjacent HPR-capable nodes. NCP determines values for these based on values coded by the user or exchanged with the adjacent node during link station activation. However, this does not always work correctly because the correct CAPACITY value cannot always be worked out by the NCP.

The ARB mechanism determines, at Route Setup time, what traffic rate it will use when the RTP connection is initialized. The Route Setup flows will contain the lowest CAPACITY value of any link in the RTP path, and the RTP endpoint will use a figure of 10% of that value as the ARB starting point. As the Route Setup traverses the RTP path, each node checks the CAPACITY value of the next link and substitutes that value in the Route Setup if it is found to be smaller than the previous value; thus the Route Setup reaches its destination with the minimum CAPACITY value.

If an NCP is on the Route Setup path, it cannot use the APPN CAPACITY value as it does not know it. There is no CP in an NCP and therefore no knowledge of the TG characteristics. Therefore, until V7R5 the NCP used the link SPEED value (if a FID-2 connection) or the HPRMLC value from the BUILD (if a VR-TG connection). Note that Route Setup does not flow between CPs; like a BIND, it flows on what will become the session path.

The issue arises because the SPEED keyword is not coded correctly, if at all, on many NCP definitions. There is usually no need to code it unless internal clocking is used or it is required for performance monitoring.

NCP V7R5 learns the true CAPACITY values of its FID-2 links from its owning VTAM (at VTAM V4R4 or above), and therefore uses the correct values in the Route Setup for the ARB algorithm. This function was implemented by APAR IR33946.

7.1.2 Controlling the Flow of HPR Data across the CNN

The HPRQLIM keyword on the BF-TG link station (PU) statement lets you control the amount of APPN HPR data queued for transmission on link stations in the subarea network. If an NLP arriving at a link station would cause the transmission queue to exceed its limit, the frame is discarded. HPRQLIM can also be coded on the BUILD statement to make it apply to all BF-TGs. The default is zero, meaning no limit.

The corresponding thresholds for subarea TGs are defined in the PATH statements. The first ERn keyword on the PATH statement specifies the subarea flow control thresholds for the three subarea transmission priorities as well as a total threshold for all priorities combined. The sixth and last operand, the total threshold, is also used by NCP to limit the the amount of HPR data that can be queued on the TG at any time. When the threshold is exceeded, additional data is discarded.

While RTP will recover lost data packets, you should try to ensure that packet loss is a rare occurrence. Lost data will cause drastic cuts in the ARB flow control rates, thus affecting throughput.

7.1.3 Link-Level Error Recovery

HPR, which is designed for high-speed highly reliable networks, does not require each individual link to perform error recovery; the RTP endpoints can do this. Therefore, on many links the user has the choice of defining whether or not link-level error recovery is used. Link-level error recovery is advisable if the error rate is high, but on reliable connections it is better to turn it off and let RTP handle the rare occurrences.

NCP supports the use of link-level error recovery on all its HPR-capable connections, but allows it to be turned off only on the following:

- Frame relay (subarea or peripheral)
- Token-ring (peripheral, and subarea only on TIC-3)
- ISDN (subarea or peripheral)

Link-level error recovery is controlled by the LLERP keyword. LLERP may be coded on the physical port (PU) definition for the above connections, and overridden individually on the logical link station (PU) definitions. You should code LLERP=NOTPREF unless the connection is unreliable. NOTPREF means bypass link-level error recovery unless the adjacent node demands it. LLERP=REQUIRED is the only other valid value. This means use link-level error recovery unless the adjacent node forbids it, in which case do not use HPR at all.

7.1.4 Session Control Block Requirements

As you move LU-LU sessions to the APPN/HPR network, you may find that you can lower the values you generate for the following keywords, which define session control block requirements:

- ADDSESS (BUILD)
- AUXADDR (BUILD)
- SESSACC (BUILD)
- NUMILU (LUDRPOOL)

This is because, as you move from LEN or APPN to HPR, session awareness moves from the NCP boundary function to the RTP endpoint (which is never NCP).

You can monitor your session control block usage with NTuneMON.

7.2 HPR across Channel Links

To test HPR between CNNs, we added two NCPs to our environment as shown in Figure 41 on page 77. RAA owns NCP RA6NCR0 while RAK owns NCP RA7NCPB.

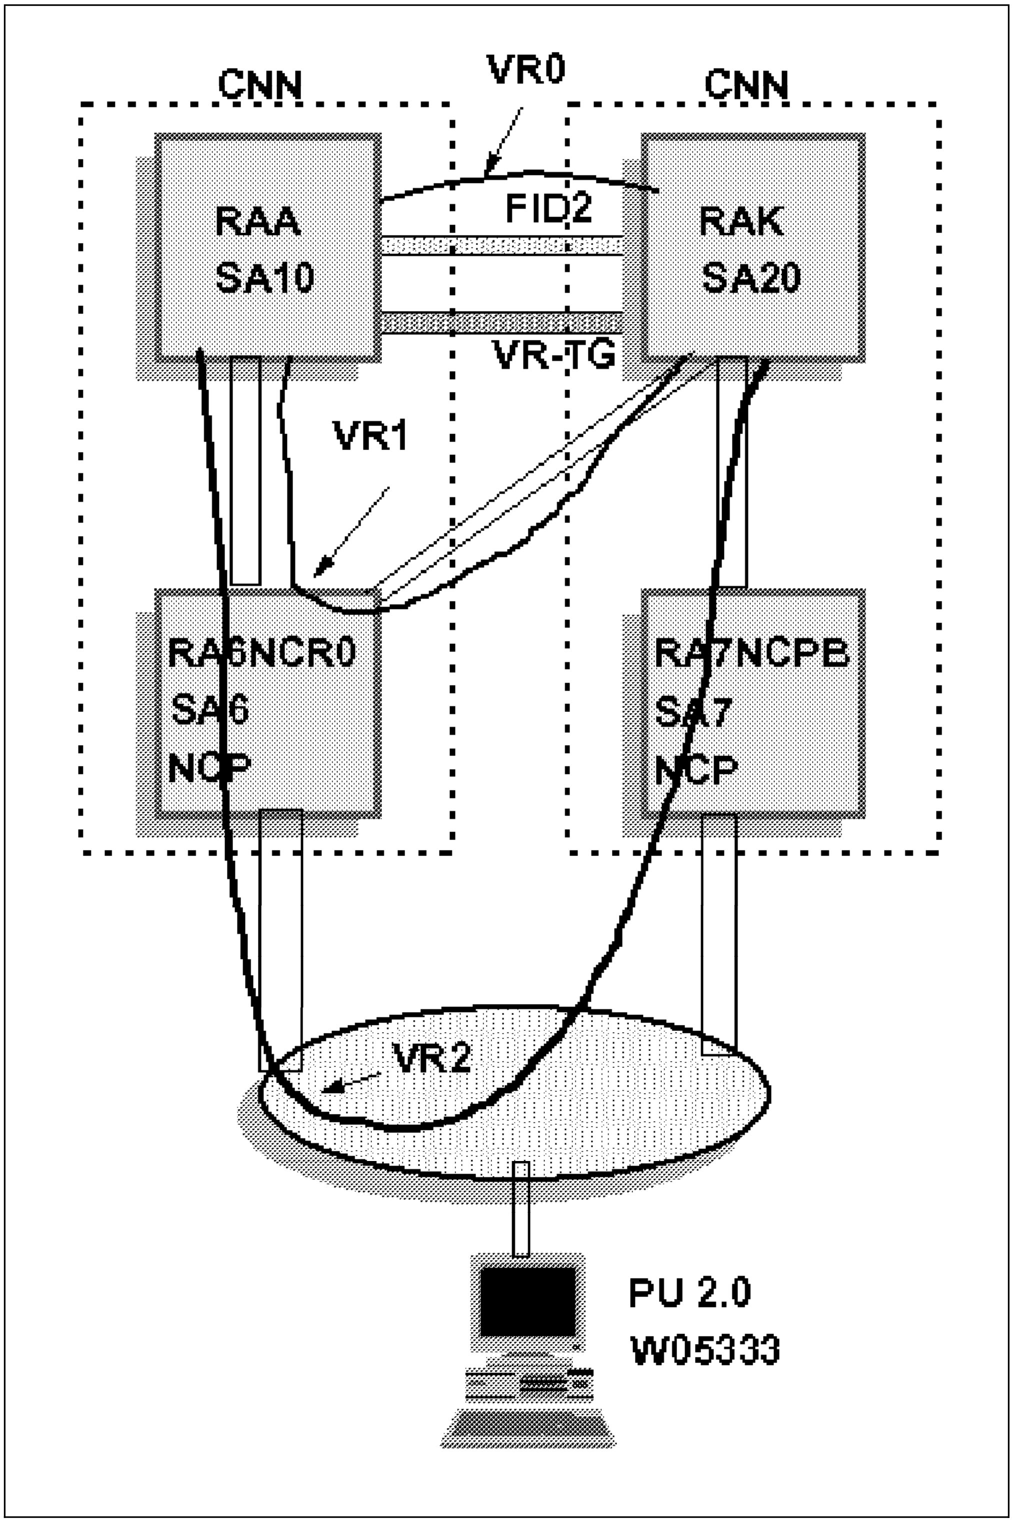

Figure 41. Network Configuration

The two CNNs have three subarea connections, on which a VR-TG is defined: a CTC, a channel link from RAK to RA6NCR0, and a token-ring connection between the NCPs. These connections form part of VR0, VR1 and VR2 respectively between the host subareas. We also have an APPN connection over MPC between the two VTAMs; thus there are four physical paths between them. In APPN terms, as shown in Figure 42 on page 78 (the topology as seen from

RAA), we have two one-hop connections between RAA and RAK. There is a BF-TG over the MPC link **1** and a VR-TG **2**.

```
D NET,TOPO,ID=RAA,LIST=ALL
IST097I  DISPLAY  ACCEPTED
IST350I  DISPLAY TYPE = TOPOLOGY
IST1295I  CP NAME           NODETYPE ROUTERES CONGESTION  CP-CP WEIGHT
IST1296I  USIBMRA.RAA       NN       128      NONE        *NA*  *NA*
IST1579I                    ------------------------------------------
IST1297I                    ICN/MDH  CDSERVR  RSN         HPR
IST1298I                    YES      YES      115782      RTP
IST1579I                    ------------------------------------------
IST1223I                    BN       NATIVE   TIME LEFT
IST1224I                    NO       YES      15
IST1299I  TRANSMISSION GROUPS ORIGINATING AT CP USIBMRA.RAA
IST1357I                                            CPCP
IST1300I  DESTINATION CP  TGN  STATUS   TGTYPE      VALUE WEIGHT
IST1301I  USIBMRA.RAK     255  OPER     INTERM VRTG YES   *NA*  2
IST1301I  USIBMRA.RAK     21   OPER     INTERM      YES   *NA*  1
```

Figure 42. APPN View of the Configuration

We established a session between an LU local to RAK and TSO running on RAA. Two RTP pipes were immediately created as demonstrated in Figure 43 on page 79.

```
IST1488I  ACTIVATION   FOR RTP CNR00006 AS ACTIVE  PARTNER COMPLETED
IST1488I  ACTIVATION   FOR RTP CNR00007 AS ACTIVE  PARTNER COMPLETED
*
D NET,ID=ISTRTPMN,E
IST097I  DISPLAY  ACCEPTED
IST075I  NAME = ISTRTPMN         , TYPE = RTP MAJOR NODE
IST486I  STATUS= ACTIV     , DESIRED STATE= ACTIV
IST1486I  RTP NAME   STATE         DESTINATION CP   MNPS  TYPE
IST1487I  CNR00007   CONNECTED     USIBMRA.RAK      NO    LULU  4
IST1487I  CNR00006   CONNECTED     USIBMRA.RAK      NO    RSTP  4
IST314I  END
*
D NET,ID=CNR00007,E
IST097I  DISPLAY  ACCEPTED
IST075I  NAME = CNR00007         , TYPE = PU_T2.1
IST1392I  DISCNTIM = 00010 DEFINED AT PU   FOR DISCONNECT
IST486I  STATUS= ACTIV--LX-, DESIRED STATE= ACTIV
IST1043I  CP NAME = RAK     , CP NETID = USIBMRA , DYNAMIC LU = YES
IST1589I  XNETALS = YES
IST933I  LOGMODE=***NA***, COS=#CONNECT
IST1476I  TCID X'2E277490000000A5' - REMOTE TCID X'1EC14F41000001F5'
IST1481I  DESTINATION CP USIBMRA.RAK       - NCE X'D000000000000000'
IST1587I  ORIGIN NCE X'D000000000000000'
IST1477I  ALLOWED DATA FLOW RATE =   3200 KBITS/SEC
IST1516I  INITIAL DATA FLOW RATE =   3200 KBITS/SEC
IST1511I  MAXIMUM NETWORK LAYER PACKET SIZE =  20476 BYTES
IST1478I  NUMBER OF UNACKNOWLEDGED BUFFERS =     0
IST1479I  RTP CONNECTION STATE = CONNECTED             - MNPS = NO
IST1480I  RTP END TO END ROUTE - PHYSICAL PATH
IST1460I  TGN  CPNAME             TG TYPE      HPR
IST1461I   21  USIBMRA.RAK        APPN         RTP  3
IST231I  RTP      MAJOR NODE = ISTRTPMN
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I  STATE TRACE = OFF
IST355I  LOGICAL UNITS:
IST080I  RAKTX055 ACT/S----Y
IST314I  END
```

Figure 43. Network Log and Displays Issued on RAA

The RTP connections CNR00007 and CNR00006 **4** are the session pipe and the long-lived Route Setup pipe respectively.

CNR00007 is carrying our 3270 session and is using the APPN BF-TG **3**. As the subarea CTC link was activated before the local SNA major node for RAK, the CP-CP sessions were started over the VR-TG and therefore there is no CP-CP pipe in the ISTRTPMN major node **4**.

Next, we brought down the MPC connection between RAA and RAK, as seen in Figure 44 on page 80.

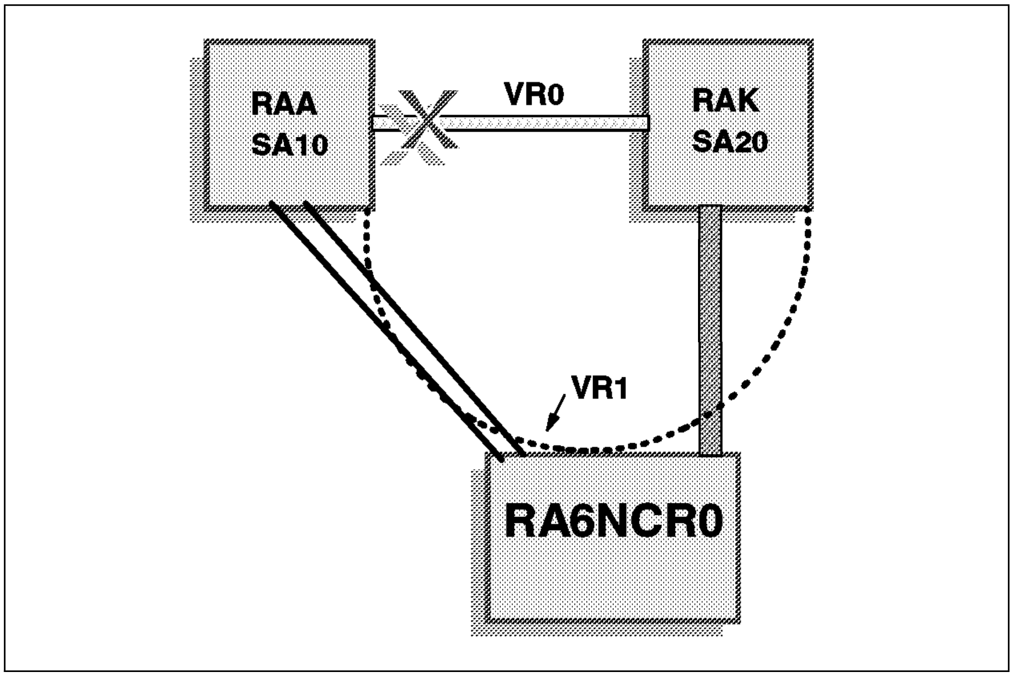

Figure 44. New Path

The related RAA network log is in Figure 45.

```
V NET,INACT,ID=RAAAHHK
IST097I  VARY     ACCEPTED
IST1196I  APPN CONNECTION FOR USIBMRA.RAK  INACTIVE - TGN =  21
IST1494I  PATH SWITCH STARTED   FOR RTP CNR00007
IST1133I  RAAIRAK           IS NOW INACTIVE, TYPE = PU_T2
IST1488I  INACTIVATION FOR RTP CNR00006 AS PASSIVE PARTNER COMPLETED
IST1416I  ID = CNR00006         FAILED - RECOVERY IN PROGRESS
IST1133I  RAAAHHK       IS NOW INACTIVE, TYPE = LCL SNA MAJ NODE
IST1136I  VARY INACT CNR00006    SCHEDULED - UNRECOVERABLE ERROR
IST1133I  CNR00006          IS NOW INACTIVE, TYPE = PU_T2.1
IST871I  RESOURCE CNR00006 DELETED
IST1494I  PATH SWITCH COMPLETED FOR RTP CNR00007
IST1480I  RTP END TO END ROUTE - PHYSICAL PATH
IST1460I  TGN  CPNAME             TG TYPE      HPR
IST1461I  255  USIBMRA.RAK        VRTG         RTP  5
IST314I  END
```

Figure 45. Network Log on RAA

As we expected, the long-lived pipe CNR00006 was deactivated while the session pipe was switched to the VR-TG connection **5**. The CP-CP sessions were unaffected by the failure so there is no message related to them.

Now we verified the active virtual routes between RAA and RAK, as seen in Figure 46 on page 81 on the RAA side.

```
D NET,ROUTE,DESTSUB=20
IST097I  DISPLAY  ACCEPTED
IST535I  ROUTE DISPLAY   7 FROM SA         10 TO SA         20
IST808I  ORIGIN PU = ISTPUSA0 DEST PU = ***NA*** NETID = USIBMRA
IST536I  VR  TP    STATUS   ER       ADJSUB  TGN  STATUS  CUR MIN MAX
IST537I   0   0    ACTIV     0           20    1  ACTIV3   11  10  30
IST537I   0   1    INACT     0           20    1  ACTIV3
IST537I   0   2    ACTIV     0           20    1  ACTIV3   11  10  30
IST537I   1   0    INACT     1            6    1  INACT
IST537I   1   1    INACT     1            6    1  INACT   11
IST537I   1   2    INACT     1            6    1  INACT
```

Figure 46. Display Routes on RAA

Here we come to an important consideration. HPR over VR-TG only uses a virtual route for the Route Setup process. Once the RTP connection has been established, the VR is no longer needed and it may be deactivated if no other sessions are using it. HPR only uses the explicit route (ER) to forward the NLP packets over a FID-4 link.

In our case we have VR0 active because the SSCP and CP-CP sessions are flowing over it. VR1 is inactive, but this does not prove that the HPR connection is also flowing across VR0. In fact, the only alternative path at the time is VR1, whose ER has not yet been activated **11**; VR2 via the token-ring has not yet been defined. The fact that ER1 was never activated indicates that ER0 is, indeed, being used for the HPR connection. However, in the current implementation, the only way to be sure of the actual link being used by HPR over a VR-TG is to trace the physical link and verify the presence of NLPs.

Next, we tried another path switch by breaking the channel-to-channel connection (on which ER0 is mapped) to verify that the CNR00007 pipe stayed up and used ER1 to reach RAK (see Figure 47).

```
V NET,INACT,ID=RAACTCA
IST097I  VARY     ACCEPTED
IST526I  ROUTE FAILED FROM         10 TO         20 - DSA
IST526I  ROUTE FAILED FROM         10 TO         20 - DSA
IST1196I APPN CONNECTION FOR USIBMRA.RAK INACTIVE - TGN = 255  6
IST819I  CDRM RAK   COMMUNICATION LOST - RECOVERY IN PROGRESS
IST1110I  ACTIVATION OF CP-CP SESSION WITH USIBMRA.RAK FAILED
IST1280I  SESSION TYPE = CONWINNER - SENSE = 80200007
IST1002I  RCPRI=0004 RCSEC=0000
IST314I  END
IST1133I  RAACTCA       IS NOW INACTIVE, TYPE = CA MAJOR NODE
IST1086I APPN CONNECTION FOR USIBMRA.RAK  IS ACTIVE - TGN = 255  7
IST1132I  USIBMRA.RAK       IS ACTIVE, TYPE = CDRM
IST1096I  CP-CP SESSIONS WITH USIBMRA.RAK    ACTIVATED
IST1494I  PATH SWITCH STARTED   FOR RTP CNR00007
IST1494I  PATH SWITCH COMPLETED FOR RTP CNR00007
IST1480I  RTP END TO END ROUTE - PHYSICAL PATH
IST1460I  TGN  CPNAME            TG TYPE      HPR
IST1461I  255  USIBMRA.RAK        VRTG         RTP  8
IST314I  END
```

Figure 47. Network Log on RAA

You can see that the TG 255 between RAA and RAK is deactivated **6** but it is then opened again **7**. Of course, it now traverses a different subarea VR. The

CP-CP sessions have been recycled; they were on the CTC connection with the SSCP-SSCP session. This is correct as we did not have a CPCP RTP pipe. The SSCP-SSCP session has been similarly terminated and restored on the alternate path.

You can see that CNR00007 has been switched correctly, as expected, to the same VR-TG **8**. You are not told the new subarea route after the path switch message; TG 255 is always shown.

We then verified what happened to the subarea routes from RAA, as shown in Figure 48.

```
D NET,ROUTE,DESTSUB=20
IST097I  DISPLAY  ACCEPTED
IST535I  ROUTE DISPLAY   8 FROM SA         10 TO SA
IST808I  ORIGIN PU = ISTPUSA0 DEST PU = ***NA*** NETID = USIBMRA
IST536I  VR  TP    STATUS   ER       ADJSUB  TGN  STATUS
IST537I   0   0    INACT     0           20    1  INOP
IST537I   0   1    INACT     0           20    1  INOP
IST537I   0   2    INACT     0           20    1  INOP
IST537I   1   0    ACTIV     1            6    1  ACTIV3
IST537I   1   1    INACT     1            6    1  ACTIV3
IST537I   1   2    ACTIV     1            6    1  ACTIV3
```

Figure 48. Active Subarea Routes on RAA

You can see that the only active explicit route now is ER1, which goes through NCP RA6NCR0. All sessions are flowing over VR1, while the RTP pipe CNR00007 must also flow across ER1 as ER0 is inoperative.

7.3 HPR across Token-Ring

This test demonstrates the ability of HPR to switch the path to the token-ring subarea connection (please refer to Figure 41 on page 77).

We started by disabling the NCP connection between RA6NCR0 and RAK, then established a session, and an HPR connection, betwen RAK and RAA through the VR-TG as shown in Figure 49 on page 83.

```
D NET,ID=CNR00008,E
IST097I  DISPLAY  ACCEPTED
IST075I  NAME = CNR00008         , TYPE = PU_T2.1
IST1392I  DISCNTIM = 00010 DEFINED AT PU   FOR DISCONNECT
IST486I  STATUS= ACTIV--LX-, DESIRED STATE= ACTIV
IST1043I  CP NAME = RAK     , CP NETID = USIBMRA , DYNAMIC LU = YES
IST1589I  XNETALS = YES
IST933I  LOGMODE=***NA***, COS=#CONNECT
IST1476I  TCID X'2E2774910000008D' - REMOTE TCID X'1EC14F42000001F2'
IST1481I  DESTINATION CP USIBMRA.RAK       - NCE X'D000000000000000'
IST1587I  ORIGIN NCE X'D000000000000000'
IST1477I  ALLOWED DATA FLOW RATE =     51 KBITS/SEC
IST1516I  INITIAL DATA FLOW RATE =   6400  BITS/SEC
IST1511I  MAXIMUM NETWORK LAYER PACKET SIZE =   4046 BYTES
IST1478I  NUMBER OF UNACKNOWLEDGED BUFFERS =     1
IST1479I  RTP CONNECTION STATE = CONNECTED            - MNPS = NO
IST1480I  RTP END TO END ROUTE - PHYSICAL PATH
IST1460I  TGN  CPNAME             TG TYPE      HPR
IST1461I  255  USIBMRA.RAK        VRTG         RTP
IST231I  RTP      MAJOR NODE = ISTRTPMN
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I  STATE TRACE = OFF
IST355I  LOGICAL UNITS:
IST080I  RAKTX082 ACT/S----Y RAKTX089 ACT/S----Y
IST314I  END
```

Figure 49. Display of RTP Connection

To ensure that the session was established on the route we anticipated, a display of the route was done as shown in Figure 50.

```
IST535I  ROUTE DISPLAY  11 FROM SA         10 TO SA         20
IST808I  ORIGIN PU = ISTPUSA0 DEST PU = ***NA*** NETID = USIBMRA
IST536I  VR  TP    STATUS   ER       ADJSUB  TGN  STATUS  CUR MIN MAX
IST537I   0   0    ACTIV     0           20    1  ACTIV3   11  10  30
IST537I   0   1    INACT     0           20    1  ACTIV3
IST537I   0   2    ACTIV     0           20    1  ACTIV3   11  10  30
IST537I   1   0    INACT     1            6    1  INOP
IST537I   1   1    INACT     1            6    1  INOP
IST537I   1   2    INACT     1            6    1  INOP
IST537I   2   0    INACT     2            6    1  INACT
IST537I   2   1    INACT     2            6    1  INACT
IST537I   2   2    INACT     2            6    1  INACT
```

Figure 50. Routes Showing ER0 Active

The above display verifies that we had only two routes available for the VR-TG. ER0 was the CTC connection between RAA and RAK, and ER2 was through the token-ring connection. Since ER0 was the only active ER, the RTP pipe must have been flowing through the CTC.

We then continued by inactivating the CTC connection. As we expected a path switch occurred, as shown in Figure 51 on page 84.

```
IST1494I  PATH SWITCH STARTED   FOR RTP CNR00008
IST1494I  PATH SWITCH COMPLETED FOR RTP CNR00008
IST1480I  RTP END TO END ROUTE - PHYSICAL PATH
IST1460I  TGN  CPNAME            TG TYPE      HPR
IST1461I  255  USIBMRA.RAK       VRTG         RTP
IST314I  END
```

Figure 51. Path Switch Messages

The above display shows that the connection is still going through TG255, so we did another display of the route (see Figure 52) to check where the RTP pipe now was. Since ER2 was the only active ER, CNR00008 must have taken that path.

```
D NET,ROUTE,DESTSUB=20
IST097I  DISPLAY  ACCEPTED
IST535I  ROUTE DISPLAY  12 FROM SA         10 TO SA         20
IST808I  ORIGIN PU = ISTPUSA0 DEST PU = ***NA*** NETID = USIBMRA
IST536I  VR  TP    STATUS   ER      ADJSUB  TGN  STATUS  CUR MIN MAX
IST537I   0   0    INACT     0          20    1  INOP
IST537I   0   1    INACT     0          20    1  INOP
IST537I   0   2    INACT     0          20    1  INOP
IST537I   1   0    INACT     1           6    1  INOP
IST537I   1   1    INACT     1           6    1  INOP
IST537I   1   2    INACT     1           6    1  INOP
IST537I   2   0    INACT     2           6    1  ACTIV3
IST537I   2   1    INACT     2           6    1  ACTIV3
IST537I   2   2    ACTIV     2           6    1  ACTIV3   11  10  30
```

Figure 52. Route Showing ER2 Active

7.4 Using HPRNCPBF

In this section we demonstrate the HPRNCPBF VTAM start option that was introduced in VTAM V4R4 with APAR OW25950. This allows a customer to decide whether HPR is to be used even if it will cause session data to travel through an NCP twice. To illustrate why this might be desirable, consider the network shown in Figure 41 on page 77.

Suppose the PC labelled W05333 is owned by RAA, and a dependent LU on it logs on to an application on RAK. Such a session will use the path RA6NCR0 - RAK if the network is configured to choose the shortest route. This session will not be able to use HPR, because the NCP cannot be an RTP endpoint.

If you want such a session to use an RTP connection once it leaves RAA′s domain, the session path must traverse RAA itself because that is the only RTP-capable subarea node in the domain. Now if the connection from RA6NCR0 to RAK is the only one available, the session will have to go back on itself through RA6NCR0 to reach RAK by means of an RTP pipe. Thus the path will be (W05333 - RA6NCR0 - RAA) as the subarea portion and (RAA - RA6NCR0 - RAK) as the HPR portion.

If HPRNCPBF=NO, RAA will not permit such a route to be used. If HPRNCPBF=YES, it will be permitted. The purpose of the HPRNCPBF option is to allow you to trade performance (fewest hops on a path) against availability (maximum HPR portion of the route).

HPRNCPBF defaults to NO, and is modifiable. If the value is changed from YES to NO, new sessions can use existing RTP pipes, but no new RTP pipes will be set up that cause the traffic to go through an NCP twice. Sessions will still be established but will use base APPN until they happen to reach an RTP-capable node.

Referring again to Figure 41 on page 77, we deactivated both the MPC connection and the CTC link between RAA and RAK, as well as the token-ring subarea connection that we only used in the previous test. Thus the only operative path between RAA and RAK was the VR-TG mapped on to ER 1. Figure 53 showed that the only active route was through RA6NCR0.

```
D NET,ROUTE,DESTSUB=20
IST097I  DISPLAY  ACCEPTED
IST535I  ROUTE DISPLAY  22 FROM SA         10 TO SA
IST808I  ORIGIN PU = ISTPUSA0 DEST PU = ***NA*** NETID = USIBMRA
IST536I  VR  TP    STATUS   ER       ADJSUB  TGN  STATUS
IST537I   0   0    INACT     0           20    1  INOP
IST537I   0   1    INACT     0           20    1  INOP
IST537I   0   2    INACT     0           20    1  INOP
IST537I   1   0    ACTIV     1            6    1  ACTIV3
IST537I   1   1    INACT     1            6    1  ACTIV3
IST537I   1   2    ACTIV     1            6    1  ACTIV3
```

Figure 53. Subarea Routes between RAA and RAK

From an APPN point of view, as seen in Figure 54, there was only one active TG.

```
D NET,TOPO,ID=RAA,LIST=ALL
IST097I  DISPLAY  ACCEPTED
IST350I  DISPLAY TYPE = TOPOLOGY
IST1295I  CP NAME           NODETYPE ROUTERES CONGESTION  CP-CP WEIGHT
IST1296I  USIBMRA.RAA       NN       128      NONE        *NA*  *NA*
IST1579I                    ------------------------------------------
IST1297I                    ICN/MDH  CDSERVR  RSN         HPR
IST1298I                    YES      YES      115782      RTP
IST1579I                    ------------------------------------------
IST1223I                    BN       NATIVE   TIME LEFT
IST1224I                    NO       YES      14
IST1299I  TRANSMISSION GROUPS ORIGINATING AT CP USIBMRA.RAA
IST1357I                                                  CPCP
IST1300I  DESTINATION CP    TGN      STATUS   TGTYPE      VALUE WEIGHT
IST1301I  USIBMRA.RAK       255      OPER     INTERM VRTG YES   *NA*
IST1301I  USIBMRA.RAK       21       INOP     INTERM      YES   *NA*
```

Figure 54. Display of Topology Database

Next we displayed the ISTRTPMN major node on RAA to verify the current number of active RTP connections to RAK and which sessions they carried (see Figure 55 on page 86).

```
DISPLAY NET,ID=ISTRTPMN,SCOPE=ALL
IST097I  DISPLAY  ACCEPTED
IST075I  NAME = ISTRTPMN         , TYPE = RTP MAJOR NODE
IST486I  STATUS= ACTIV     , DESIRED STATE= ACTIV
IST1486I    RTP NAME   STATE             DESTINATION CP     MNPS  TYPE
IST1487I    CNR0000E   CONNECTED         USIBMRA.RAK        NO    LULU
IST314I  END
*
DISPLAY NET,ID=CNR0000E,SCOPE=ALL
IST097I  DISPLAY  ACCEPTED
IST075I  NAME = CNR0000E         , TYPE = PU_T2.1
IST1392I  DISCNTIM = 00010 DEFINED AT PU   FOR DISCONNECT
IST486I  STATUS= ACTIV--LX-, DESIRED STATE= ACTIV
IST1043I  CP NAME = RAK     , CP NETID = USIBMRA , DYNAMIC LU = YES
IST1589I  XNETALS = YES
IST933I  LOGMODE=***NA***, COS=#CONNECT
IST1476I  TCID X'2E277497000000BD' - REMOTE TCID X'1EC14F49000001F6'
IST1481I  DESTINATION CP USIBMRA.RAK       - NCE X'D000000000000000'
IST1587I  ORIGIN NCE X'D000000000000000'
IST1477I  ALLOWED DATA FLOW RATE =     10 KBITS/SEC
IST1516I  INITIAL DATA FLOW RATE =   6400  BITS/SEC
IST1511I  MAXIMUM NETWORK LAYER PACKET SIZE =   4276 BYTES
IST1478I  NUMBER OF UNACKNOWLEDGED BUFFERS =     0
IST1479I  RTP CONNECTION STATE = CONNECTED            - MNPS = NO
IST1480I  RTP END TO END ROUTE - PHYSICAL PATH
IST1460I  TGN  CPNAME            TG TYPE      HPR
IST1461I  255  USIBMRA.RAK        VRTG         RTP
IST231I  RTP      MAJOR NODE = ISTRTPMN
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I  STATE TRACE = OFF
IST355I  LOGICAL UNITS:
IST080I  RAKTX025 ACT/S----Y RAKTX034 ACT/S----Y RAKTX047 ACT/S----Y
IST314I  END
```

Figure 55. Displays Issued on RAA

We only had one active pipe, and it was carrying three sessions between RAA and RAK. These sessions had been established before we broke the MPC and the CTC. They were still there because HPR always managed to find an alternative path for them.

Then we activated the workstation W05333, with its dependent LUs, connected to NCP subarea 6. From the USS10 message displayed by RAA we logged on to the TSO subsystem on RAK.

The session path for the TSO session, as expected, went via RA6NCR0 directly to RAK. The display in Figure 56 on page 87 confirms that no new RTP connections have been set up, so our new session uses base APPN.

```
13:49:16 IST590I  CONNECTIN  ESTABLISHED FOR PU W05333   ON LINE J0006011
*
DISPLAY NET,ID=ISTRTPMN,SCOPE=ALL
IST097I  DISPLAY  ACCEPTED
IST075I  NAME = ISTRTPMN         , TYPE = RTP MAJOR NODE
IST486I  STATUS= ACTIV     , DESIRED STATE= ACTIV
IST1486I    RTP NAME   STATE             DESTINATION CP     MNPS  TYPE
IST1487I    CNR0000E   CONNECTED         USIBMRA.RAK        NO    LULU
IST314I  END
*
DISPLAY NET,ID=CNR0000E,SCOPE=ALL
IST097I  DISPLAY  ACCEPTED
IST075I  NAME = CNR0000E         , TYPE = PU_T2.1
IST1392I  DISCNTIM = 00010 DEFINED AT PU   FOR DISCONNECT
IST486I  STATUS= ACTIV--LX-, DESIRED STATE= ACTIV
IST1043I  CP NAME = RAK     , CP NETID = USIBMRA , DYNAMIC LU = YES
IST1589I  XNETALS = YES
IST933I  LOGMODE=***NA***, COS=#CONNECT
IST1476I  TCID X'2E277497000000BD' - REMOTE TCID X'1EC14F49000001F6'
IST1481I  DESTINATION CP USIBMRA.RAK       - NCE X'D000000000000000'
IST1587I  ORIGIN NCE X'D000000000000000'
IST1477I  ALLOWED DATA FLOW RATE =     10 KBITS/SEC
IST1516I  INITIAL DATA FLOW RATE =   6400  BITS/SEC
IST1511I  MAXIMUM NETWORK LAYER PACKET SIZE =   4276 BYTES
IST1478I  NUMBER OF UNACKNOWLEDGED BUFFERS =     0
IST1479I  RTP CONNECTION STATE = CONNECTED            - MNPS = NO
IST1480I  RTP END TO END ROUTE - PHYSICAL PATH
IST1460I  TGN  CPNAME             TG TYPE      HPR
IST1461I  255  USIBMRA.RAK        VRTG         RTP
IST231I  RTP      MAJOR NODE = ISTRTPMN
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I  STATE TRACE = OFF
IST355I  LOGICAL UNITS:
IST080I  RAKTX025 ACT/S----Y RAKTX034 ACT/S----Y RAKTX047 ACT/S----Y
IST314I  END
```

Figure 56. Display of CNR0000E and ISTRTPMN on RAA after the PU Connection

To see where the session went, we displayed the workstation dependent LU from RAA, as shown in Figure 57 on page 88.

```
DISPLAY NET,ID=W0533302,SCOPE=ALL
IST097I  DISPLAY  ACCEPTED
IST075I  NAME = USIBMRA.W0533302 , TYPE = LOGICAL UNIT
IST486I  STATUS= ACT/S---X-, DESIRED STATE= ACTIV
IST1447I  REGISTRATION TYPE = NETSRVR
IST977I  MDLTAB=***NA*** ASLTAB=***NA***
IST861I  MODETAB=ISTINCLM USSTAB=US327X   LOGTAB=***NA***
IST934I  DLOGMOD=D4C32XX3 USS LANGTAB=***NA***
IST597I  CAPABILITY-PLU INHIBITED,SLU ENABLED  ,SESSION LIMIT 00000001
IST136I  SWITCHED SNA MAJOR NODE = ISTDSWMN
IST081I  LINE NAME = J0006011, LINE GROUP = EG06L01 , MAJNOD = RA6NCR0
IST135I  PHYSICAL UNIT = W05333
IST1131I  DEVICE = LU
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I  STATE TRACE = OFF
IST228I  ENCRYPTION = NONE
IST1563I  CKEYNAME = LUMOD05D CKEY = PRIMARY   CERTIFY = NO
IST1552I  MAC = NONE        MACTYPE = NONE
IST171I  ACTIVE SESSIONS = 0000000001, SESSION REQUESTS = 0000000000
IST206I  SESSIONS:
IST634I  NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I  RAKAT06  ACTIV-P    F8D3D164311BD942           0  0 USIBMRA
IST314I  END
```

Figure 57. Display of Workstation LU on RAA

In fact this display does not prove whether the session used base APPN or HPR; all it shows is that W0533302 was an LU owned by RAA and connected to RA6NCR0, and that it was in session with RAK's TSO. What *did* prove that the session was base APPN was a similar display from RAK. This showed that the link station being used by the CDRSC named W0533302 was the FID-2 connection to RA6NCR0 and not a PU of the form CNRxxxxx.

Now we enabled the HPRNCPBF VTAM start option as you can see in Figure 58.

```
F NETA0,VTAMOPTS,HPRNCPBF=YES
IST097I  MODIFY   ACCEPTED
IST223I  MODIFY          COMMAND COMPLETED
D NET,VTAMOPTS,OPT=HPRNCPBF
IST097I  DISPLAY  ACCEPTED
IST1188I  ACF/VTAM V4R4   STARTED AT 12:46:02 ON 01/28/98
IST1349I  COMPONENT ID IS 5695-11701-401
IST1348I  VTAM STARTED AS INTERCHANGE NODE
IST1189I  HPRNCPBF = YES
IST314I  END
```

Figure 58. Modify Start Option Command Issued on RAA

From another emulator session on the same workstation, we logged on to TSO on RAK again. This time a display of ISTRTPMN showed that a new RTP pipe had been established, as in Figure 59 on page 89.

```
14:13:10 IST1488I  ACTIVATION   FOR RTP CNR00045 AS PASSIVE PARTNER COMPLETED
*
DISPLAY NET,ID=ISTRTPMN,SCOPE=ALL
IST097I  DISPLAY  ACCEPTED
IST075I  NAME = ISTRTPMN         , TYPE = RTP MAJOR NODE
IST486I  STATUS= ACTIV     , DESIRED STATE= ACTIV
IST1486I    RTP NAME   STATE             DESTINATION CP     MNPS  TYPE
IST1487I    CNR00045   CONNECTED         USIBMRA.RAK        NO    LULU
IST1487I    CNR0000E   CONNECTED         USIBMRA.RAK        NO    LULU
IST314I  END
*
DISPLAY NET,ID=CNR00045,SCOPE=ALL
IST097I  DISPLAY  ACCEPTED
IST075I  NAME = CNR00045         , TYPE = PU_T2.1
IST1392I  DISCNTIM = 00010 DEFINED AT PU   FOR DISCONNECT
IST486I  STATUS= ACTIV--LX-, DESIRED STATE= ACTIV
IST1043I  CP NAME = RAK     , CP NETID = USIBMRA , DYNAMIC LU = YES
IST1589I  XNETALS = YES
IST933I  LOGMODE=***NA***, COS=#CONNECT
IST1476I  TCID X'2E27749C0000008D' - REMOTE TCID X'1EC14F53000001EB'
IST1481I  DESTINATION CP USIBMRA.RAK       - NCE X'D000000000000000'
IST1587I  ORIGIN NCE X'D100000000000000'
IST1477I  ALLOWED DATA FLOW RATE =     14 KBITS/SEC
IST1516I  INITIAL DATA FLOW RATE =   6400  BITS/SEC
IST1511I  MAXIMUM NETWORK LAYER PACKET SIZE =   4276 BYTES
IST1478I  NUMBER OF UNACKNOWLEDGED BUFFERS =     0
IST1479I  RTP CONNECTION STATE = CONNECTED            - MNPS = NO
IST1480I  RTP END TO END ROUTE - PHYSICAL PATH
IST1460I  TGN  CPNAME             TG TYPE      HPR
IST1461I  255  USIBMRA.RAK        VRTG         RTP  10
IST231I  RTP      MAJOR NODE = ISTRTPMN
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I  STATE TRACE = OFF
IST355I  LOGICAL UNITS:
IST080I  RAKAT06  ACT/S----Y   9
IST314I  END
```

Figure 59. Displays Issued on RAA after the New Logon

We can see that the new connection CNR00045 carries the session to TSO **9** and uses the VR-TG **10**. The session path is shown in Figure 60 on page 90, and comprises (W05333 - RA6NCR0 - RAA - RA6NCR0 - RAK). The RTP connection forms only part of this route.

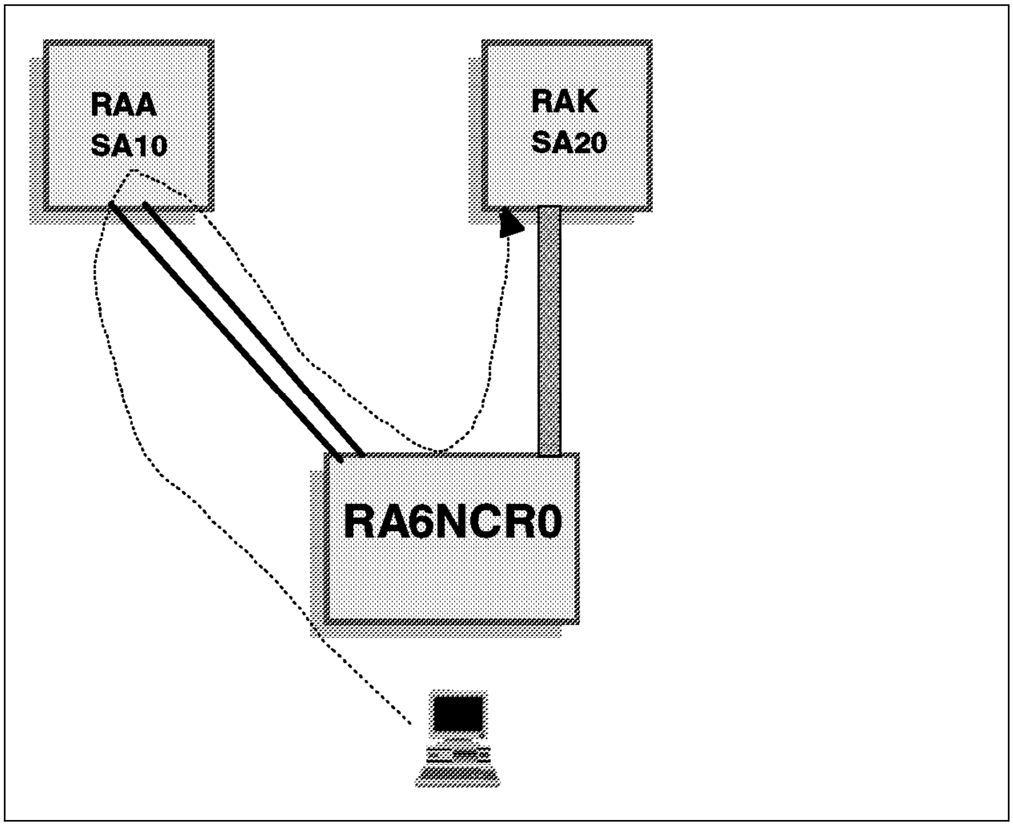

Figure 60. Session Establishment Path

7.5 HPR on Communications Server/2

The purpose of this final test was to extend HPR support all the way to the workstation. In the first part of this chapter HPR has been restricted to the connection between the two VTAM CNNs. HPR was available for cross-domain sessions only, and even then for just that portion of the session between the two VTAMs.

Allowing the CS/2 workstation to be an RTP endpoint allows any session to use HPR all the way. However, if dependent LU sessions are to traverse an HPR path all the way they must first be capable of APPN. This means using DLUR, because only DLUR allows that vital first hop (CS/2 to NCP) to be APPN and therefore HPR. Figure 61 on page 91 illustrates this setup.

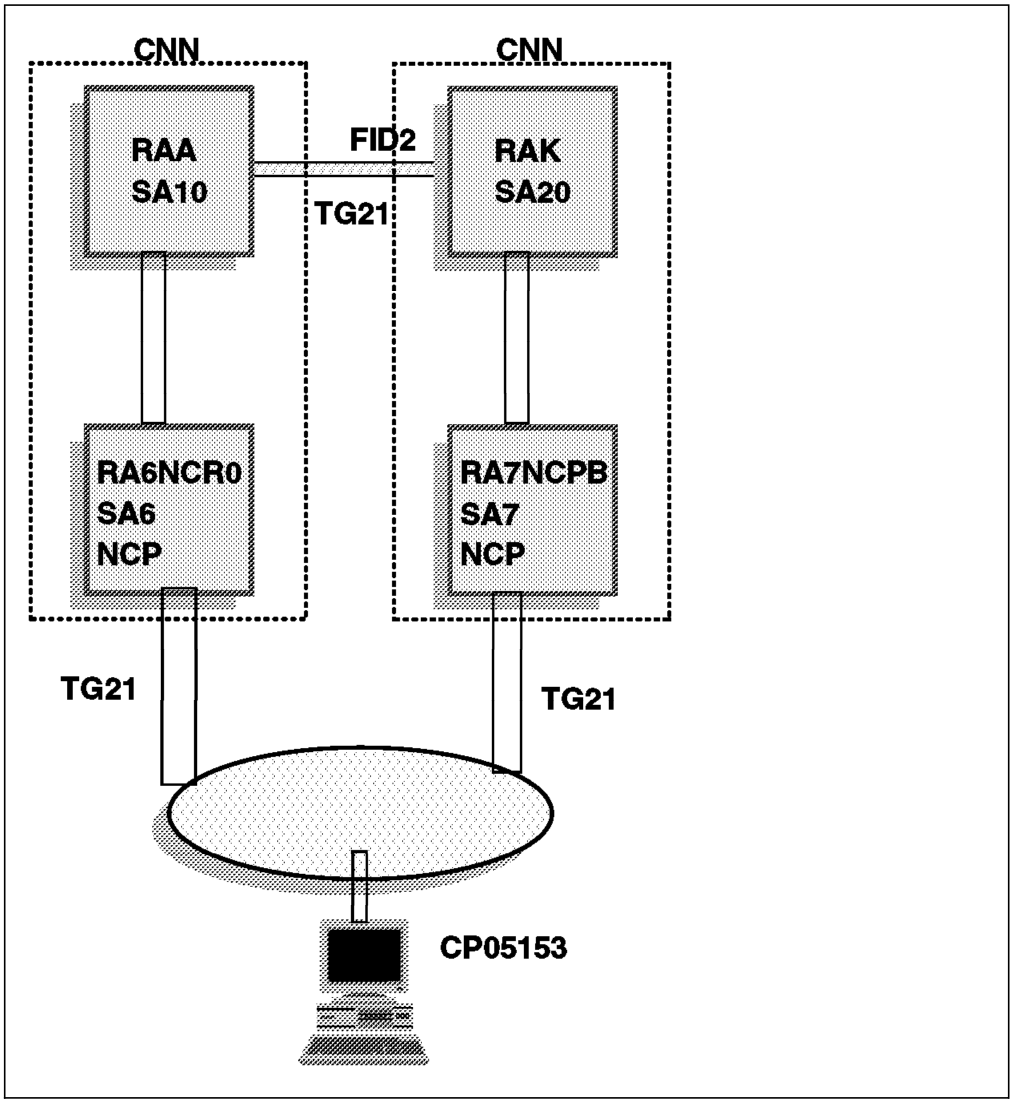

Figure 61. Network Scenario

The CS/2 machine CP05153 was configured as a network node, although the node type did not affect the HPR or DLUR configuration. CP05153 was connected to both NCPs (and therefore to both VTAM APPN nodes), and we wanted it to be able to use both VTAMs as dependent LU servers at the same time. This was not necessary; we could have defined one VTAM as the primary DLU server and the other as backup. However, we wanted to demonstrate the full flexibility of the DLUR configuration on CS/2.

7.5.1 Configuring HPR and DLUR on CS/2

To enable an RTP connection between CS/2 and VTAM no definitions are necessary. Both VTAM′s and CS/2′s configuration parameters default to RTP support. We checked the customization by invoking the Token-ring DLC Adapter Parameters panel. Selecting **HPR Parameters** then gives you another panel which allows you to select:

- Whether HPR is enabled on this adapter (YES by default)
- Whether HPR MLTG is allowed on this adapter (NO by default)
- The degree of link-level error recovery required on this adapter (not preferred by default)

In our environment the customization of CS/2 for DLUR support involved four link definitions, two being real connections (one to each NCP) and two being DLUR/S logical connections (one to each VTAM). The four connections were:

1. The real connection between CP05153 and RAA via RA6NCR0. The identification of the adjacent link station to VTAM could be either by CP name (CP05153) or the node ID (05D-05153, mapped to switched major node keywords IDBLK=05D and IDNUM=05153).
2. The real connection between CP05153 and RAK via RA7NCPB. The identification of the adjacent link station to VTAM is exactly the same as that to RAA. The CP name presented to each VTAM can, and must, be the same. The node ID could be different (CS/2 can send a unique node ID on each link), but this not necessary.
3. The logical connection used by the internal PU and dependent LUs for which CP05153 is DLUR and RAA is DLUS. In order to ensure that this logical connection can be distinguished from the real connection, a different identification is required which must rely on the node ID, as the internal PU has no CP name of its own. Therefore, the node ID was configured as 05D-05154, mapped to switched major node keywords IDBLK=05D and IDNUM=05154.
4. The logical connection used by the internal PU and dependent LUs for which CP05153 is DLUR and RAK is DLUS. The identification of this connection must differ from that of the real connection, and we configured it as 05D-DDD54. This maps to switched major node definitions IDBLK=05D and IDNUM=DDD54. This particular node ID could have been the same as that presented to RAA for its DLUR PU, because the logical connection was to a different VTAM. However, we chose to give the two internal PUs different node IDs to minimize confusion.

Since both our VTAMs were NNs, no further definitions were required to give them DLUS capability. All VTAM NNs from V4R2 onwards are automatically able to perform DLUS functions.

On the real link stations, we used no VTAM definitions and allowed them to be created dynamically using the configuration services exit ISTEXCCS. These link stations had no dependent LUs, so each VTAM acquired a link station represented by the PU W05153, as named by ISTEXCCS from the node ID.

As to the DLUR PUs, we allowed the configuration services exit to define the one on RAK but we coded a manual definition on RAA as seen in Figure 62 on page 93. All DLUR/S connections appear to VTAM as switched connections; the actual physical connectivity is irrelevant.

```
LO5153 VBUILD TYPE=SWNET,
PU05153 PU     ADDR=02,
               IDBLK=05D,
               IDNUM=05154,  1
               MAXDATA=1033,      MAXIMUM AMOUNT OF DATA
               PUTYPE=2,
               SSCPFM=USSSCS,     (V) VTAM
               ISTATUS=ACTIVE
RA5153L1 LU LOCADDR=2,            FIRST LU MUST BE LOCADDR=2
               MODETAB=AMODETAB,DLOGMOD=M2SDLCQ,
               ISTATUS=ACTIVE     (V) VTAM
```

Figure 62. Switched Major Node in RAA

The IDNUM **1** in this definition matched the IDNUM we defined for the type 2 PU which used RAA as its DLUS. Note that although MAXDATA was coded, it is irrelevant for an *internal* DLUR PU. Sessions using DLUR resources flow using APPN protocols and the maximum PIU size is negotiated by the DLUR node with its adjacent nodes. However, an external PU has a real link to the DLUR and on this link MAXDATA (and other parameters such as MAXOUT) can be used to configure the link.

We then customized the CS/2 machine (running CS/2 Version 5) to correspond. We defined the DLC characteristics, the local node characteristics and the two physical connections in the usual way. The companion volume, *Subarea to APPN Migration: VTAM and APPN Implementation* has examples of CS/2 customization.

To define our DLUR/S setup, we went to the Communications Manager Configuration List, as shown in Figure 63 on page 94. There is an option here for **SNA Dependent LU Server definitions**, which we selected.

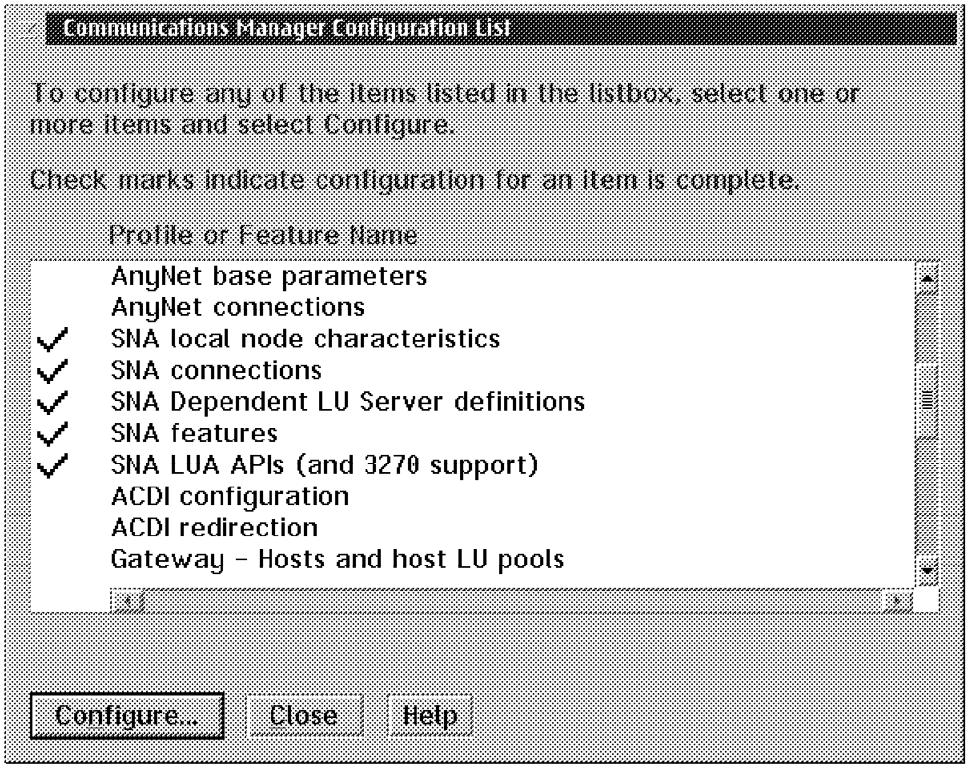

Figure 63. Communications Manager Configuration List

The Dependent LU Server Definitions panel is used to define the DLUR/S connection pipes, which appear to Communications Server as another type of link, albeit logical rather than physical. The normal flows for SSCP-LU and SSCP-PU sessions are used across this link, except that they are carried on a pair of LU 6.2 sessions.

Figure 64 on page 95 shows our configuration for the two DLUR/S connections. Each logical connection has the DLUS name and the node ID defined. CS/2 needs no more information than this to get the dependent LUs activated. APPN is used to locate the defined DLU server, and the node ID is used to identify the appropriate PU from the IDBLK and IDNUM keywords in a switched major node.

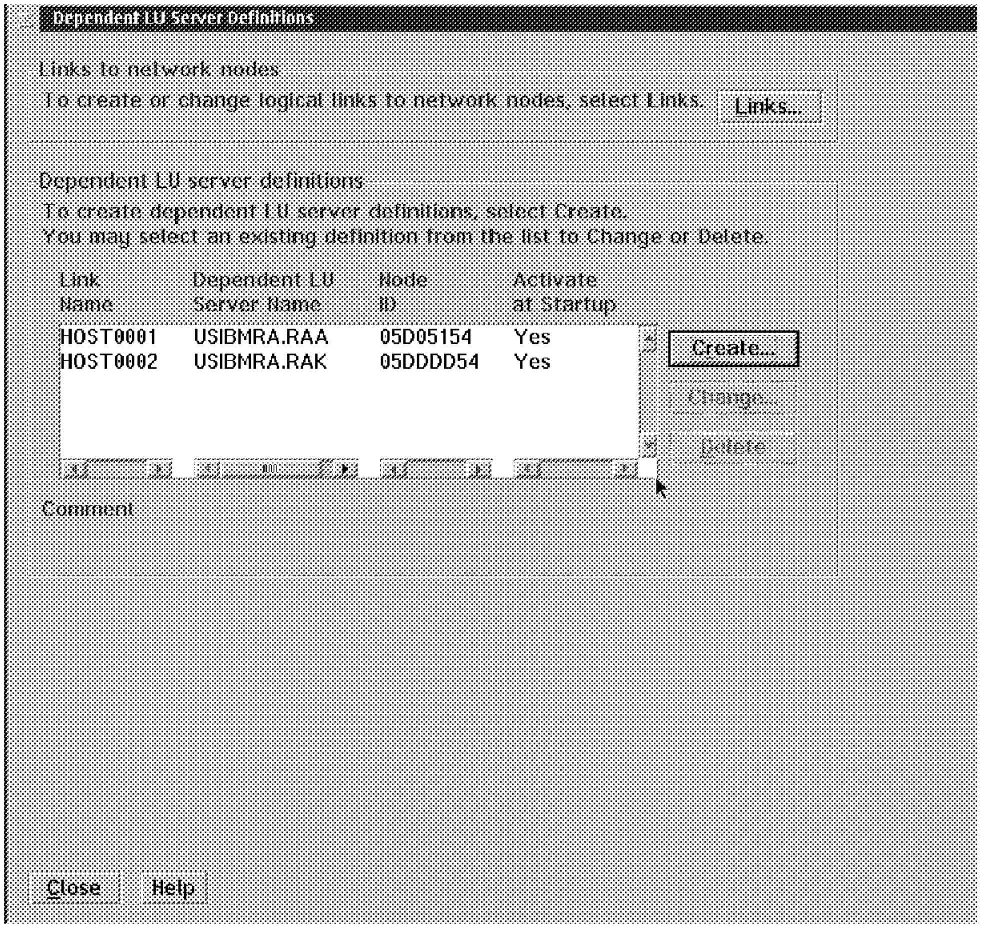

Figure 64. Dependent LU Server Definitions

If you select Create or Change, you get the Dependent LU Server Parameters panel, as in Figure 65 on page 96. This demonstrates what you can define for a DLUR/S connection.

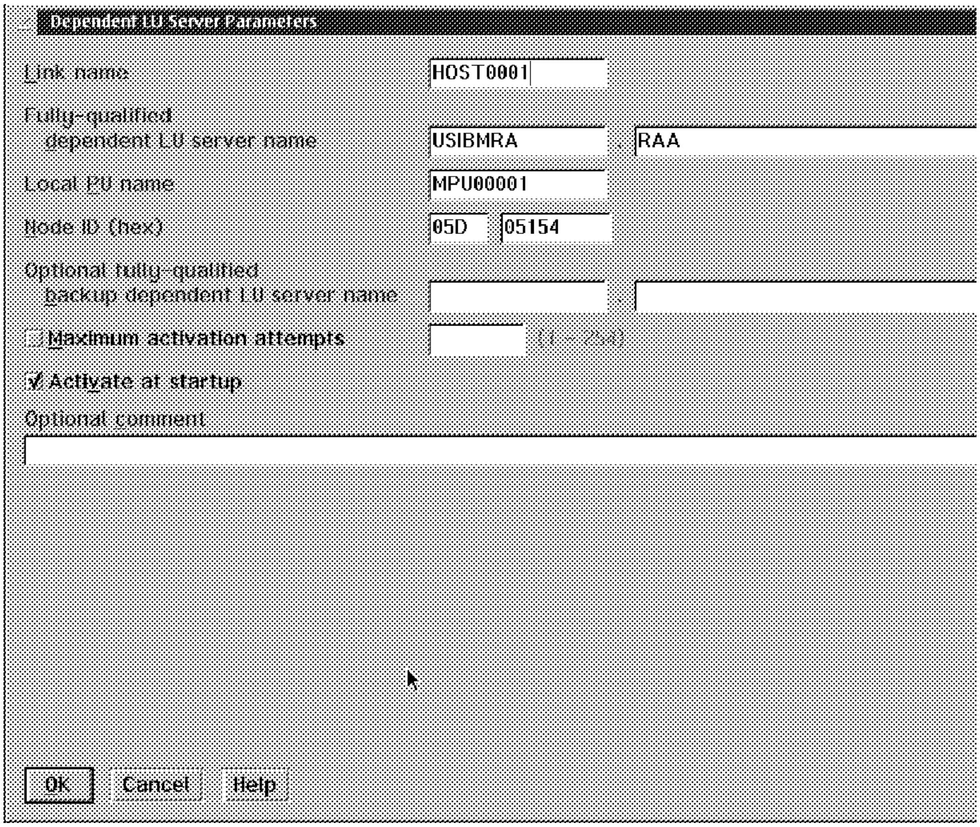

Figure 65. Dependent LU Server Parameters

The link name and local PU name are known only to this CS/2 node, and need not match anything outside the node. The backup DLUS name is optional. If the DLUR/S pipe to the primary DLU server breaks, CS/2 is able to connect to a backup DLU server without disrupting existing dependent LU sessions. This is comparable to the existing support for SSCP takeover of real link stations defined in an NCP.

Once the DLUR/S connection(s) have been defined, the dependent LUs themselves are defined on the appropriate logical links. Figure 66 on page 97 shows that we defined two LUs on each DLUS. The LU names specified here are *not* the LU names known to VTAM. These (local to CS/2) names are used to connect the LU definitions to a product that uses the LUA API, PComm being the prime example. When you define SNA LUs to PComm you define only high-level parameters such as screen sizes and graphics capability. All the rest comes from the CS/2 definitions which you refer to using this LU name.

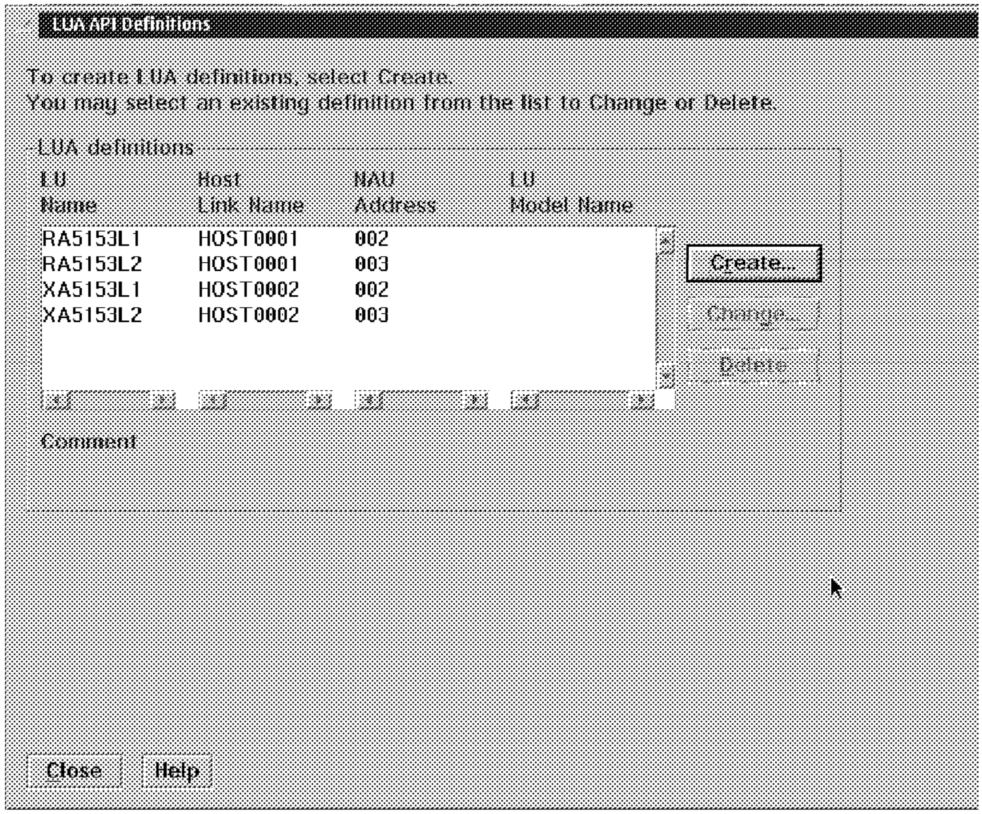

Figure 66. Dependent LU Definitions

When you select Create or Change you get the LUA API Parameters panel as shown in Figure 67 on page 98. This panel allows you to enter the LU name (locally known only), the host link name, and the NAU address which must, of course, correspond to the VTAM LOCADDR keyword on the LU definition. The LU model name, an optional parameter, allows you to supply a name to the ISTEXCSD exit which can define the dependent LUs dynamically. ISTEXCSD and ISTEXCCS can be used together, to provide the most flexible way to define resources dynamically. ISTEXCCS can define adjacent link stations (and PUs) based on XID fields, while ISTEXCSD can define dependent LUs based on NMVT requests sent on the SSCP-PU session.

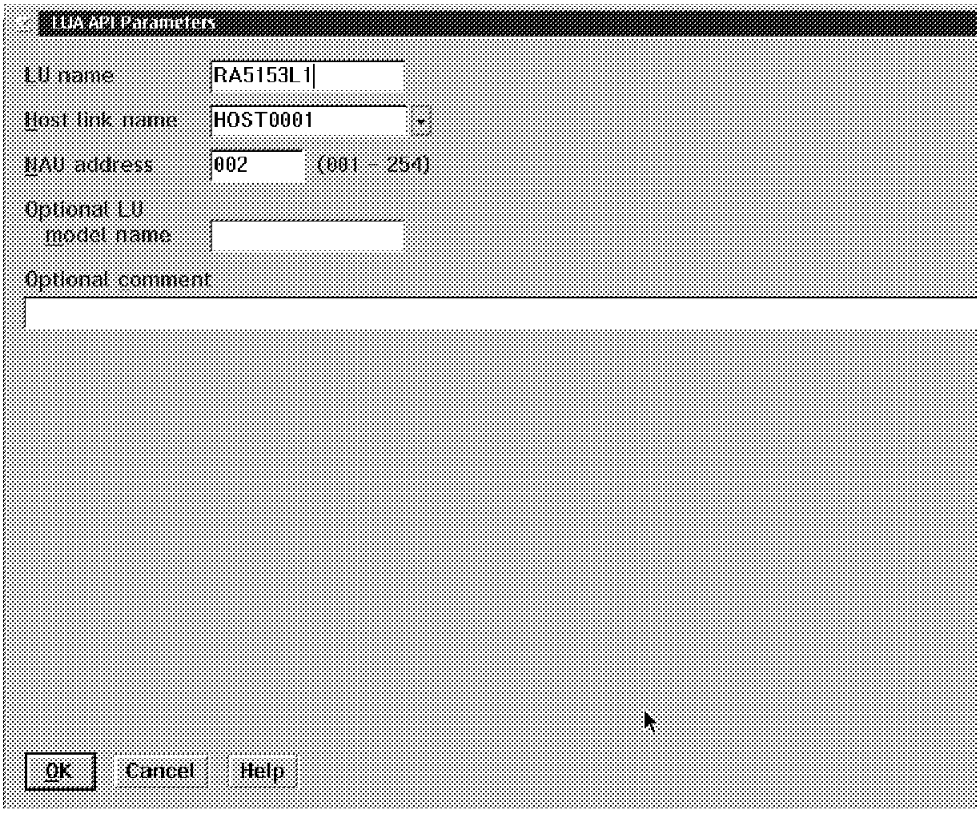

Figure 67. Dependent LU Parameters

7.5.2 Using HPR and DLUR on CS/2

The administration facilities of CS/2 provide you with a very good graphical interface to display and manage your APPN/HPR network. You can obtain from display panels the same information on HPR pipes, and other logical connections, as you can from VTAM. Some CS/2 displays are shown in later chapters of this book, but here we used the VTAM displays to demonstrate what went on.

After starting CS/2, we first verified the connectivity. Figure 68 shows the messages issued by RAA when the CS/2 node was connected. Figure 69 on page 99 shows the corresponding messages on RAK.

```
17:30:43 IST1488I ACTIVATION FOR RTP CNR00710 AS PASSIVE PARTNER COMPLETED   6
17:30:43 IST1488I ACTIVATION FOR RTP CNR00711 AS ACTIVE  PARTNER COMPLETED   7
17:30:45 IST1132I PU05153  IS ACTIVE, TYPE = PU_T2    8

17:31:30 IST590I  CONNECTIN  ESTABLISHED FOR PU W05153 ON LINE J0006027   9
17:31:30 IST1086I APPN CONNECTION FOR USIBMRA.CP05153 IS ACTIVE   10
17:31:30 IST1096I CP-CP SESSIONS WITH USIBMRA.CP05153 ACTIVATED   11
```

Figure 68. CP05153 Connects to RAA

```
17:30:40 IST590I  CONNECTIN  ESTABLISHED FOR PU W05153 ON LINE J0007025   1
17:30:40 IST1086I  APPN CONNECTION FOR USIBMRA.CP05153 IS ACTIVE    2
17:30:40 IST1096I  CP-CP SESSIONS WITH USIBMRA.CP05153 ACTIVATED    3
17:30:43 IST1488I  ACTIVATION   FOR RTP CNR009D4 AS PASSIVE PARTNER COMPLETED   4
17:30:43 IST1488I  ACTIVATION   FOR RTP CNR009D5 AS ACTIVE  PARTNER COMPLETED   5
```

Figure 69. CP05153 Connects to RAK

What happened here was as follows:

- The first link activated by CP05153 was that to RAK **1**. This was immediately followed by the APPN connection **2** and the CP-CP sessions **3**. Since CP05153 is an NN we allow it to establish CP-CP sessions to both VTAMs. Because VTAM does not support APPN Control Flows over RTP on an NCP-attached link, the CP-CP sessions flow as base APPN and no RTP pipes are set up.
- Next, CP05153 activated its DLUR/S connection to RAK. The DLUR/S sessions are not CP-CP sessions in the strict sense, so they can flow over RTP connections. In this case two separate RTP pipes were set up **4** and **5**, one for each of the two LU 6.2 sessions.
- CP05153 then established its other DLUR/S connection to RAA. Since at this time the direct link to RAA was not there, the RTP pipes **6** and **7** must have gone via RA7NCPB (RAK′s NCP). The PU type 2 named PU05153 **8** was the statically defined PU shown in Figure 62 on page 93.
- Finally, CP05153 activated its direct connection to RAA via RA6NCR0 **9**. This, again, resulted in an APPN connection **10** and a pair of CP-CP sessions **11**.

We displayed one of the new RTP pipes from RAA, as shown in Figure 70 on page 100.

```
D NET,ID=CNR00710,E
IST097I DISPLAY  ACCEPTED
IST075I    NAME = CNR00710   , TYPE = PU_T2.1
IST1392I   DISCNTIM = 00010 DEFINED AT PU   FOR DISCONNECT
IST486I  STATUS= ACTIV--LX-, DESIRED STATE= ACTIV
IST1043I   CP NAME = CP05153 , CP NETID = USIBMRA , DYNAMIC LU = YES
IST1589I   XNETALS = YES
IST933I  LOGMODE=***NA***, COS=SNASVCMG   12
IST1476I   TCID X'310BB4C6000000C2' - REMOTE TCID X'000000000000005D'
IST1481I   DESTINATION CP USIBMRA.CP05153 - NCE X'80'
IST1587I   ORIGIN NCE X'D000000000000000'
IST1477I   ALLOWED DATA FLOW RATE =   9000  BITS/SEC   15
IST1516I   INITIAL DATA FLOW RATE =   1000  BITS/SEC
IST1511I   MAXIMUM NETWORK LAYER PACKET SIZE = 2224 BYTES
IST1478I   NUMBER OF UNACKNOWLEDGED BUFFERS =   0
IST1479I   RTP CONNECTION STATE = CONNECTED        - MNPS = NO
IST1480I   RTP END TO END ROUTE - PHYSICAL PATH
IST1460I   TGN  CPNAME                   TG TYPE      HPR
IST1461I    21  USIBMRA.RAK              APPN         RTP   13
IST1461I    21  USIBMRA.CP05153          APPN         RTP
IST231I  RTP      MAJOR NODE = ISTRTPMN
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I  STATE TRACE = OFF
IST355I  LOGICAL UNITS:
IST080I  CP05153  ACT/S----Y   14
```

Figure 70. DLUR/S RTP Pipe from RAA

Note that:

- This pipe connects VTAM with remote LU CP05153 **14** using COS SNASVCMG **12**, as we expect for a DLUR/S connection.
- The route is via the MPC connection to RAK **13**, as we deduced from the order of activation of the various connections.
- We have not checked our TG characteristics **15**. Somewhere on this RTP pipe there is a connection whose capacity VTAM does not know. We did not investigate the cause of this because our purpose was to demonstrate function, not performance.

One of the corresponding DLUR/S pipes from RAK to CP05153 is displayed in Figure 71 on page 101. This shows the session taking the direct route from RAK to CP05153 **16**.

```
D NET,ID=CNR009D5,E
IST097I DISPLAY  ACCEPTED
IST075I NAME = CNR009D5         , TYPE = PU_T2.1
IST1392I DISCNTIM = 00010 DEFINED AT PU   FOR DISCONNECT
IST486I STATUS= ACTIV--LX-, DESIRED STATE= ACTIV
IST1043I CP NAME = CP05153 , CP NETID = USIBMRA , DYNAMIC LU = YES
IST1589I XNETALS = YES
IST933I LOGMODE=***NA***, COS=SNASVCMG
IST1476I TCID X'1EC14FE50000023B' - REMOTE TCID X'000000000000005E'
IST1481I DESTINATION CP USIBMRA.CP05153   - NCE X'80'
IST1587I ORIGIN NCE X'D000000000000000'
IST1477I ALLOWED DATA FLOW RATE =     13 KBITS/SEC
IST1516I INITIAL DATA FLOW RATE =   1000  BITS/SEC
IST1511I MAXIMUM NETWORK LAYER PACKET SIZE =   2224 BYTES
IST1478I NUMBER OF UNACKNOWLEDGED BUFFERS =     0
IST1479I RTP CONNECTION STATE = CONNECTED             - MNPS = NO
IST1480I RTP END TO END ROUTE - PHYSICAL PATH
IST1460I TGN  CPNAME             TG TYPE      HPR
IST1461I  21  USIBMRA.CP05153    APPN         RTP   16
IST231I RTP      MAJOR NODE = ISTRTPMN
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST355I LOGICAL UNITS:
IST080I CP05153  ACT/S----Y
IST314I END
```

Figure 71. RTP Connection Detail

Next, we established a session from RA5153L1 to an application on RAA. RA5153L1 is a dependent LU owned by RAA, but on the DLUR PU PU05153 which is located in CP05153. For this session a new RTP connection was set up, as shown in Figure 72.

```
D NET,ID=CNR00714,E
IST097I DISPLAY  ACCEPTED
IST075I NAME = CNR00714         , TYPE = PU_T2.1
IST1392I DISCNTIM = 00010 DEFINED AT PU   FOR DISCONNECT
IST486I STATUS= ACTIV--LX-, DESIRED STATE= ACTIV
IST1043I CP NAME = CP05153 , CP NETID = USIBMRA , DYNAMIC LU = YES
IST1589I XNETALS = YES
IST933I  LOGMODE=***NA***, COS=#CONNECT
IST1476I TCID X'310BB4CA000000C3' - REMOTE TCID X'000000000000005A'
IST1481I DESTINATION CP USIBMRA.CP05153   - NCE X'80'
IST1587I ORIGIN NCE X'D000000000000000'
IST1477I ALLOWED DATA FLOW RATE = 9000 BITS/SEC
IST1516I INITIAL DATA FLOW RATE = 1000 BITS/SEC
IST1511I MAXIMUM NETWORK LAYER PACKET SIZE = 2224 BYTES
IST1478I NUMBER OF UNACKNOWLEDGED BUFFERS =   0
IST1479I RTP CONNECTION STATE = CONNECTED
IST1480I  RTP END TO END ROUTE - PHYSICAL PATH
IST1460I   TGN  CPNAME                TG TYPE      HPR
IST1461I    21  USIBMRA.CP05153       APPN         RTP
IST231I  RTP      MAJOR NODE = ISTRTPMN
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I  STATE TRACE = OFF
IST355I  LOGICAL UNITS:
IST080I  RA5153L1 ACT/S
IST314I  END
```

Figure 72. New RTP Connection

We can see that this session takes the direct path; exactly the same route, in fact, as it would have taken without DLUR. Only when we broke the link later did we see how DLUR/HPR differs from the old subarea way of working.

We also performed an APING transaction between CP05153 and RAK. This resulted in the setting up of a new RTP connection, CNR009D6. A display of CP05153 from RAK then showed all the sessions active at the time between the two nodes (refer to Figure 73).

```
D NET,ID=CP05153,E
IST097I  DISPLAY  ACCEPTED
IST075I NAME = USIBMRA.CP05153  , 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.CP05153   - 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 = 0000000006, SESSION REQUESTS = 0000000000
IST206I SESSIONS:
IST1081I  ADJACENT LINK STATION = CNR009D6   17
IST634I NAME     STATUS         SID          SEND RECV VR TP
IST635I RAK      ACTIV-S    ECBF5D353E4D269F 000A 000D  0  0
IST1081I  ADJACENT LINK STATION = CNR009D4   18
IST634I NAME     STATUS         SID          SEND RECV VR TP
IST635I RAK      ACTIV/SV-S ECBF5D353D4D269F 0001 0001  0  0
IST635I RAK      ACTIV/DL-S ECBF5D35394D269F 0016 0000  0  0
IST1081I  ADJACENT LINK STATION = W05153     19
IST634I NAME     STATUS         SID          SEND RECV VR TP
IST635I RAK      ACTIV/CP-S ECBF5D35374D269F 0055 0001  0  0
IST635I RAK      ACTIV/CP-P F8D3D164311C85EA 0001 0056  0  0
IST1081I  ADJACENT LINK STATION = CNR009D5   20
IST634I NAME     STATUS         SID          SEND RECV VR TP
IST635I RAK      ACTIV/DL-P F8D3D164311C85EB 0000 0016  0  0
IST1355I PHYSICAL UNITS SUPPORTED BY DLUR USIBMRA.CP05153
IST089I WDDD54   TYPE = PU_T2      21
IST924I -----------------------------------------------------
IST075I NAME = USIBMRA.CP05153  , TYPE = DIRECTORY ENTRY
IST1186I DIRECTORY ENTRY = DYNAMIC    NN
IST1184I CPNAME = USIBMRA.CP05153   - NETSRVR = ***NA***
IST1402I SRTIMER =   120  SRCOUNT =    60
IST314I  END
```

Figure 73. Display of CS/2 CP and Its DLUR PU

There were actually four logical links between RAK and CP05153:

- The new RTP pipe CNR009D6 **17** carries the new APING LU 6.2 session. Because this session is the only one that uses the APPN COS #INTER, it has an RTP pipe to itself.
- The RTP pipe CNR009D4 **18** carries one of the two DLUR/S sessions. The notation ACTIV/DL indicates a DLUR/S session. The APPN COS used by DLUR/S pipes is SNASVCMG, thus an RTP connection carrying a DLUR/S

session can be shared by other sessions using this COS. Indeed, there is another session on this RTP pipe. The session labelled ACTIV/SV is the CNOS session from the APING transaction, which also uses SNASVCMG.

- The CP-CP sessions flow on the real link W05153 **19**. The name W05153 was defined by the ISTEXCCS exit in response to the receipt of IDNUM 05153 on the XID. Because VTAM does not support Control Flows on an NCP connection, the CP-CP sessions cannot flow on an RTP connection.
- The RTP pipe CNR009D5 **20** carries the second DLUR/S LU 6.2 session. Although it uses APPN COS SNASVCMG, it does not share the pipe CNR009D4. This is probably because the two sessions were started within a short time of each other, and the RTP pipe setup overlapped.

Note also that CP05153 is identified **21** as the DLU requester that looks after the type 2 PU WDDD54. This name was also created by ISTEXCCS from the IDNUM we gave it in our CS/2 setup. Figure 74 illustrates the structure of the sessions and RTP pipes between RAK and CP05153 at this time.

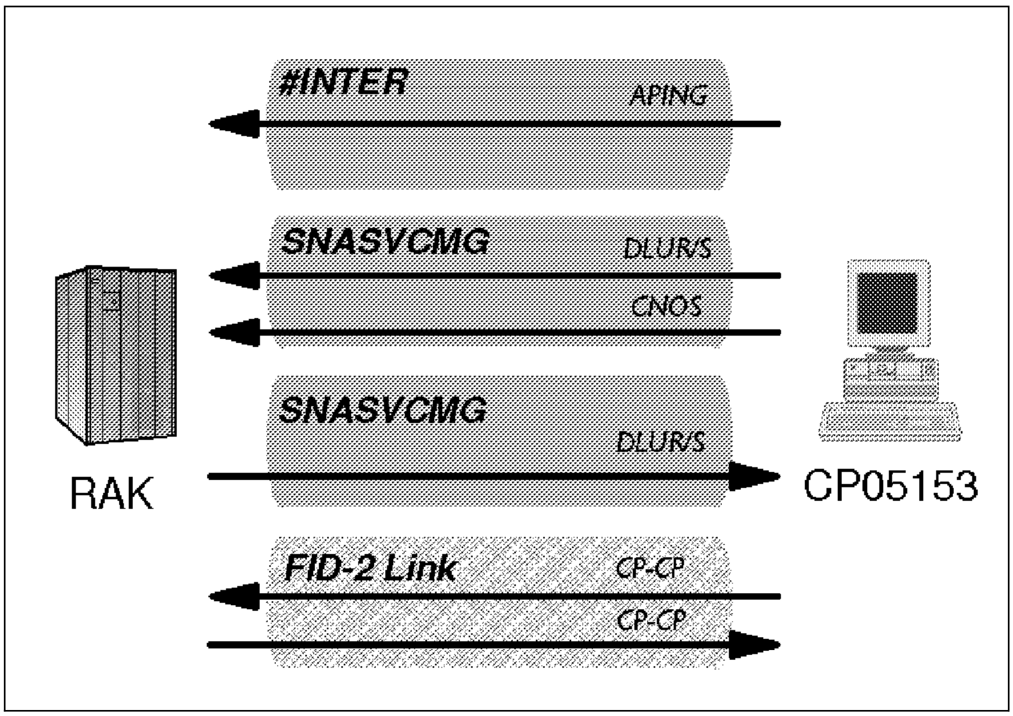

Figure 74. RTP Pipes and Sessions

Having established all the sessions and displayed all the details of the RTP and DLUR connections, we deactivated the connection between RAA and CP05153 across the token-ring. Figure 75 on page 104 shows the APPN view of what happened. Figure 76 on page 104 shows the resulting display.

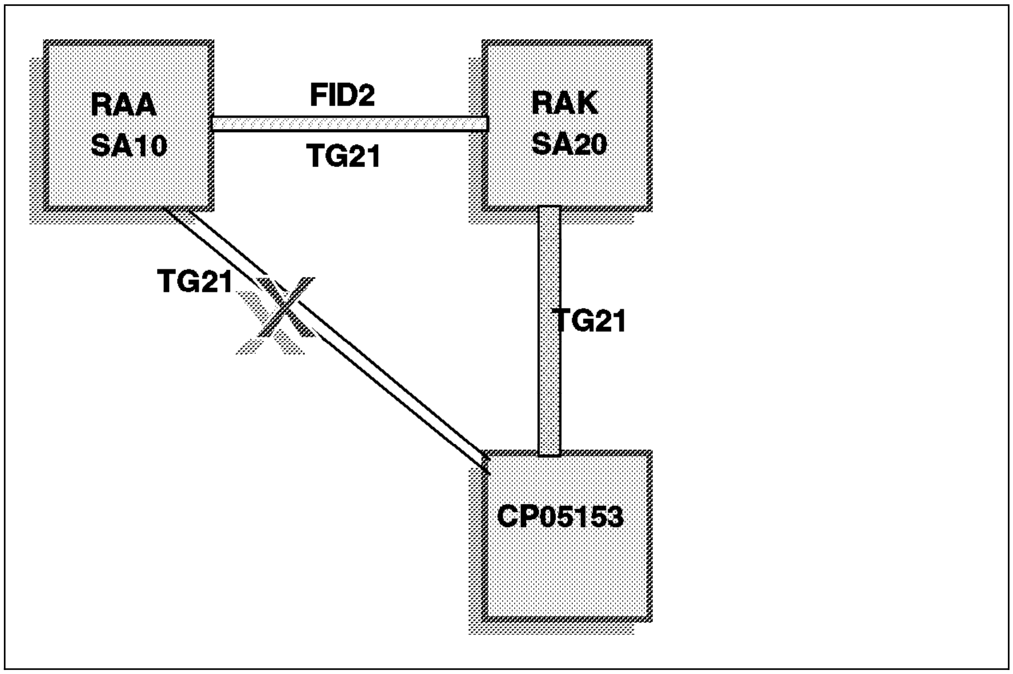

Figure 75. Failing Link and Alternative Path

```
IST259I  INOP RECEIVED FOR W05153  CODE = 04
IST1416I  ID = W05153     FAILED - RECOVERY IN PROGRESS
IST1136I  VARY INACT W05153        SCHEDULED - UNRECOVERABLE ERROR
IST1097I  CP-CP SESSION WITH USIBMRA.CP05153   TERMINATED
IST1280I  SESSION TYPE = CONLOSER - SENSE = 80030004
IST314I  END
IST1196I  APPN CONNECTION FOR USIBMRA.CP05153 INACTIVE - TG
IST590I  CONNECTION TERMINATED  FOR PU W05153 ON LINE J0006
IST1097I  CP-CP SESSION WITH USIBMRA.CP05153  TERMINATED
IST1280I  SESSION TYPE = CONWINNER - SENSE = 08420001
IST314I  END
IST1133I  W05153      IS NOW INACTIVE, TYPE = PU_T2
```

Figure 76. Network Log on RAA during the Link Failure

We see that the APPN connection is broken, the CP-CP sessions are terminated and the PU W05153 is deactivated. But a display of the RTP pipe carrying the dependent LU session (Figure 77 on page 105) shows that the session is still alive.

```
D NET,ID=CNR00714,E
IST097I DISPLAY  ACCEPTED
IST075I NAME = CNR00714         , TYPE = PU_T2.1
IST1392I DISCNTIM = 00010 DEFINED AT PU   FOR DISCONNECT
IST486I STATUS= ACTIV--LX-, DESIRED STATE= ACTIV
IST1043I CP NAME = CP05153 , CP NETID = USIBMRA , DYNAMIC LU = YES
IST1589I XNETALS = YES
IST933I  LOGMODE=***NA***, COS=#CONNECT
IST1476I TCID X'310BB4CA000000C3' - REMOTE TCID X'000000000000005A'
IST1481I DESTINATION CP USIBMRA.CP05153   - NCE X'80
IST1587I ORIGIN NCE X'D000000000000000'
IST1477I ALLOWED DATA FLOW RATE = 9000 BITS/SEC
IST1516IINITIAL DATA FLOW RATE = 1000BITS/SEC
IST1511I MAXIMUM NETWORK LAYER PACKET SIZE = 2224 BYTES
IST1478I NUMBER OF UNACKNOWLEDGED BUFFERS =   0
IST1479I RTP CONNECTION STATE = CONNECTED
IST1480I  RTP END TO END ROUTE - PHYSICAL PATH
IST1460I  TGN  CPNAME           TG TYPE    HPR
IST1461I   21  USIBMRA.RAK      APPN       RTP
IST1461I   21  USIBMRA.CP05153  APPN       RTP
IST231I  RTP      MAJOR NODE = ISTRTPMN
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I  STATE TRACE = OFF
IST355I  LOGICAL UNITS:
IST080I  RA5153L1 ACT/S
IST314I  END
```

Figure 77. Session Display

The session has moved to the path via RAK and RA7NCPB.

This shows that, with HPR and DLUR, dependent LU sessions can receive exactly the same availability benefits from end to end as independent LU sessions. Without DLUR the HPR resilience only works between RTP-capable nodes, which means that the portion of the session between the workstation and the nearest RTP node on the path is not recoverable in case of failure.

Chapter 8. HPR and DLUR on the 3746

In this chapter we have replaced the NCPs with 3746 network node processors, and we describe how HPR and DLUR may be implemented on the 3746-9X0 platform. Before we discuss how the APPN functions are configured in the 3746, we show in Figure 78 the test scenario that we used.

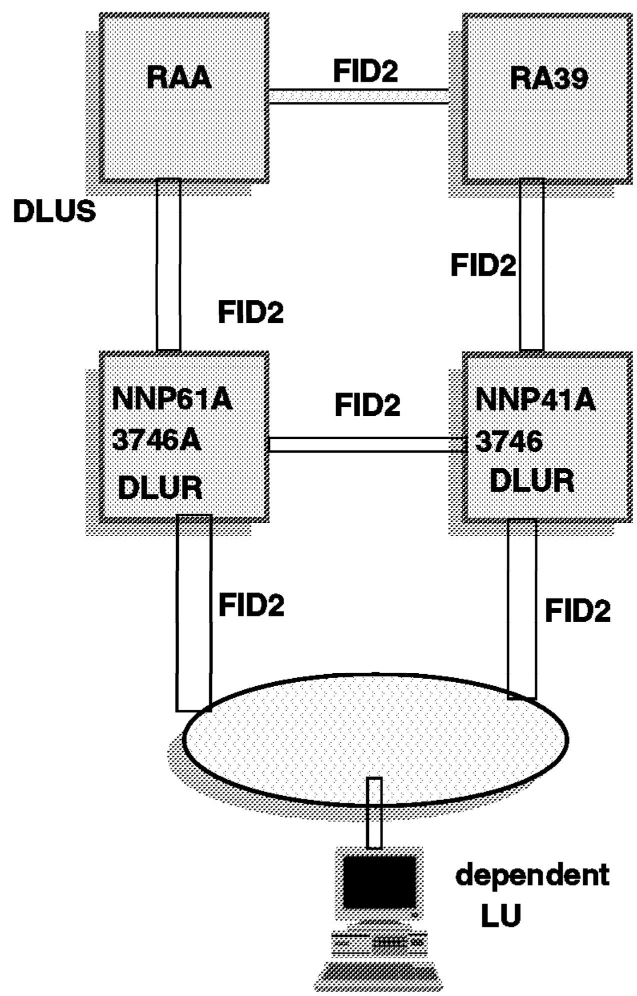

Figure 78. Network Configuration

Here there is no need for VR-TG connection as most of the reasons for it are no longer valid. There are no NCPs, no subarea MLTGs, no SSCP takeover and no BSC 3270 sessions. The two VTAM NNs we used, RAA and RA39, are connected by an APPN MPC link. Two 3746 NNPs, NNP61A and NNP41A, have replaced the NCPs and are connected to the VTAMs in a square as shown. All connections are FID-2. Strictly speaking, they should be called APPN rather than FID-2 because they can all carry NLPs as well as FID-2 traffic. The connection between the two 3746s is a token-ring, while the connections between 3746s and VTAMs are ESCON channels.

The workstation we used, with its dependent LUs, was configured in two ways to illustrate various methods of connecting such a device:

1. In the first case, the workstation was configured without DLUR, and the dependent LUs used peripheral subarea protocols to the 3746s. The 3746s therefore had to perform the DLUR function. This is an example of DLUR passthrough support, where the DLUR node preserves the identity of the attached nodes to the DLUS. RTP endpoints are in the 3746s, which gives greater resilience than the NCP configuration but not as much as the full end-to-end DLUR/S path.
2. In the second case, we defined the DLUR function within the workstation, allowing APPN and HPR all the way from the dependent LUs to the application. This is similar to our last scenario in 7.5.2, "Using HPR and DLUR on CS/2" on page 98, except that the 3746s provide the ANR routing instead of the NCPs.

We used token-ring and ESCON connections throughout for simplicity. We realize that in the real world there would be frame relay, SDLC and X.25 connections in such a network. However, the differences between these protocols are at the DLC level and have little effect on the operation of APPN and HPR. Please refer to Appendix F, "Related Publications" on page 325 for a list of publications that describe how to configure other DLCs on the platforms we used.

The 3746s we used were both Model 900s with NNPs attached to 3745s. As there were no NCP-owned resources in use, Model 950 stand-alone machines would have worked in exactly the same way.

8.1 3746 Configuration

In the next few sections we detail the connection scenario that was used to attach APPN and non-APPN nodes to our 3746 NNs. The required parameter settings for the 3746s and their adjacent nodes have been included. It should be remembered that at the time of testing the emphasis was on connectivity rather than on optimizing performance.

Note: Unless specifically mentioned, similar definitions apply for attachments to a 3746 Nways Controller (3746-950). The generic term 3746 NN is used if definitions apply to both 3746-900 and 3746-950. We refer to either 3746-900 or 3746-950 when definitions apply to one of these only.

Each of the configurations depicted in this section requires network node, port and link station definitions on the 3746 NN. A port provides access to the physical medium (ESCON, SDLC, frame relay, token-ring) enabling communication to a destination node. A link station represents the adjacent node on an APPN link and specifies the characteristics of the connection. Apart from VTAM, all the products we used (3746, 2216, 2210, CS/2) have a similar pattern when it comes to defining link resources.

First you define a port, which corresponds to an adapter, a protocol and possibly a local SAP used by that protocol. Next, in a node that will initiate contact, you define the adjacent link station that will be contacted. Except for non-switched SDLC connections, the adjacent link station that initiates the contact is not defined at the node that receives the contact. In general, an adjacent link station representing a connection between a workstation EN and an NN is defined only in the EN.

Normally, multiple APPN link stations can be defined on a physical port. The major exception is SDLC, which allows only a multipoint primary station to share secondary link stations.

The configuration definition is shown on the basis of the type of attachment used to connect our equipment to the 3746 Model 900:

- ESCON coupler:
  - Connection between a VTAM network node and a 3746 NN using a single ESCON port
- Token-ring coupler:
  - TIC3 connection between two 3746-900 machines
  - Connection between a PS/2 (CS/2) node supporting dependent LUs and a 3746 NN
  - Connection between a PS/2 (CS/2) network node with DLUR support and a 3746 NN

8.2 Controller Configuration and Management

Controller Configuration and Management (CCM) is an application that resides on the network node processor and is accessed through the MOSS-E interface of the service processor. Optionally, the CCM tool can also operate on a stand-alone PC. The prerequisites for stand-alone operation are:

- OS/2 V2.1 or later
- A minimum of 30 MB disk space for the CCM application
- A minimum of 20 MB disk space for the swapper file

The two main components of the CCM application are the configurator and the management interface. The management interface allows the user to manage the APPN NN, for example starting/stopping the APPN node, ports and link stations. The configurator is used to define all resources used by the 3746-9x0. It enables the user to customize the APPN NN, define APPN and non-APPN attachments, and configure DLUR functions. In addition, there is support for IP and associated routing protocols, and for frame relay as a frame handler (FRFH).

The configurator produces a 3746 NN configuration file that is used by the APPN control point (CP). It also produces files to define all supported interfaces, such as ESCON, serial and token-ring.

The CCM program runs like any other application. To start CCM double-click on the CCM icon that is automatically created when the CCM program is installed. Before opening a configuration, a list of available configuration data files is presented. After selecting the appropriate configuration, the topology is displayed on the primary configuration window. This reflects the hardware configuration of the 3746-9x0; different configurations can be customized to your particular requirements.

8.2.1 Configuration Files

CCM provides a graphical user interface (GUI) to define the APPN NN and DLUR parameters, plus the ESCON, token-ring, IP, frame relay, and SDLC connections. The GUI enables the user to set default values for a great number of configuration parameters. The ESCON Generation Assistant, a tool formerly used to define the ESCON attachments, has been integrated into CCM.

CCM creates a set of output files. The files of a given configuration are compressed into a unique configuration file, which will be referred to as the 3746 NN configuration file. Several configurations may be defined by the user but only one can be active (in use by the 3746) at a time. Configuration files cannot be edited by the customer other than using CCM facilities. Facilities exist to export the configuration file to disk, or import the configuration file from disk. Import and export functions are especially important when creating a configuration on a stand-alone PC.

8.2.2 CCM Environments

CCM can operate on both the service processor and on a stand-alone PC. However, CCM′s method of operation, and the functions available, are different.

8.2.2.1 CCM on the Service Processor

All the configuration, operation, and management functions are available on the service processor. The CCM program is accessible via the service processor, either locally or remotely using a DCAF station.

Once CCM is active it can be used to configure and manage the 3746 NN, start and stop the APPN control point, and perform adapter traces.

8.2.2.2 CCM on a Stand-Alone PC

Although the CCM application can run on a stand-alone PC, its functions are limited. Most configuration options are available. However, as the stand-alone PC has no access to the resources critical for APPN operation, CCM operation and management functions are excluded. A configuration file made on the stand-alone PC can be exported to diskette and imported by the CCM running on the service processor.

8.3 Importing and Exporting a Configuration

3746 APPN NN configuration files can be imported and exported to allow activation of the configuration file produced on a stand-alone PC.

To export a configuration from your stand-alone PC to a diskette, select **File** on the CCM primary menu. From the pull-down menu click on **Open**. The system will prompt you with a window listing the configurations. By selecting a configuration and clicking on **Export**, the configuration is copied to one or more diskettes, depending on the configuration size (see Figure 80 on page 111).

To import a configuration from a diskette, select **File** on the CCM primary menu. From the pull-down menu click on **Import a configuration**. You will then be prompted to insert a diskette in the drive. Once you click on **Yes** the configuration file will be copied to the hard disk of the network node processor(s) and the service processor. If the configuration file already exists, you have the option to overwrite the file already stored.

8.4 Activating a Configuration

To activate a configuration, click on the **File** button and from the pull-down menu select **Open** as in Figure 79.

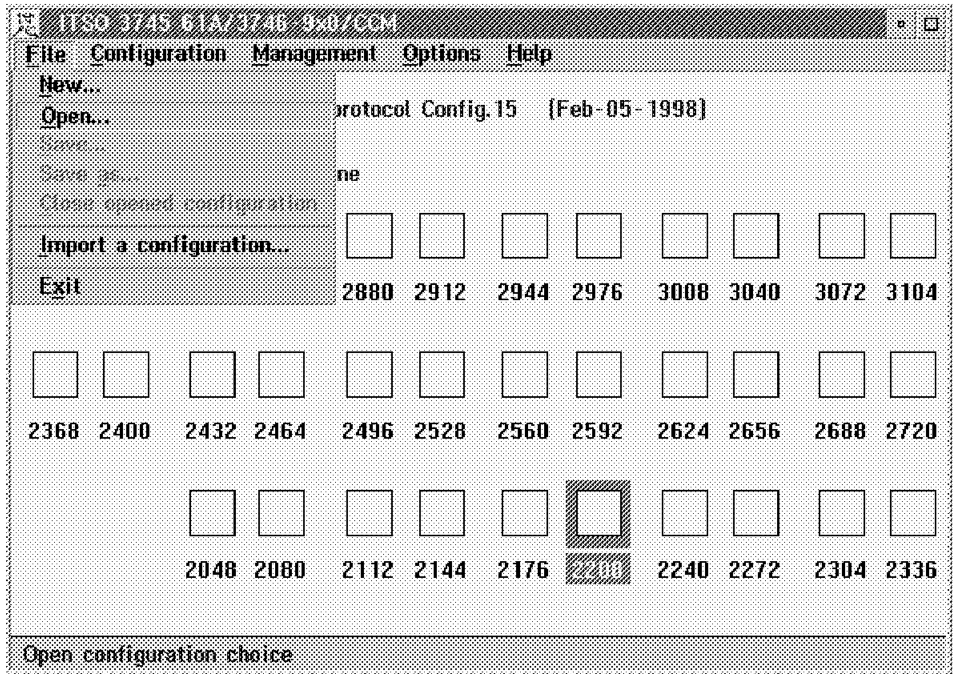

Figure 79. Open Configuration

A configuration list will result as shown in Figure 80.

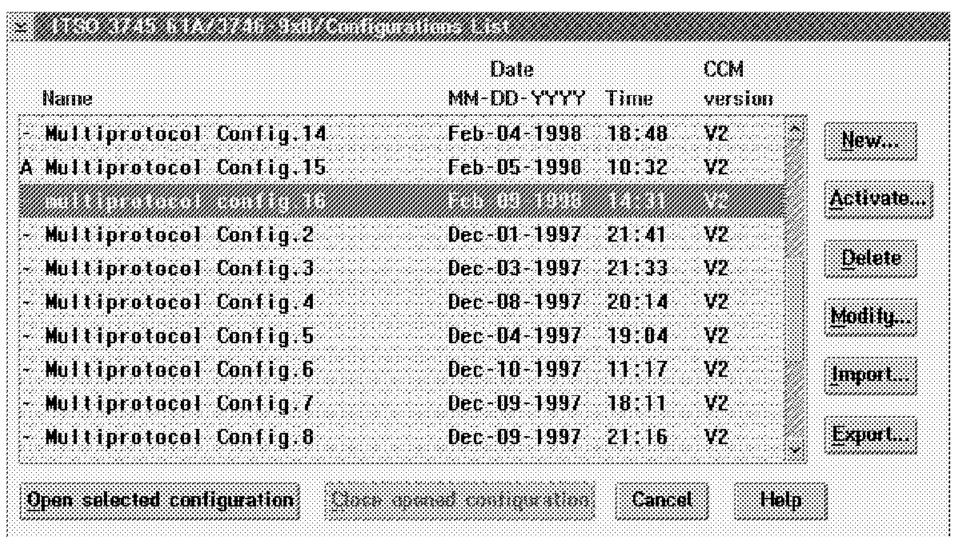

Figure 80. Display of Already Defined Configurations

From the configuration files displayed, select the one you want to activate and click on **Activate**. You will then be prompted to confirm your activation request, as seen in Figure 81 on page 112. This is because some functions (such as CP restart) are disruptive; activation of a new configuration file can result in a reset of the entire NNP and most of the 3746 adapters.

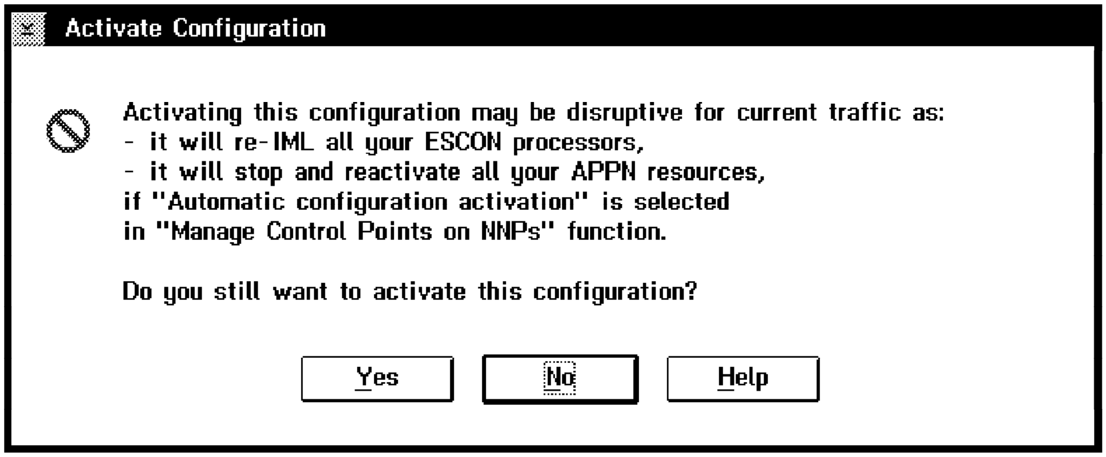

Figure 81. New Configuration Activated

8.5 Creating or Modifying a Configuration

To create or change a configuration, start CCM to obtain the primary CCM window. As shown in Figure 82 on page 113, along the top of the CCM panel are pull-down menus from which you can select the specific tasks you want to perform. The configuration pull-down menu will initiate a series of screens that will prompt the user for the required information.

The CCM window will show the active APPN configuration that the network node node processor is using and the configuration opened for customization.

The CCM primary screen also displays the installation and configuration status of each coupler for the opened configuration. Each coupler is represented by an icon with the coupler address and coupler type (if installed and known) indicated just below it.

The basic arrangement of the 3746 adapters is:

- The 3746 I/O processors (ESCON, serial, token-ring) are connected to the bus, and each processor is assigned 64 addresses.
- The physical attachments to the media are handled by the couplers (ESCON, token-ring, serial). One or two couplers can be attached to each processor. Each coupler is assigned 32 addresses.
- Token-ring and serial processors can handle two couplers, but ESCON processors have only one. Therefore each ESCON processor always has at least one spare coupler slot.
- A pair of serial couplers can be attached to an adjacent processor, whether for production or for backup purposes. This results in a processor looking after 128 addresses. However, it also restricts the amount of traffic that the processor can handle.
- The capacity of a serial processor also depends on the speeds of the lines that are connected to it.

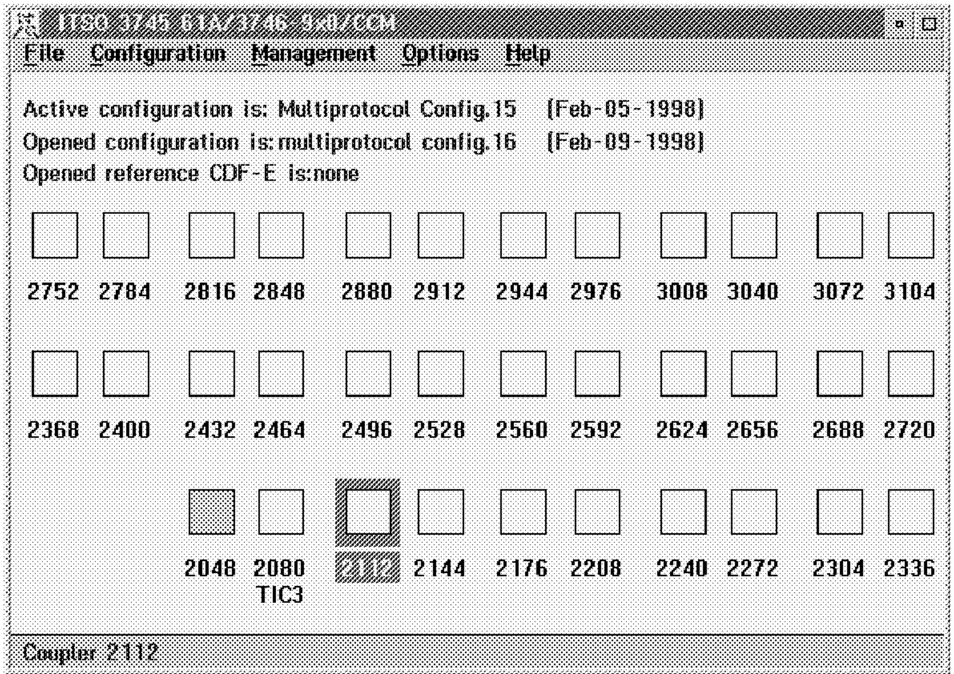

Figure 82. CCM Configuration, Nothing Configured

Figure 82 shows a pristine 3746 configuration with nothing defined as yet. All the couplers are shown as unshaded, with no coupler description, except for three that CCM knows about even without any predefinition:

- Addresses 2048 and 2112 are the CBSPs (controller bus and service couplers) that connect the two 3745 CCUs to the 3746.
- Address 2080 is the token-ring coupler that connects the NNP to the service processor.

For our configurations, we selected each coupler to be configured, selected the configuration pull-down, and filled out the fields for the various coupler screens. Help is available for each field by pressing the F1 key or clicking on the **Help** button. Note that when you try to configure the first coupler, CCM asks you for the 3746 model and the 3745 operating mode if not already configured.

After completing the task of configuring all our adapters, our CCM primary display for this particular 3746 was updated to Figure 83 on page 114. A transparent icon with a check mark indicates a coupler that has been configured (ports 2144 and 2176). These couplers now have their correct coupler type shown beneath the coupler icon. At the same time, coupler 2208 has been shaded to denote that it is not available to be configured. CCM knows that an ESCON processor can have only one coupler attached to it, so the second coupler slot of the pair has been shaded.

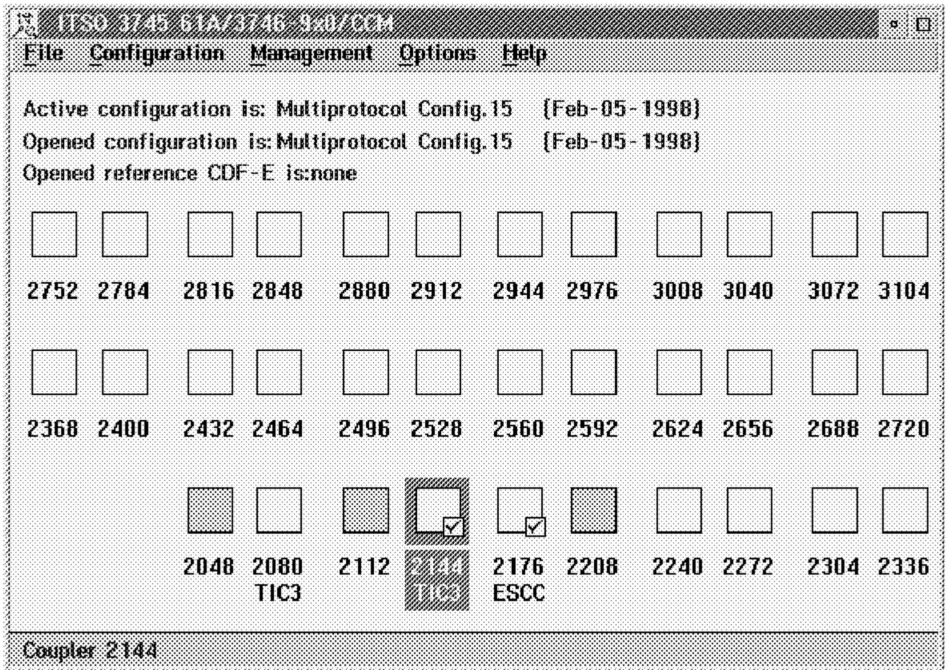

Figure 83. CCM Configuration, TR and ESCON Configured

To modify an existing configuration, click on **File**, **Open**, then **Open Selected Configuration** as shown in Figure 80 on page 111. To create a new configuration select **File** then **New**, in which case the panel shown in Figure 84 will appear. Fill in the required information, select **OK** and you will be returned to the primary screen to customize your configuration.

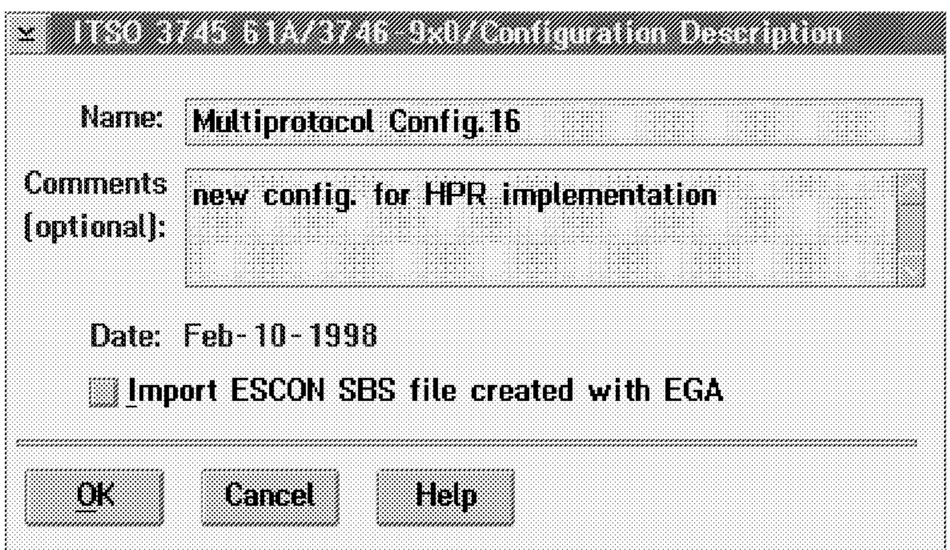

Figure 84. Configuration Description

The file name you filled in will appear as the opened APPN configuration file. You can now start to configure the node, DLUR and the couplers as needed.

Note: The import ESCON file check box visible in Figure 84 allows migration from the old EGA environment. This can be of great assistance in converting ESCON configurations.

8.6 3746 APPN Network Node Definition

Once a new configuration file has been opened, you need to specify the 3746 NN control point name, the management focal point, and the DLU server. For the focal point and the DLU server, backup nodes can be configured. To configure the 3746 NN, focal point and DLUR parameters, select **Configuration** on the primary CCM window. On the pull-down menu click on **APPN NN/FP/DLUR parameters** and you will be prompted with the window as shown in Figure 85.

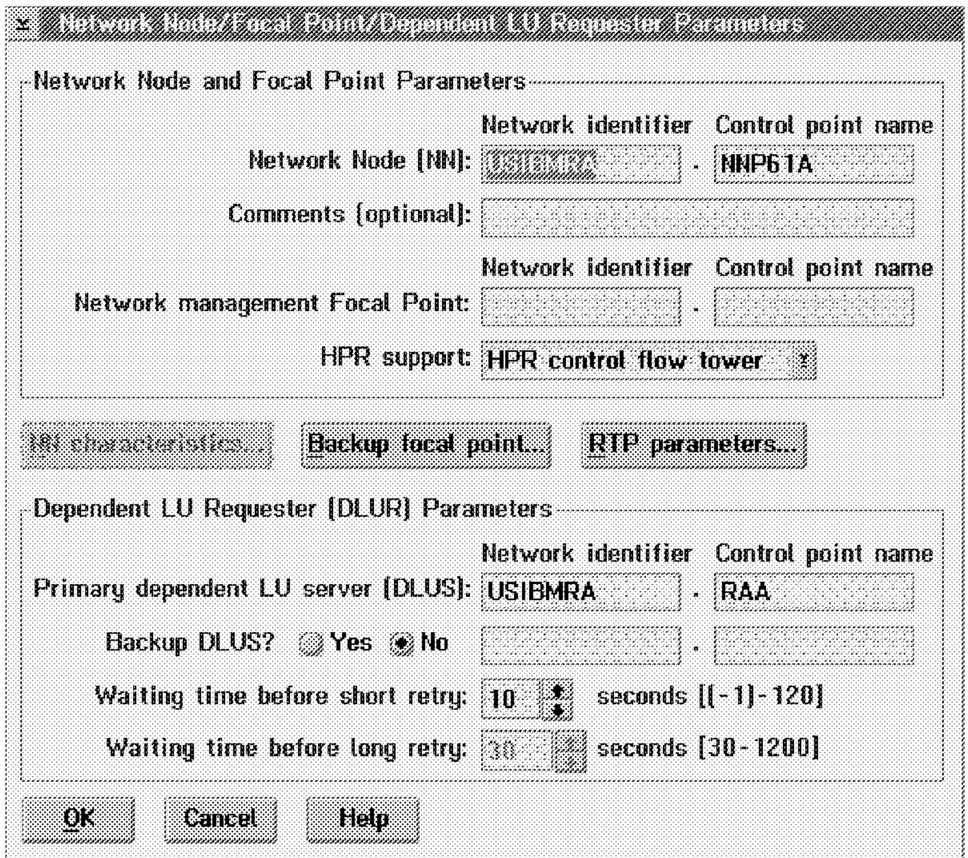

Figure 85. APPN NN, FP and DLUR Parameters

Figure 85 shows how the NN, DLUR, and focal point definitions have been specified for our test environment.

The CP name of the 3746 NN is USIBMRA.NNP61A. The primary DLU server is USIBMRA.RAA; if the 3746 NN is not able to contact this host then a back up DLU server can be specified. We did not choose this option, nor did we define a network management focal point.

The 3746 NN CP name and network ID are mandatory; a CCM configuration cannot be saved before filling in this information. The primary and backup dependent LU servers configured on the APPN NN/FP/DLUR Parameters screen can be considered as the default DLU servers. These parameters will be used for all link stations for which no explicit DLU server is defined.

The level of HPR support can be specified on the scroll bar labelled HPR support. Figure 86 shows that the options available range from no HPR support to Control Flows over RTP. The default value is ANR support. HPR capability can be turned on or off for each individual port or link station.

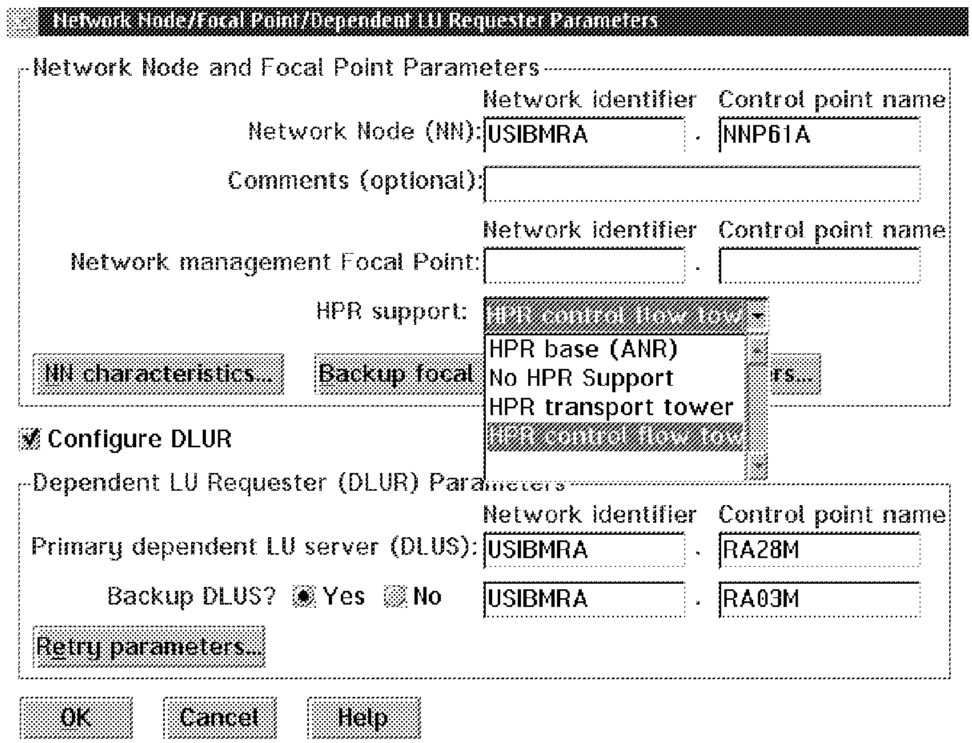

Figure 86. HPR Levels Supported by 3746

The observant reader will see that the panel in Figure 86 differs slightly from that used in Figure 85 on page 115. We displayed a similar configuration using the two latest releases of CCM, to show that the panels you see may not be identical to ours. Figure 86 is the later (F12380). If you select NN characteristics from this panel you will see Figure 87 on page 117, whereas selecting DLUR Retry parameters will give you Figure 88 on page 117.

The DLUR retry algorithm on the 3746 is the same as that on the 2216/2210. If the cause of failure of the DLUR/S pipe is a non-disruptive UNBIND, the 3746 attempts to contact the DLUS or the backup DLUS at intervals determined by the *long retries* setting. With other failures, it tries more often; it performs a sequence of attempts based on the *short retries* timer and count, this sequence being repeated at intervals based on the *long retries* setting.

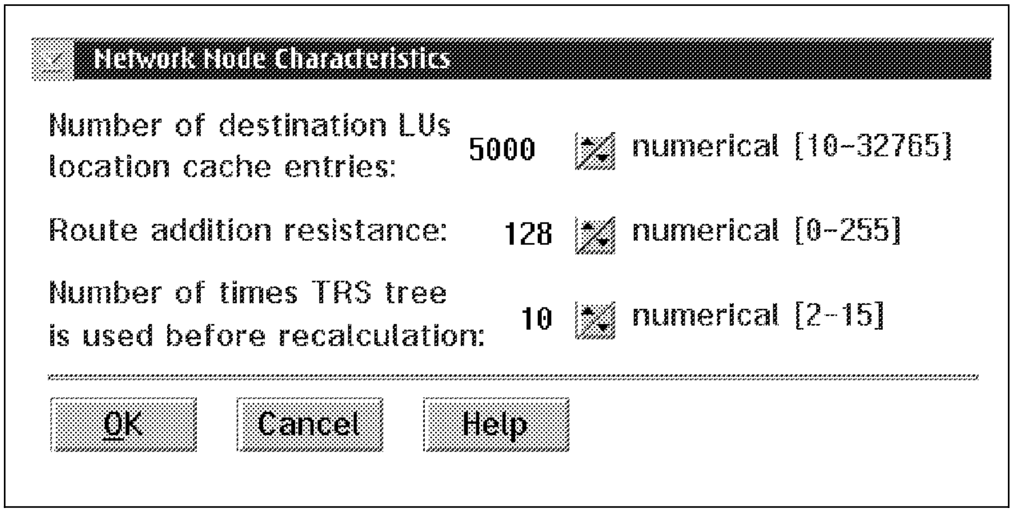

Figure 87. 3746 NN Characteristics

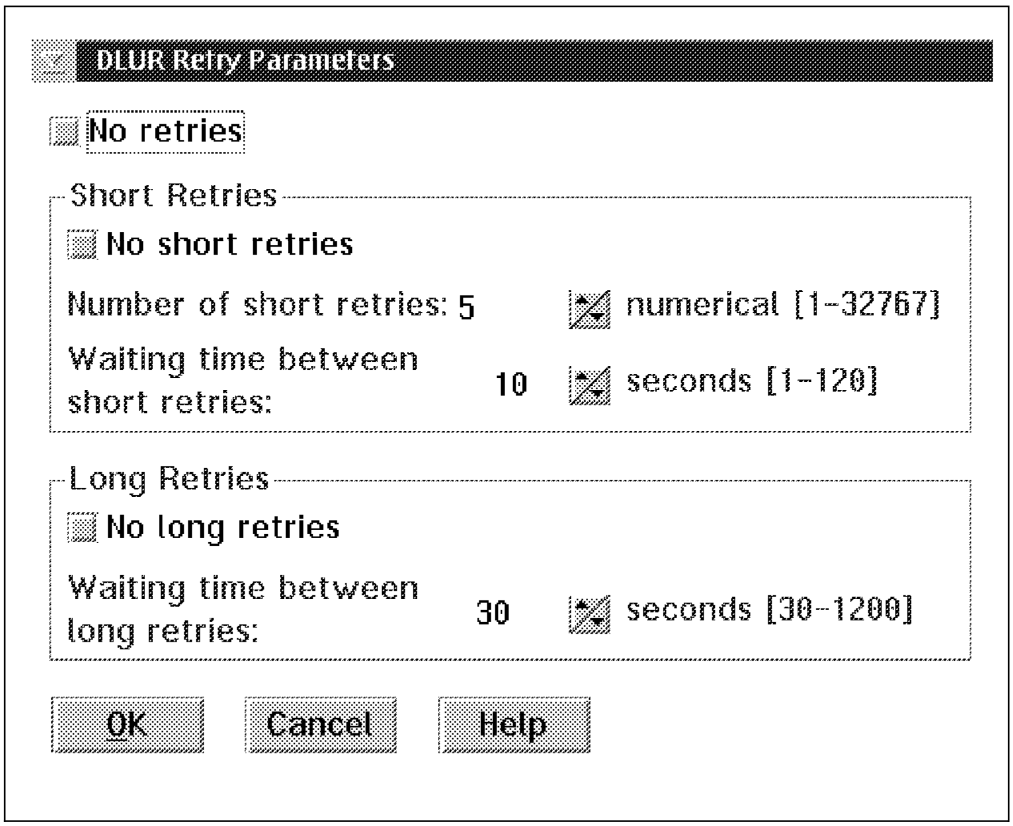

Figure 88. 3746 DLUR Retry Parameters

You can also define the RTP parameters such as the path switch timers, by means of the **RTP Parameters** button. Figure 89 on page 118 shows the choices you have.

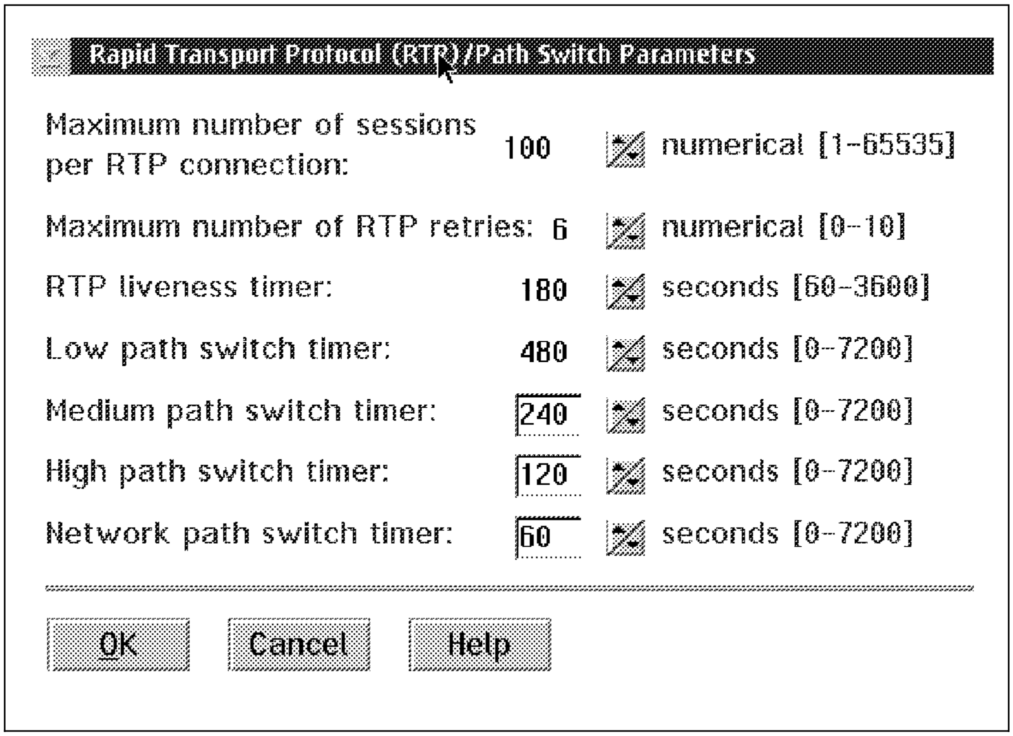

Figure 89. RTP Parameters

8.7 Configuring an ESCON Connection

In this section we detail how we configured our ESCON connection for the 3746 Model 61A. The same steps were done for the 3746 Model 41A, so we omit them here. We show how on a single ESCON port, one host link with one link station can be defined.

We identify the required configuration information in the following series of windows:

- Coupler type
- Environment parameters
- Port configuration
  - APPN parameters
  - DLC parameters
- Host link configuration
- Station configuration
  - APPN parameters
  - DLC parameters

The ESCON configuration has three levels of hierarchy rather than the two normally associated with APPN links. The physical port can have multiple host links, which are effectively logical ports connecting the 3746 with different hosts. Each host link can have multiple link stations for various purposes (for example, APPN parallel TGs and NCP connections).

Coupler type, environment, and port configuration parameters are only entered once per ESCON port. The host link configuration needs to be entered for each host link and the station configuration parameters repeated for each link station on each of the host links.

In the following sections the relevant configuration screens to define an ESCON attachment to the 3746 NN are shown. Only a limited number of the parameter fields are discussed. For an overview of all parameters refer to Table 2 on page 313 or the CCM help screens. The configuration depicted in Figure 90 shows a single ESCON connection between the 3746 NN and our VTAM RAA. RAA, along with our other VTAM hosts RA39 and RAK, runs in a virtual machine under VM/ESA.

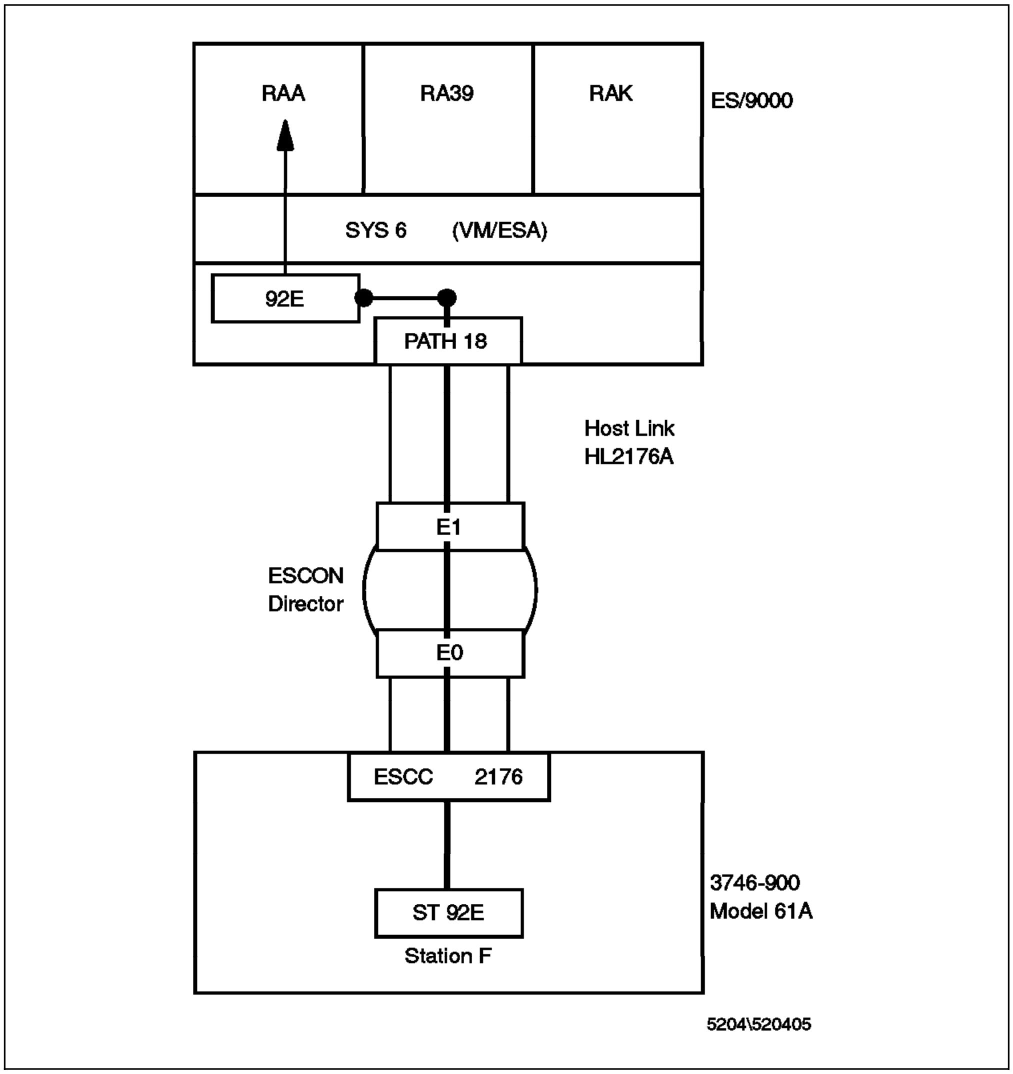

Figure 90. 3746 NN ESCON Connection

Port 2176 contains an ESCON coupler that is connected to an ESCON director. Without an ESCON director each coupler is permanently connected to one host by one piece of fiber optic cable, but an ESCON director can dynamically switch host links (logical ports) between physical cables. This gives ESCON the appearance of a LAN rather than a series of point-to-point connections, so you can define multiple host links for connection to different hosts.

The ESCON host link corresponds to the CHPID on the host, whereas the link station corresponds to the channel address.

We needed only one connection (to RAA), so we defined a single host link with a single link station. The host link was CHPID 18 and the link station was address 92E on RAA.

The names assigned in the 3746 were APPN2176 for the port, HL2176A for the host link and ST92E for the link station.

In the next section we show how port 2176, its host link and the link station have been defined on the 3746 NN.

8.7.1 Coupler Type

Starting from the primary CCM panel (Figure 82 on page 113), we selected coupler 2176, then **Configuration**. As shown in Figure 91, we chose ESCON for the coupler type.

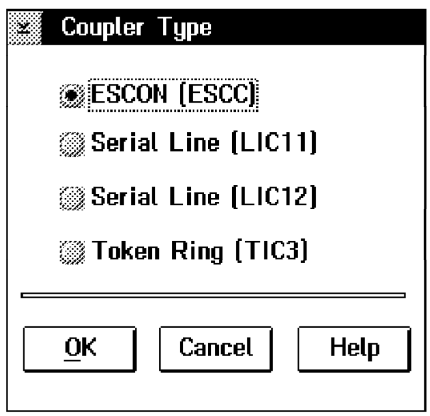

Figure 91. Coupler Type

Once a coupler has been configured for ESCON, this window will be skipped the next time the ESCON port, host links, or link stations are customized.

8.7.2 Port Configuration

Once we clicked on OK we were presented with the Port Configuration panel in Figure 92 on page 121. This panel contains fields in which you can specify configuration parameters for the ESCON port and the ESCON Director (ESCD). The ESCD number must correspond to the definitions in the ESCON Director, whereas the Control Unit Link Address must correspond to the LINK keyword on the CNTLUNIT statement in the IOCP definition. For additional information regarding the parameters, refer to Table 2 on page 313.

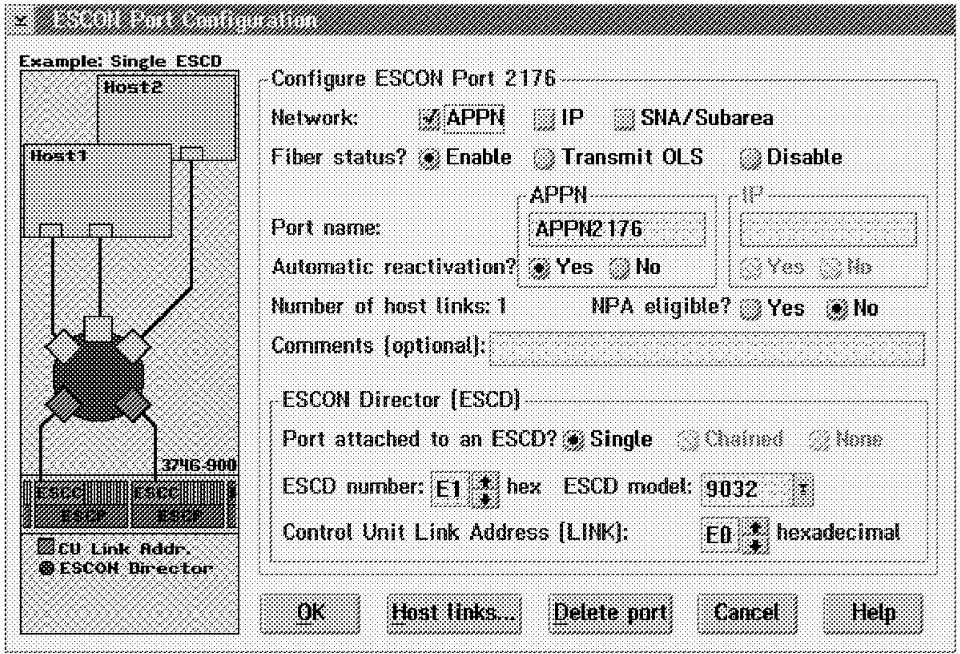

Figure 92. ESCON Port Configuration

8.7.3 Host Link Configuration

On the ESCON Port Configuration panel we selected **Host links** and used Figure 95 on page 124 to configure a host link for the ESCON port.

The ESCON coupler can be configured in one of six ways:

1. Basic mode, no ESCON Director. Each ESCON coupler is connected to one host, thus there is only one host link and one CHPID per port. Multiple link stations can be defined, each of which corresponds to a host channel address.

2. LPAR mode, no ESCON Director. Each ESCON coupler is connected to one LPAR in one host. There is one host link and one CHPID per port. Multiple link stations can be defined, each of which corresponds to a host channel address.

3. EMIF mode, no ESCON Director. Each ESCON coupler is connected to one host. Multiple host links can be defined, each of which is connected to a different LPAR. There is a single CHPID on each host which supports all the LPARs. Multiple link stations can be defined on each host link, each of which corresponds to a host channel address.

4. Basic mode, ESCON Director. Each ESCON coupler can be connected to multiple hosts using a host link each. There is one CHPID per host. Multiple link stations can be defined on each host; and each link station has its own channel address.

5. LPAR mode, ESCON Director. Each ESCON coupler can be connected to multiple LPARs using a host link each. There is one CHPID per LPAR. Multiple link stations can be defined on each LPAR, each link station having one channel address.

6. EMIF mode, ESCON Director. Each ESCON coupler can be connected to multiple LPARs using one host link per physical host. There is one CHPID

per host, which is shared between the host links connected to the LPARs on that host. Multiple link stations can be defined on each LPAR, each link station having its own channel address.

Figure 93 illustrates the ESCON configurations available without an ESCON Director. Figure 94 on page 123 shows how an ESCON coupler can be configured for the three corresponding options with an ESCON Director.

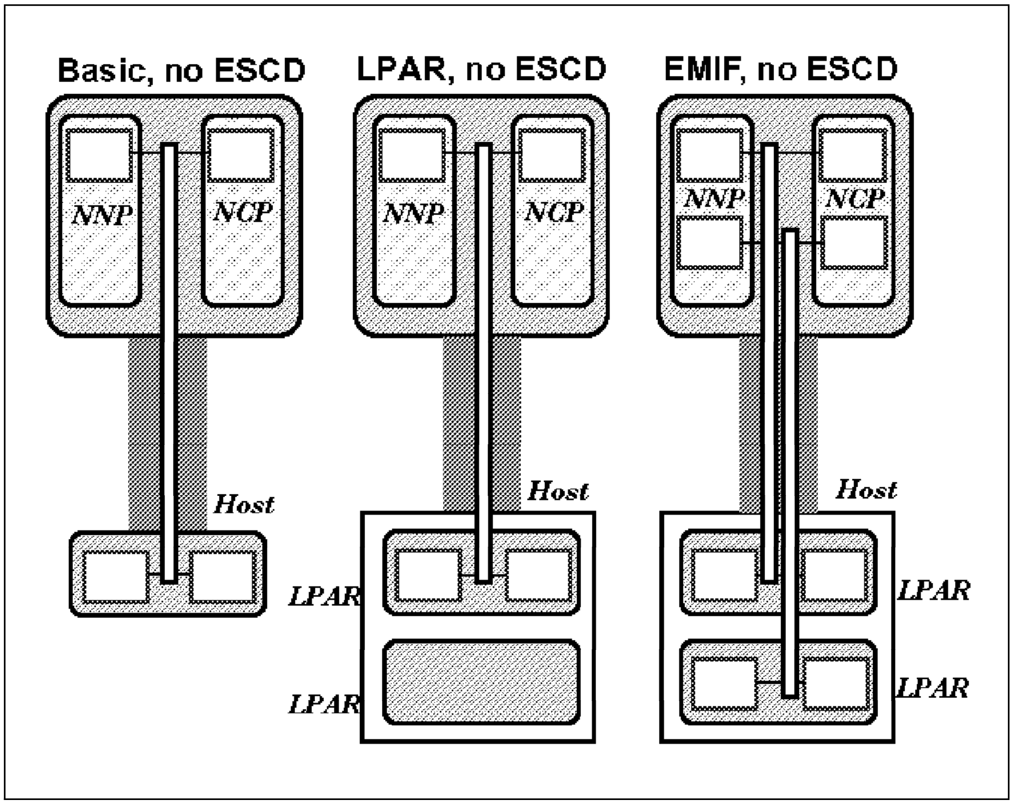

Figure 93. ESCON Configurations, No Director

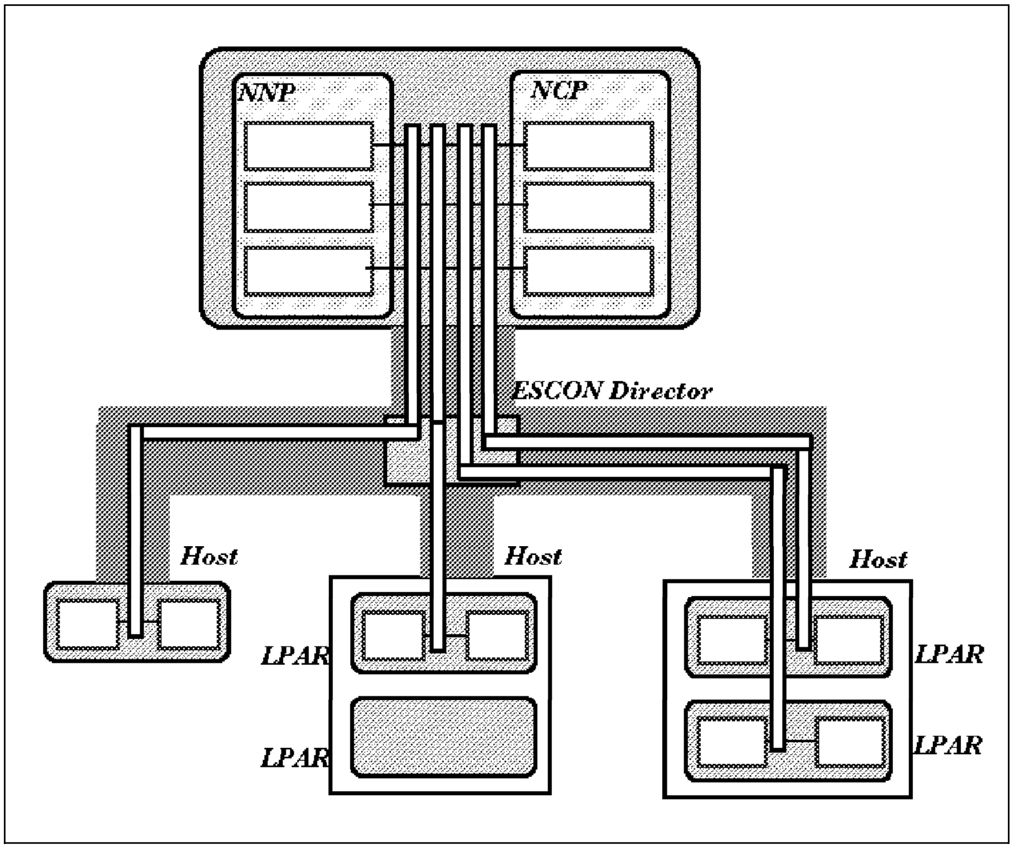

Figure 94. ESCON Configurations, with Director

The critical fields that we had to define correctly on the ESCON panels were:

- We needed APPN but not TCP/IP.
- The ESCON connection was running in basic mode (no LPARs, no EMIF).
- The CHPID must correspond to the PATH keyword on the CHPID statement in the IOCP definition.
- The host link address can be safely left to the ESCON subsystem to be defined dynamically, although we chose to enter it manually. Only if you plan to IPL an NCP over this connection must you specify a real address. This field defines the port in the ESCON director to which the host is connected.

  Note that, in this example, the ESCD number is the same as the host link address. This is purely coincidental.

We selected the **Add** option to complete this panel. Once a host link is defined (and selected in the Already Configured window) then you can proceed to define APPN parameters and link stations. Table 2 on page 313 gives more details on the parameters on this panel.

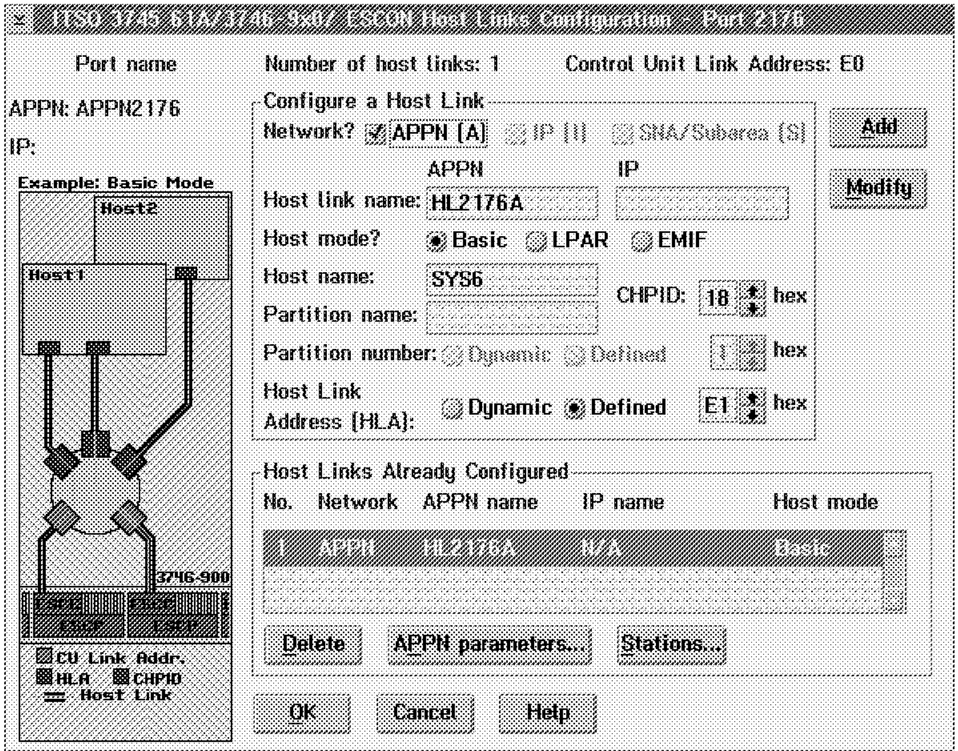

Figure 95. ESCON Host Links Configuration

We clicked on **APPN parameters** to update the additional ESCON port parameters as shown in Figure 96 on page 125. Note in particular:

- The TG characteristics can be specified here, and overridden if necessary on the link station definition.
- The maximum PIU sizes are the values used during Route Setup to compute the maximum permitted NLP size.
- The only HPR Support values permitted are ERP Required and No HPR. ESCON channels always require link-level error recovery. Once again, these values sift down to the link stations unless they are overridden.

We took the default for ERP support (required), enabling HPR support over the channel.

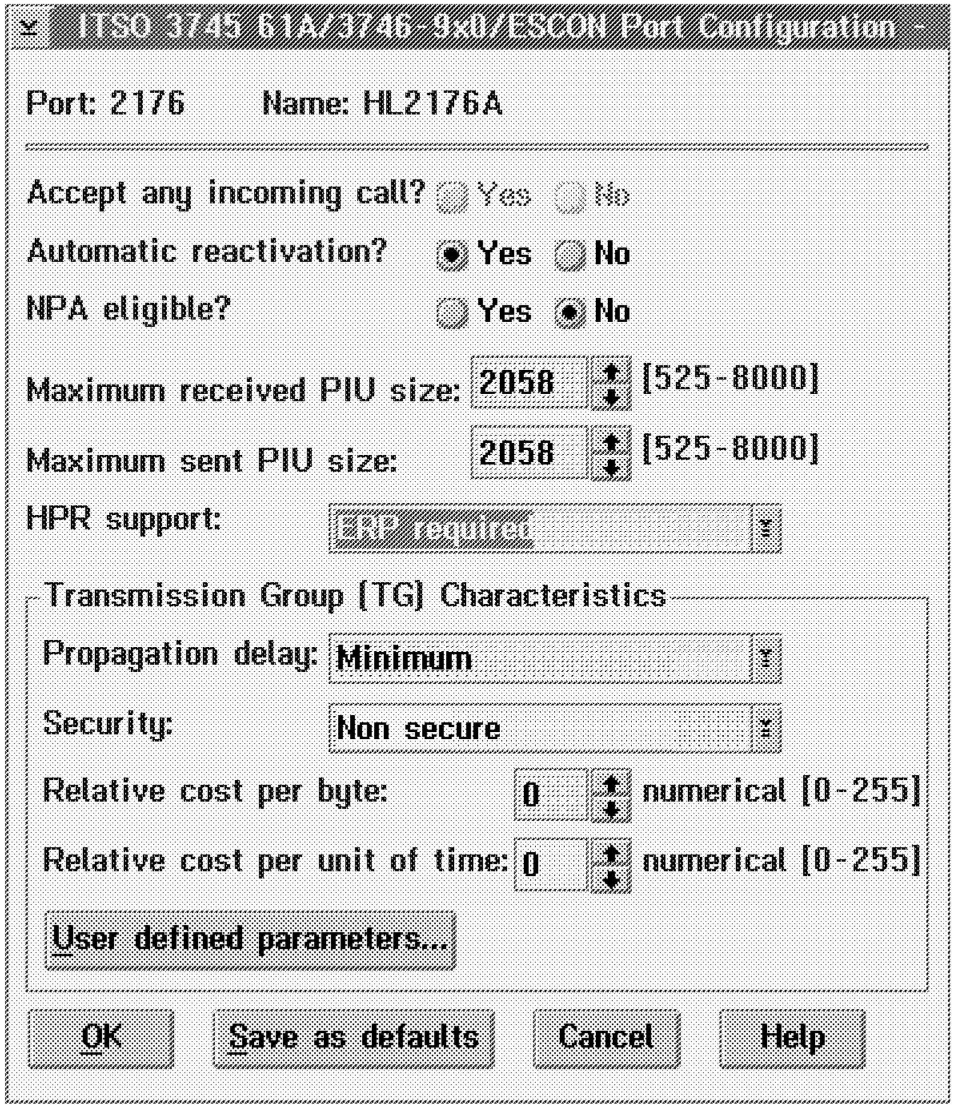

Figure 96. Port Configuration - APPN Parameters

8.7.4 Link Station Configuration

Returning to the ESCON Host Links Configuration panel, we clicked on **Stations** to get to the screen shown in Figure 97 on page 126. This window is used to define each individual ESCON station and will either have an APPN or SNA/subarea type of connection. Because we specified APPN in Figure 95 on page 124, only APPN is highlighted. On the host link HL2176A we only defined one link station. The address (F) of this station must correspond to the UNITADD keyword on the IODEVICE statement in the IOCP definition.

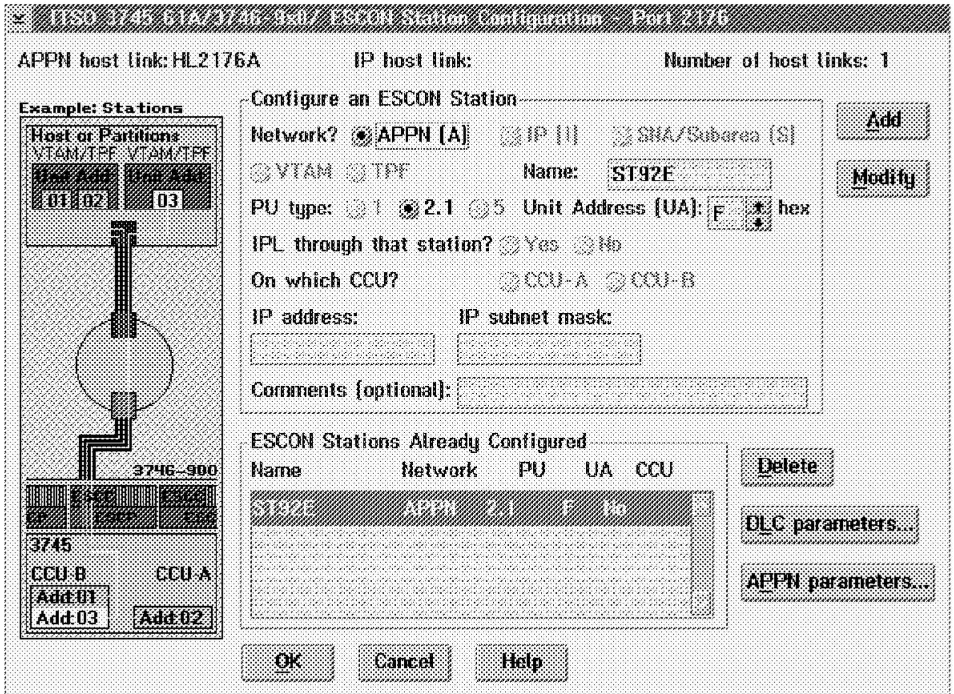

Figure 97. ESCON Station Configuration

8.7.5 Station - APPN Parameters

We selected the station we had just added and then chose **APPN Parameters** as seen in Figure 98 on page 127. Here, aside from the HPR and TG characteristics data, you have the choice of defining an MLTG on this connection. You can also override the global DLUR parameters for the PU that this link station represents. The DLUR parameters are not available for selection or modification on ESCON stations, since such stations never support dependent resources using peripheral subarea protocols.

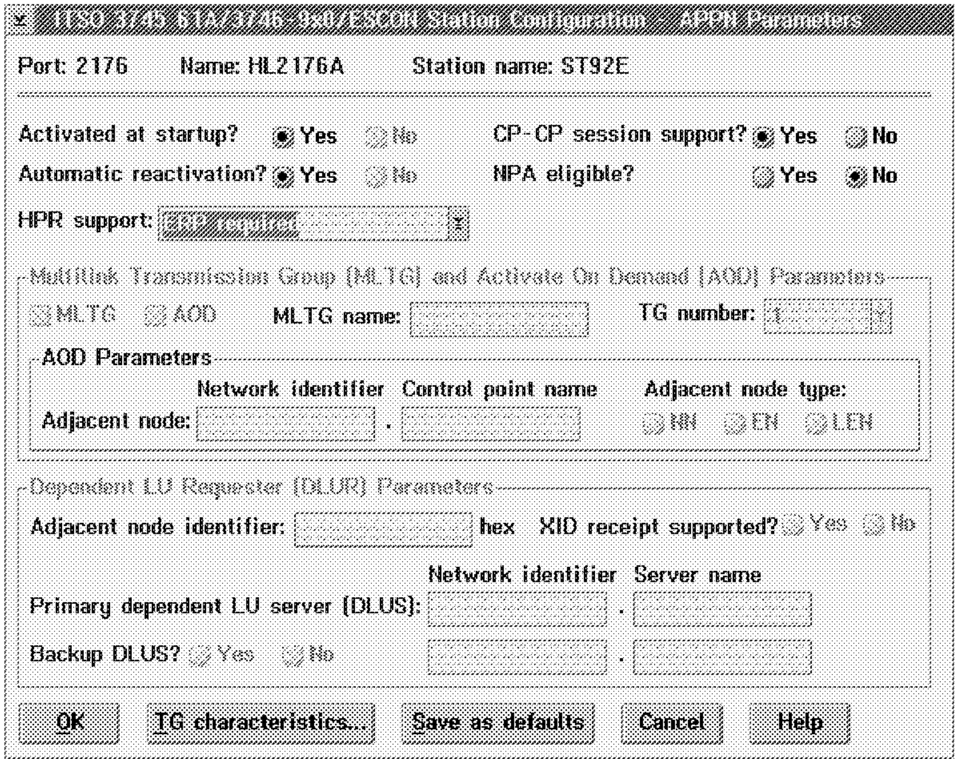

Figure 98. Station Configuration - APPN Parameters

8.7.6 IOCP Definition for ESCON Channel

The CCM application also gives you some guidelines for coding the IOCP definitions in your host. From the CCM main panel, select **Options** and then **View** in the related pull-down menu. You can select **IOCP** to see an example of IOCP definitions consistent with what you defined in the 3746 ESCON definition panels.

Figure 99 shows our own IOCP configuration.

```
*********************************************************************
* IOCP ESCON GENERATION SUBSET: Multipro
*                    FOR HOST: SYS6
*     COMMUNICATION CONTROLLER: 3745-61A
*********************************************************************
         CHPID PATH=(18),SWITCH=E1,TYPE=CNC
         CNTLUNIT CUNUMBR=92E,                                         X
               UNITADD=((01,16)),                                      X
               PATH=(18),                                              X
               LINK=E0,                                                X
               UNIT=3745
         IODEVICE CUNUMBR=92E,         * Same as CNTLUNIT CUNUMBR *    X
               UNIT=3745,                                              X
               ADAPTER=TYPE7,                                          X
               UNITADD=0F,                                             X
               ADDRESS=(92E,2) *xxx must match the CUADDR of a PCCU
*                              * in NCP , when requested
```

Figure 99. CCM IOCP File Updated for RAA

Note the following:

- The SWITCH keyword identifies the ESCON director to which this device is connected. This may not be the same ESCD as the one to which the 3746 is connected, since they may be cascaded.
- The PATH keyword corresponds to the CHPID definition in the CCM panels.
- The LINK keyword corresponds to the Control Unit Link Address definition in the CCM panels.
- The UNITADD keyword corresponds to the station address in the CCM panels. This links the definitions to the channel address known by VTAM.

8.7.7 Station - DLC Parameters

The final panel in the ESCON definitions is the DLC Parameters panel as shown in Figure 100. This contains tuning parameters which we left alone. Please see Table 2 on page 313 for an explanation of the fields.

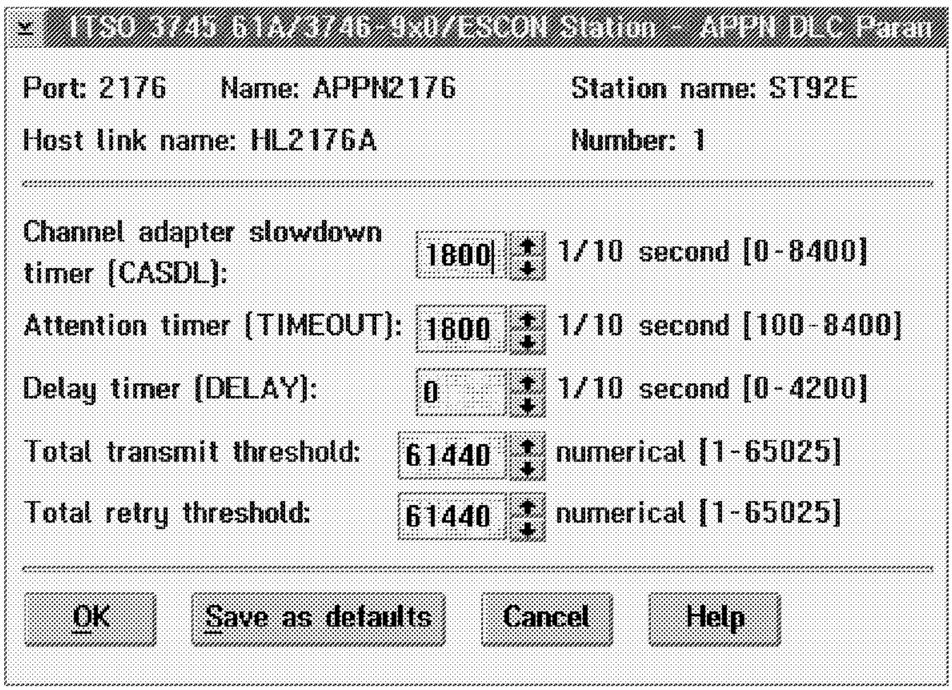

Figure 100. ESCON Station - DLC Parameters

8.7.8 VTAM Definitions

Lastly, we created some VTAM definitions for our ESCON connection. These are quite straightforward. We just defined one local SNA major node for each 3746, as shown in Figure 101 on page 129 (used in RAA for the 3745-61A) and Figure 102 on page 129 (used in RA39 for the 3745-41A).

```
**********************************************************************
*   LOCAL    MAJOR NODE FOR CP-CP LINK TO 3746-900 NNP61A            *
**********************************************************************
CP90061A VBUILD TYPE=LOCAL
**
**     USING 92E TO CONNECT TO RAA.
**
CP900PU1 PU    CUADDR=92E,XID=YES,                                    X
               CPCP=YES,MAXBFRU=15,                                   X
               CONNTYPE=APPN,                                         X
               PUTYPE=2,DYNLU=YES,DYNADJCP=YES,                       X
               ISTATUS=ACTIVE
```

Figure 101. Local Major Node on RAA for NNP61A

```
**********************************************************************
*   LOCAL MAJOR NODE FROM RA39 TO 3746 NNP41A                        *
**********************************************************************
CP900D   VBUILD TYPE=LOCAL
**
*-------CHANNEL PU
**
CP9DD41A PU    PUTYPE=2,CUADDR=90F,XID=YES,                           X
               CONNTYPE=APPN,MAXBFRU=15,                              X
               CPCP=YES
```

Figure 102. Local Major Node on RA39 for NNP41A

Note that we needed to code XID=YES because the default is NO on a leased connection. We have shown DYNLU=YES in one of these examples; however, it is normally specified as a VTAM start option. We allowed HPR= to take the defaults (full RTP support) from the VTAM start options.

Note also that the CUADDR keyword in the VTAM PU definition must correspond with the ADDRESS keyword in the IODEVICE statement in the IOCP. The CUNUMBR keywords in the IOCP are used to match the CNTLUNIT to the IODEVICE; they have *no* relationship with the CUADDR and it is merely a coincidence that they are the same.

8.8 Configuring a Token-Ring Connection

In this section we show an example of how to configure a token-ring connection for the 3746 NN. Using CCM we entered our configuration information into the following series of windows:

- Coupler type
- Port configuration:
  - APPN parameters
  - Default station parameters
- Station configuration:
  - APPN parameters
  - DLC parameters

Station configuration is only required for connections established by the 3746 NN (dial-out) or for attachments where the default station parameters do not apply.

- Connection network

The coupler type needs to be entered only once per coupler. The port configuration parameters are required for each token-ring (TIC3) port. Default token-ring station parameters can be used, or they can be individually configured for each station. Stations must be defined for dial-out. Dynamically defined link stations automatically use the default station parameters. In cases where non-default parameters are to be used, a station must be predefined, and the appropriate parameters set. Multiple connection networks can be defined per token-ring port; connection networks can span a single or multiple ports.

In the following sections the relevant configuration screens to define a token-ring attachment to the 3746 NN are shown. Only a limited number of the parameter fields will be discussed. For an overview of all parameters refer to Table 3 on page 314 or the CCM help screens.

8.8.1 Coupler Type

We started by opening our configuration file and selecting coupler address 2144. The Coupler Type panel (Figure 91 on page 120) appeared, but the only available choice was to click **OK** because the token-ring coupler had been selected. On the 61A and 41A models of 3745, the coupler at address 2144 can only be a token-ring coupler.

8.8.2 Port Configuration

On the Token-Ring Port Configuration panel we specified that our port name was APPN2144. As shown in Figure 103 on page 131, we filled in the local MAC address and our local SAP. These must correspond with the definitions on remote stations that will initiate a connection to this port. For additional information regarding the parameters, refer to Table 3 on page 314.

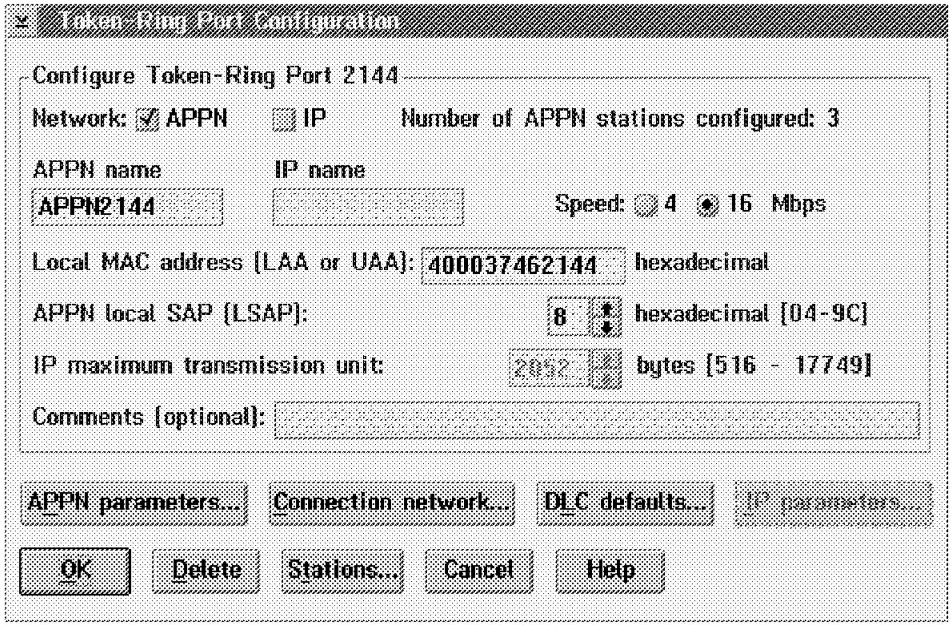

Figure 103. Token-Ring Port Configuration

By clicking on **APPN parameters** we were presented with the panel in Figure 104 on page 132. This panel is very similar to the ESCON Host Link parameters panel (Figure 96 on page 125), except that the TG characteristics are more suitable to a LAN connection and there are more choices in the HPR scroll bar. Figure 105 on page 133 shows that link-level ERP can be prohibited or allowed to be negotiated with the adjacent node. You must make sure that the adjacent node definitions are consistent. We allowed all the parameters to remain at their default values.

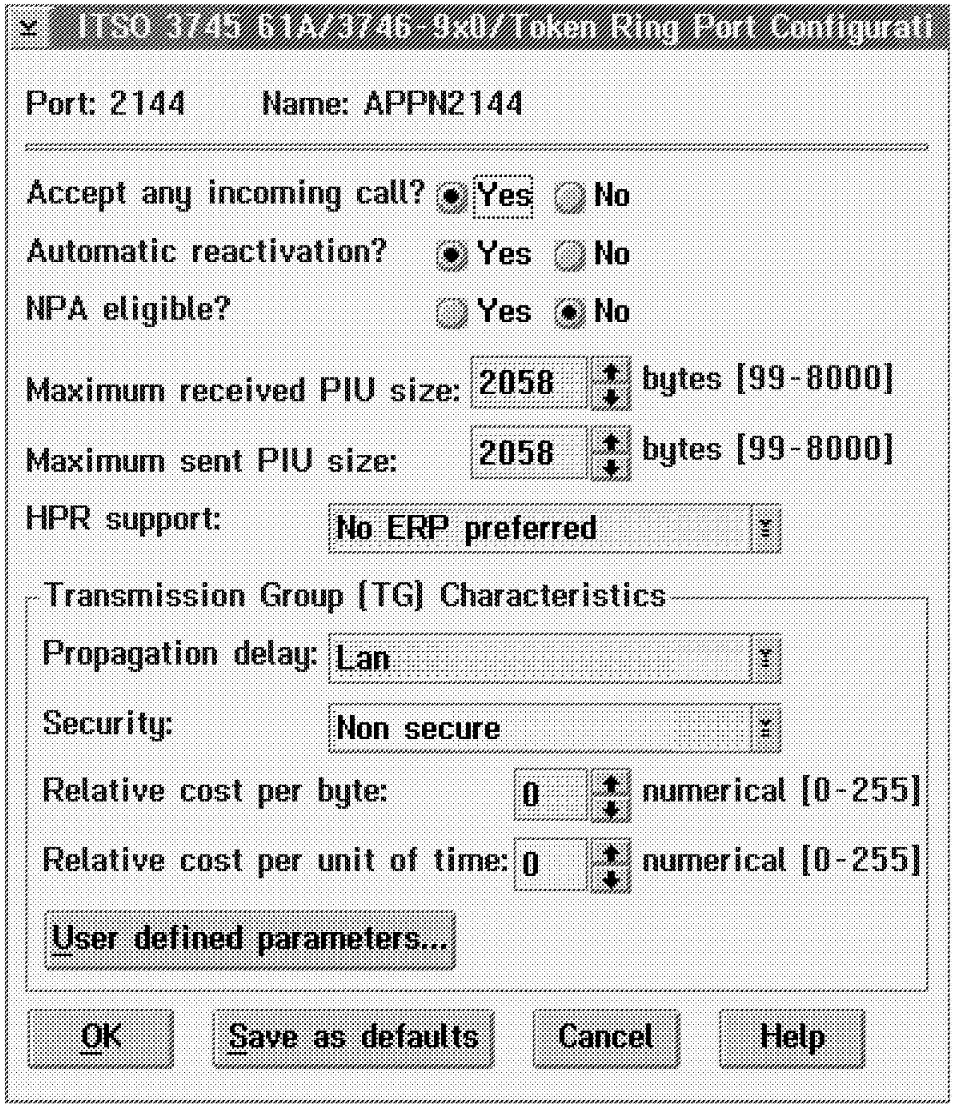

Figure 104. Port Configuration - APPN Parameters

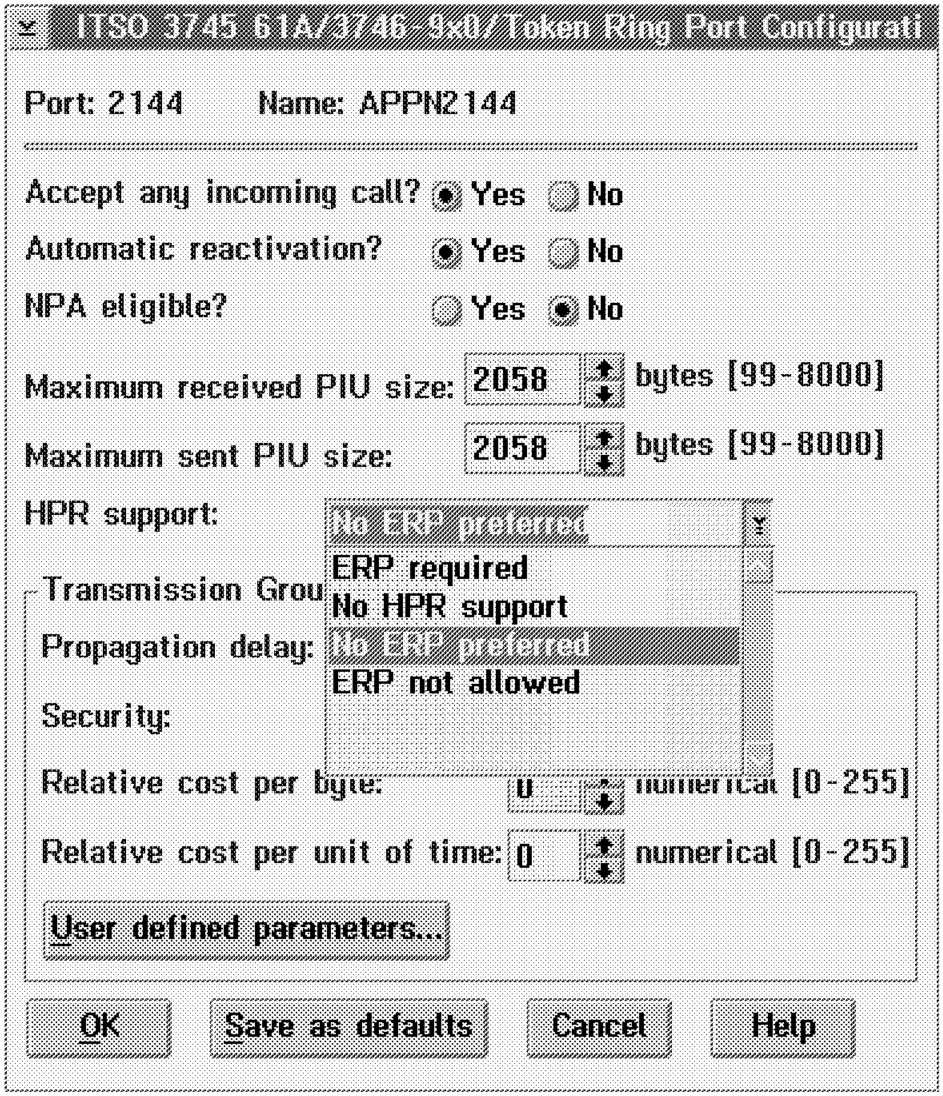

Figure 105. HPR Support on 3746 Token-Ring Lines

8.8.3 Default Station Parameters

Returning to our initial Port Configuration panel, we clicked on **DLC defaults** and checked the values for the timers, MAXOUT, MAXIN, and the retries as shown in Figure 106 on page 134. The values shown in Figure 106 on page 134 are in fact the supplied defaults. Whatever changes you make will be used in all subsequent station definitions on this port, unless specifically overridden. For additional information regarding the parameters, refer to Table 3 on page 314.

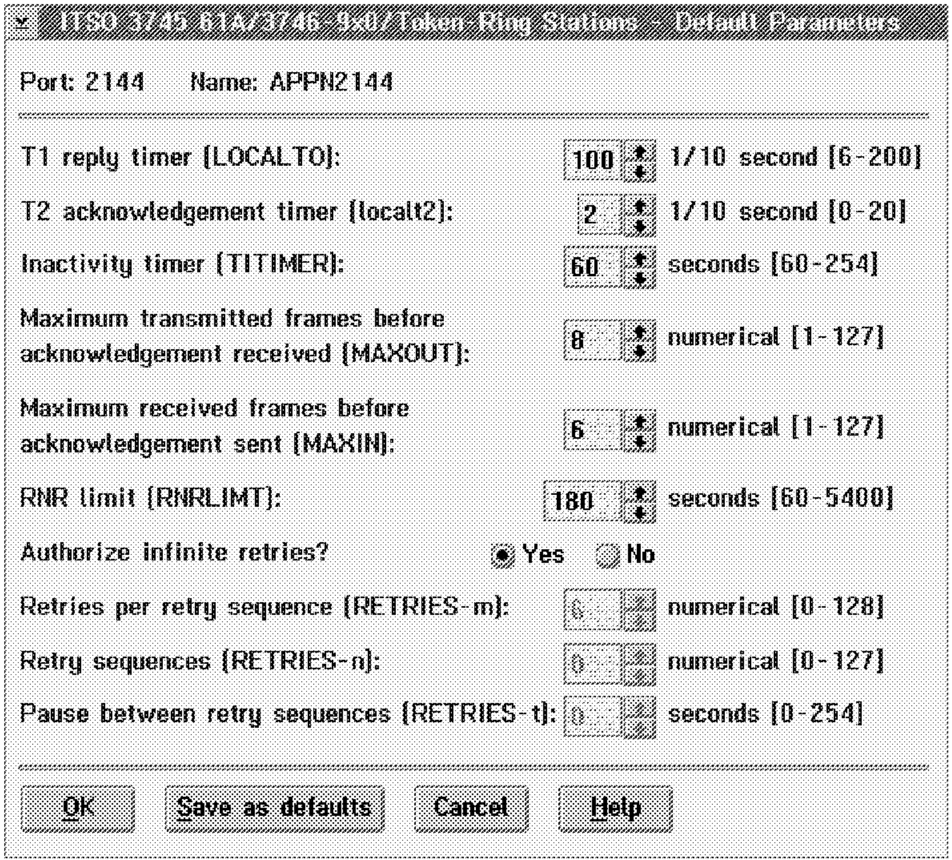

Figure 106. Token-Ring Stations - Default Parameters

We did not define a connection network because we only ever had a small number of nodes on our LAN, mostly network nodes. The Connection Network panel accessible from the Port Configuration panel (Figure 103 on page 131) prompts you for the fully qualified names of all the virtual nodes you wish to define on this port.

8.8.4 Station Configuration

Once the port definitions were set, we once again returned to the Port Configuration screen and selected **Stations**. As shown in Figure 107 on page 135, the station configuration is used to add stations to the token-ring port. You fill in all the fields and click **Add** to define each station. For our connection to the 3746 Model 41A we added station P2144AP.

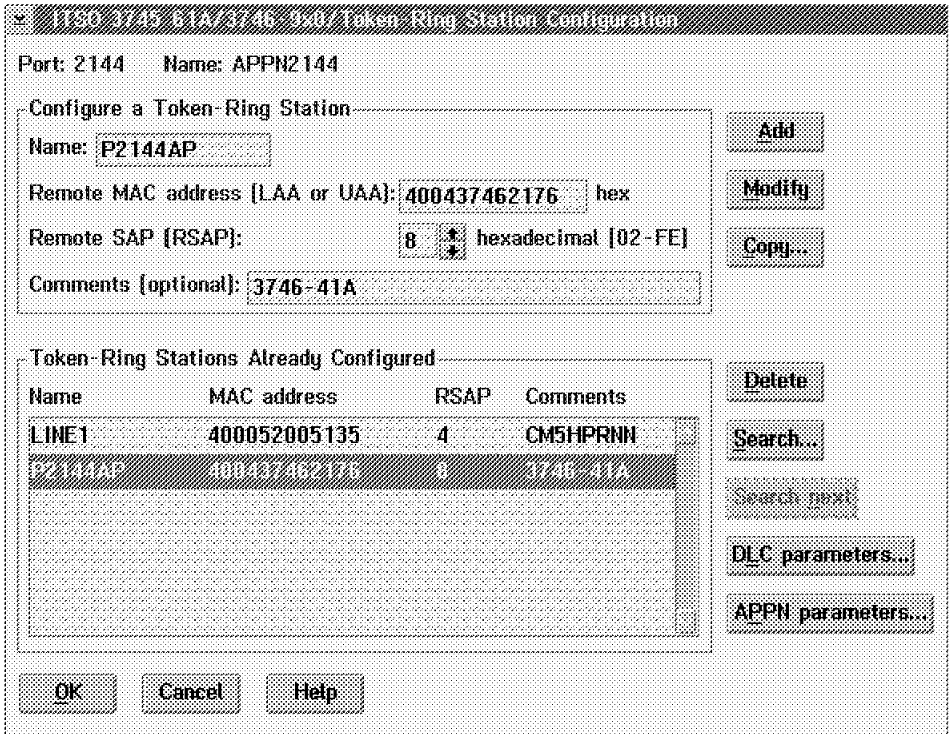

Figure 107. NNP41A Token-Ring Station Configuration on NNP61A

Note: The remote SAP defaults to 08. You have to make sure that the MAC address and the SAP correspond with the remote station.

The LAN-attached CS/2 node (PU05170) was not explicitly defined as a token-ring station in the 3746-61A in the first test, so the 3746 dynamically created a link station when the CS/2 node connected to it. We just defined a link in CS/2 pointing to the 3746 MAC/SAP address.

We defined in the 3746 the network node used for the second test, CM5HPRNN. As stated previously, the link station can be defined at either partner node. So, if you are defining the link station at the 3746 NN you do not have to define a link station at the PC and vice versa. The side that does not have explicit definitions for its partner cannot initiate the connection itself and must wait for the other side to do so.

Note: If the link station in CS/2 is dynamically defined, care must be taken not to specify a value of zero in the Percent of Incoming Calls field in the DLC Adapters definition.

8.8.5 Station Parameters

Once the station was added, we selected **APPN parameters** and verified the appropriate parameters as shown in Figure 108 on page 136. This is very similar to the equivalent ESCON panel (Figure 98 on page 127). Note the Activate On Demand (AOD) fields where you can define auto-active connections. Because such connections must be registered to the APPN topology database before activation, you must pre-define the adjacent node name and node type.

The panel also allows you to enter DLUR parameters. This information is applicable only if the remote station contains dependent LUs, and the default parameters in Figure 85 on page 115 are not appropriate. For additional information regarding the parameters, refer to Table 3 on page 314.

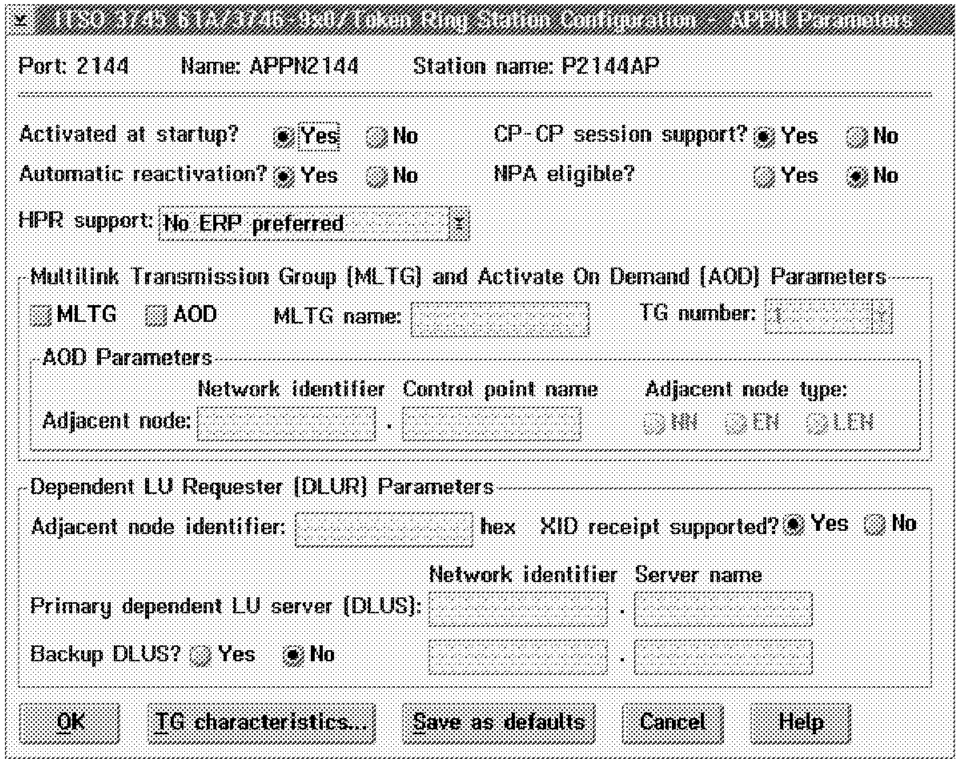

Figure 108. Token-Ring Configuration - APPN Parameters

Once the station APPN parameters were set we selected the **DLC Parameters** to verify them. We left them alone because the defaults (shown in Figure 106 on page 134) were acceptable.

8.9 Example with 3746 As DLUR Node

The first configuration we implemented had the DLUR function within the 3746s, so that the workstations on the LAN were acting as peripheral subarea nodes. Thus VTAM's RTP partner was always one or other 3746. Please see Figure 78 on page 107 for the network diagram.

8.9.1 Activate VTAM-to-VTAM Connection

We first established an ANNC (MPC+) connection between RAA and RA39. Figure 109 shows the messages that appeared on RA39.

```
V NET,ID=RACAHHC,SCOPE=ALL,ACT
IST097I VARY ACCEPTED
IST093I RACAHHC ACTIVE
IST1086I APPN CONNECTION FOR USIBMRA.RAA IS ACTIVE - TGN = 21
IST093I RACHRAA ACTIVE
IST1488I  ACTIVATION   FOR  RTP CNR00001 AS PASSIVE PARTNER COMPLETED   1
IST1096I CP-CP SESSIONS WITH USIBMRA.RAA ACTIVATED
```

Figure 109. VTAM-to-VTAM HPDT Connection

Note **1** that the CP-CP RTP connection was established before the CP-CP sessions were activated. This MPC connection supports control flows over RTP. The corresponding PU name given to the CP-CP pipe on RAA was CNR00761.

A display of ISTRTPMN confirmed that the RTP pipe was indeed for CP-CP sessions. There was as yet no LU-LU pipe because there were no LU-LU sessions between the VTAMs. Nor was there an RSTP pipe because there were no LU-LU sessions using this link (please see Figure 110).

```
D NET,ID=ISTRTPMN,E
IST097I DISPLAY ACCEPTED
IST075I NAME = ISTRTPMN, TYPE = RTP MAJOR NODE 584
IST486I STATUS= ACTIV, DESIRED STATE= ACTIV
IST1486I   RTP NAME   STATE             DESTINATION CP     MNPS  TYPE
IST1487I   CNR00001   CONNECTED         USIBMRA.RAA        NO    CPCP
IST314I END
```

Figure 110. RTP Major Node on RA39

A detailed display of the CNR00001 connection (Figure 111) showed that:

- The APPN COS was CPSVCMG **2**, which is reserved for CP-CP sessions.
- The NCE used by VTAM for CP-CP RTP connections starts with X′D4′ **3**, as opposed to the X′D0′ for normal LU-LU pipes. RSTP pipes use X′D2′.
- The ARB flow control algorithm had been initialized to reasonable values **4**, indicating that VTAM can select a suitable CAPACITY for its channel connections.
- The partner CP **5** was the only LU using this pipe.

```
D NET,ID=CNR00001,E
IST097I DISPLAY ACCEPTED
IST075I NAME = CNR00001, TYPE = PU_T2.1 587
IST1392I DISCNTIM = 00010 DEFINED AT PU FOR DISCONNECT
IST486I STATUS= ACTIV--LX-, DESIRED STATE= ACTIV
IST1043I CP NAME = RAA, CP NETID = USIBMRA, DYNAMIC LU = YES
IST1589I XNETALS = YES
IST933I LOGMODE=***NA***, COS=CPSVCMG   2
IST1476I TCID X'38273EAD0000009D' - REMOTE TCID X'310BB5170000009C'
IST1481I DESTINATION CP USIBMRA.RAA - NCE X'D400000000000000'  3
IST1587I ORIGIN NCE X'D400000000000000'   3
IST1477I ALLOWED DATA FLOW RATE = 1876 KBITS/SEC   4
IST1516I INITIAL DATA FLOW RATE = 3200 KBITS/SEC
IST1511I MAXIMUM NETWORK LAYER PACKET SIZE = 20479 BYTES
IST1478I NUMBER OF UNACKNOWLEDGED BUFFERS = 0
IST1479I RTP CONNECTION STATE = CONNECTED - MNPS = NO
IST1480I RTP END TO END ROUTE - PHYSICAL PATH
IST1460I TGN  CPNAME             TG TYPE      HPR
IST1461I  21  USIBMRA.RAA        APPN         RTP
IST231I RTP MAJOR NODE = ISTRTPMN
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST355I LOGICAL UNITS:
IST080I RAA      ACT/S----Y   5
IST314I END
```

Figure 111. Details of CP-CP RTP Pipe

8.9.2 3746 Activation

We now activated the connections to the 3746 NNs: the link to NNP61A from RAA and the link to NNP41A from RA39. Figure 112 shows what happened on RAA.

```
V NET,ACT,ID=NNP61A
IST097I VARY ACCEPTED
IST1488I  ACTIVATION   FOR RTP CNR0076F AS ACTIVE  PARTNER COMPLETED
IST1096I CP-CP SESSIONS WITH USIBMRA.NNP61A ACTIVATED
```

Figure 112. Activation of Link to 3746 NN

Once again, the HPR connection CNR0076F was activated before the CP-CP sessions were established. Both VTAM and the 3746 support Control Flows over RTP over a channel connection.

A display of the RTP connection showed (Figure 113) very similar characteristics to those of the VTAM-VTAM connection. Note that the 3746 uses different formats for TCIDs and NCEs.

```
D NET,ID=CNR0076F,E
IST097I DISPLAY ACCEPTED
IST075I NAME = CNR0076F, TYPE = PU_T2.1 393
IST1392I DISCNTIM = 00010 DEFINED AT PU FOR DISCONNECT
IST486I STATUS= ACTIV--LX-, DESIRED STATE= ACTIV
IST1043I CP NAME = NNP61A, CP NETID = USIBMRA, DYNAMIC LU = YES
IST1589I XNETALS = YES
IST933I LOGMODE=***NA***, COS=CPSVCMG
IST1476I TCID X'310BB525000000AA' - REMOTE TCID X'000000000B447050'
IST1481I DESTINATION CP USIBMRA.NNP61A - NCE X'D0201025'
IST1587I ORIGIN NCE X'D400000000000000'
IST1477I ALLOWED DATA FLOW RATE = 1876 KBITS/SEC
IST1516I INITIAL DATA FLOW RATE = 3200 KBITS/SEC
IST1511I MAXIMUM NETWORK LAYER PACKET SIZE = 2074 BYTES
IST1478I NUMBER OF UNACKNOWLEDGED BUFFERS = 0
IST1479I RTP CONNECTION STATE = CONNECTED - MNPS = NO
IST1480I RTP END TO END ROUTE - PHYSICAL PATH
IST1460I TGN  CPNAME             TG TYPE      HPR
IST1461I  21  USIBMRA.NNP61A     APPN         RTP
IST231I RTP MAJOR NODE = ISTRTPMN
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST355I LOGICAL UNITS:
IST080I NNP61A   ACT/S----Y
IST314I END
```

Figure 113. RTP Pipe to 3746

Displaying the 3746 control point yielded Figure 114 on page 139.

```
D NET,ID=NNP61A,E
IST097I DISPLAY ACCEPTED
IST075I NAME = USIBMRA.NNP61A, TYPE = ADJACENT CP 398
IST486I STATUS= ACT/S----Y, DESIRED STATE= ACTIV
IST1402I SRTIMER =    30  SRCOUNT =    10
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.NNP61A - NETSRVR = ***NA***
IST1044I ALSLIST = ISTAPNPU
IST1131I DEVICE = ILU/CDRSC
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST228I ENCRYPTION = NONE
IST1563I CKEYNAME = NNP61A CKEY = PRIMARY CERTIFY = NO
IST1552I MAC = NONE MACTYPE = NONE
IST171I ACTIVE SESSIONS = 0000000002, SESSION REQUESTS = 0000000000
IST206I SESSIONS:
IST1081I ADJACENT LINK STATION = CNR0076F   6
IST634I NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I RAA      ACTIV/CP-S D3974BDF909413B6 00F4 0001       USIBMRA
IST635I RAA      ACTIV/CP-P F7FF6164529F529D 0001 00F1  0  0 USIBMRA
IST924I -------------------------------------------------------------
IST075I NAME = USIBMRA.NNP61A, TYPE = DIRECTORY ENTRY
IST1186I DIRECTORY ENTRY = DYNAMIC NN
IST1184I CPNAME = USIBMRA.NNP61A - NETSRVR = ***NA***
IST1402I SRTIMER =    30  SRCOUNT =    10
IST314I END
```

Figure 114. Display of 3746 NN CP

Note **6** that both CP-CP sessions used the same RTP pipe. On NNP41A, as it happened, the CP-CP sessions used separate pipes. This depends purely on the timing of the activation requests from the partner nodes. The APPN COS and the route taken for this pipe is always the same.

8.9.3 Displays on the 3746

The management interface of CCM allows you to display and manage the APPN connections to the 3746. In this section we illustrate some of the displays we were able to take on NNP61A. To access this information we selected the **Management** option of the action bar.

To display the link stations known to the 3746, we clicked on **Stations** to see Figure 115 on page 140.

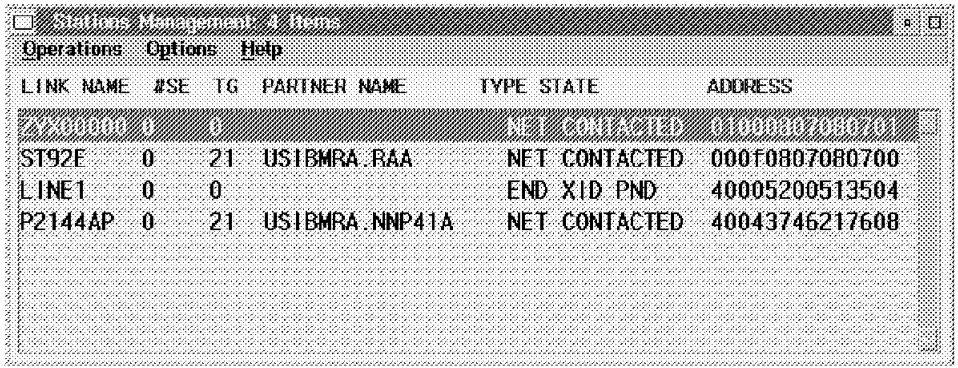

Figure 115. Display of Active Stations on NNP61A

The first connection on the list was not relevant to our tests, but:

- ST92E is the ESCON link station to RAA. RAA is a network node, and is in contact with the 3746.
- LINE1 is the connection we defined to the CS/2 network node. As this station has not been contacted, its details (except for the defined MAC address) are not known.
- P2144AP is the connection we defined to NNP41A. This is also in contact with NNP61A. The MAC address was predefined but the CP name has been discovered by the 3746 during XID exchange. We did not define it.

Now we checked the active HPR connections ending in this NNP by selecting **Management**, then **APPN Specifics** followed by **HPR Connections**. This resulted in the panel in Figure 116.

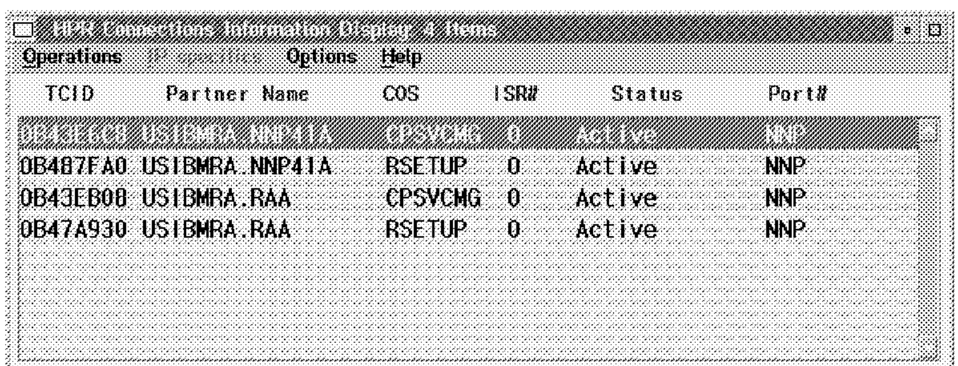

Figure 116. Active HPR Pipes on NNP61A

We can see two CP-CP session pipes (to each adjacent APPN node that supports Control Flows) and two Route Setup pipes (because by this time we had established some LU-LU sessions across those links).

CCM also allows you to verify how many LU 6.2 sessions starting or ending in this NNP are active. From the CCM main panel select **Management** and then **Non Intermediate Sessions** to view the display as in Figure 117 on page 141.

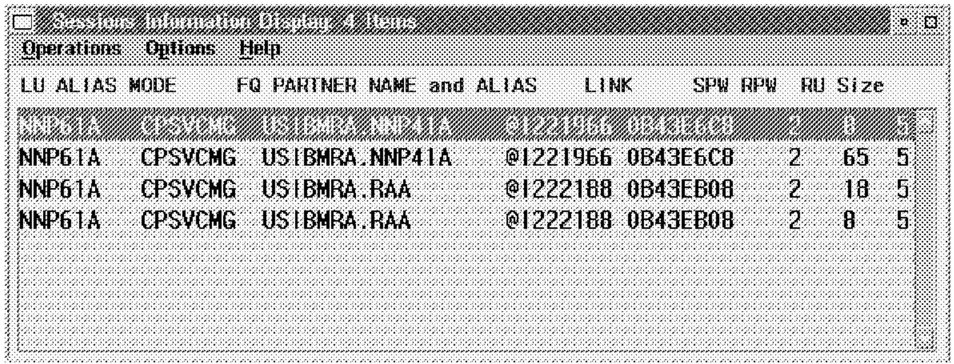

Figure 117. Active LU 6.2 Sessions from NNP61A

The only sessions of which the 3746 was aware were the CP-CP sessions to adjacent nodes. As yet there were no DLUR sessions; the LU-LU sessions that caused the Route Setup pipes to be established are transparent to the 3746, which only performs ANR routing for them.

8.9.4 Activation of Dependent LU Workstation

Next, we started CS/2 on the workstation that was configured with dependent LUs. As soon as we started CS/2 we did some displays from CCM to check what new connections there may have been. See Figure 118 for the refreshed Stations panel.

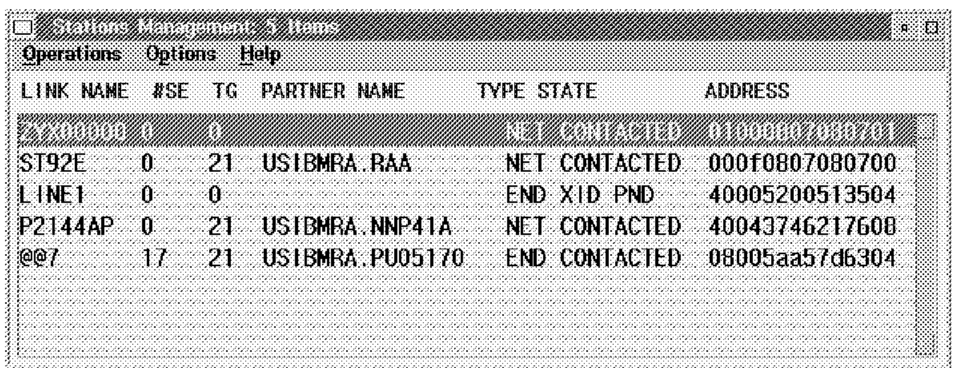

Figure 118. Display of Active Stations on NNP61A

The node PU05170 was not predefined on the 3746, so the link station was implicitly defined as soon as the XIDs were exchanged. We had defined in CS/2 a link to the 3746, giving the TIC3 MAC and SAP address as the destination. PU05170 is in fact an end node with a LEN connection to NNP61A.

Figure 119 on page 142 shows what happened to the active HPR connections.

HPR Connections Information Display: 5 Items

Operations IP specifics Options Help

TCID	Partner Name	COS	ISR#	Status	Port#
0B43E6C8	USIBMRA.NNP41A	CPSVCMG	0	Active	NNP
0B487FA0	USIBMRA.NNP41A	RSETUP	0	Active	NNP
0B43EB08	USIBMRA.RAA	CPSVCMG	0	Active	NNP
0B47A930	USIBMRA.RAA	RSETUP	0	Active	NNP
0B47BA30	USIBMRA.RAA	SNASVCMG	0	Active	NNP

Figure 119. Active HPR Pipes on NNP61A

The last RTP connection on the display, with APPN COS SNASVCMG, has been created to carry the DLUR/S LU 6.2 sessions. Of course, we immediately displayed the Non Intermediate Sessions panel again (Figure 120) to see these.

Sessions Information Display: 6 Items

Operations Options Help

LU ALIAS	MODE	FQ PARTNER NAME and ALIAS	LINK		SPW	RPW	RU Size
NNP61A	CPSVCMG	USIBMRA.NNP41A	@1221966	0B43E6C8	2	8	5
NNP61A	CPSVCMG	USIBMRA.NNP41A	@1221966	0B43E6C8	2	65	5
NNP61A	CPSVCMG	USIBMRA.RAA	@1222188	0B43EB08	2	20	5
NNP61A	CPSVCMG	USIBMRA.RAA	@1222188	0B43EB08	2	8	5
NNP61A	CPSVRMGR	USIBMRA.RAA	@1222188	0B47BA30	2	1	5
NNP61A	CPSVRMGR	USIBMRA.RAA	@1222188	0B47BA30	2	8	5

Figure 120. Active LU 6.2 Sessions From NNP61A

We now had four sessions between NNP61A and RAA. The two new ones (the last two) were the two DLUR/S sessions using mode CPSVRMGR.

8.9.5 Dependent LU Sessions

Next, we logged on from a dependent LU on PU05170 to NetView on RA39. A display of the RTP major node on RAA (Figure 121) showed a new RTP connection, CNR0077C.

```
D NET,ID=ISTRTPMN,E
IST097I DISPLAY ACCEPTED
IST075I NAME = ISTRTPMN, TYPE = RTP MAJOR NODE 636
IST486I STATUS= ACTIV, DESIRED STATE= ACTIV
IST1486I   RTP NAME   STATE            DESTINATION CP     MNPS  TYPE
IST1487I   CNR0077C   CONNECTED        USIBMRA.NNP61A     NO    LULU
IST1487I   CNR0076F   CONNECTED        USIBMRA.NNP61A     NO    CPCP
IST1487I   CNR0076C   CONNECTED        USIBMRA.NNP61A     NO    RSTP
```

Figure 121. RTP Pipes from DLU Server

Remember that RAA was the DLU server for the dependent LU, but not the application owner. We therefore expected an LU-LU session pipe for the DLUR/S sessions, but not for the dependent LU session. Figure 122 on page 143 indicated (but did not prove) that CNR0077C was indeed the DLUR/S pipe.

```
D NET,ID=CNR0077C,E
IST097I DISPLAY ACCEPTED
IST075I NAME = CNR0077C, TYPE = PU_T2.1 645
IST1392I DISCNTIM = 00010 DEFINED AT PU FOR DISCONNECT
IST486I STATUS= ACTIV--LX-, DESIRED STATE= ACTIV
IST1043I CP NAME = NNP61A, CP NETID = USIBMRA, DYNAMIC LU = YES
IST1589I XNETALS = YES
IST933I LOGMODE=***NA***, COS=SNASVCMG  7
IST1476I TCID X'310BB532000000A3' - REMOTE TCID X'000000000B46BBA8'
IST1481I DESTINATION CP USIBMRA.NNP61A - NCE X'D0201025'
IST1587I ORIGIN NCE X'D000000000000000'
IST1477I ALLOWED DATA FLOW RATE = 2800 KBITS/SEC
IST1516I INITIAL DATA FLOW RATE = 3200 KBITS/SEC
IST1511I MAXIMUM NETWORK LAYER PACKET SIZE = 2074 BYTES
IST1478I NUMBER OF UNACKNOWLEDGED BUFFERS = 0
IST1479I RTP CONNECTION STATE = CONNECTED - MNPS = NO
IST1480I RTP END TO END ROUTE - PHYSICAL PATH
IST1460I TGN  CPNAME             TG TYPE      HPR
IST1461I  21  USIBMRA.NNP61A     APPN         RTP   8
IST231I RTP MAJOR NODE = ISTRTPMN
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST355I LOGICAL UNITS:
IST080I NNP61A   ACT/S----Y
IST314I END
```

Figure 122. RTP Pipe for DLUR/S

The APPN COS was SNASVCMG **7** and the partner node was NNP61A **8**. To prove that this was the DLUR/S HPR pipe, we displayed the CP name of the partner node as in Figure 123 on page 144.

```
D NET,ID=NNP61A,E
IST097I DISPLAY ACCEPTED
IST075I NAME = USIBMRA.NNP61A, TYPE = ADJACENT CP 654
IST486I STATUS= ACT/S----Y, DESIRED STATE= ACTIV
IST1402I SRTIMER =    30  SRCOUNT =    10
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.NNP61A - NETSRVR = ***NA***
IST1044I ALSLIST = ISTAPNPU
IST1131I DEVICE = ILU/CDRSC
IST654I I/O TRACE = OFF, BUFFER TRACE = ON - AMOUNT = FULL
IST1500I STATE TRACE = OFF
IST228I ENCRYPTION = NONE
IST1563I CKEYNAME = NNP61A CKEY = PRIMARY CERTIFY = NO
IST1552I MAC = NONE MACTYPE = NONE
IST171I ACTIVE SESSIONS = 0000000004, SESSION REQUESTS = 0000000000
IST206I SESSIONS:
IST1081I ADJACENT LINK STATION = CNR0077C   9
IST634I NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I RAA      ACTIV/DL-S D3974BDF909413C3 0016 0000  0  0 USIBMRA
IST635I RAA      ACTIV/DL-P F7FF6164529F533F 0000 0018  0  0 USIBMRA
IST1081I ADJACENT LINK STATION = CNR0076F   10
IST634I NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I RAA      ACTIV/CP-S D3974BDF909413B6 016E 0001       USIBMRA
IST635I RAA      ACTIV/CP-P F7FF6164529F529D 0001 016B  0  0 USIBMRA
IST1355I PHYSICAL UNITS SUPPORTED BY DLUR USIBMRA.NNP61A   11
IST089I W05170   TYPE = PU_T2.1          , ACTIV---X-
IST924I -------------------------------------------------------------
IST075I NAME = USIBMRA.NNP61A, TYPE = DIRECTORY ENTRY
IST1186I DIRECTORY ENTRY = DYNAMIC NN
IST1184I CPNAME = USIBMRA.NNP61A - NETSRVR = ***NA***
IST1402I SRTIMER =    30  SRCOUNT =    10
IST314I END
```

Figure 123. DLUR Node Display

There were four sessions between RAA and NNP61A. Two of them, using the RTP pipe CNR0077C **9**, were labelled ACTIV/DL meaning that they were DLUR/S sessions. The other two, using RTP pipe CNR0076F **10**, were the CP-CP sessions. The message IST1355I **11** shows that this node is a served DLUR that is acting on behalf of the type 2 PU W05170.

A display of the PU W05170 (Figure 124 on page 145) showed that a DLUR PU is seen by VTAM almost as any other peripheral type 2 link station. The only difference is the presence of the message IST1354I **12** showing that the PU is on a served DLUR node.

```
D NET,ID=W05170,E
IST097I DISPLAY ACCEPTED
IST075I NAME = W05170, TYPE = PU_T2.1 666
IST486I STATUS= ACTIV---X-, DESIRED STATE= ACTIV
IST1043I CP NAME = PU05170, CP NETID = USIBMRA, DYNAMIC LU = YES
IST1589I XNETALS = YES
IST1354I DLUR NAME = NNP61A             MAJNODE = ISTDSWMN    12
IST136I SWITCHED SNA MAJOR NODE = ISTDSWMN
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST355I LOGICAL UNITS:
IST080I W0517002 ACT/S---X- W0517003 ACTIV---X- W0517004 ACTIV---X-
IST080I W0517005 ACTIV---X- W0517006 ACTIV---X- W0517007 ACTIV---X-
IST080I W0517008 ACTIV---X- W0517009 ACTIV---X- W051700A ACTIV---X-
IST080I W051700E ACTIV---X- W051700F ACTIV---X- W0517010 ACTIV---X-
IST080I W0517011 ACTIV---X-
IST314I END
```

Figure 124. DLUR Dependent PU

The display of the LU, in fact, shows absolutely nothing that indicates it is on a DLUR node.

When we displayed the RTP connections from RA39 we saw Figure 125.

```
 D NET,ID=ISTRTPMN,E
IST097I DISPLAY ACCEPTED
IST075I NAME = ISTRTPMN, TYPE = RTP MAJOR NODE 360
IST486I STATUS= ACTIV, DESIRED STATE= ACTIV
IST1486I   RTP NAME   STATE            DESTINATION CP     MNPS  TYPE
IST1487I   CNR0000F   CONNECTED        USIBMRA.NNP61A     NO    LULU   12
IST1487I   CNR0000E   CONNECTED        USIBMRA.NNP41A     NO    RSTP
IST1487I   CNR0000D   CONNECTED        USIBMRA.NNP41A     NO    CPCP
IST1487I   CNR0000C   CONNECTED        USIBMRA.NNP41A     NO    CPCP
IST1487I   CNR00003   CONNECTED        USIBMRA.RAK        NO    LULU
IST1487I   CNR00002   CONNECTED        USIBMRA.RAA        NO    RSTP
IST1487I   CNR00001   CONNECTED        USIBMRA.RAA        NO    CPCP
```

Figure 125. RTP Connection to DLUR LU

We saw that a new LU-LU session pipe, CNR0000F, has been created **12**. There is, of course, no Route Setup pipe to NNP61A because NNP61A is not adjacent to RA39. The RSTP pipe CNR00002 was probably used to set up CNR0000F. A detailed display of CNR0000F gave us Figure 126 on page 146.

```
D NET,ID=CNR0000F,E
IST097I DISPLAY ACCEPTED
IST075I NAME = CNR0000F, TYPE = PU_T2.1 366
IST1392I DISCNTIM = 00010 DEFINED AT PU FOR DISCONNECT
IST486I STATUS= ACTIV--LX-, DESIRED STATE= ACTIV
IST1043I CP NAME = NNP61A, CP NETID = USIBMRA, DYNAMIC LU = YES
IST1589I XNETALS = YES
IST933I LOGMODE=***NA***, COS=#CONNECT    13
IST1476I TCID X'38273EBB000000BA' - REMOTE TCID X'00000000FF267B40'
IST1481I DESTINATION CP USIBMRA.NNP61A - NCE X'D020200E'   14
IST1587I ORIGIN NCE X'D000000000000000'
IST1477I ALLOWED DATA FLOW RATE = 3200 KBITS/SEC
IST1516I INITIAL DATA FLOW RATE = 3200 KBITS/SEC
IST1511I MAXIMUM NETWORK LAYER PACKET SIZE = 2074 BYTES
IST1478I NUMBER OF UNACKNOWLEDGED BUFFERS = 0
IST1479I RTP CONNECTION STATE = CONNECTED - MNPS = NO
IST1480I RTP END TO END ROUTE - PHYSICAL PATH
IST1460I TGN  CPNAME             TG TYPE      HPR
IST1461I  21  USIBMRA.RAA        APPN         RTP   15
IST1461I  21  USIBMRA.NNP61A     APPN         RTP
IST231I RTP MAJOR NODE = ISTRTPMN
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST355I LOGICAL UNITS:
IST080I W0517002 ACT/S----Y   16
IST314I END
```

Figure 126. DLUR LU RTP Pipes

This pipe used APPN COS #CONNECT **13** and contained a session to the dependent LU W0517002 **16**. However, the RTP partner was NNP61A **14**, the DLUR node. The path taken by the pipe (and therefore the session) **15** is shown graphically in Figure 127 on page 147.

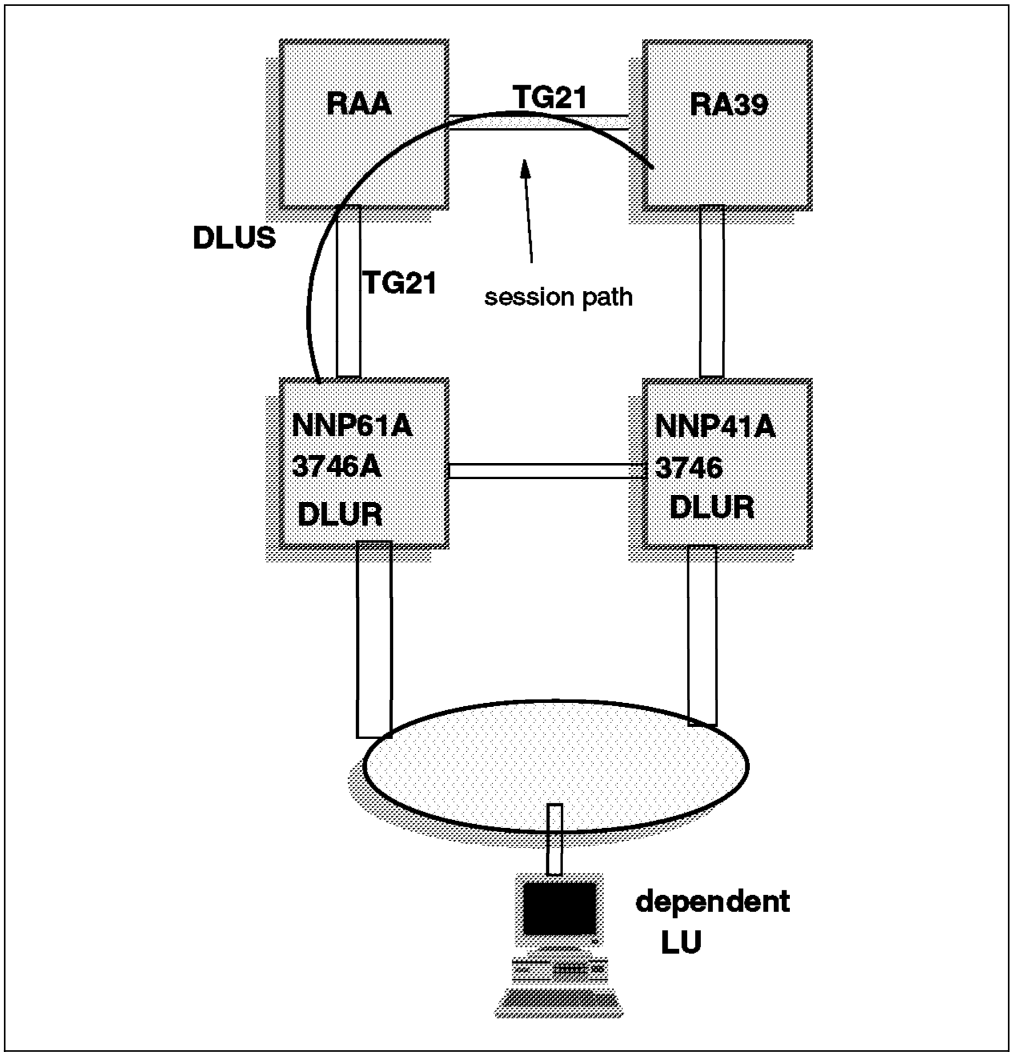

Figure 127. LU-LU Session Path

The session path was calculated by RA39 and took the optimum route through the APPN network. The fact that it happened to pass through its DLU server (RAA) was purely coincidental.

Next, we displayed the dependent LU from RA39. Figure 128 on page 148 showed that it used RTP pipe CNR0000F.

```
D NET,ID=W0517002,E
IST097I DISPLAY ACCEPTED
IST075I NAME = USIBMRA.W0517002, TYPE = CDRSC 371
IST486I STATUS= ACT/S----Y, DESIRED STATE= ACTIV
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 = RA39, VERIFY OWNER = NO
IST1184I CPNAME = USIBMRA.NNP61A - NETSRVR = ***NA***
IST082I DEVTYPE = INDEPENDENT LU / CDRSC
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST228I ENCRYPTION = NONE
IST1563I CKEYNAME = W0517002 CKEY = PRIMARY CERTIFY = NO
IST1552I MAC = NONE MACTYPE = NONE
IST171I ACTIVE SESSIONS = 0000000001, SESSION REQUESTS = 0000000000
IST206I SESSIONS:
IST1081I ADJACENT LINK STATION = CNR0000F
IST634I NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I RA39N008 ACTIV-P    F70794547C2C20D9 0001 0002  0  0 USIBMRA
IST924I -------------------------------------------------------------
IST075I NAME = USIBMRA.W0517002, TYPE = DIRECTORY ENTRY
IST1186I DIRECTORY ENTRY = DYNAMIC LU
IST1184I CPNAME = USIBMRA.NNP61A - NETSRVR = USIBMRA.RAA
IST314I END
```

Figure 128. Dependent LU Session on RTP Pipe to DLUR

Now we checked the 3746 HPR display again to see how the new pipe showed up, as in Figure 129.

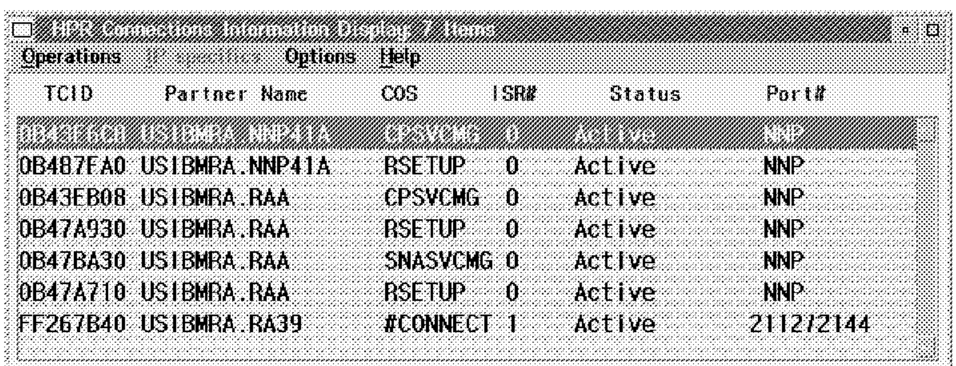

TCID	Partner Name	COS	ISR#	Status	Port#
0B43E6C8	USIBMRA.NNP41A	CPSVCMG	0	Active	NNP
0B487FA0	USIBMRA.NNP41A	RSETUP	0	Active	NNP
0B43EB08	USIBMRA.RAA	CPSVCMG	0	Active	NNP
0B47A930	USIBMRA.RAA	RSETUP	0	Active	NNP
0B47BA30	USIBMRA.RAA	SNASVCMG	0	Active	NNP
0B47A710	USIBMRA.RAA	RSETUP	0	Active	NNP
FF267B40	USIBMRA.RA39	#CONNECT	1	Active	2112/2144

Figure 129. Active RTP Connections after Establishing an LU-LU Session

The last line of the display shows the RTP connection carrying the CS/2 session to RA39′s NetView, using APPN COS #CONNECT. It is interesting to note that this RTP pipe originates in the TIC (Port 2112/2144), rather than in the NNP as do the control sessions.

8.9.6 Path Switch for DLUR Session

The MPC connection between the VTAMs, which was on the session route, was now deactivated. Figure 130 shows what happened on RAA.

```
V NET,INACT,ID=RAAAHHD
IST097I VARY ACCEPTED
IST1196I APPN CONNECTION FOR USIBMRA.RA39 INACTIVE - TGN = 21
IST1494I  PATH SWITCH STARTED   FOR RTP CNR00761    18
IST1133I RAAHRAC IS NOW INACTIVE, TYPE = PU_T2
IST1133I RAAAHHD IS NOW INACTIVE, TYPE = LCL SNA MAJ NODE
IST1488I  INACTIVATION FOR RTP CNR00763 AS PASSIVE PARTNER COMPLETED
IST1416I  ID = CNR00763          FAILED - RECOVERY IN PROGRESS   17
IST1136I  VARY INACT CNR00763          SCHEDULED - UNRECOVERABLE ERROR
...     ...     ...    ...     ...
IST1097I  CP-CP SESSION WITH USIBMRA.RA39      TERMINATED
```

Figure 130. Deactivate MPC Connection

The Route Setup pipe to RA39, CNR00763, was deactivated immediately **17**. A path switch was initiated for the CP-CP pipe CNR00761 **18**, but failed after the path switch timer expired because there were no alternative routes (let alone Control Flows-capable routes) to RA39. When the pipe was deactivated the CP-CP sessions also failed.

The DLUR/S pipe, CNR0077C, was not switched because it did not cross the failing connection. No messages were issued regarding the dependent LU session, because RAA (as an ANR node on the session path) was not aware that the session passed through it.

RA39, as the RTP endpoint, detected the failure. In fact it was the first RTP endpoint to detect the failure because its adjacent link failed. Figure 131 shows what happened on RA39.

```
IST259I  INOP RECEIVED FOR RACHRAA  CODE = 01
IST619I  ID = RACHRAA  FAILED - RECOVERY IN PROGRESS
IST1196I  APPN CONNECTION FOR USIBMRA.RAA       INACTIVE - TGN =  21
IST1494I  PATH SWITCH STARTED   FOR RTP CNR0000F
IST1494I  PATH SWITCH COMPLETED FOR RTP CNR0000F
IST1480I  RTP END TO END ROUTE - PHYSICAL PATH
IST1460I  TGN  CPNAME            TG TYPE      HPR
IST1461I   21  USIBMRA.NNP41A     APPN         RTP
IST1461I   21  USIBMRA.NNP61A     APPN         RTP
IST314I  END
```

Figure 131. RTP Path Switch for DLUR Session

RA39 initiated a path switch, and the dependent LU session was moved to a new route. Figure 132 on page 150 shows the new route, now passing through both 3746s. The session no longer even passes through the VTAM that owns the SLU. However, it must pass through its DLUR (NNP61A), because the DLUR provides the subarea boundary function which VTAM or NCP would have provided without DLUR. To obtain the maximum benefit from APPN routing you must go one stage further and put the DLUR function in the workstation itself.

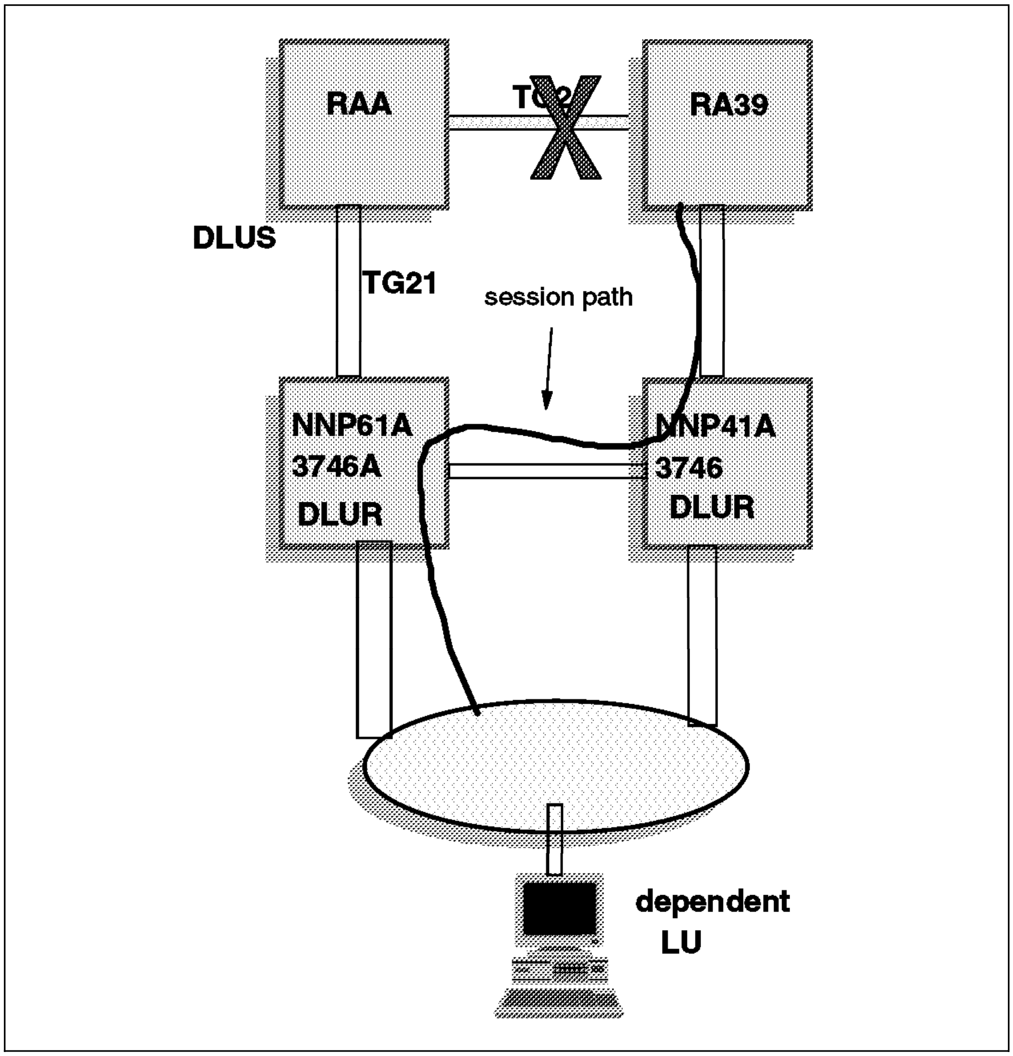

Figure 132. Path after Path Switch

A display of the switched pipe from RA39 confirms the new route, as seen in Figure 133 on page 151.

```
D NET,ID=CNR0000F,E
IST097I DISPLAY ACCEPTED
IST075I NAME = CNR0000F, TYPE = PU_T2.1 435
IST1392I DISCNTIM = 00010 DEFINED AT PU FOR DISCONNECT
IST486I STATUS= ACTIV--LX-, DESIRED STATE= ACTIV
IST1043I CP NAME = NNP61A, CP NETID = USIBMRA, DYNAMIC LU = YES
IST1043I CP NAME = NNP61A, CP NETID = USIBMRA, DYNAMIC LU = YES
IST1589I XNETALS = YES
IST933I LOGMODE=***NA***, COS=#CONNECT
IST1476I TCID X'38273EBB000000BA' - REMOTE TCID X'00000000FF267B40'
IST1481I DESTINATION CP USIBMRA.NNP61A - NCE X'D020200E'
IST1587I ORIGIN NCE X'D000000000000000'
IST1477I ALLOWED DATA FLOW RATE = 1597 KBITS/SEC
IST1516I INITIAL DATA FLOW RATE = 1597 KBITS/SEC
IST1511I MAXIMUM NETWORK LAYER PACKET SIZE = 2058 BYTES
IST1478I NUMBER OF UNACKNOWLEDGED BUFFERS = 0
IST1479I RTP CONNECTION STATE = CONNECTED - MNPS = NO
IST1480I RTP END TO END ROUTE - PHYSICAL PATH
IST1460I TGN  CPNAME            TG TYPE      HPR
IST1461I  21  USIBMRA.NNP41A     APPN         RTP
IST1461I  21  USIBMRA.NNP61A     APPN         RTP
IST231I RTP MAJOR NODE = ISTRTPMN
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST355I LOGICAL UNITS:
IST080I W0517002 ACT/S----Y
IST314I END
```

Figure 133. RTP Pipe after Path Switch

We took some traces on the 3746 and the CS/2 node during these tests, which showed the subarea session setup protocols flowing natively between the CS/2 and the 3746, but being encapsulated in LU 6.2 sessions between the 3746 DLUR and VTAM DLUS. Appendix B, “A Complete Scenario” on page 247 shows extracts from these traces.

8.10 Example with 3746 As ANR Node

For our second test with the 3746s, we defined a CS/2 node as a DLUR, so that the 3746s would act purely as ANR nodes. Please refer again to Figure 78 on page 107 for the network setup.

The configuration of the VTAMs and the 3746s was identical to the previous configuration; only the workstation was changed. 7.5, “HPR on Communications Server/2” on page 90 shows how to set up DLUR on a CS/2 node.

Our setup in this case was as follows:

- The CS/2 was defined as a network node with CP name CM5HPRNN.
- CM5HPRNN had APPN connections, with CP-CP sessions, to both NNP61A and NNP41A.
- A dependent PU with IDNUM AAA61 was defined in the CS/2 node, using a DLUR connection as its host link.
- That DLUR logical link was specified as being to RAA. Thus RAA was the DLUS for the dependent resources on CM5HPRNN.

8.10.1 CS/2 As DLUR Node

When the CS/2 node was started, it immediately set up several RTP pipes. Because both CS/2 and 3746 support Control Flows over RTP, there were CP-CP session pipes to each 3746 (as can be seen in later displays). There was also a Route Setup pipe (used to establish the DLUR/S sessions) to one 3746 (in fact NNP61A). Finally there was the DLUR/S RTP connection itself, whose VTAM display we show in Figure 134.

```
D NET,ID=CNR00790,E
IST097I DISPLAY ACCEPTED
IST075I NAME = CNR00790, TYPE = PU_T2.1 454
IST1392I DISCNTIM = 00010 DEFINED AT PU FOR DISCONNECT
IST486I STATUS= ACTIV--LX-, DESIRED STATE= ACTIV
IST1043I CP NAME = CM5HPRNN, CP NETID = USIBMRA, DYNAMIC LU = YES
IST1589I XNETALS = YES
IST933I LOGMODE=***NA***, COS=SNASVCMG
IST1476I TCID X'310BB546000000CE' - REMOTE TCID X'000000000000001A'
IST1481I DESTINATION CP USIBMRA.CM5HPRNN - NCE X'80'
IST1587I ORIGIN NCE X'D000000000000000'
IST1477I ALLOWED DATA FLOW RATE = 1597 KBITS/SEC
IST1516I INITIAL DATA FLOW RATE = 1597 KBITS/SEC
IST1511I MAXIMUM NETWORK LAYER PACKET SIZE = 2058 BYTES
IST1478I NUMBER OF UNACKNOWLEDGED BUFFERS = 0
IST1479I RTP CONNECTION STATE = CONNECTED - MNPS = NO
IST1480I RTP END TO END ROUTE - PHYSICAL PATH
IST1460I TGN  CPNAME            TG TYPE      HPR
IST1461I  21  USIBMRA.NNP61A     APPN         RTP
IST1461I  25  USIBMRA.CM5HPRNN   APPN         RTP
IST231I RTP MAJOR NODE = ISTRTPMN
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST355I LOGICAL UNITS:
IST080I CM5HPRNN ACT/S----Y
```

Figure 134. DLUR/S RTP Pipe from CS/2

The path used for this pipe was as shown in Figure 135.

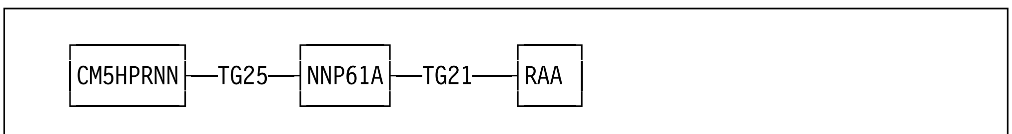

Figure 135. DLUR/S Path

A display of CM5HPRNN from RAA shows that it is a DLUR node serving one type 2 PU WAA61A **1**. The only sessions between CM5HPRNN and RAA are the two DLUR/S (ACTIV/DL) sessions on RTP connection CNR00790 **2** (please see Figure 136 on page 153).

```
D NET,ID=CM5HPRNN,E
IST097I DISPLAY ACCEPTED
IST075I NAME = USIBMRA.CM5HPRNN, TYPE = ADJACENT CP 461
IST486I STATUS= ACT/S----Y, DESIRED STATE= ACTIV
IST1402I SRTIMER =    30  SRCOUNT =    10
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 = RAA, VERIFY OWNER = NO
IST1184I CPNAME = USIBMRA.CM5HPRNN - NETSRVR = ***NA***
IST1044I ALSLIST = ISTAPNPU
IST1131I DEVICE = ILU/CDRSC
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST228I ENCRYPTION = NONE
IST1563I CKEYNAME = CM5HPRNN CKEY = PRIMARY CERTIFY = NO
IST1552I MAC = NONE MACTYPE = NONE
IST171I ACTIVE SESSIONS = 0000000002, SESSION REQUESTS = 0000000000
IST206I SESSIONS:
IST1081I ADJACENT LINK STATION = CNR00790   2
IST634I NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I RAA      ACTIV/DL-S C96F2A09108FC1DF 0013 0000  0  0 USIBMRA
IST635I RAA      ACTIV/DL-P F7FF6164529F5423 0000 0014  0  0 USIBMRA
IST1355I PHYSICAL UNITS SUPPORTED BY DLUR USIBMRA.CM5HPRNN
IST089I WAA61A   TYPE = PU_T2            , ACTIV---X-    1
IST924I -------------------------------------------------------------
IST075I NAME = USIBMRA.CM5HPRNN, TYPE = DIRECTORY ENTRY
IST1186I DIRECTORY ENTRY = DYNAMIC NN
IST1184I CPNAME = USIBMRA.CM5HPRNN - NETSRVR = ***NA***
IST1402I SRTIMER =    30  SRCOUNT =    10
IST314I END
```

Figure 136. DLUR CP of CS/2

A display of the PU WAA61A (Figure 137) confirms that it is on a DLUR node **3**.

```
D NET,ID=WAA61A,E
IST097I DISPLAY ACCEPTED
IST075I NAME = WAA61A, TYPE = PU_T2 653
IST486I STATUS= ACTIV---X-, DESIRED STATE= ACTIV
IST1043I CP NAME = ***NA***, CP NETID = USIBMRA, DYNAMIC LU = YES
IST1589I XNETALS = YES
IST1354I DLUR NAME = CM5HPRNN           MAJNODE = ISTDSWMN   3
IST136I SWITCHED SNA MAJOR NODE = ISTDSWMN
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST355I LOGICAL UNITS:
IST080I WAA61A02 ACTIV---X- WAA61A03 ACTIV---X- WAA61A04 ACTIV---X-
IST080I WAA61A05 ACTIV---X- WAA61A06 ACTIV---X- WAA61A07 ACTIV---X-
IST080I WAA61A08 ACTIV---X- WAA61A09 ACTIV---X- WAA61A0A ACTIV---X-
IST080I WAA61A0B ACTIV---X- WAA61A0C ACTIV---X- WAA61A0D ACTIV---X-
IST080I WAA61A0E ACTIV---X- WAA61A0F ACTIV---X- WAA61A10 ACTIV---X-
IST080I WAA61A11 ACTIV---X-
IST314I END
```

Figure 137. DLUR PU on CS/2 Node

We displayed again the adjacent link stations from CCM on the 3746 NNP61A, as seen in Figure 138 on page 154.

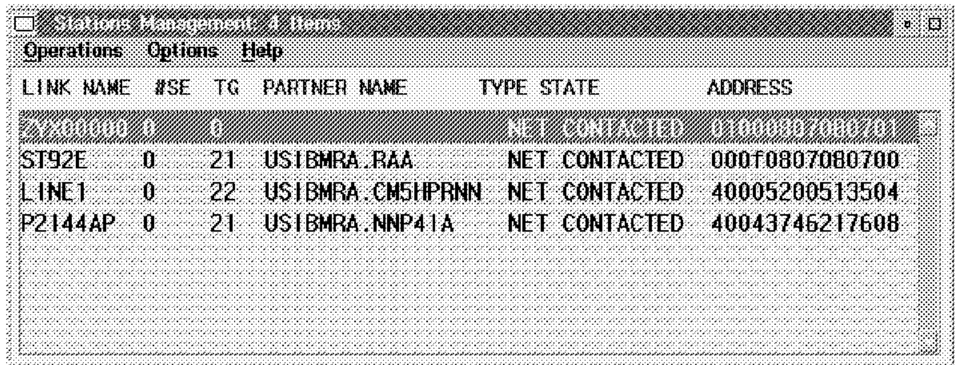

Figure 138. Active Stations on NNP61A Node

We can see now that the predefined LINE1 station has been connected to the network node CM5HPRNN. Figure 139 then shows the active RTP connections.

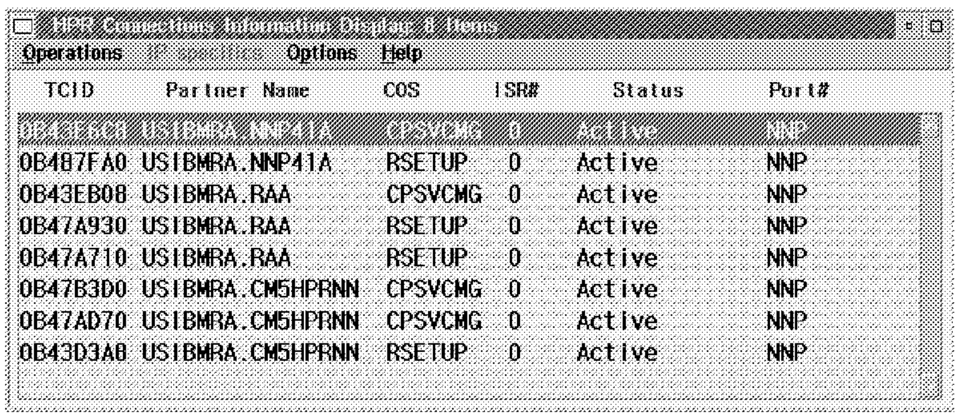

Figure 139. Display of Active RTP Pipes on NNP61A

You can see the two CP-CP pipes (the two CP-CP pipes were set up independently between these two NNs), and the single Route Setup pipe used to establish the DLUR/S sessions. Those sessions themselves, as well as their RTP pipe, are not known to NNP61A. Figure 140 shows details of the sessions ending in NNP61A.

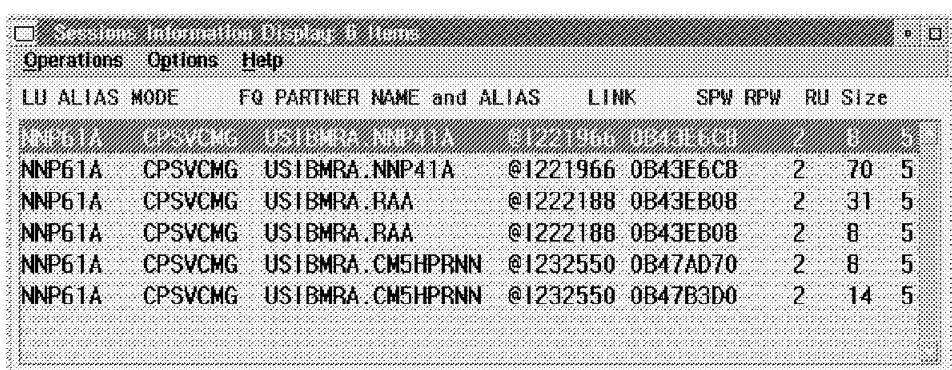

Figure 140. Display of Non-Intermediate Sessions on NNP61A

You will notice that we have no information related to the CS/2 sessions flowing across NNP61A because the 3746 is acting as an ANR node for them now. Previously they were visible in these panels; Figure 129 on page 148, for example, shows the RTP connections carrying the LU-LU and DLUR/S sessions.

Next, we logged to TSO on RA39 from one of the dependent LUs on CM5HPRNN. We saw no new RTP pipes on RAA (the DLUS) but one new one, CNR00021, was established on RA39. Figure 141 shows what we saw when we displayed it.

```
D NET,ID=CNR00021,E
IST097I DISPLAY ACCEPTED
IST075I NAME = CNR00021, TYPE = PU_T2.1 862
IST1392I DISCNTIM = 00010 DEFINED AT PU FOR DISCONNECT
IST486I STATUS= ACTIV--LX-, DESIRED STATE= ACTIV
IST1043I CP NAME = CM5HPRNN, CP NETID = USIBMRA, DYNAMIC LU = YES
IST1589I XNETALS = YES
IST933I LOGMODE=***NA***, COS=#CONNECT
IST1476I TCID X'38273ECD000000AC' - REMOTE TCID X'000000000000001B
IST1481I DESTINATION CP USIBMRA.CM5HPRNN - NCE X'80'
IST1587I ORIGIN NCE X'D000000000000000'
IST1477I ALLOWED DATA FLOW RATE = 1597 KBITS/SEC
IST1516I INITIAL DATA FLOW RATE = 1597 KBITS/SEC
IST1511I MAXIMUM NETWORK LAYER PACKET SIZE = 2058 BYTES
IST1478I NUMBER OF UNACKNOWLEDGED BUFFERS = 0
IST1479I RTP CONNECTION STATE = CONNECTED - MNPS = NO
IST1480I RTP END TO END ROUTE - PHYSICAL PATH
IST1460I TGN  CPNAME             TG TYPE      HPR
IST1461I  21  USIBMRA.NNP41A     APPN         RTP    4
IST1461I  23  USIBMRA.CM5HPRNN   APPN         RTP
IST231I RTP MAJOR NODE = ISTRTPMN
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST355I LOGICAL UNITS:
IST080I WAA61A02 ACT/S----Y
IST314I END
```

Figure 141. RTP Pipe to CS/2 Dependent LU

You can see that this particular pipe went all the way from RA39 to CM5HPRNN, via NNP41A **4**. Figure 142 on page 156 shows the route the session took. It went nowhere near the DLUS VTAM, RAA.

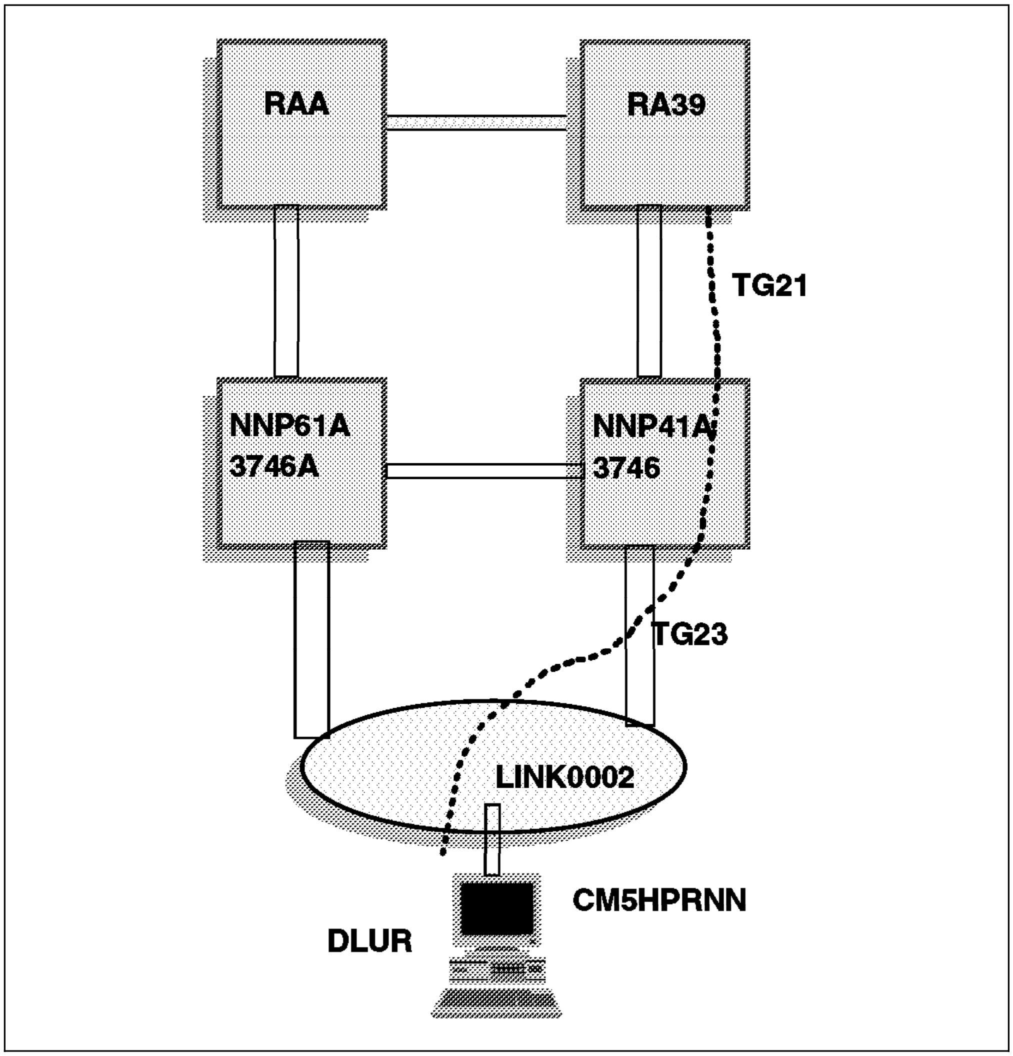

Figure 142. Session Path for DLUR on CS/2

A display of the LU, Figure 143 on page 157, from RA39 shows it simply as an APPN LU (a CDRSC **5** with owner RA39 **6**) using the link station CNR00021 **7** (an HPR connection).

```
D NET,ID=WAA61A02,E
IST097I DISPLAY ACCEPTED
IST075I NAME = USIBMRA.WAA61A02, TYPE = CDRSC 867
IST486I STATUS= ACT/S----Y, DESIRED STATE= ACTIV
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 = RA39, VERIFY OWNER = NO  6
IST1184I CPNAME = USIBMRA.CM5HPRNN - NETSRVR = ***NA***
IST082I DEVTYPE = INDEPENDENT LU / CDRSC  5
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST228I ENCRYPTION = NONE
IST1563I CKEYNAME = WAA61A02 CKEY = PRIMARY CERTIFY = NO
IST1552I MAC = NONE MACTYPE = NONE
IST171I ACTIVE SESSIONS = 0000000001, SESSION REQUESTS = 0000000000
IST206I SESSIONS:
IT1081I ADJACENT LINK STATION = CNR00021   7
IST634I NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I RA39T03  ACTIV-P    F70794547C2C213C 0003 000A  0  0 USIBMRA
IST924I -------------------------------------------------------------
IST075I NAME = USIBMRA.WAA61A02, TYPE = DIRECTORY ENTRY
IST1186I DIRECTORY ENTRY = DYNAMIC LU
IST1184I CPNAME = USIBMRA.CM5HPRNN - NETSRVR = USIBMRA.RAA
IST314I END
```

Figure 143. DLUR LU from RTP Partner

8.10.2 Path Switch for CS/2 DLUR Session

Finally, we deactivated from the CS/2 node the connection to NNP41A. Displays on RAA showed that nothing had changed. A display of ISTRTPMN on RA39 showed exactly the same RTP pipes, but a detailed display of CNR00021 (Figure 144 on page 158) tells us more.

```
D NET,ID=CNR00021,E
IST097I DISPLAY ACCEPTED
IST075I NAME = CNR00021, TYPE = PU_T2.1 950
IST1392I DISCNTIM = 00010 DEFINED AT PU FOR DISCONNECT
IST486I STATUS= ACTIV--LX-, DESIRED STATE= ACTIV
IST1043I CP NAME = CM5HPRNN, CP NETID = USIBMRA, DYNAMIC LU = YES
IST1589I XNETALS = YES
IST933I LOGMODE=***NA***, COS=#CONNECT
IST1476I TCID X'38273ECD000000AC' - REMOTE TCID X'000000000000001B'
IST1481I DESTINATION CP USIBMRA.CM5HPRNN - NCE X'80'
IST1587I ORIGIN NCE X'D000000000000000'
IST1477I ALLOWED DATA FLOW RATE = 1597 KBITS/SEC
IST1516I INITIAL DATA FLOW RATE = 1597 KBITS/SEC
IST1511I MAXIMUM NETWORK LAYER PACKET SIZE = 2058 BYTES
IST1478I NUMBER OF UNACKNOWLEDGED BUFFERS = 0
IST1479I RTP CONNECTION STATE = CONNECTED - MNPS = NO
IST1480I RTP END TO END ROUTE - PHYSICAL PATH
IST1460I TGN  CPNAME            TG TYPE      HPR
IST1461I  21  USIBMRA.NNP41A     APPN         RTP
IST1461I  21  USIBMRA.NNP61A     APPN         RTP
IST1461I  25  USIBMRA.CM5HPRNN   APPN         RTP
IST231I RTP MAJOR NODE = ISTRTPMN
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST355I LOGICAL UNITS:
IST080I WAA61A02 ACT/S----Y
IST314I END
```

Figure 144. RTP Pipe after Switch

The session path has moved to the route shown in Figure 145 on page 159. In this case there were no path switch messages on VTAM as the CS/2 node detected the failure (an adjacent link deactivation) long before RA39 could do so (by waiting for a timer to expire). Therefore, CM5HPRNN started and completed the path switch before VTAM's timer expired. VTAM only issues a path switch message if it initiates the switch itself.

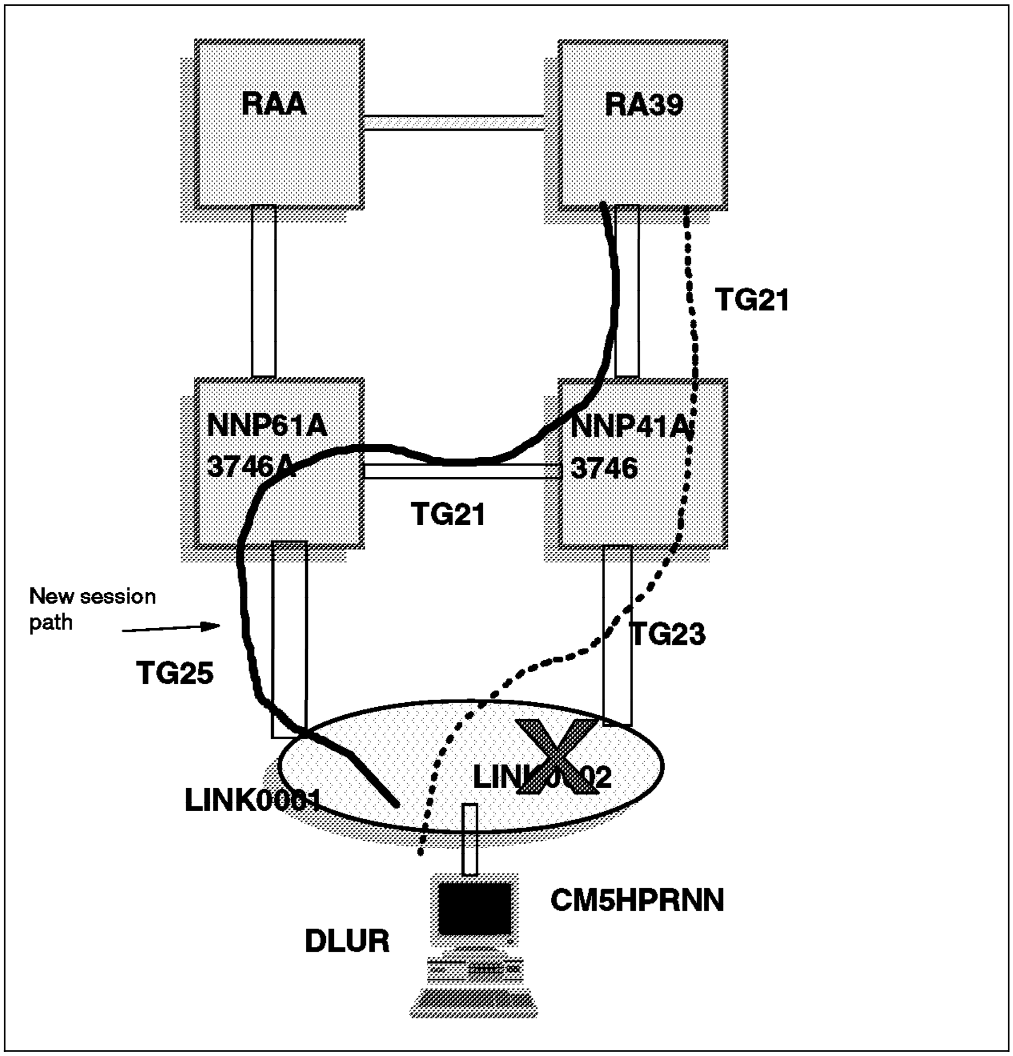

Figure 145. Session Route after CS/2 Link Failure

After that, we displayed the dependent LU from RA39. There was no change from the display shown in Figure 143 on page 157.

For this test, we performed a comprehensive audit (traces and displays) of what went on in the CS/2 node. The results of this exercise are documented in Appendix B, "A Complete Scenario" on page 247.

Chapter 9. HPR and DLUR on the 2216

In this chapter we extend our network to a remote site in which a 2216 router has been installed. We used token-ring connections to represent the wide area network between the 3746s and the 2216. The actual DLC used has only a marginal effect on the HPR and DLUR customization and operation in such a configuration.

Figure 146 shows the network we used in these scenarios.

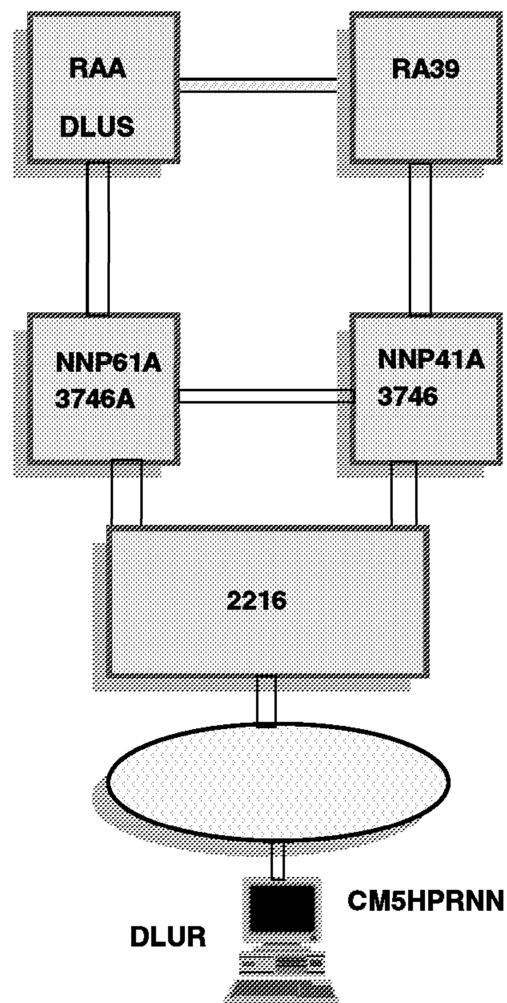

Figure 146. 2216 Test Scenario

Once again, the VTAM nodes RA39 and RAA are connected by an MPC link. As in the previous chapter, each VTAM is connected to one 3746 NNP, the NNPs being joined together via a token-ring. This time the 2216 has a connection to each 3746. The 2216 is configured as a network node (there is no choice in the matter), and it has a separate downstream token-ring to which our workstation is attached. We ran two distinct tests, as in the 3746-only case:

1. First, we defined DLUR support in the CS/2 workstation and used the 2216 and the 3746 purely for ANR switching.

2. Second, we defined DLUR support in the 2216 and used peripheral subarea protocols between the 2216 and the CS/2 PC.

9.1 Configuration with 2216 As ANR Node

In our first scenario we used the 2216 as an HPR router and the CS/2 node as the DLUR and the RTP endpoint. The 2216 was at Version 2, Release 2 of its software, Multiprotocol Access Services (MAS). Please refer to Figure 146 on page 161 for the diagram of the network setup.

9.1.1 APPN/HPR Configuration for 2216 Router

The 2216 configuration interface can be accessed directly via an ASCII terminal connected to its service port, or across the TCP/IP network via TELNET. Clearly the initial configuration must be done by means of the direct attachment, since the IP address is not known at that time. For our tests we always used an ASCII emulator on a PC connected to the service port.

The 2216 also has a GUI method of configuration. Using this, the configuration files are created on a PC without requiring online access to the 2216. When complete, the files are downloaded to the 2216 using either SNMP over the IP network, or directly through the service port.

When the ASCII terminal is connected to the 2216, the command prompt displayed is an asterisk (*). To enter configuration mode, type `talk 6` at the prompt, as in Figure 147. The basic 2216 parameters such as the token-ring MAC addresses had already been configured, so we invoked the APPN configuration function by entering `protocol APPN` at the configuration prompt (`Config>`), as shown in Figure 147.

```
*talk 6
Config>protocol appn
APPN user configuration
APPN config>
```

Figure 147. Invoking APPN Configuration on 2216

At the `APPN Config>` prompt, you can enter the configuration commands as shown below. To return to the base prompt, type `exit`.

The 2216′s configuration has basically the same structure as that of the 3746. First you define the node, then each port, then the link stations on the ports. Once again, only those link stations where the 2216 will make the connection need to be defined.

Each configuration question is presented with the default value in brackets immediately before the question mark. If you wish to accept the defaults, then simply press the Enter key, otherwise type in the value you want and press Enter. Some questions ask whether you wish to edit certain subgroups of parameters; we recommend that you answer YES to all of these until you are familiar with 2216 configuration. We found that some parameters are quite well concealed within these subgroups, and are not presented unless you answer YES to the question that reveals them.

Figure 148 on page 163 shows how we configured the APPN node parameters on our 2216. Type set node at the APPN Configuration prompt to do this.

```
Config>protocol appn
APPN config>set node
Enable APPN (Y)es (N)o (Y)? y
Network ID (Max 8 characters) (  )? USIBMRA   1
Control point name (Max 8 characters) (  )? NN2216A 1
Enable branch extender (Y)es (N)o (N)?   2
Route addition resistance(0-255) (128)?
XID ID number for subarea connection (5 hex digits) (00000)?   3
Use enhanced #BATCH COS (Y)es (N)o (Y)?   4
Use enhanced #BATCHSC COS (Y)es (N)o (Y)?
Use enhanced #INTER COS (Y)es (N)o (Y)?
Use enhanced #INTERSC COS (Y)es (N)o (Y)?
Write this record? (Y)? y
```

Figure 148. 2216 Node Definition

In the node configuration:

- We entered the CP name of the 2216 NN at **1**.
- Branch extender would be tested at a later stage **2** so we allowed this answer to default to N.
- The subarea link question **3** provides a node ID for use if VTAM cannot identify the 2216 by CP name.
- The 2216 offers the option **4** of using the four new APPN COS entries designed for ATM networks. These same entries were introduced in VTAM V4R4, and provide the ability to discriminate more easily between various high-speed links. It is recommended that if you use these in one APPN NN, then you should use them in all NNs.

At the end of the node configuration we entered y to the last question so the 2216 would store the details for later use. Next, we entered the port configurations. Figure 149 on page 164 shows the definitions for the token-ring port which we called trn1. This was the port we used to connect the 2216 to both 3746s. The other port (trn0) was used for the separate LAN on which the workstation lived. Enter add port at the APPN Configuration prompt to define a port.

```
APPN config>add port
APPN Port
Link Type: (P)PP, (FR)AME RELAY, (E)THERNET, (T)OKEN RING,
(S)DLC, (X)25, (D)LSw, (A)TM ( )? t
Interface number(Default 0): (0)? 1
Port name (Max 8 characters) (TR000)? trn1
Enable APPN on this port (Y)es (N)o (Y)? y
Port Definition
Service any node: (Y)es (N)o (Y)? y   5
High performance routing: (Y)es (N)o (Y)?   6
Maximum BTU size (768-17745) (2048)?
Maximum number of link stations (1-976) (512)?
Percent of link stations reserved for incoming calls (0-100) (0)
Percent of link stations reserved for outgoing calls (0-100) (0)
Local SAP address (04-EC) (4)?
Local HPR SAP address (04-EC) (C8)?
Edit TG Characteristics: (Y)es (N)o (N)?
Edit LLC Characteristics: (Y)es (N)o (N)? y   7
Remote SAP(04-EC) (4)? 8   8
Maximum number of outstanding I-format LPDUs (1-127) (26)?
Receive window size (1-127) (26)?
Inactivity timer(1-254 seconds) (30)?
Reply timer (1-254 half seconds) (2)?
Maximum number of retransmissions(1-254) (8)?
Receive acknowledgement timer (1-254 half seconds) (1)?
Acknowledgements needed to increment working window(0-127) (1)?
Edit HPR defaults: (Y)es (N)o (N)?
Write this record? (Y)? y
```

Figure 149. 2216 Port Definition

In this figure:

- We specified **5** that this port could accept any incoming connection requests.
- HPR is enabled on the port by default **6**.
- We needed to edit the LLC characteristics **7** in order to change the remote SAP **8**. You may remember that the 3746 local SAPs default to 8, while the 2216's remote SAP defaults to 4. The remote SAP, of course, applies to the *station* rather than the port, but (as with the 3746) the station defaults may be set at the port level.

On the port trn1 we defined the two link stations connecting the 2216 to the 3746s. You do this by typing `add link` at the APPN Configuration prompt. The 2216 asks you on which port you want to define this station. Figure 150 on page 165 shows the definition of the connection to NNP61A.

```
APPN config>add link
APPN Station
Port name for the link station ( )? trn1     9
Station name (Max 8 characters) ( )? t61a    10
Activate link automatically (Y)es (N)o (Y)?
MAC address of adjacent node (000000000000)? 400037462144   11
Adjacent node type: 0 = APPN network node,
1 = APPN end node or Unknown node type,
2 = LEN end node (0)?    12
High performance routing: (Y)es (N)o (Y)?    13
Allow CP-CP sessions on this link (Y)es (N)o (Y)?
CP-CP session level security (Y)es (N)o (N)?
Configure CP name of adjacent node: (Y)es (N)o (N)?
Edit TG Characteristics: (Y)es (N)o (N)?
Edit LLC Characteristics: (Y)es (N)o (N)? y
Remote SAP(04-EC) (4)?
Maximum number of outstanding I-format LPDUs (1-127) (26)?
Receive window size (1-127) (26)?
Inactivity timer(1-254 seconds) (30)?
Reply timer (1-254 half seconds) (2)?
Maximum number of retransmissions(1-254) (254)?
Receive acknowledgement timer (1-254 half seconds) (1)?
Acknowledgements needed to increment working window(0-127) (1)?
Edit HPR defaults: (Y)es (N)o (N)?
Write this record? (Y)? y
The record has been written.
```

Figure 150. Definition of Link to NN61A

In our station definitions for NNP61A:

- This link station was defined on port trn1 **9**.
- The name we assigned to the station **10** will be used in subsequent displays from the 2216.
- We specified the MAC address of the destination TIC3 on the 3746 **11**. The SAP address is not entered at this time, but needs to be unearthed by replying y to the question about LLC characteristics. We had, of course, already specified it as 8 on the port defaults question.
- As with VTAM, you can check the type of the adjacent node or allow the 2216 to accept any node type as its partner **12**.
- HPR is enabled on this link by default **13**, because the port has HPR support.

We entered a very similar definition for the other 3746, NNP41A. We then defined a second port, trn0, for the downstream token-ring connection to CM5HPRNN. We did not define a link station on port trn0 because the CS/2 node would be making the connection.

Once we had entered the whole APPN configuration, we were able to display a summary of it by typing `list all` at the APPN Configuration prompt, as shown in Figure 151 on page 166.

```
APPN config>list all
NODE:
NETWORK ID: USIBMRA
CONTROL POINT NAME: NN2216A
XID: 00000
APPN ENABLED: YES
MAX SHARED MEMORY: 5108
MAX CACHED: 4000
DLUR:
DLUR ENABLED: NO
PRIMARY DLUS NAME:
CONNECTION NETWORK:
      CN NAME       LINK TYPE  PORT INTERFACES
-------------------------------------------------------------
COS:
COS NAME
--------
#BATCH
#BATCHSC
#CONNECT
  #INTER
#INTERSC
 CPSVCMG
SNASVCMG

 MODE:
 MODE NAME  COS NAME
---------------------
PORT:
  INTF     PORT      LINK      HPR    SERVICE   PORT
 NUMBER    NAME      TYPE    ENABLED    ANY    ENABLED
------------------------------------------------------
    0       TRN0   IBMTRNET    YES      YES      YES
    1       TRN1   IBMTRNET    YES      YES      YES
STATION:
STATION    PORT       DESTINATION      HPR    ALLOW  ADJ NODE
 NAME      NAME         ADDRESS      ENABLED  CP-CP    TYPE
 ------------------------------------------------------------
  T61A     TRN1      400037462144     YES     YES       0
  T41A     TRN1      400437462176     YES     YES       0
 LU NAME:
      LU NAME        STATION NAME          CP NAME
 ------------------------------------------------------------
```

Figure 151. Listing of APPN/HPR Configuration

There were several other options we could have defined: DLUR, additional modes and COS entries, connection networks, and LUs. If the 2216 is serving a LEN node, independent LUs on that node need to be defined to the 2216 so that it can respond to search requests for them.

9.1.2 APPN/HPR Configuration for CS/2 DLUR Node

We configured our CS/2 network node CM5HPRNN with a new connection to the 2216, as well as the DLUR function. This time we show the node definition file (NDF) rather than the GUI panels. For minor changes it is often easier to edit the NDF and issue the CMVERIFY command than to go through the GUI panels. It is also easier to make an error. See Figure 152 on page 167 for the NDF file we used in this test.

```
DEFINE_LOCAL_CP  FQ_CP_NAME(USIBMRA.CM5HPRNN )
                 CP_ALIAS(CM5HPRNN)
                 NAU_ADDRESS(INDEPENDENT_LU)
                 NODE_TYPE(NN)
                 NODE_ID(X'05D00000')
                 NW_FP_SUPPORT(NONE)
                 HOST_FP_SUPPORT(YES)
                 SEARCH_REQUIRED(NO)
                 BRANCH_EXTENDER_SUPPORT(NO)
                 FREE_UNUSED_SESSIONS(NO)
                 FREE_UNUSED_SESSIONS_TIME(10)
                 MAX_COMP_LEVEL(NONE)
                 MAX_COMP_TOKENS(0);

DEFINE_LOGICAL_LINK  LINK_NAME(LINK0001)      1
                     ADJACENT_NODE_TYPE(NN)
                     PREFERRED_NN_SERVER(NO)
                     DLC_NAME(IBMTRNET)
                     ADAPTER_NUMBER(0)
                     DESTINATION_ADDRESS(X'400022160100004')   2
                     ETHERNET_FORMAT(NO)
                     CP_CP_SESSION_SUPPORT(YES)
                     SOLICIT_SSCP_SESSION(NO)
                     USE_PUNAME_AS_CPNAME(NO)
                     MAX_ACTIVATION_ATTEMPTS(USE_ADAPTER_DEFINITION)
                     AUTO_REACTIVATE(NO_RETRY)
                     ACTIVATE_AT_STARTUP(YES)
                     LIMITED_RESOURCE(NO)
                     LINK_STATION_ROLE(USE_ADAPTER_DEFINITION)
                     EFFECTIVE_CAPACITY(USE_ADAPTER_DEFINITION)
                     COST_PER_CONNECT_TIME(USE_ADAPTER_DEFINITION)
                     COST_PER_BYTE(USE_ADAPTER_DEFINITION)
                     SECURITY(USE_ADAPTER_DEFINITION)
                     PROPAGATION_DELAY(USE_ADAPTER_DEFINITION)
                     HPR_SUPPORT(USE_ADAPTER_DEFINITION)
                     HPR_LLERP_SUPPORT(USE_ADAPTER_DEFINITION)
                     HPR_MLTG_NUMBER(0)
                     BRANCH_EXTENDER_UPLINK(USE_ADAPTER_DEFINITION)
                     MAX_I_FIELD_SIZE(USE_ADAPTER_DEFINITION)
                     LIMITED_RESOURCE_TIMEOUT(USE_ADAPTER_DEFINITION)
                     INACTIVITY_TIMEOUT(USE_ADAPTER_DEFINITION)
                     USER_DEFINED_1(USE_ADAPTER_DEFINITION)
                     USER_DEFINED_2(USE_ADAPTER_DEFINITION)
                     USER_DEFINED_3(USE_ADAPTER_DEFINITION);
```

Figure 152 (Part 1 of 2). CS/2 NDF Listing

```
DEFINE_DEPENDENT_LU_SERVER  LINK_NAME(HOST0001)    3
                            FQ_DLUS_NAME(USIBMRA.RAA )    4
                            PU_NAME(MPU00001)
                            NODE_ID(X'05DAA61A')
                            MAX_ACTIVATION_ATTEMPTS(INFINITE)
                            ACTIVATE_AT_STARTUP(YES);

DEFINE_LUA  LU_NAME(@LUA0001)     5
            HOST_LINK_NAME(HOST0001)
            NAU_ADDRESS(2);

DEFINE_LUA  LU_NAME(@LUA0002)     5
            HOST_LINK_NAME(HOST0001)
            NAU_ADDRESS(3);

DEFINE_LUA  LU_NAME(@LUA0003)     5
            HOST_LINK_NAME(HOST0001)
            NAU_ADDRESS(4);

DEFINE_DEFAULTS  IMPLICIT_INBOUND_PLU_SUPPORT(YES)
                 DEFAULT_MODE_NAME(BLANK)
                 MAX_MC_LL_SEND_SIZE(32767)
                 DIRECTORY_FOR_INBOUND_ATTACHES(*)
                 DEFAULT_TP_OPERATION(NONQUEUED_AM_STARTED)
                 DEFAULT_TP_PROGRAM_TYPE(BACKGROUND)
                 DEFAULT_TP_CONV_SECURITY_RQD(NO)
                 MAX_HELD_ALERTS(10)
                 DEFAULT_ROUTING_PREFERENCE(NATIVE_FIRST)
                 RETRY_COUNT(6)
                 ALIVE_TIMER(60)
                 PATH_SWITCH_TIMER_LOW(480)
                 PATH_SWITCH_TIMER_MEDIUM(240)
                 PATH_SWITCH_TIMER_HIGH(120)
                 PATH_SWITCH_TIMER_NET(60)
                 ROUTE_SETUP_TIMEOUT(10)
                 MOBILE(NO)
                 TN3270E_PORT(23)
                 TN3270E_KEEPALIVE_TYPE(NONE)
                 TN3270E_AUTOMATIC_LOGOFF(0)
                 DISABLE_DLUR_REGISTRATION(NO);

START_ATTACH_MANAGER;

SET_DISCOVERY_SERVER  ADAPTER_NUMBER(0)
                      GROUP_NAMES(IROUTSNA)
                      ROUTING_CAPABILITIES(NN);
```

Figure 152 (Part 2 of 2). CS/2 NDF Listing

In this file we can see a real logical link **1** to the 2216, specifying the MAC and SAP address **2** of the 2216′s adapter known as trn0. There is also a DLUR logical link **3** specifying the CP name of the desired DLUS **4**. Defined on the DLUR logical link are three dependent LUs **5**.

9.2 Example with 2216 As ANR Node

Before we started the CS/2 node, we used the ASCII emulator on the 2216's service port to display some information about its connections. To get to the APPN display prompt, type `talk 5` at the main (*) prompt, then `protocol APPN`. Figure 153 shows the list of available APPN display commands, which we were able to produce by typing `list ?` as shown. The commands can all be abbreviated as the examples show.

```
2216T3 APPN >list ?
CP-CP_SESSIONS
ISR_SESSIONS
SESSION_INFORMATION
RTP_CONNECTIONS
PORT_INFORMATION
LINK_INFORMATION
FOCAL_POINT
APPC_SESSIONS
DUMPS
```

Figure 153. List of Available APPN Displays on 2216

We issued `list link` to see the active APPN links, as shown in Figure 154.

```
2216T3 APPN >list link
    Name   Port Name  Intf       Adj CP Name  Type       HPR      State
=========================================================================
    T61A       TRN1    1     USIBMRA.NNP61A    NN    ACTIVE     ACT_LS
    T41A       TRN1    1     USIBMRA.NNP41A    NN    ACTIVE     ACT_LS
```

Figure 154. Active Links on 2216 before CS/2 Started

The 2216, as expected, had direct connections to both 3746s. To display the active CP-CP sessions, we issued `list cp-cp_sessions` as seen in Figure 155.

```
2216T3 APPN >list cp-cp sessions
           CP Name      Type    Status  Connwinner ID    Conloser ID
========================================================================
    USIBMRA.NNP41A       NN     Active    34E22E1C        34E22E1B
    USIBMRA.NNP61A       NN     Active    34E1DA4E        34E1DA4D
```

Figure 155. Active CP-CP Sessions on 2216

Similarly, we saw the active RTP connections using `list rtp_connections`, as in Figure 156 on page 170.

```
2216T3 APPN >list rtp connections
RTP CONNECTION TABLE:
  TCID            CP Name  ISR  APPC  Pathswitch    Alive    COS TPF   T
=========================================================================
1E2B510    USIBMRA.NNP61A    0    1        180       180    CPSVCMG
1E2C2A8    USIBMRA.NNP61A    0    1        180        90    CPSVCMG
1E149F8    USIBMRA.NNP41A    0    1        180       180    CPSVCMG
1E2B998    USIBMRA.NNP41A    0    1        180        90    CPSVCMG
```

Figure 156. Active RTP Connections on 2216

Both the 2216 and the 3746 support Control Flows over RTP, and each adjacent pair of nodes has initiated CP-CP session activation from each side concurrently. Therefore, we have four RTP pipes for four sessions.

Another display (`list appc_sessions`) shows us all the LU 6.2 sessions (Figure 157), which are in fact all the CP-CP sessions at this point.

```
2216T3 APPN >list appc sessions
LU Name             Mode  Type  FSM
======================================
USIBMRA.NNP61A    CPSVCMG  Pri  ACT
USIBMRA.NNP61A    CPSVCMG  Sec  ACT
USIBMRA.NNP41A    CPSVCMG  Pri  ACT
USIBMRA.NNP41A    CPSVCMG  Sec  ACT
```

Figure 157. Active LU 6.2 Sessions before CS/2 Connection

Next, we verified the same information from one of the 3746s, NNP41A. Figure 158 shows the list of link stations known to NNP41A.

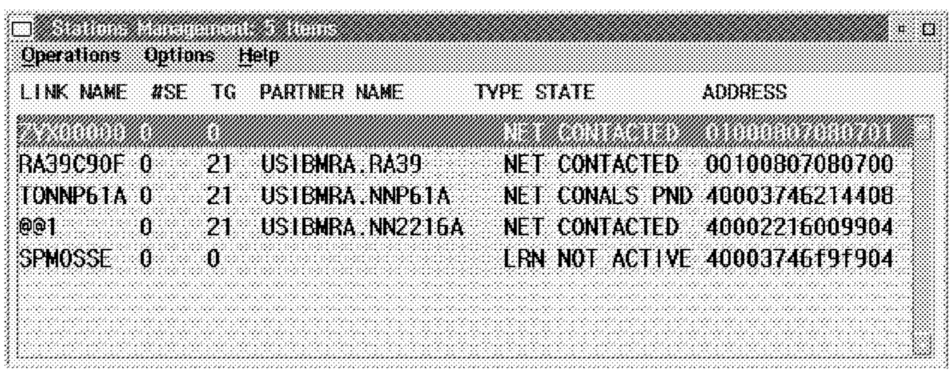

Stations Management: 5 Items

Operations Options Help

LINK NAME	#SE	TG	PARTNER NAME	TYPE	STATE	ADDRESS
ZYX00000	0	0		NET	CONTACTED	01000807080701
RA39C90F	0	21	USIBMRA.RA39	NET	CONTACTED	00100807080700
TONNP61A	0	21	USIBMRA.NNP61A	NET	CONALS PND	40003746214408
@@1	0	21	USIBMRA.NN2216A	NET	CONTACTED	40002216009904
SPMOSSE	0	0		LRN	NOT ACTIVE	40003746f9f904

Figure 158. Active Stations on NNP41A

We can see that the connection to NNP61A is not active at this time, but the connections to VTAM RA39 and to the 2216 are active. The 2216 connection has been defined dynamically by the 3746 (the 2216 initiated it and it was not predefined in the 3746), so it has a 3746-assigned name of @@1.

9.2.1 CS/2 As DLUR and RTP Endpoint

Now we started the CS/2 node without (yet) starting any dependent LU sessions.

Using the CS/2 Subsystem Management facility, we displayed the active logical links from CM5HPRNN as shown in Figure 159.

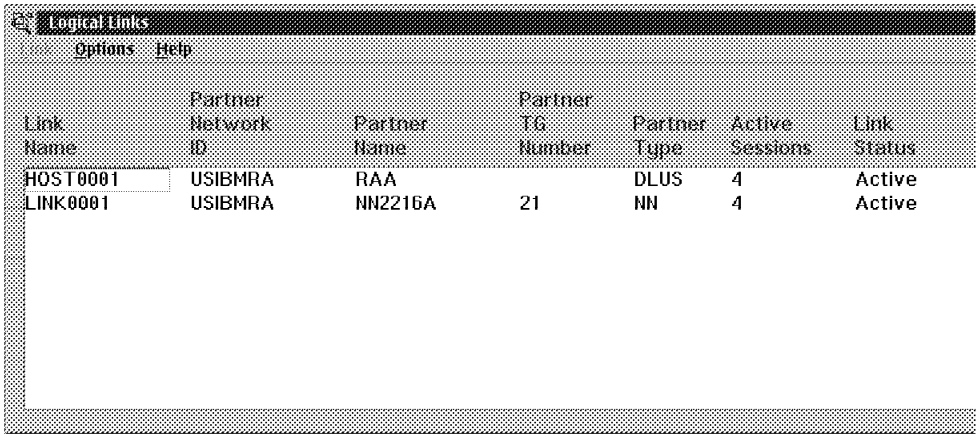
Logical Links

Options Help

Link Name	Partner Network ID	Partner Name	Partner TG Number	Partner Type	Active Sessions	Link Status
HOST0001	USIBMRA	RAA		DLUS	4	Active
LINK0001	USIBMRA	NN2216A	21	NN	4	Active

Figure 159. Logical Link Display

This shows that the CS/2 node has an active real link to the 2216 (TG 21) and an active DLUR/S pipe to RAA.

A display of the LU 6.2 sessions known to CM5HPRNN shows (Figure 160) that it has a pair of CP-CP sessions with NN2216A and a pair of DLUR/S sessions with RAA.

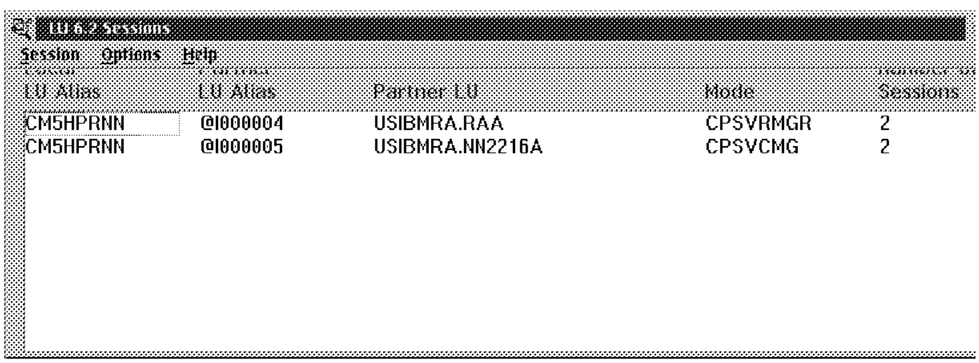
LU 6.2 Sessions

Session Options Help

LU Alias	LU Alias	Partner LU	Mode	Sessions
CM5HPRNN	@I000004	USIBMRA.RAA	CPSVRMGR	2
CM5HPRNN	@I000005	USIBMRA.NN2216A	CPSVCMG	2

Figure 160. LU 6.2 Sessions

Another display, this time of the RTP connections (Figure 161 on page 172) shows two independent CP-CP pipes to the adjacent 2216. This is usual for adjacent NNs. The DLUR/S sessions are carried on an RTP pipe with APPN COS SNASVCMG. The long-lived pipe (TCID 12F) was used to set up the DLUR/S sessions across the link to NN2216A.

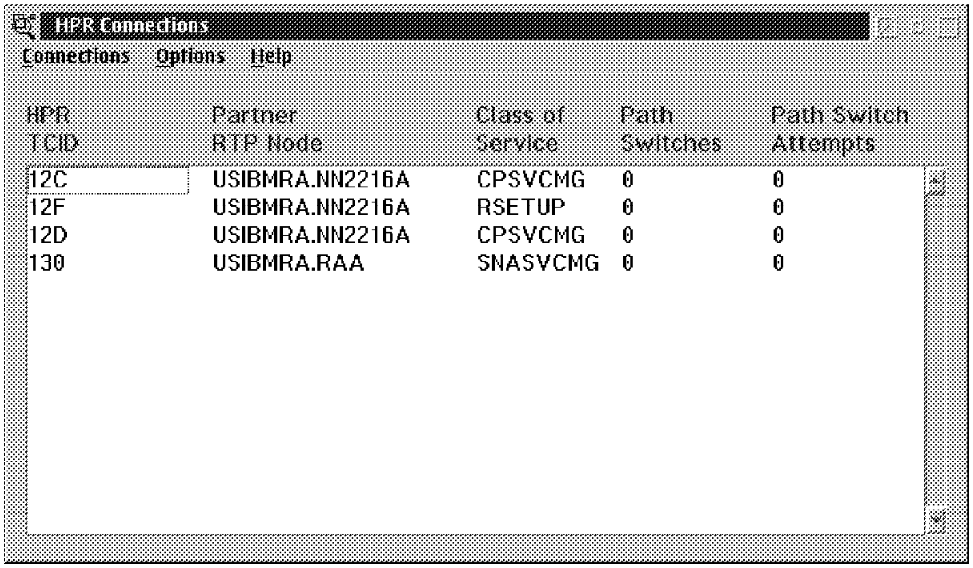

Figure 161. HPR Connections

Next, we established a dependent LU session from CM5HPRNN to an application on RAA, and another to an application on RA39. Figure 162 shows the RTP connections from RAA at this stage.

```
10:11:14 DISPLAY NET,ID=ISTRTPMN,SCOPE=ALL
10:11:14 IST097I  DISPLAY  ACCEPTED
10:11:14 IST075I  NAME = ISTRTPMN         , TYPE = RTP MAJOR NODE
10:11:14 IST486I  STATUS= ACTIV     , DESIRED STATE= ACTIV
10:11:14 IST1486I  RTP NAME   STATE         DESTINATION CP   MNPS  TYPE
10:11:14 IST1487I  CNR007FF   CONNECTED     USIBMRA.CM5HPRNN NO    LULU   6
10:11:14 IST1487I  CNR007FE   CONNECTED     USIBMRA.CM5HPRNN NO    LULU   7
10:11:14 IST1487I  CNR007F5   CONNECTED     USIBMRA.NNP61A   NO    RSTP   8
10:11:14 IST1487I  CNR007F2   CONNECTED     USIBMRA.NNP61A   NO    RSTP   8
10:11:14 IST1487I  CNR007F1   CONNECTED     USIBMRA.NNP61A   NO    CPCP
10:11:14 IST1487I  CNR00795   CONNECTED     USIBMRA.RAK      NO    LULU
10:11:14 IST1487I  CNR00780   CONNECTED     USIBMRA.RA39     NO    RSTP
10:11:14 IST1487I  CNR0077F   CONNECTED     USIBMRA.RA39     NO    CPCP
10:11:14 IST1487I  CNR0075E   CONNECTED     USIBMRA.RAK      NO    RSTP
10:11:14 IST1487I  CNR0075D   CONNECTED     USIBMRA.RAK      NO    CPCP
10:11:14 IST314I  END
```

Figure 162. RTP Connections from RAA

This shows two HPR connections **6** and **7** to CM5HPRNN. One is the DLUR/S pipe and the other is the pipe carrying the dependent LU session, which has a different APPN COS. Note also the two Route Setup pipes connecting RAA with NNP61A **8**. This is unusual because there is only one physical link between them, but it could happen if each end needs to establish an LU-LU session at the same time.

To see which HPR pipe to CM5HPRNN was which, we displayed CNR007FF from RAA (see Figure 163 on page 173).

```
DISPLAY NET,ID=CNR007FF,SCOPE=ALL
IST097I  DISPLAY  ACCEPTED
IST075I  NAME = CNR007FF         , TYPE = PU_T2.1
IST1392I  DISCNTIM = 00010 DEFINED AT PU   FOR DISCONNECT
IST486I  STATUS= ACTIV--LX-, DESIRED STATE= ACTIV
IST1043I  CP NAME = CM5HPRNN, CP NETID = USIBMRA , DYNAMIC LU = YES
IST1589I  XNETALS = YES
IST933I  LOGMODE=***NA***, COS=#CONNECT
IST1476I  TCID X'310BB5B5000000BC' - REMOTE TCID X'000000000000012B'
IST1481I  DESTINATION CP USIBMRA.CM5HPRNN  - NCE X'80             '
IST1587I  ORIGIN NCE X'D000000000000000'
IST1477I  ALLOWED DATA FLOW RATE =    399 KBITS/SEC
IST1516I  INITIAL DATA FLOW RATE =    399 KBITS/SEC
IST1511I  MAXIMUM NETWORK LAYER PACKET SIZE =   2048 BYTES
IST1478I  NUMBER OF UNACKNOWLEDGED BUFFERS =     0
IST1479I  RTP CONNECTION STATE = CONNECTED            - MNPS = NO
IST1480I  RTP END TO END ROUTE - PHYSICAL PATH
IST1460I  TGN  CPNAME             TG TYPE      HPR
IST1461I   21  USIBMRA.NNP61A     APPN         RTP
IST1461I   21  USIBMRA.NN2216A    APPN         RTP
IST1461I   21  USIBMRA.CM5HPRNN   APPN         RTP
IST231I  RTP      MAJOR NODE = ISTRTPMN
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I  STATE TRACE = OFF
IST355I  LOGICAL UNITS:
IST080I  WAA61A04 ACT/S---X-
IST314I  END
```

Figure 163. Dependent LU RTP Pipe

This pipe represents the dependent LU session, and takes the following path:

```
┌──────────────────────┐          ┌────────┐        ┌───────┐       ┌─────┐
│(WAA61A04)   CM5HPRNN ├───TG21───┤NN2216A ├──TG21──┤NNP61A ├─TG21──┤ RAA │
└──────────────────────┘          └────────┘        └───────┘       └─────┘
```

Displaying the same connection from the CS/2 node shows the same path, as in Figure 164 on page 174.

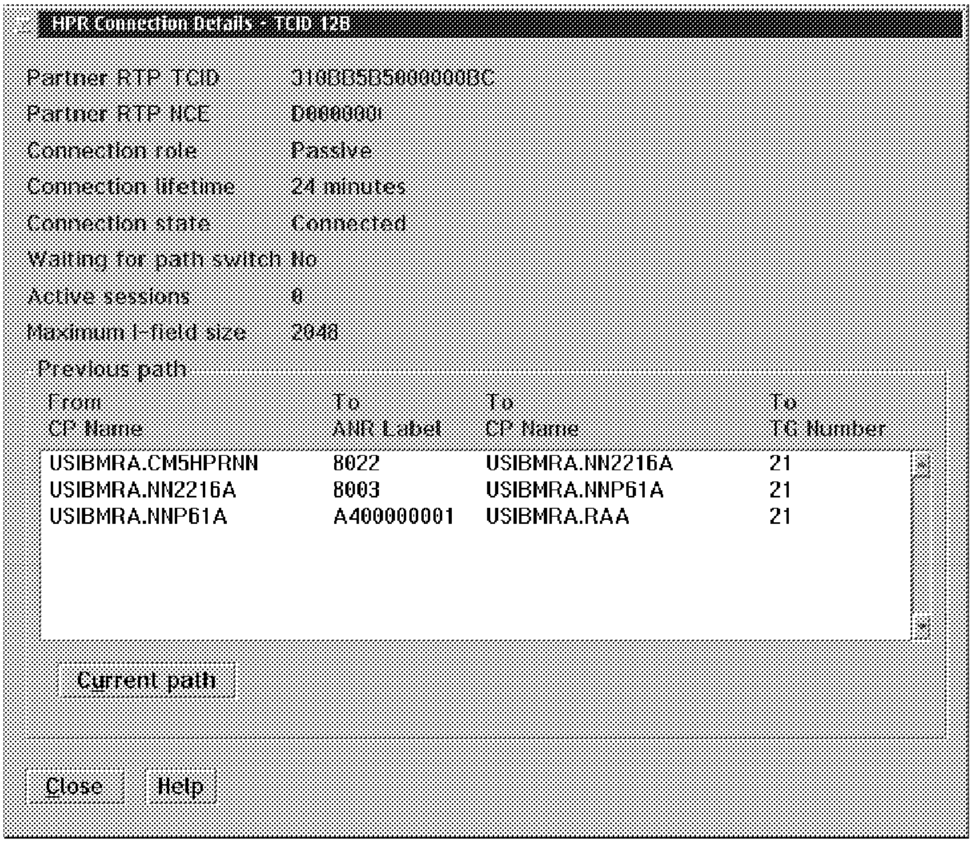

Figure 164. Dependent LU Pipe from CS/2

The other HPR pipe between RAA and CM5HPRNN is confirmed as the DLUR/S pipe by Figure 165 on page 175. It has APPN COS SNASVCMG and carries one or more sessions to the CS/2 CP itself, over the same route as the other pipe.

```
DISPLAY NET,ID=CNR007FE,SCOPE=ALL
IST097I  DISPLAY  ACCEPTED
IST075I  NAME = CNR007FE         , TYPE = PU_T2.1
IST1392I  DISCNTIM = 00010 DEFINED AT PU   FOR DISCONNECT
IST486I  STATUS= ACTIV--LX-, DESIRED STATE= ACTIV
IST1043I  CP NAME = CM5HPRNN, CP NETID = USIBMRA , DYNAMIC LU = YES
IST1589I  XNETALS = YES
IST933I  LOGMODE=***NA***, COS=SNASVCMG
IST1476I  TCID X'310BB5B4000000CE' - REMOTE TCID X'0000000000000130'
IST1481I  DESTINATION CP USIBMRA.CM5HPRNN  - NCE X'80               '
IST1587I  ORIGIN NCE X'D000000000000000'
IST1477I  ALLOWED DATA FLOW RATE =    399 KBITS/SEC
IST1516I  INITIAL DATA FLOW RATE =    399 KBITS/SEC
IST1511I  MAXIMUM NETWORK LAYER PACKET SIZE =   2048 BYTES
IST1478I  NUMBER OF UNACKNOWLEDGED BUFFERS =     0
IST1479I  RTP CONNECTION STATE = CONNECTED             - MNPS = NO
IST1480I  RTP END TO END ROUTE - PHYSICAL PATH
IST1460I  TGN  CPNAME             TG TYPE      HPR
IST1461I   21  USIBMRA.NNP61A     APPN         RTP
IST1461I   21  USIBMRA.NN2216A    APPN         RTP
IST1461I   21  USIBMRA.CM5HPRNN   APPN         RTP
IST231I  RTP      MAJOR NODE = ISTRTPMN
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I  STATE TRACE = OFF
IST355I  LOGICAL UNITS:
IST080I  CM5HPRNN ACT/S----Y
```

Figure 165. DLUR/S RTP Pipe

Similarly, a display from the Subsystem Management panel of CS/2 shows us the DLUR/S pipe from the other end, as in Figure 166 on page 176.

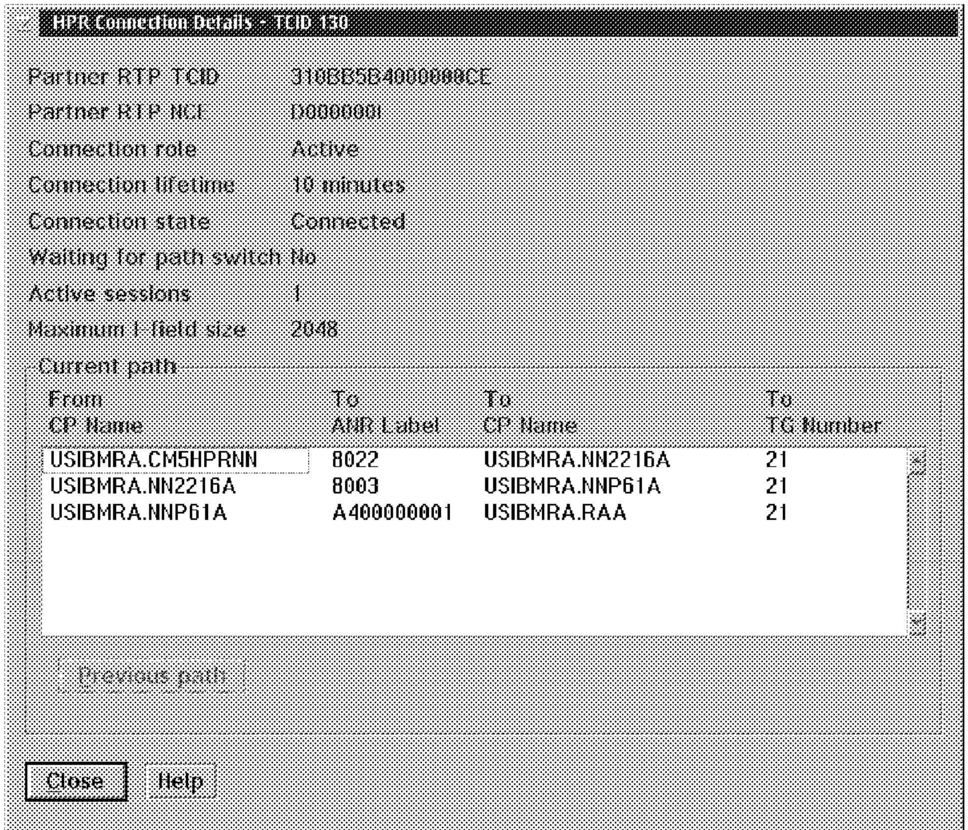

Figure 166. DLUR/S RTP Connection

We investigated the other dependent LU session (from CM5HPRNN) to RA39 in a similar fashion, but the displays are not shown here because they show nothing new. The session was on an RTP pipe, known as CNR0005B by RA39, which took the following path:

```
┌───────────────────────┐       ┌────────┐     ┌───────┐    ┌──────┐
│(WAA61A02)    CM5HPRNN ├─TG21──┤NN2216A ├─TG21┤NNP41A ├TG21┤RA39  │
└───────────────────────┘       └────────┘     └───────┘    └──────┘
```

At this time we also took displays in the 2216 to see how the picture had changed. Figure 167 on page 177 showed the connections that were created after the CS/2 node was linked into the network.

```
2216T3 APPN >list cp-cp sessions
          CP Name      Type    Status  Connwinner ID    Conloser ID
========================================================================
    USIBMRA.NNP41A      NN     Active    34E22E1C        34E22E1B
  USIBMRA.CM5HPRNN      NN     Active    34E230A7        34E230A5
    USIBMRA.NNP61A      NN     Active    34E1DA4E        34E1DA4D
*
2216T3 APPN >list rtp connections
RTP CONNECTION TABLE:
  TCID             CP Name  ISR  APPC  Pathswitch    Alive    COS TPF   T
=========================================================================
1E2B510    USIBMRA.NNP61A    0    1        180       180    CPSVCMG
1E2C2A8    USIBMRA.NNP61A    0    1        180        90    CPSVCMG
1E149F8    USIBMRA.NNP41A    0    1        180       180    CPSVCMG
1E2B998    USIBMRA.NNP41A    0    1        180        90    CPSVCMG
1E2C730  USIBMRA.CM5HPRNN    0    1        180        61    CPSVCMG  9
1E47B10  USIBMRA.CM5HPRNN    0    1        180       180    CPSVCMG  9
1E47200  USIBMRA.CM5HPRNN    0    0          0        61     RSETUP  10
1E47688    USIBMRA.NNP61A    0    0          0        90     RSETUP
*
2216T3 APPN >list appc sessions
LU Name            Mode  Type  FSM
======================================
USIBMRA.NNP61A    CPSVCMG  Pri  ACT
USIBMRA.NNP61A    CPSVCMG  Sec  ACT
USIBMRA.NNP41A    CPSVCMG  Pri  ACT
USIBMRA.NNP41A    CPSVCMG  Sec  ACT
USIBMRA.CM5HPRNN  CPSVCMG  Pri  ACT   11
USIBMRA.CM5HPRNN  CPSVCMG  Sec  ACT   11
```

Figure 167. Displays Issued on 2216 after Session Establishment

The only new logical link was to the CS/2 node CM5HPRNN. Using this link were the CP-CP **9** and long-lived RTP pipes **10**, and the CP-CP sessions **11**. The 2216 was not aware of the dependent LU sessions or of the RTP connections that they were using.

9.2.2 Path Switch

We physically unplugged the connection between the 2216 and NNP61A to see what happened to the sessions using it. Those sessions included the DLUR/S pipe from RAA to CM5HPRNN as well as the dependent LU session going to RAA. In Figure 168 you can see what NNP61A thought of the active connections after this event.

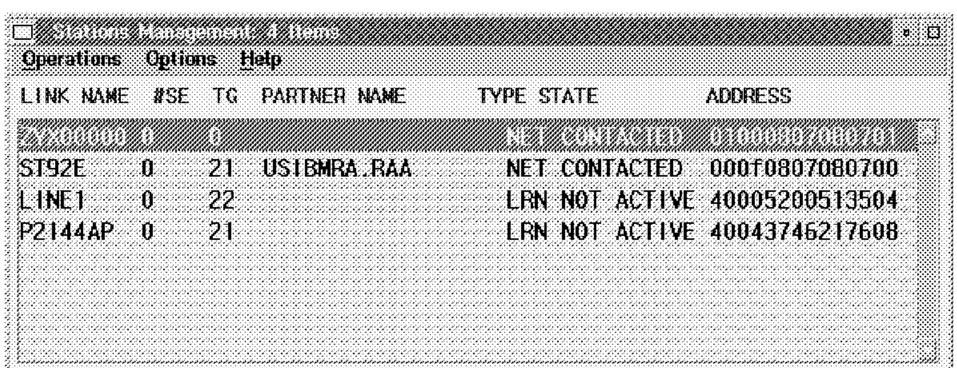

Figure 168. Active Stations on NNP61A after Breaking the Token-Ring

A display of the HPR connections from CS/2 showed that a path switch had occurred for the RTP pipes with TCID 130 and TCID 12B (Figure 169 on page 178).

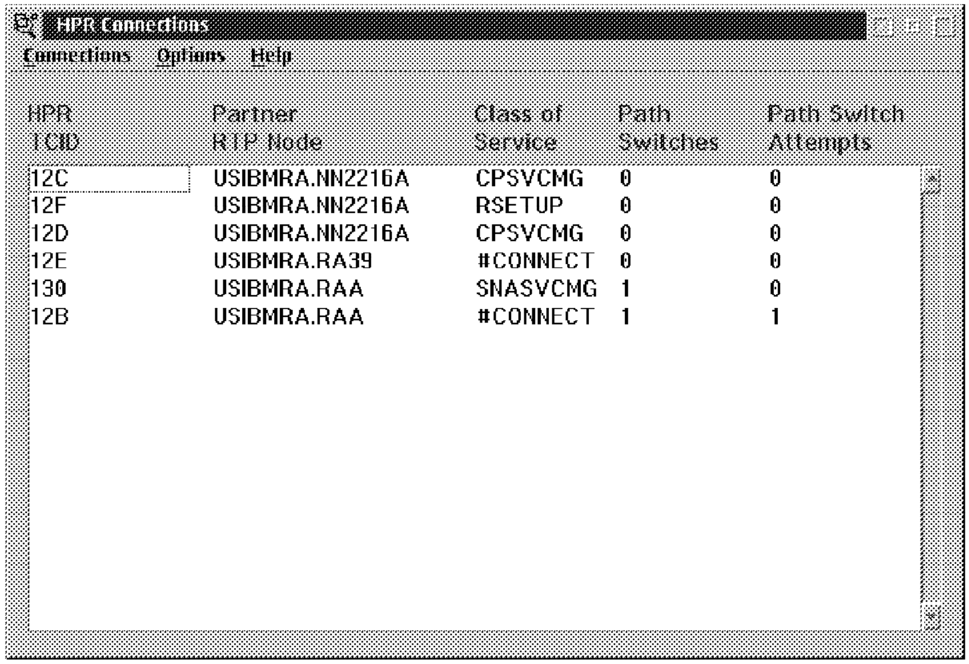

Figure 169. HPR Connections after Path Switch

RAA′s log told a similar tale (Figure 170).

```
10:21:26   IST1494I  PATH SWITCH STARTED   FOR RTP CNR007FE
10:21:26   IST1494I  PATH SWITCH COMPLETED FOR RTP CNR007FE
10:21:26   IST1480I  RTP END TO END ROUTE - PHYSICAL PATH
10:21:26   IST1460I  TGN  CPNAME             TG TYPE      HPR
10:21:26   IST1461I   21  USIBMRA.RA39       APPN         RTP
10:21:26   IST1461I   21  USIBMRA.NNP41A     APPN         RTP
10:21:26   IST1461I   21  USIBMRA.NN2216A    APPN         RTP
10:21:26   IST1461I   21  USIBMRA.CM5HPRNN   APPN         RTP
10:21:26   IST314I  END
```

Figure 170. Path Switch on RAA Log

A display of the newly switched LU-LU session pipe, CNR007FF, from RAA shows the new route (Figure 171 on page 179).

```
DIS CNR007FF
DISPLAY NET,ID=CNR007FF,SCOPE=ALL
IST097I  DISPLAY  ACCEPTED
IST075I  NAME = CNR007FF         , TYPE = PU_T2.1
IST1392I  DISCNTIM = 00010 DEFINED AT PU   FOR DISCONNECT
IST486I  STATUS= ACTIV--LX-, DESIRED STATE= ACTIV
IST1043I  CP NAME = CM5HPRNN, CP NETID = USIBMRA , DYNAMIC LU = YES
IST1589I  XNETALS = YES
IST933I  LOGMODE=***NA***, COS=#CONNECT
IST1476I  TCID X'310BB5B5000000BC' - REMOTE TCID X'000000000000012B'
IST1481I  DESTINATION CP USIBMRA.CM5HPRNN  - NCE X'80              '
IST1587I  ORIGIN NCE X'D000000000000000'
IST1477I  ALLOWED DATA FLOW RATE =    399 KBITS/SEC
IST1516I  INITIAL DATA FLOW RATE =    399 KBITS/SEC
IST1511I  MAXIMUM NETWORK LAYER PACKET SIZE =   2048 BYTES
IST1478I  NUMBER OF UNACKNOWLEDGED BUFFERS =     0
IST1479I  RTP CONNECTION STATE = CONNECTED             - MNPS = NO
IST1480I  RTP END TO END ROUTE - PHYSICAL PATH
IST1460I  TGN  CPNAME             TG TYPE      HPR
IST1461I   21  USIBMRA.RA39       APPN         RTP
IST1461I   21  USIBMRA.NNP41A     APPN         RTP
IST1461I   21  USIBMRA.NN2216A    APPN         RTP
IST1461I   21  USIBMRA.CM5HPRNN   APPN         RTP
IST231I  RTP      MAJOR NODE = ISTRTPMN
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I  STATE TRACE = OFF
IST355I  LOGICAL UNITS:
IST080I  WAA61A04 ACT/S---X-
IST314I  END
```

Figure 171. Newly Switched RTP Pipe

The new route for both the DLUR/S sessions and the dependent LU session is shown in Figure 172 on page 180.

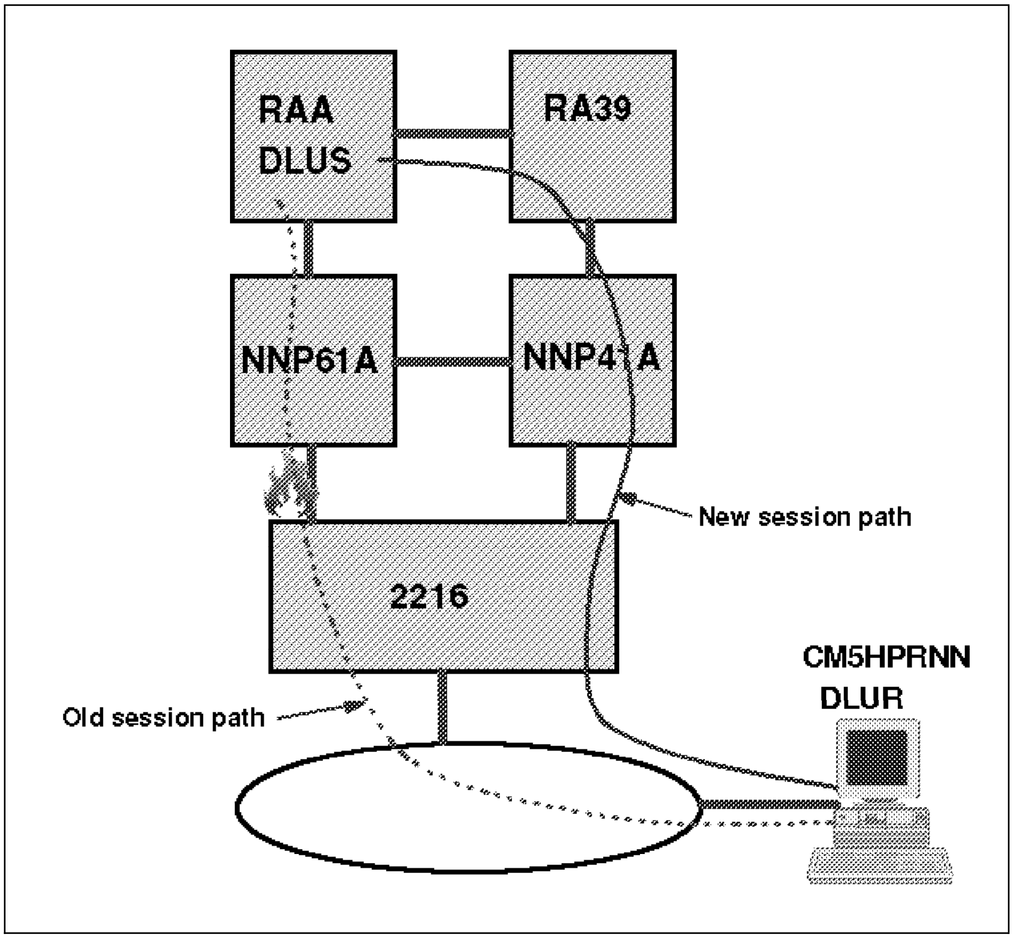

Figure 172. New Path after Token-Ring Break

The HPR Connection Details display from CS/2 (Figure 173 on page 181) confirms this route.

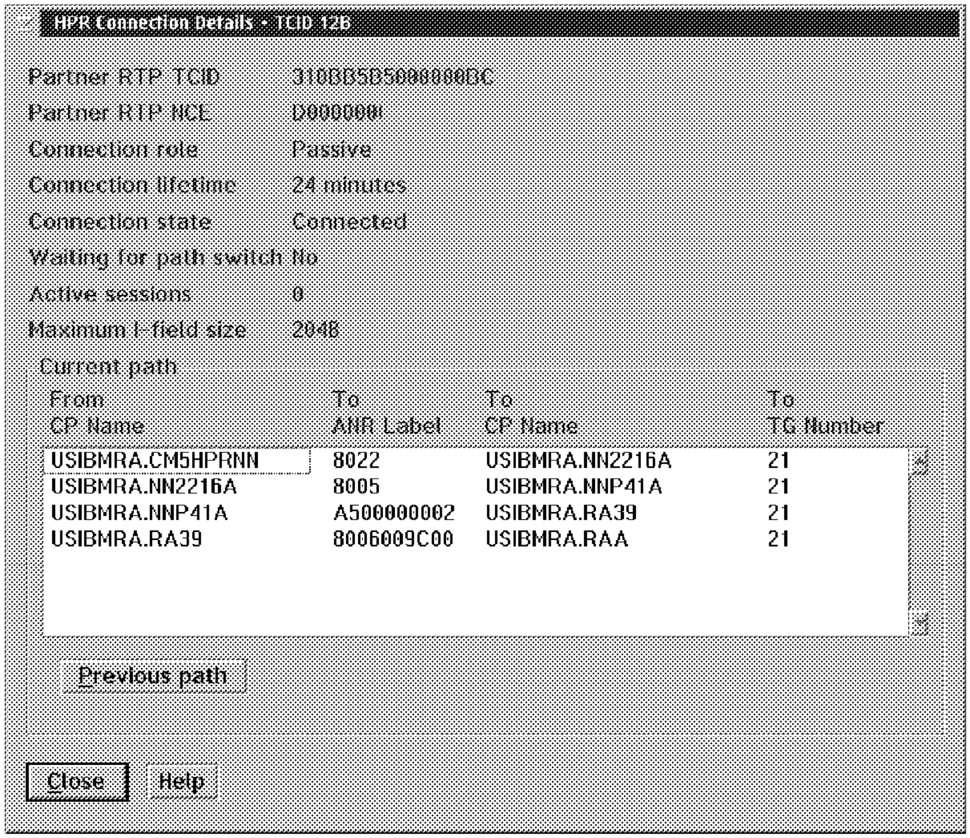

Figure 173. HPR Connection after Path Switch

Finally, we displayed the RTP connections once again from the 2216 (Figure 174).

```
2216T3 APPN >list rtp connections
RTP CONNECTION TABLE:
  TCID            CP Name  ISR  APPC  Pathswitch    Alive    COS TPF   T
========================================================================
1E149F8    USIBMRA.NNP41A    0     1         180       180    CPSVCMG
1E2B998    USIBMRA.NNP41A    0     1         180        90    CPSVCMG
1E2C730  USIBMRA.CM5HPRNN    0     1         180        61    CPSVCMG
1E47B10  USIBMRA.CM5HPRNN    0     1         180       180    CPSVCMG
1E47200  USIBMRA.CM5HPRNN    0     0           0        61     RSETUP
1E46D78    USIBMRA.NNP41A    0     0           0       180     RSETUP
```

Figure 174. Active RTP Connections on 2216 after TG to NNP61A Failed

As the 2216 was an ANR node for the switched sessions, it recorded no information about those sessions or their RTP pipes. The only difference now is that the CP-CP and long-lived pipes to NNP61A have disappeared.

9.3 Configuration with 2216 As DLUR Node

The configuration for this test is depicted in Figure 175 on page 182. This time the CS/2 node is configured as it would be for a direct NCP attachment, with dependent LUs and no APPN support. The 2216 is providing the DLUR function, and appears as a subarea node to CS/2.

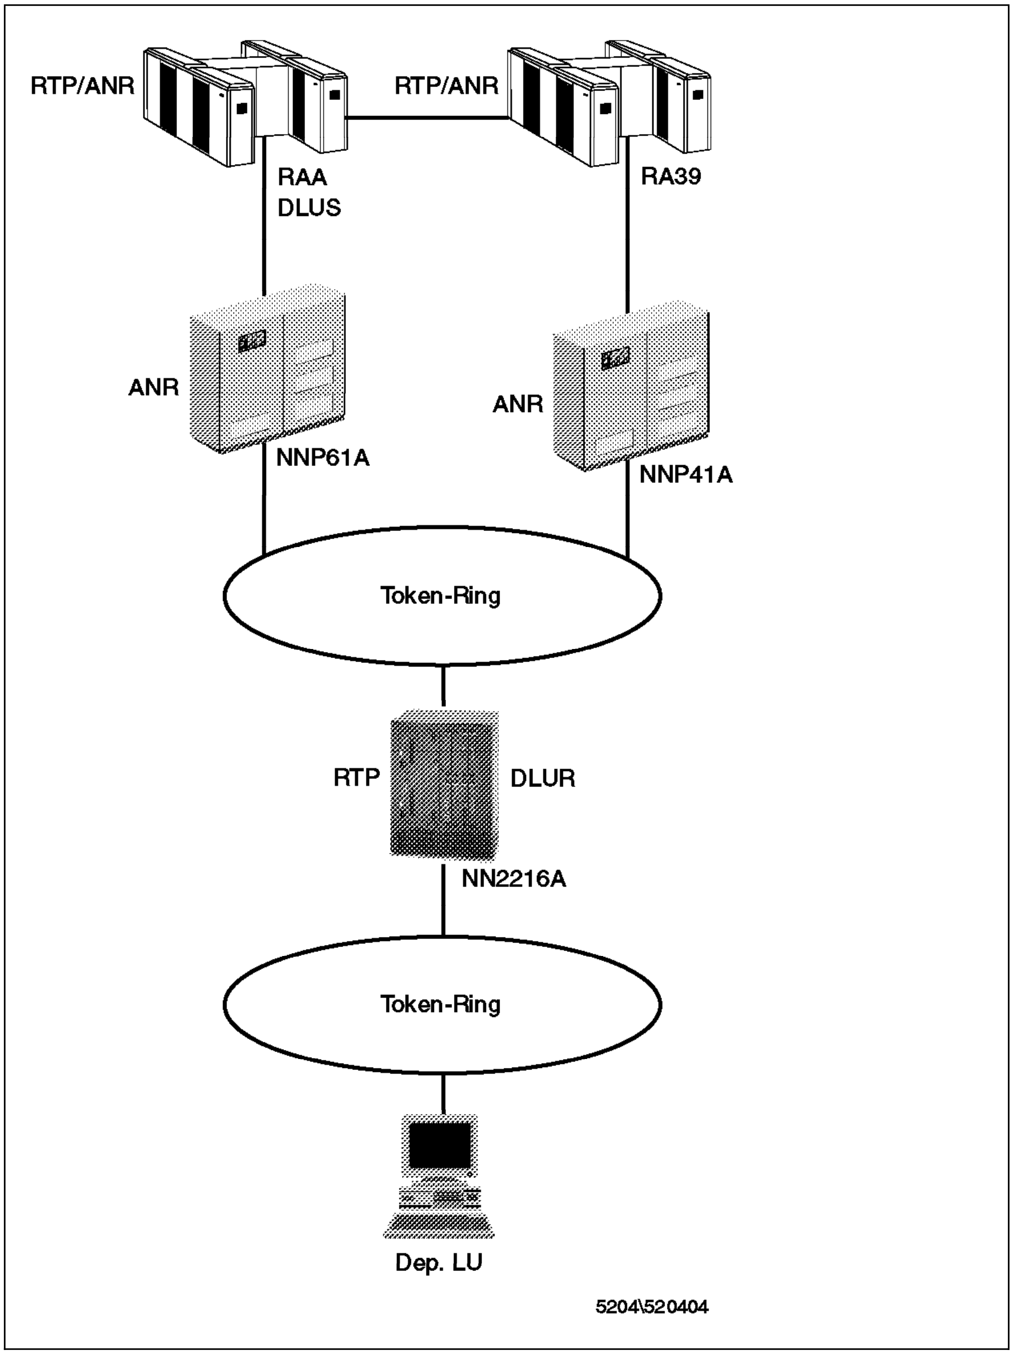

Figure 175. 2216 As DLUR and RTP Endpoint

9.3.1 DLUR Configuration for 2216 Router

DLUR capability is not enabled by default on the 2216. To implement DLUR, you must enable it at the node level. You then have the same choices as in the 3746: you can define a primary (and backup, if needed) DLUS at the node level and override this information for each attached node if you so wish. If your DLUS is to initiate the DLUR/S connections, then you do not have to define the DLUS name(s) in the 2216, but you always have to enable the function at the node level.

See Figure 176 on page 183 for details of how we enabled DLUR for our 2216.

```
APPN config>set dlur
Enable DLUR (Y)es (N)o [N]? y   1
Fully-qualified CP name of primary DLUS []?usibmra.raa 2
Fully-qualified CP name of backup DLUS []? usibmra.ra39  3
Perform retries to restore disrupted pipe [Y]?
Delay before initiating retries(0-2756000 seconds) [120]?
Perform short retries to restore disrupted pipe [Y]?
Short retry timer(0-2756000 seconds) [120]?
Short retry count(0-65535) [5]?
Perform long retry to restore disrupted pipe [Y]?
Long retry timer(0-2756000 seconds) [300]?
Write this record? [Y]?
The record has been written.
APPN config>
```

Figure 176. Enabling DLUR on 2216

1 is the minimum you need to enter. We also entered RAA as a primary 2216-wide DLU server **2** and RA39 as a backup server **3**.

The Perform retries settings determine how the 2216 attempts to recover a failed DLUR/S pipe. A full description of the algorithm is in *Multiprotocol Access Services Protocol Configuration and Monitoring Reference*. The action taken by the 2216 differs depending on the cause of the failure: a nondisruptive UNBIND from the DLUS (less frequent attempts at recovery) or any other cause (more frequent attempts at recovery).

Because we only used one downstream node that required DLUR services from the 2216, these definitions were enough for our purposes. Figure 177 shows how you might override the DLUR/S configuration parameters for a particular link.

```
APPN config>add link-station
APPN Station
Port name for the link station [ ]? trn1
Station name (Max 8 characters) [ ]? t41a
Activate link automatically (Y)es (N)o [Y]?
MAC address of adjacent node [000000000000]? 400052005115
Adjacent node type: 0 = APPN network node,
1 = APPN end node or Unknown node type                    4
2 = LEN end node, 3 = PU 2.0 node [0]?
High performance routing: (Y)es (N)o [Y]?
Edit Dependent LU Server: (Y)es (N)o [N]? y          5
Fully-qualified CP name of primary DLUS [USIBMRA.RAA]? usibmra.RA39   6
Fully-qualified CP name of backup DLUS [USIBMRA.RA39]? usibmra.raa    7
Allow CP-CP sessions on this link (Y)es (N)o [Y]?
CP-CP session level security (Y)es (N)o [N]?
Configure CP name of adjacent node: (Y)es (N)o [N]?
Edit TG Characteristics: (Y)es (N)o [N]?
Edit LLC Characteristics: (Y)es (N)o [N]?
Edit HPR defaults: (Y)es (N)o [N]?
Write this record? [Y]?
The record has been written.
APPN config>
```

Figure 177. Specifying a Different DLUS for a Station

Here we take the station whose MAC address is 400052005115 and assign it a primary DLUS of RA39 **6** and a backup DLUS of RAA **7**. To get these questions asked, you need to respond y to the question **5** about editing the

DLUS parameters. Note also that the node with the dependent LUs can be NN, EN, LEN or simply type 2.0 **4**.

Apart from recustomizing the 2216, we configured our PC as a LEN node with CP name PU05170 and four dependent LUs on a single logical link to the 2216. In fact, CS/2 cannot be configured specifically as a LEN node; what you have to do is configure it as an APPN node but turn off APPN support on the host (2216 in this case) link.

9.4 Example with 2216 As DLUR Node

Before we started the CS/2 node, we performed some displays on the 2216, as shown in Figure 178.

```
2216T3 APPN >list cp-cp
           CP Name      Type    Status  Connwinner ID    Conloser ID
========================================================================
    USIBMRA.NNP61A       NN     Active    34E2D7F7        34E2D7F9   8
    USIBMRA.NNP41A       NN     Active    34E2D7FC        34E2D7FB
*
2216T3 APPN >list isr sessions
No ISR information available    9
*
2216T3 APPN >list rtp connections
RTP CONNECTION TABLE:
  TCID            CP Name  ISR  APPC  Pathswitch    Alive    COS TPF   T
=========================================================================
1E14570    USIBMRA.NNP61A    0    1         180       180    CPSVCMG
1E14E80    USIBMRA.NNP61A    0    1         180        90    CPSVCMG
1E2D9A8    USIBMRA.NNP41A    0    1         180       180    CPSVCMG    10
1E2E2B8    USIBMRA.NNP41A    0    1         180        90    CPSVCMG
*
2216T3 APPN >list link
    Name   Port Name  Intf       Adj CP Name  Type       HPR      State
=========================================================================
    T61A       TRN1    1     USIBMRA.NNP61A    NN    ACTIVE     ACT_LS   11
    T41A       TRN1    1     USIBMRA.NNP41A    NN    ACTIVE     ACT_LS
*
2216T3 APPN >list appc sessions
LU Name             Mode  Type  FSM
======================================
USIBMRA.NNP61A    CPSVCMG  Pri  ACT   12
USIBMRA.NNP41A    CPSVCMG  Pri  ACT
USIBMRA.NNP61A    CPSVCMG  Sec  ACT
USIBMRA.NNP41A    CPSVCMG  Sec  ACT
*
```

Figure 178. Displays Issued on 2216 before CS/2 Activation

At this stage the 2216 has CP-CP sessions with the 3746s **8**. There are no ISR sessions (meaning sessions passing through the 2216) **9**. The only RTP connections active are the CP-CP pipes **10**, because no LU-LU sessions have been started yet. The DLUR/S pipe will not be set up until some dependent LUs arrive on the scene. The only LU 6.2 sessions are the four CP-CP sessions **12**, and there is not yet a connection to CS/2 **11**.

We then started the CS/2 node and performed some displays using the Subsystem Management Facility.

The only link active on the CS/2 node, as shown in Figure 179 on page 185, is that to the 2216 NN. The five sessions shown are the SSCP-PU and four SSCP-LU sessions.

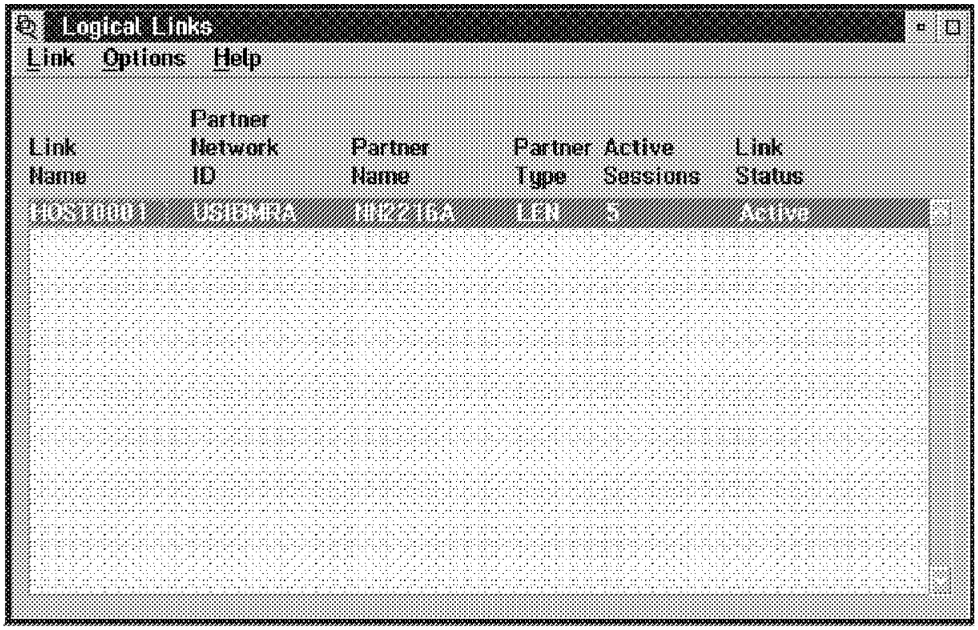

Figure 179. Logical Links on LEN Node

Displaying the details of this link (Figure 180 on page 186) confirms the MAC/SAP address of the 2216′s downstream port and the fact that CP-CP sessions are not supported.

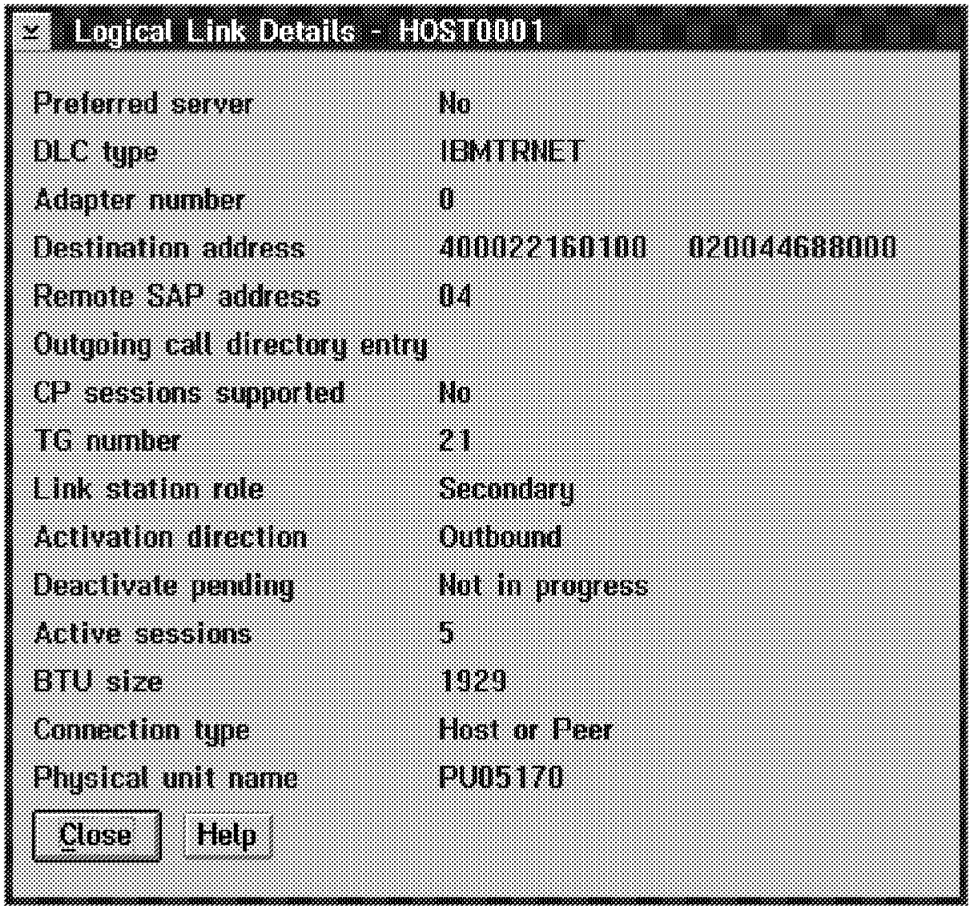

Figure 180. Logical Link Details

As soon as the link from CS/2 to 2216 was started, we displayed the 2216 connection status as seen in Figure 181 on page 187.

```
2216T3 APPN >list link
    Name   Port Name  Intf       Adj CP Name  Type       HPR      State
=========================================================================
    T61A        TRN1    1     USIBMRA.NNP61A    NN    ACTIVE     ACT_LS
    T41A        TRN1    1     USIBMRA.NNP41A    NN    ACTIVE     ACT_LS
     @@0  1     TRN0    0    USIBMRA.PU05170    EN  INACTIVE     ACT_LS
*
2216T3 APPN >list cp-cp  2
           CP Name      Type    Status  Connwinner ID    Conloser ID
========================================================================
    USIBMRA.NNP61A       NN     Active    34E2D7F7        34E2D7F9
    USIBMRA.NNP41A       NN     Active    34E2D7FC        34E2D7FB
*
2216T3 APPN >list rtp connections
RTP PARTNER TABLE:
 Remote Partner Name  Remote Boundary Name  TG Number
======================================================
         USIBMRA.RAA           USIBMRA.RAA     -1
RTP CONNECTION TABLE:
  TCID             CP Name  ISR  APPC  Pathswitch    Alive    COS TPF   T
=========================================================================
1E14570    USIBMRA.NNP61A    0    1        180       180    CPSVCMG
1E14E80    USIBMRA.NNP61A    0    1        180        90    CPSVCMG
1E2D9A8    USIBMRA.NNP41A    0    1        180       180    CPSVCMG
1E2E2B8    USIBMRA.NNP41A    0    1        180        90    CPSVCMG
1E2E740    USIBMRA.NNP61A    0    0          0        90     RSETUP
1E2EBC8       USIBMRA.RAA    0    2        180  3     90   SNASVCMG
*
2216T3 APPN >list appc sessions
LU Name             Mode  Type  FSM
======================================
USIBMRA.NNP61A    CPSVCMG  Pri  ACT
USIBMRA.NNP41A    CPSVCMG  Pri  ACT
USIBMRA.NNP61A    CPSVCMG  Sec  ACT
USIBMRA.NNP41A    CPSVCMG  Sec  ACT
USIBMRA.RAA       CPSVRMGR Pri  ACT  4
USIBMRA.RAA       CPSVRMGR Sec  ACT
```

Figure 181. Displays on 2216 after CS/2 Activation

This shows that the CS/2 connection has been dynamically defined **1**; we defined this only on the CS/2 node. HPR is not available because APPN is not available. Indeed, no CP-CP sessions are shown **2** between the 2216 and the CS/2.

The SNASVCMG RTP connection **3** has been established to RAA, and carries the two DLUR/S sessions **4**. As soon as the CS/2 node contacted the 2216, the 2216 knew (from the absence of the ACTPU not supported bit in the XID) that DLUR services were required. Therefore, the 2216 looked up its default DLUS name (USIBMRA.RAA), set up the RTP pipe, established the two DLUR/S sessions, and forwarded the XID information from CS/2 in a REQACTPU request unit. VTAM RAA did the rest.

Note also the presence of the Route Setup pipe to NNP61A; this indicates that the DLUR/S RTP pipe and its sessions are routed via NNP61A.

The console log on RAA (Figure 182 on page 188) shows what happened on the DLUS at this time.

```
IST1488I  ACTIVATION   FOR RTP CNR00801 AS PASSIVE PARTNER COMPLETED 5
IST1576I  DYNAMIC SWITCHED MAJOR NODE ISTDSWMN CREATED  6
```

Figure 182. DLUR/S Pipe Activation

The RTP pipe **5** to carry the DLUR/S sessions was established first. When the REQACTPU with the XID information hit RAA, it invoked the configuration services exit ISTEXCCS to define the type 2 PU and its LUs, since all DLUR/S flows use switched procedures. Because the dynamically defined PU was the first one on this VTAM node, the major node ISTDSWMN was created for it **6**. This feature (creation of ISTDSWMN on demand) is new in VTAM V4R4, and also allows ISTEXCCS to specify an alternative major node name for dynamically created PUs. Previous releases of VTAM created ISTDSWMN at startup time and placed all dynamic PUs in the one major node.

A display of ISTRTPMN, the RTP major node, from RAA confirmed that CNR00801 linked RAA with NN2216A. We then took a detailed display as in Figure 183.

```
DISPLAY NET,ID=CNR00801,SCOPE=ALL
IST097I  DISPLAY  ACCEPTED
IST075I  NAME = CNR00801         , TYPE = PU_T2.1
IST1392I  DISCNTIM = 00010 DEFINED AT PU   FOR DISCONNECT
IST486I  STATUS= ACTIV--LX-, DESIRED STATE= ACTIV
IST1043I  CP NAME = NN2216A , CP NETID = USIBMRA , DYNAMIC LU = YES
IST1589I  XNETALS = YES
IST933I  LOGMODE=***NA***, COS=SNASVCMG
IST1476I  TCID X'310BB5B7000000C3' - REMOTE TCID X'0000000001E2EBC8'
IST1481I  DESTINATION CP USIBMRA.NN2216A   - NCE X'8280           '
IST1587I  ORIGIN NCE X'D000000000000000'
IST1477I  ALLOWED DATA FLOW RATE =   1597 KBITS/SEC
IST1516I  INITIAL DATA FLOW RATE =   1597 KBITS/SEC
IST1511I  MAXIMUM NETWORK LAYER PACKET SIZE =   2048 BYTES
IST1478I  NUMBER OF UNACKNOWLEDGED BUFFERS =     0
IST1479I  RTP CONNECTION STATE = CONNECTED            - MNPS = NO
IST1480I  RTP END TO END ROUTE - PHYSICAL PATH
IST1460I  TGN  CPNAME             TG TYPE      HPR
IST1461I   21  USIBMRA.NNP61A     APPN         RTP
IST1461I   21  USIBMRA.NN2216A    APPN         RTP
IST231I  RTP      MAJOR NODE = ISTRTPMN
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I  STATE TRACE = OFF
IST355I  LOGICAL UNITS:
IST080I  NN2216A  ACT/S----Y
IST314I  END
```

Figure 183. New DLUR/S Pipe from RAA

The path for CNR00801 was as follows:

```
┌───────┐       ┌───────┐       ┌─────┐
│NN2216A├─TG21──┤NNP61A ├─TG21──┤RAA  │
└───────┘       └───────┘       └─────┘
```

Next, we displayed NN2216A from RAA. NN2216A is a remote NN in the APPN network, as well as (now) the served DLUR node that supports PU05170 (see Figure 184 on page 189).

```
DISPLAY NET,ID=NN2216A,SCOPE=ALL
IST097I  DISPLAY  ACCEPTED
IST075I  NAME = USIBMRA.NN2216A  , TYPE = ADJACENT CP
IST486I  STATUS= ACT/S----Y, DESIRED STATE= ACTIV
IST1402I  SRTIMER =    30  SRCOUNT =    10
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 = RAA     , VERIFY OWNER = NO
IST1184I  CPNAME = USIBMRA.NN2216A   - NETSRVR = ***NA***
IST1044I  ALSLIST = ISTAPNPU
IST1131I  DEVICE = ILU/CDRSC
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I  STATE TRACE = OFF
IST228I  ENCRYPTION = NONE
IST1563I  CKEYNAME = NN2216A  CKEY = PRIMARY   CERTIFY = NO
IST1552I  MAC = NONE        MACTYPE = NONE
IST171I  ACTIVE SESSIONS = 0000000002, SESSION REQUESTS = 0000000000
IST206I  SESSIONS:
IST1081I  ADJACENT LINK STATION = CNR00801
IST634I  NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I  RAA      ACTIV/DL-S D1C38C9F6F91897E 0016 0000  0  0 USIBMRA
IST635I  RAA      ACTIV/DL-P F7FF6164529F95C8 0000 0018  0  0 USIBMRA
IST1355I  PHYSICAL UNITS SUPPORTED BY DLUR USIBMRA.NN2216A
IST089I  W05170   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 =    30  SRCOUNT =    10
IST314I  END
```

Figure 184. NN and DLUR Node Display

The PU name W05170, of course, is that dynamically created by ISTEXCCS. The actual CP name (PU05170) of the CS/2 node is not known to VTAM. As far as RAA is concerned, the PU and the LUs are all associated with the CP NN2216A, and PU05170 does not exist. The only time RAA would become aware of PU05170 is if an APPN session was established between an RAA-owned LU and an LU on PU05170. PU05170 appears as a LEN node served by NN2216A, and is not known to the APPN network until it needs a session.

A display of W05170 itself from RAA showed (Figure 185 on page 190) that it was a switched PU supported by DLUR NN2216A **7**.

```
DISPLAY NET,ID=W05170,SCOPE=ALL
IST097I  DISPLAY  ACCEPTED
IST075I  NAME = W05170            , TYPE = PU_T2.1
IST486I  STATUS= ACTIV---X-, DESIRED STATE= ACTIV
IST1043I  CP NAME = PU05170 , CP NETID = USIBMRA , DYNAMIC LU = YES
IST1589I  XNETALS = YES
IST1354I  DLUR NAME = NN2216A            MAJNODE = ISTDSWMN   7
IST136I  SWITCHED SNA MAJOR NODE = ISTDSWMN
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I  STATE TRACE = OFF
IST355I  LOGICAL UNITS:
IST080I  W0517002 ACTIV---X- W0517003 ACTIV---X- W0517004 ACTIV---X-
IST080I  W0517005 ACTIV---X- W0517006 ACTIV---X- W0517007 ACTIV---X-
IST080I  W0517008 ACTIV---X- W0517009 ACTIV---X- W051700A ACTIV---X-
IST080I  W051700B ACTIV---X- W051700C ACTIV---X- W051700D ACTIV---X-
IST080I  W051700E ACTIV---X- W051700F ACTIV---X- W0517010 ACTIV---X-
IST080I  W0517011 ACTIV---X-
IST314I  END
```

Figure 185. DLUR-Owned PU

Note that no dependent LUs on W05170 are yet in session.

9.4.1 Session Establishment with 2216 As DLUR

Now we logged on from the CS/2 node to NetView on RAA (the DLUS) and TSO on RA39 (cross-domain). A display of the LUs on PU05170 (Figure 186) showed that two LUs were active on the link (HOST0001) to the 2216.

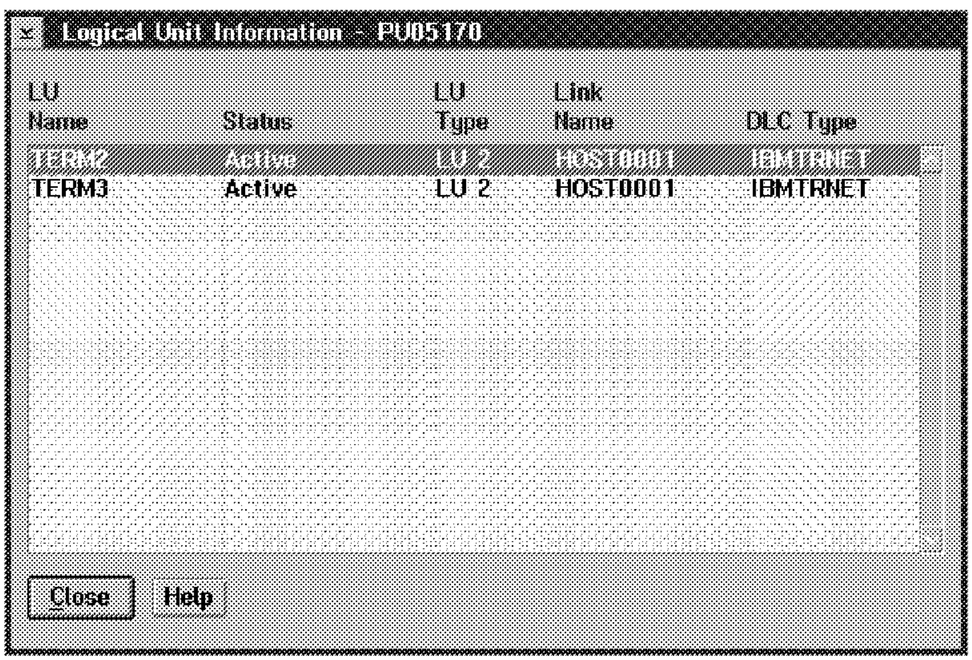

Figure 186. Logical Unit Information

Both RAA and RA39 immediately set up RTP connections to NN2216A for the new sessions. These were CNR00802 on RAA and CNR0005C ON RA39.

Displays on RAA and RA39 showed the paths for the sessions were:

```
W0517003 ─── NN2216A ─TG21─ NNP61A ─TG21─ RAA

W0517002 ─── NN2216A ─TG21─ NNP41A ─TG21─ RA39
```

We displayed the dependent LUs from both RAA and RA39. Figure 187 shows RAA′s view of the LU in session with its own NetView.

```
DISPLAY NET,ID=W0517003,SCOPE=ALL
IST097I  DISPLAY  ACCEPTED
IST075I  NAME = USIBMRA.W0517003 , TYPE = LOGICAL UNIT
IST486I  STATUS= ACT/S---X-, DESIRED STATE= ACTIV
IST1447I  REGISTRATION TYPE = NETSRVR
IST977I  MDLTAB=***NA*** ASLTAB=***NA***
IST861I  MODETAB=ISTINCLM USSTAB=US327X   LOGTAB=***NA***
IST934I  DLOGMOD=D4C32XX3 USS LANGTAB=***NA***
IST597I  CAPABILITY-PLU INHIBITED,SLU ENABLED  ,SESSION LIMIT 00000001
IST136I  SWITCHED SNA MAJOR NODE = ISTDSWMN
IST135I  PHYSICAL UNIT = W05170
IST1131I  DEVICE = LU
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I  STATE TRACE = OFF
IST228I  ENCRYPTION = NONE
IST1563I  CKEYNAME = LUMOD05D CKEY = PRIMARY   CERTIFY = NO
IST1552I  MAC = NONE        MACTYPE = NONE
IST171I  ACTIVE SESSIONS = 0000000001, SESSION REQUESTS = 0000000000
IST206I  SESSIONS:
IST634I  NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I  RAAAN008 ACTIV-P    F7FF6164529F95DE 0004 0005       USIBMRA
IST314I  END
```

Figure 187. Owned LU in Session with Application

This LU (W0517003) is shown as an owned LU in session with a local application. There is no indication here that it is a DLUR LU, or that this session actually traverses an HPR connection. For that information you must perform other displays, such as that of its RTP PU.

Figure 188 on page 192 shows the corresponding LU displayed from RA39.

```
DISPLAY NET,ID=W0517002,SCOPE=ALL
IST097I  DISPLAY  ACCEPTED
IST075I  NAME = USIBMRA.W0517002 , TYPE = CDRSC
IST486I  STATUS= ACT/S----Y, DESIRED STATE= ACTIV
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 = RA39    , VERIFY OWNER = NO
IST1184I  CPNAME = USIBMRA.NN2216A   - NETSRVR = ***NA***
IST082I  DEVTYPE = INDEPENDENT LU / CDRSC
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I  STATE TRACE = OFF
IST228I  ENCRYPTION = NONE
IST1563I  CKEYNAME = W0517002 CKEY = PRIMARY   CERTIFY = NO
IST1552I  MAC = NONE        MACTYPE = NONE
IST171I  ACTIVE SESSIONS = 0000000001, SESSION REQUESTS = 0000000000
IST206I  SESSIONS:
IST1081I  ADJACENT LINK STATION = CNR0005C
IST634I  NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I  RA39T04  ACTIV-P    F70794547C2C2281 0001 0003  0  0 USIBMRA
IST924I  -------------------------------------------------------------
IST075I  NAME = USIBMRA.W0517002 , TYPE = DIRECTORY ENTRY
IST1186I  DIRECTORY ENTRY = DYNAMIC    LU
IST1184I  CPNAME = USIBMRA.NN2216A   - NETSRVR = USIBMRA.RAA    8
IST314I  END
```

Figure 188. Cross Domain LU in Session with Application

This LU (W0517002) is seen by RA39 as outside RA39′s domain, yet in session with an LU owned by RA39. It is therefore represented by a CDRSC. The RTP connection (CNR0005C) is displayed as the link station through which this session is flowing. Note the APPN directory information **8**. W0517002 is owned by NN2216A (the DLUR) but served by RAA (the DLUS, not the NN server of the owning CP as you might expect). This is all in accordance with the DLUR/S architecture, because RAA as the DLUS initiates and responds to session requests on behalf of the DLUR LUs.

On the 2216, we displayed the connection information again as shown in Figure 189 on page 193.

```
2216T3 APPN >list rtp connections
RTP PARTNER TABLE:
 Remote Partner Name  Remote Boundary Name  TG Number
======================================================
         USIBMRA.RAA           USIBMRA.RAA     -1
        USIBMRA.RA39                           -1
RTP CONNECTION TABLE:
  TCID             CP Name  ISR  APPC  Pathswitch    Alive    COS TPF   T
=========================================================================
1E14570    USIBMRA.NNP61A    0    1         180       180    CPSVCMG
1E14E80    USIBMRA.NNP61A    0    1         180        90    CPSVCMG
1E2D9A8    USIBMRA.NNP41A    0    1         180       180    CPSVCMG
1E2E2B8    USIBMRA.NNP41A    0    1         180        90    CPSVCMG
1E2E740    USIBMRA.NNP61A    0    0           0        90     RSETUP
1E2EBC8       USIBMRA.RAA    0    2         180        90   SNASVCMG
1E56C80    USIBMRA.NNP41A    0    0           0       180     RSETUP
1E567F8      USIBMRA.RA39    1    0         240        90   #CONNECT  9
1E57108       USIBMRA.RAA    1    0         240        90   #CONNECT  9
*
2216T3 APPN >list isr sessions
  Adjacent CP Name  TG Number   ISR Sessions
===============================================
   USIBMRA.PU05170      21            2
*
2216T3 APPN >list sessions
     Origin CP Name          Primary LU        Secondary LU  Mode Name
=========================================================================
       USIBMRA.RA39
        USIBMRA.RAA
```

Figure 189. Displays on 2216 after Opening Two Sessions

You can see that two RTP pipes with APPN COS #CONNECT have been set up **9**. Each links this DLUR node with one of the session partners of the dependent LUs.

The session count is also interesting. On the 2216:

- Sessions that originate on the 2216 itself are called APPC.
- Sessions that pass through the 2216 are called ISR because this is an intermediate node on the session route and not an endpoint.

This is not the same terminology as used in APPN, where APPC means LU 6.2 and ISR means not ANR. The sessions listed on the various RTP pipes are:

- One APPC to NNP61A (CP-CP)
- One APPC to NNP61A (CP-CP)
- One APPC to NNP41A (CP-CP)
- One APPC to NNP41A (CP-CP)
- Two APPC to RAA (DLUR/S pair)
- One ISR to RA39 (W0517002 to TSO)
- One ISR to RAA (W0517003 to NetView)

9.4.2 Path Switch with 2216 As DLUR

Now we broke the link between the 2216 and the 3746 NNP61A (please refer to Figure 175 on page 182). The log on RAA showed Figure 190.

```
11:42:48   IST1494I  PATH SWITCH STARTED   FOR RTP CNR00801
11:42:48   IST1494I  PATH SWITCH COMPLETED FOR RTP CNR00801
11:42:48   IST1480I  RTP END TO END ROUTE - PHYSICAL PATH
11:42:48   IST1460I  TGN  CPNAME            TG TYPE      HPR
11:42:48   IST1461I   21  USIBMRA.RA39      APPN         RTP
11:42:48   IST1461I   21  USIBMRA.NNP41A    APPN         RTP
11:42:48   IST1461I   21  USIBMRA.NN2216A   APPN         RTP
11:42:48   IST314I  END
```

Figure 190. Path Switch of DLUR/S Pipe

There were no such messages for the LU-LU RTP connection, CNR00802. Presumably NN2216A initiated the path switch, in which case VTAM would issue no message. We displayed CNR00802 (Figure 191) to confirm that it had indeed been switched.

```
DISPLAY NET,ID=CNR00802,SCOPE=ALL
IST097I  DISPLAY  ACCEPTED
IST075I  NAME = CNR00802         , TYPE = PU_T2.1
IST1392I  DISCNTIM = 00010 DEFINED AT PU   FOR DISCONNECT
IST486I  STATUS= ACTIV--LX-, DESIRED STATE= ACTIV
IST1043I  CP NAME = NN2216A , CP NETID = USIBMRA , DYNAMIC LU = YES
IST1589I  XNETALS = YES
IST933I  LOGMODE=***NA***, COS=#CONNECT
IST1476I  TCID X'310BB5B8000000D0' - REMOTE TCID X'0000000001E57108'
IST1481I  DESTINATION CP USIBMRA.NN2216A   - NCE X'8280           '
IST1587I  ORIGIN NCE X'D000000000000000'
IST1477I  ALLOWED DATA FLOW RATE =   1597 KBITS/SEC
IST1516I  INITIAL DATA FLOW RATE =   1597 KBITS/SEC
IST1511I  MAXIMUM NETWORK LAYER PACKET SIZE =   2048 BYTES
IST1478I  NUMBER OF UNACKNOWLEDGED BUFFERS =     0
IST1479I  RTP CONNECTION STATE = CONNECTED            - MNPS = NO
IST1480I  RTP END TO END ROUTE - PHYSICAL PATH
IST1460I  TGN  CPNAME            TG TYPE      HPR
IST1461I   21  USIBMRA.RA39      APPN         RTP
IST1461I   21  USIBMRA.NNP41A    APPN         RTP
IST1461I   21  USIBMRA.NN2216A   APPN         RTP
IST231I  RTP      MAJOR NODE = ISTRTPMN
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I  STATE TRACE = OFF
IST355I  LOGICAL UNITS:
IST080I  W0517003 ACT/S---X-
IST314I  END
```

Figure 191. Path Switch for LU-LU Pipe

Both RTP connections are now using the same route:

```
┌─────────┐           ┌──────┐       ┌────┐       ┌───┐
│ NN2216A ├───TG21────┤NNP41A├─TG21──┤RA39├─TG21──┤RAA│
└─────────┘           └──────┘       └────┘       └───┘
```

Nothing was observed on RA39, and nothing changed as far as it was aware. Its sessions were not affected by the failure.

On the 2216, we verified the new connection state as in Figure 192 on page 195.

```
2216T3 APPN >list link
    Name   Port Name  Intf       Adj CP Name  Type       HPR      State
==========================================================================
    T61A       TRN1    1  1  USIBMRA.NNP61A    NN   ENABLED   RESET_LS
    T41A       TRN1    1     USIBMRA.NNP41A    NN    ACTIVE     ACT_LS
     @@0       TRN0    0    USIBMRA.PU05170    EN  INACTIVE     ACT_LS
*
2216T3 APPN >list appc
LU Name             Mode  Type  FSM
======================================
USIBMRA.NNP41A    CPSVCMG  Pri  ACT
USIBMRA.NNP41A    CPSVCMG  Sec  ACT
USIBMRA.RAA       CPSVRMGR Pri  ACT
USIBMRA.RAA       CPSVRMGR Sec  ACT
*
2216T3 APPN >list cp-cp        2
           CP Name      Type    Status  Connwinner ID    Conloser ID
=========================================================================
    USIBMRA.NNP61A       NN   Inactive    00000000       00000000
    USIBMRA.NNP41A       NN     Active    34E2D7FC       34E2D7FB
*
2216T3 APPN >list rtp connections
RTP PARTNER TABLE:
 Remote Partner Name  Remote Boundary Name  TG Number
======================================================
         USIBMRA.RAA           USIBMRA.RAA     -1
        USIBMRA.RA39                           -1
RTP CONNECTION TABLE:  3
  TCID             CP Name  ISR  APPC  Pathswitch    Alive    COS TPF   T
==========================================================================
1E2D9A8    USIBMRA.NNP41A    0    1         180       180    CPSVCMG
1E2E2B8    USIBMRA.NNP41A    0    1         180        90    CPSVCMG
1E2EBC8       USIBMRA.RAA    0    2         180        90   SNASVCMG
1E56C80    USIBMRA.NNP41A    0    0  4        0       180     RSETUP
1E567F8      USIBMRA.RA39    1    0         240        90   #CONNECT
1E57108       USIBMRA.RAA    2    0         180        90   #CONNECT
```

Figure 192. Displays on 2216 After Link Failure

This display shows that the link to NNP61A is in RESET status **1**, and the CP-CP sessions no longer exist **2**. The RTP connections **3** with their LU-LU sessions are all still there, but the only remaining Route Setup pipe is **4** to NNP41A, where the newly switched sessions now go.

Figure 193 on page 196 summarizes this section, where we set up and switched LU-LU sessions with the 2216 acting as the DLUR node on behalf of a node acting as a peripheral subarea node.

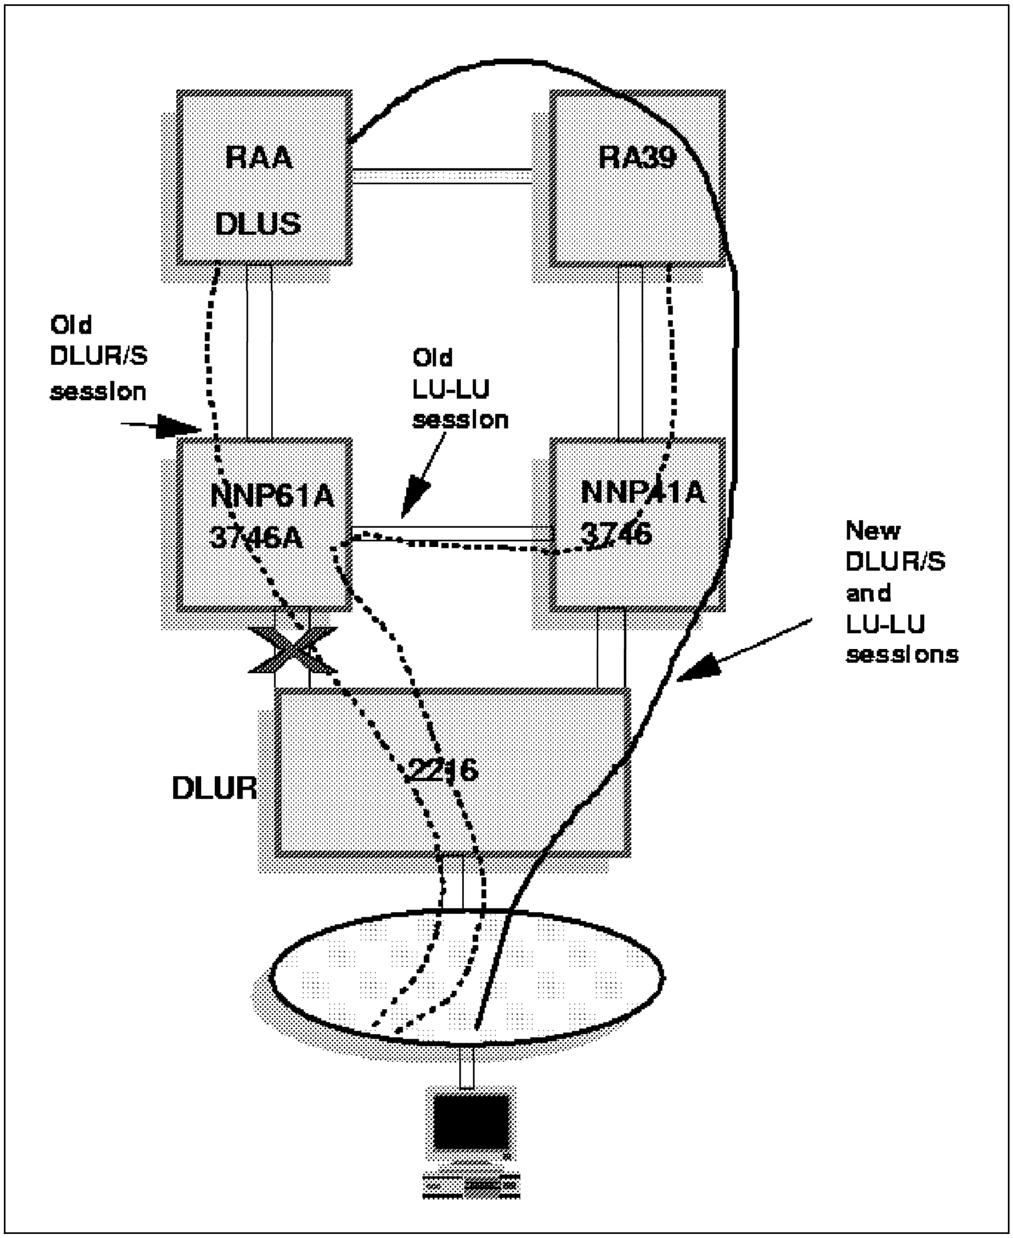

Figure 193. Summary of DLUR Test

Chapter 10. HPR and DLUR on the 2210

In this chapter we extend the HPR and DLUR implementation to the 2210 platform. Figure 194 shows the network we built for this configuration.

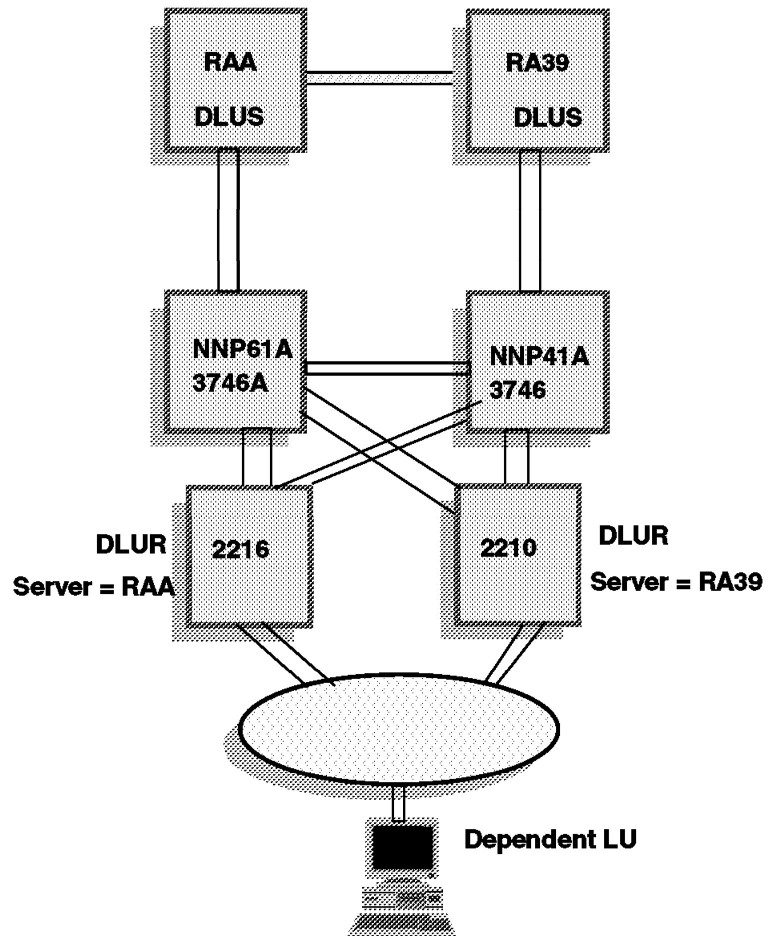

Figure 194. Network Configuration

We installed the 2210 router in parallel with the 2216, to simulate an alternative gateway to a remote site. We cross-connected the 2216 and the 2210 to each of the central 3746s, as would probably be the case if the wide area network was a shared facility such as frame relay. The VTAM and 3746 configurations were the same as in the previous chapter. Both 2216 and 2210 were configured as DLUR nodes, and a CS/2 node PU05170, configured for LEN attachment, served as the remote workstation as in the previous chapter. All connections in the diagram are in fact token-ring except the ESCON channels to the hosts. The CS/2 node is on its own separate LAN, to conform more closely with the typical customer environment. The backbone LAN, which connects the 3746s and the remote routers, represents the wide area network.

10.1 Configuration with 2210 As DLUR Node

The purpose of this section is to show the configuration of the 2210 and the steps taken to activate it. We enabled the node on the 2210 to support APPN, then we defined two ports, TRN0 to serve the downstream LAN (with the CS/2 node) and TRN5 to provide connectivity to the mainframe through its links to the 3746-900 network nodes. Three links were defined: two on the upstream port to connect to NNP41A and NNP61A, and one link to the 2216 through the downstream port. We finished off by enabling the DLUR requester function of the 2210 and defining RA39 as the DLU server.

10.1.1 HPR/DLUR Configuration for the 2210

The process of configuring the 2210 is remarkably similar to that of the 2216, since they share a common code base. The steps outlined below show how we configured our machine, CP2210T, for HPR and DLUR. As with the 2216, the base configuration (before you get to the APPN bits) allows you to define the MAC addresses of the physical ports. In our case they were 400022100099 (upstream) and 40002210A000 (downstream).

As with the 2216, the major steps in the 2210 configuration are:

1. APPN/HPR node configuration
2. APPN/HPR port configuration
3. APPN/HPR link configuration
4. APPN/HPR DLUR/DLUS configuration

Figure 195 shows that invoking the APPN configuration for the 2210 is exactly the same as for the 2216. You enter `talk 6` at the basic "*" prompt, then `protocol appn`. As with the 2216, the commands can be abbreviated (p for protocol, for example) but in general we used longer forms for clarity.

```
*talk 6
Config>protocol appn
APPN user configuration
APPN config>
```

Figure 195. Invoking 2210 APPN Configuration

First, we defined the APPN node-level parameters as shown in Figure 196 on page 199.

```
APPN config>set node
Enable APPN (Y)es (N)o (Y)? y
Network ID (Max 8 characters) (  )? USIBMRA
Control point name (Max 8 characters) (  )? CP2210T
Enable branch extender (Y)es (N)o (N)?
Route addition resistance(0-255) (128)?
XID ID number for subarea connection (5 hex digits) (00000)?
Use enhanced #BATCH COS (Y)es (N)o (Y)?
Use enhanced #BATCHSC COS (Y)es (N)o (Y)?
Use enhanced #INTER COS (Y)es (N)o (Y)?
Use enhanced #INTERSC COS (Y)es (N)o (Y)?
Write this record? (Y)? y
The record has been written.
```

Figure 196. 2210 Node Definition

Next, we defined the two ports trn0 and trn5, the downstream and upstream ports respectively. Figure 197 shows the definitions for trn0. The definitions for trn5 are identical except for (of course) the port name and interface number.

```
APPN config>add port
APPN Port
Link Type: (P)PP, (FR)AME RELAY, (E)THERNET, (T)OKEN RING,
(S)DLC, (X)25, (D)LSw, (A)TM ( )? t
Interface number(Default 0): (0)?
Port name (Max 8 characters) (TR000)? trn0
Enable APPN on this port (Y)es (N)o (Y)? y
Port Definition
Service any node: (Y)es (N)o (Y)? y
High performance routing: (Y)es (N)o (Y)?
Maximum BTU size (768-17745) (2048)?
Maximum number of link stations (1-976) (512)?
Percent of link stations reserved for incoming calls (0-100) (0)?
Percent of link stations reserved for outgoing calls (0-100) (0)?
Local SAP address (04-EC) (4)?
Local HPR SAP address (04-EC) (C8)?
Edit TG Characteristics: (Y)es (N)o (N)?
Edit LLC Characteristics: (Y)es (N)o (N)? y
Remote SAP(04-EC) (4)?
Maximum number of outstanding I-format LPDUs (1-127) (26)?
Receive window size (1-127) (26)?
Inactivity timer(1-254 seconds) (30)?
Reply timer (1-254 half seconds) (2)?
Maximum number of retransmissions(1-254) (8)? 254
Receive acknowledgement timer (1-254 half seconds) (1)?
Acknowledgements needed to increment working window(0-127) (1)?
Edit HPR defaults: (Y)es (N)o (N)?
Write this record? (Y)? y
The record has been written.
```

Figure 197. Downstream Port of 2210

Next, we configured the link stations. We defined both upstream stations to the 3746s, as we wanted the remote routers to initiate these connections. We also defined the connection to the 2216 on the downstream port, as we had not defined this previously on the 2216. Figure 198 on page 200 shows the connection to NNP41A on trn5.

```
APPN config>add link
APPN Station
Port name for the link station ( )? trn5
Station name (Max 8 characters) ( )? st41a
Activate link automatically (Y)es (N)o (Y)?
MAC address of adjacent node (000000000000)? 400437462176
Adjacent node type: 0 = APPN network node,
1 = APPN end node or Unknown node type,
2 = LEN end node (0)?
High performance routing: (Y)es (N)o (Y)?
Allow CP-CP sessions on this link (Y)es (N)o (Y)?
CP-CP session level security (Y)es (N)o (N)?
Configure CP name of adjacent node: (Y)es (N)o (N)?
Edit TG Characteristics: (Y)es (N)o (N)?
Edit LLC Characteristics: (Y)es (N)o (N)? Y
Remote SAP(04-EC) (4)? 8   1
Maximum number of outstanding I-format LPDUs (1-127) (26)?
Receive window size (1-127) (26)?
Inactivity timer(1-254 seconds) (30)?
Reply timer (1-254 half seconds) (2)?
Maximum number of retransmissions(1-254) (8)? 254
Receive acknowledgement timer (1-254 half seconds) (1)?
Acknowledgements needed to increment working window(0-127) (1)?
Edit HPR defaults: (Y)es (N)o (N)?
Write this record? (Y)? y
The record has been written.
```

Figure 198. Definition of Link to NNP41A

Note the remote SAP of 8 **1**, necessary because the 3746s used their own defaults instead of the 2210's default of 4.

The definition of the connection to NNP61A was the same except for the station name (st61a) and the MAC address of the remote node.

Figure 199 on page 201 shows the definition of the trn0 link station to the 2216 NN2216A.

```
APPN config>add link
APPN Station
Port name for the link station ( )? trn0
Station name (Max 8 characters) ( )? st2216
Activate link automatically (Y)es (N)o (Y)?
MAC address of adjacent node (000000000000)? 400022160100
Adjacent node type: 0 = APPN network node,
1 = APPN end node or Unknown node type,
2 = LEN end node (0)?
High performance routing: (Y)es (N)o (Y)?
Allow CP-CP sessions on this link (Y)es (N)o (Y)?
CP-CP session level security (Y)es (N)o (N)?
Configure CP name of adjacent node: (Y)es (N)o (N)?
Edit TG Characteristics: (Y)es (N)o (N)?
Edit LLC Characteristics: (Y)es (N)o (N)? y
Remote SAP(04-EC) (4)?
Maximum number of outstanding I-format LPDUs (1-127) (26)?
Receive window size (1-127) (26)?
Inactivity timer(1-254 seconds) (30)?
Reply timer (1-254 half seconds) (2)?
Maximum number of retransmissions(1-254) (254)?
Receive acknowledgement timer (1-254 half seconds) (1)?
Acknowledgements needed to increment working window(0-127) (1)?
Edit HPR defaults: (Y)es (N)o (N)?
Write this record? (Y)? y
The record has been written.
```

Figure 199. Link Definition to 2216 via Downstream Link

Having defined the connections, we next configured the DLUR/S parameters as shown in Figure 200. As with the 2216, these can be overridden for each downstream node that requests dependent LU session support. We specified only the primary DLUS (RA39), with no backup DLUS **2**.

```
APPN config>set dlur
Enable DLUR (Y)es (N)o (N)? y
Fully-qualified CP name of primary DLUS ()? usibmra.ra39
Fully-qualified CP name of backup DLUS ()?  2
Perform retries to restore disrupted pipe (N)? y
Delay before initiating retries(0-2756000 seconds) (120)?
Perform short retries to restore disrupted pipe (Y)?
Short retry timer(0-2756000 seconds) (120)?
Short retry count(0-65535) (5)?
Perform long retry to restore disrupted pipe (Y)?
Long retry timer(0-2756000 seconds) (300)?
Write this record? (Y)? y
The record has been written.
```

Figure 200. 2210 DLUR Definition

When we had finished we displayed the 2210 APPN configuration as in Figure 201 on page 202.

```
APPN config>list all
NODE:
NETWORK ID: USIBMRA
CONTROL POINT NAME: CP2210T
XID: 00000
APPN ENABLED: YES
MAX SHARED MEMORY: 5108
MAX CACHED: 4000
DLUR:
DLUR ENABLED: YES
PRIMARY DLUS NAME: USIBMRA.RA39
CONNECTION NETWORK:
      CN NAME       LINK TYPE  PORT INTERFACES
-------------------------------------------------------------
COS:
COS NAME
--------
  #BATCH
#BATCHSC
#CONNECT
  #INTER
#INTERSC
 CPSVCMG
SNASVCMG
 MODE NAME  COS NAME
---------------------
PORT:
  INTF     PORT      LINK      HPR    SERVICE   PORT
 NUMBER    NAME      TYPE    ENABLED    ANY    ENABLED
------------------------------------------------------
    0       TRN0   IBMTRNET    YES      YES      YES
    5       TRN5   IBMTRNET    YES      YES      YES
STATION:
 STATION    PORT       DESTINATION      HPR    ALLOW  ADJ NODE
  NAME      NAME         ADDRESS      ENABLED  CP-CP    TYPE
------------------------------------------------------------
  ST2216     TRN0      400022160100     YES     YES       0
   ST61A     TRN5      400037462144     YES     YES       0
   ST41A     TRN5      400437462176     YES     YES       0
LU NAME:
      LU NAME        STATION NAME           CP NAME
------------------------------------------------------------
APPN config>exit
```

Figure 201. Listing of APPN/HPR Defined Options

After the 2216 displays of earlier tests, this was quite familiar to us.

10.1.2 Configuration for CS/2 Acting As a LEN Node

The CS/2 node, PU05170, was configured exactly as in the 2216 scenario, except that this time we used the 2210 as the DLUR. We therefore defined the destination MAC address of the CS/2 node′s only link to be the appropriate port of the 2210. The fact that the DLUS (and therefore the owner) of the dependent LUs was now RA39 and not RAA was not known to CS/2. In contrast to APPN and LEN protocols, a subarea peripheral node has no need to be aware of the identity of the node managing the boundary function that supports it. Figure 202 on page 203 shows the relevant extract from the node definition file.

```
DEFINE_LOCAL_CP  FQ_CP_NAME(USIBMRA.PU05170  )
                  CP_ALIAS(THINKPAD)
                  NAU_ADDRESS(INDEPENDENT_LU)
                  NODE_TYPE(EN)
                  NODE_ID(X'05D05170')   3
                  NW_FP_SUPPORT(NONE)
                  HOST_FP_SUPPORT(YES)
                  FREE_UNUSED_SESSIONS(NO)
                  HOST_FP_LINK_NAME(HOST0001)
                  MAX_COMP_LEVEL(NONE)
                  MAX_COMP_TOKENS(0);

 DEFINE_LOGICAL_LINK  LINK_NAME(HOST0001)
                      ADJACENT_NODE_TYPE(LEN)  6
                      DLC_NAME(IBMTRNET)
                      ADAPTER_NUMBER(0)
                      DESTINATION_ADDRESS(X'40002210A00004')   4
                      ETHERNET_FORMAT(NO)
                      CP_CP_SESSION_SUPPORT(NO)   7
                      SOLICIT_SSCP_SESSION(YES)   5
                      NODE_ID(X'05D05170')   3
                      ACTIVATE_AT_STARTUP(YES)
                      USE_PUNAME_AS_CPNAME(NO)
                      LIMITED_RESOURCE(USE_ADAPTER_DEFINITION)
                      LINK_STATION_ROLE(USE_ADAPTER_DEFINITION)
                      MAX_ACTIVATION_ATTEMPTS(USE_ADAPTER_DEFINITION)
                      EFFECTIVE_CAPACITY(USE_ADAPTER_DEFINITION)
                      COST_PER_CONNECT_TIME(USE_ADAPTER_DEFINITION)
                      COST_PER_BYTE(USE_ADAPTER_DEFINITION)
                      SECURITY(USE_ADAPTER_DEFINITION)
                      PROPAGATION_DELAY(USE_ADAPTER_DEFINITION)
                      USER_DEFINED_1(USE_ADAPTER_DEFINITION)
                      USER_DEFINED_2(USE_ADAPTER_DEFINITION)
                      USER_DEFINED_3(USE_ADAPTER_DEFINITION);
```

Figure 202. NDF Listing of CS/2 Used for 3270 Sessions

Note the node ID **3**, specified at the node level and confirmed (unnecessarily) at the link level. This will be used by RA39 to identify the PU and thus to define the PU and LUs dynamically. Note also the destination MAC and SAP address of the 2210 **4** and the request for SSCP sessions **5**.

Because the adjacent node type has been specified as LEN **6** (by not requesting APPN support on the host link using the GUI configuration), CS/2 will pretend to be a LEN node at XID exchange time. The XIDs will state that parallel TGs are not supported between CS/2 and the 2210; this is all that is needed to ensure that a connection is treated as LEN, and imposes the requirement that CP-CP sessions are not supported **7**.

10.2 Example with 2210 As DLUR Node

In Figure 203 on page 204, we displayed the status of the 2210's connections before starting the downstream CS/2.

```
APPN >list port
Intf          Name       DLC Type        HPR          State
=============================================================
  5           TRN5       IBMTRNET       TRUE       ACT_PORT
  0           TRNO       IBMTRNET       TRUE       ACT_PORT
*
APPN >list link
    Name   Port Name  Intf       Adj CP Name  Type       HPR      State
=========================================================================
   ST61A        TRN5    5    USIBMRA.NNP61A    NN    ACTIVE     ACT_LS
   ST41A        TRN5    5    USIBMRA.NNP41A    NN    ACTIVE     ACT_LS
  ST2216        TRNO    0   USIBMRA.NN2216A    NN    ACTIVE     ACT_LS
```

Figure 203. Display of Active Links on 2210 before CS/2 Connection

All ports and link stations are active, both to the host gateway 3746s and to the adjoining 2216 node.

Next we displayed the RTP connections and active sessions, as in Figure 204.

```
APPN >list cp-cp sessions   8
           CP Name      Type    Status  Connwinner ID    Conloser ID
========================================================================
    USIBMRA.NNP41A       NN     Active    B88DB512        B88DB50C
   USIBMRA.NN2216A       NN     Active    B88DB4F7        B88DB4F5
    USIBMRA.NNP61A       NN     Active    B88DB514        B88DB511
*
APPN >list rtp connections   9
RTP CONNECTION TABLE:
  TCID             CP Name  ISR  APPC  Pathswitch    Alive    COS TPF   T
=========================================================================
319D2B20   USIBMRA.NN2216A    0    2         180       180    CPSVCMG
319D4650    USIBMRA.NNP41A    0    1         180       180    CPSVCMG
319DFB18    USIBMRA.NNP41A    0    1         180       180    CPSVCMG
319E0D38    USIBMRA.NNP41A    0    0           0       180     RSETUP
319DFFA0    USIBMRA.NNP61A    0    1         180       180    CPSVCMG
31A0BED8    USIBMRA.NNP61A    0    1         180       180    CPSVCMG
*
APPN >list appc    10
LU Name              Mode  Type  FSM
======================================
USIBMRA.NN2216A   CPSVCMG  Pri  ACT
USIBMRA.NN2216A   CPSVCMG  Sec  ACT
USIBMRA.NNP41A    CPSVCMG  Pri  ACT
USIBMRA.NNP61A    CPSVCMG  Pri  ACT
USIBMRA.NNP41A    CPSVCMG  Sec  ACT
USIBMRA.NNP61A    CPSVCMG  Sec  ACT
```

Figure 204. Displays Issued on 2210 before CS/2 Connection

There were pairs of CP-CP sessions with each partner network node (**8** and **10**), as we hoped. The RTP connections table **9** showed that all were using RTP pipes (all four nodes support control flows); those to the 2216 had both CP-CP sessions on a single RTP pipe but those to the 3746s had separate pipes. This is purely a timing consideration. The Route Setup pipe was from a previous LU-LU session; there were no DLUR/S sessions active at this time as **10** proved.

We also displayed the APPN connectivity from the 3746 NNP41A. Figure 205 on page 205 shows the link stations known to NNP41A.

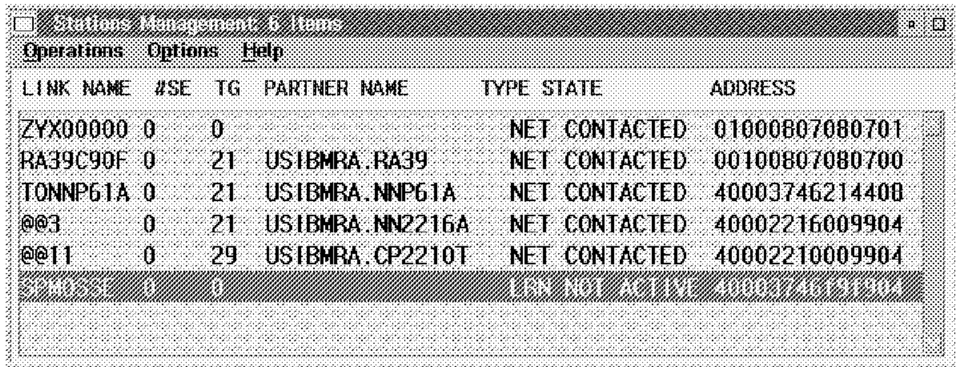

Figure 205. Active Stations on NNP41A

The links to the 2210 and the 2216 are defined dynamically on the 3746 (that is, on the partner side only), so they have 3746-defined names such as @@3 and @@11. The links to NNP61A and RA39 are explicitly defined. All are active.

Figure 206 shows the active RTP connections at the start of the test. Apart from two Route Setup pipes from previous LU-LU sessions, the only active connections are the CP-CP pipes to the four adjacent network nodes.

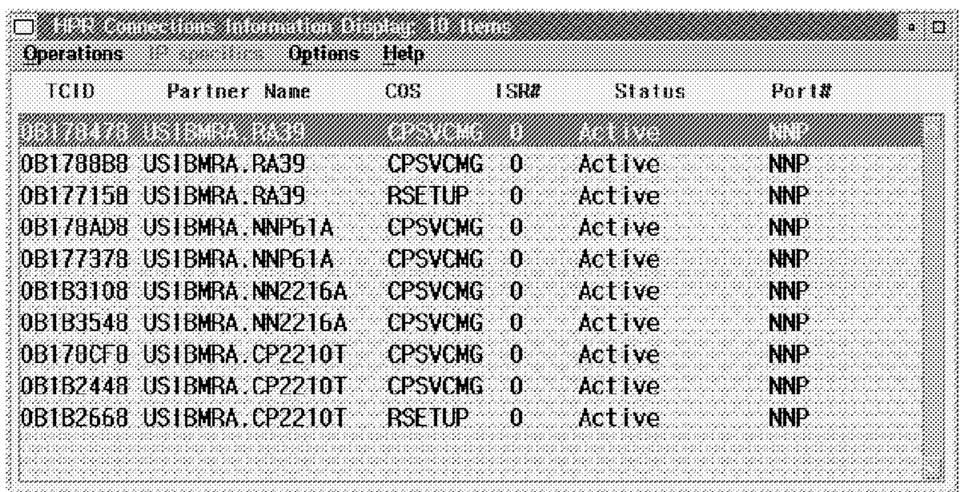

Figure 206. Active HPR Connections on NNP41A

10.2.1 CS/2 As Peripheral Node with 2210 As DLUR

Now we started the CS/2 node with its dependent LUs. We received the VTAM USS10 message from RA39 on the CS/2 emulator screens, and observed that a new RTP connection, CNR00060, was created on RA39. Figure 207 on page 206 shows a display of that RTP connection.

```
DISPLAY NET,ID=CNR00060,SCOPE=ALL
IST097I  DISPLAY  ACCEPTED
IST075I  NAME = CNR00060         , TYPE = PU_T2.1
IST1392I  DISCNTIM = 00010 DEFINED AT PU   FOR DISCONNECT
IST486I  STATUS= ACTIV--LX-, DESIRED STATE= ACTIV
IST1043I  CP NAME = CP2210T , CP NETID = USIBMRA , DYNAMIC LU = YES
IST1589I  XNETALS = YES
IST933I  LOGMODE=***NA***, COS=SNASVCMG
IST1476I  TCID X'38273F0C000000D1' - REMOTE TCID X'0000000031A0D580'
IST1481I  DESTINATION CP USIBMRA.CP2210T   - NCE X'8280            '
IST1587I  ORIGIN NCE X'D000000000000000'
IST1477I  ALLOWED DATA FLOW RATE =   1597 KBITS/SEC
IST1516I  INITIAL DATA FLOW RATE =   1597 KBITS/SEC
IST1511I  MAXIMUM NETWORK LAYER PACKET SIZE =   2048 BYTES
IST1478I  NUMBER OF UNACKNOWLEDGED BUFFERS =     0
IST1479I  RTP CONNECTION STATE = CONNECTED            - MNPS = NO
IST1480I  RTP END TO END ROUTE - PHYSICAL PATH
IST1460I  TGN  CPNAME             TG TYPE      HPR
IST1461I   21  USIBMRA.NNP41A     APPN         RTP
IST1461I   29  USIBMRA.CP2210T    APPN         RTP
IST231I  RTP      MAJOR NODE = ISTRTPMN
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I  STATE TRACE = OFF
IST355I  LOGICAL UNITS:
IST080I  CP2210T  ACT/S----Y
IST314I  END
```

Figure 207. DLUR/S Pipe from RA39

This was the DLUR/S pipe, and used the following path:

```
CP2210T ──TG29──── NNP41A ──TG21──── RA39
```

Further proof that this is the DLUR/S pipe is given by Figure 208 on page 207.

```
DISPLAY NET,ID=CP2210T,SCOPE=ALL
IST097I  DISPLAY  ACCEPTED
IST075I  NAME = USIBMRA.CP2210T  , TYPE = ADJACENT CP
IST486I  STATUS= ACT/S----Y, DESIRED STATE= ACTIV
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 = RA39    , VERIFY OWNER = NO
IST1184I  CPNAME = USIBMRA.CP2210T   - NETSRVR = ***NA***
IST1044I  ALSLIST = ISTAPNPU
IST082I  DEVTYPE = INDEPENDENT LU / CDRSC
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I  STATE TRACE = OFF
IST228I  ENCRYPTION = NONE
IST1563I  CKEYNAME = CP2210T  CKEY = PRIMARY   CERTIFY = NO
IST1552I  MAC = NONE        MACTYPE = NONE
IST171I  ACTIVE SESSIONS = 0000000002, SESSION REQUESTS = 0000000000
IST206I  SESSIONS:
IST1081I  ADJACENT LINK STATION = CNR00060   11
IST634I  NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I  RA39     ACTIV/DL-S F55F295029B5653D 000A 0000  0  0 USIBMRA
IST635I  RA39     ACTIV/DL-P F70794547C2C22A2 0000 000C  0  0 USIBMRA
IST1355I  PHYSICAL UNITS SUPPORTED BY DLUR USIBMRA.CP2210T   12
IST089I  W05170   TYPE = PU_T2.1          , ACTIV---X-
IST924I  -------------------------------------------------------------
IST075I  NAME = USIBMRA.CP2210T  , TYPE = DIRECTORY ENTRY
IST1186I  DIRECTORY ENTRY = DYNAMIC    NN
IST1184I  CPNAME = USIBMRA.CP2210T   - NETSRVR = ***NA***
IST314I  END
```

Figure 208. DLUR/S Pipe

The node CP2210T has two sessions with status ACTIV/DL on link station CNR00060 **11**, and is the DLUR for W05170 **12**, which is the type 2 PU in the CS/2 node.

Next, we issued the displays in Figure 209 on page 208 on the 2210 to see the newly created RTP connection.

```
APPN >list link
    Name   Port Name  Intf       Adj CP Name  Type       HPR      State
==========================================================================
   ST61A        TRN5    5     USIBMRA.NNP61A    NN    ACTIVE     ACT_LS
   ST41A        TRN5    5     USIBMRA.NNP41A    NN    ACTIVE     ACT_LS
  ST2216        TRN0    0    USIBMRA.NN2216A    NN    ACTIVE     ACT_LS
     @@1   13   TRN0    0    USIBMRA.PU05170    EN  INACTIVE     ACT_LS
*
APPN >list appc
LU Name             Mode  Type  FSM
======================================
USIBMRA.NN2216A   CPSVCMG  Pri  ACT
USIBMRA.NN2216A   CPSVCMG  Sec  ACT
USIBMRA.NNP41A    CPSVCMG  Pri  ACT
USIBMRA.NNP41A    CPSVCMG  Sec  ACT
USIBMRA.NNP61A    CPSVCMG  Pri  ACT
USIBMRA.NNP61A    CPSVCMG  Sec  ACT
USIBMRA.RA39      CPSVRMGR Pri  ACT  14
USIBMRA.RA39      CPSVRMGR Sec  ACT  14
*
APPN >list isr sessions
  Adjacent CP Name  TG Number   ISR Sessions
===============================================
   USIBMRA.PU05170      21             0
*
APPN >list rtp connections
RTP PARTNER TABLE:
 Remote Partner Name  Remote Boundary Name  TG Number
======================================================
        USIBMRA.RA39          USIBMRA.RA39     -1
RTP CONNECTION TABLE:
  TCID             CP Name  ISR  APPC  Pathswitch    Alive    COS TPF   T
==========================================================================
319D2B20   USIBMRA.NN2216A    0    2         180       180    CPSVCMG
319D4650    USIBMRA.NNP41A    0    1         180       180    CPSVCMG
319DFB18    USIBMRA.NNP41A    0    1         180       180    CPSVCMG
319E0D38    USIBMRA.NNP41A    0    0           0       180     RSETUP
319DFFA0    USIBMRA.NNP61A    0    1         180       180    CPSVCMG
31A0BED8    USIBMRA.NNP61A    0    1         180       180    CPSVCMG
31A0D580      USIBMRA.RA39    0    2         180  15   180   SNASVCMG
```

Figure 209. Displays Issued on 2210 after CS/2 Connection

Once again, as soon as the peripheral node sent an XID to the 2210 requesting SSCP sessions, an RTP connection was established to the DLUS defined in the 2210 node. This connection **15**, using APPN COS SNASVCMG, was to RA39. The two sessions it carried are shown at **14**. The actual link station to PU05170 **13** is dynamically defined (its name is generated by the 2210), and does not support HPR because it appears as a LEN connection.

We then logged on from a dependent LU on PU05170 to NetView on RAA. At once, RAA established a new RTP connection CNR0085D. The display of CNR0085D showed that it carried a session to W0517003, and that the session used the following path:

```
┌────────┐          ┌───────┐      ┌───────┐      ┌──────┐
│W0517003├──────────┤CP2210T├─TG34─┤NNP61A ├─TG21─┤RAA   │
└────────┘          └───────┘      └───────┘      └──────┘
```

A display of W0517003 from RAA (Figure 210 on page 209) shows that RAA sees it as an APPN LU owned by CP2210T and served by RA39.

```
DISPLAY NET,ID=W0517003,SCOPE=ALL
IST097I  DISPLAY  ACCEPTED
IST075I  NAME = USIBMRA.W0517003 , TYPE = CDRSC
IST486I  STATUS= ACT/S----Y, DESIRED STATE= ACTIV
IST1402I  SRTIMER =    30  SRCOUNT =    10
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 = RAA     , VERIFY OWNER = NO
IST1184I  CPNAME = USIBMRA.CP2210T   - NETSRVR = ***NA***
IST1131I  DEVICE = ILU/CDRSC
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I  STATE TRACE = OFF
IST228I  ENCRYPTION = NONE
IST1563I  CKEYNAME = W0517003 CKEY = PRIMARY   CERTIFY = NO
IST1552I  MAC = NONE        MACTYPE = NONE
IST171I  ACTIVE SESSIONS = 0000000001, SESSION REQUESTS = 0000000000
IST206I  SESSIONS:
IST1081I  ADJACENT LINK STATION = CNR0085D
ST634I  NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I  RAAAN008 ACTIV-P    F7FF6164529F9640 0001 0002  0  0 USIBMRA
IST924I  -------------------------------------------------------------
IST075I  NAME = USIBMRA.W0517003 , TYPE = DIRECTORY ENTRY
IST1186I  DIRECTORY ENTRY = DYNAMIC    LU
IST1184I  CPNAME = USIBMRA.CP2210T   - NETSRVR = USIBMRA.RA39
IST1402I  SRTIMER =    30  SRCOUNT =    10
IST314I  END
```

Figure 210. Cross-Domain LU Display

At this stage we looked at the 2210 displays again to see the connectivity status. The RTP connections display can be seen in Figure 211.

```
APPN >list rtp connections
RTP PARTNER TABLE:
 Remote Partner Name  Remote Boundary Name  TG Number
======================================================
        USIBMRA.RA39         USIBMRA.RA39     -1
RTP CONNECTION TABLE:
  TCID            CP Name  ISR  APPC  Pathswitch    Alive    COS TPF   T
=========================================================================
319D2B20   USIBMRA.NN2216A    0    2         180       180    CPSVCMG
319D4650    USIBMRA.NNP41A    0    1         180       180    CPSVCMG
319DFB18    USIBMRA.NNP41A    0    1         180       180    CPSVCMG
319E0D38    USIBMRA.NNP41A    0    0           0       180     RSETUP
319DFFA0    USIBMRA.NNP61A    0    1         180       180    CPSVCMG
31A0BED8    USIBMRA.NNP61A    0    1         180       180    CPSVCMG
31A0D580      USIBMRA.RA39    0    2         180       180   SNASVCMG
31A0D0F8       USIBMRA.RAA    1    0         240   16  180   #CONNECT
```

Figure 211. Active RTP Connections on 2210 after LU-LU Session Establishment

The new session is represented by the RTP pipe with APPN COS #CONNECT **16** .

The observant reader will notice the absence of a long-lived Route Setup pipe to NNP61A. If the new session was set up over this connection, where is the Route

Setup pipe? In fact, this demonstrates an aspect of HPR not easy to test in our lab. The connection between the 2210 and NNP41A was unreliable and kept breaking; you may have noticed large TG numbers in the session path displays. Every time the connection broke, the 2210 re-established it with a new TG number. All this time the LU-LU session remained active because the connection was restored before the path switch timer expired. When such a link failure happens, the Route Setup pipe is deactivated but the other RTP pipes remain active.

10.2.2 Path Switch after 2210 Is Disconnected

We broke *both* the 2210′s connections to the backbone LAN, in other words both its links to the 3746s. Now its only route to the hosts was through the partner router, the 2216. The LU-LU session remained operative so we displayed its RTP pipe from RAA as shown in Figure 212.

```
17:41:59   DISPLAY NET,ID=CNR0085D,SCOPE=ALL
IST097I  DISPLAY  ACCEPTED
IST075I  NAME = CNR0085D         , TYPE = PU_T2.1
IST1392I  DISCNTIM = 00010 DEFINED AT PU   FOR DISCONNECT
IST486I  STATUS= ACTIV--LX-, DESIRED STATE= ACTIV
IST1043I  CP NAME = CP2210T , CP NETID = USIBMRA , DYNAMIC LU = YES
IST1589I  XNETALS = YES
IST933I  LOGMODE=***NA***, COS=#CONNECT
IST1476I  TCID X'310BB6130000009E' - REMOTE TCID X'0000000031A0D0F8'
IST1481I  DESTINATION CP USIBMRA.CP2210T   - NCE X'8280           '
IST1587I  ORIGIN NCE X'D000000000000000'
IST1477I  ALLOWED DATA FLOW RATE =    399 KBITS/SEC
IST1516I  INITIAL DATA FLOW RATE =    399 KBITS/SEC
IST1511I  MAXIMUM NETWORK LAYER PACKET SIZE =   2048 BYTES
IST1478I  NUMBER OF UNACKNOWLEDGED BUFFERS =     0
IST1479I  RTP CONNECTION STATE = CONNECTED            - MNPS = NO
IST1480I  RTP END TO END ROUTE - PHYSICAL PATH
IST1460I  TGN  CPNAME             TG TYPE      HPR
IST1461I   21  USIBMRA.NNP61A     APPN         RTP
IST1461I   21  USIBMRA.NN2216A    APPN         RTP
IST1461I   29  USIBMRA.CP2210T    APPN         RTP
IST231I  RTP      MAJOR NODE = ISTRTPMN
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I  STATE TRACE = OFF
IST355I  LOGICAL UNITS:
IST080I  W0517003 ACT/S----Y
IST314I  END
```

Figure 212. LU-LU Pipe after Switch

You can see that the new path now goes through both remote routers:

```
W0517003 ──── CP2210T ──TG29── NN2216A ──TG21── NNP61A ──TG21── RAA
```

Figure 213 on page 211 shows the old and new session paths after the link failures.

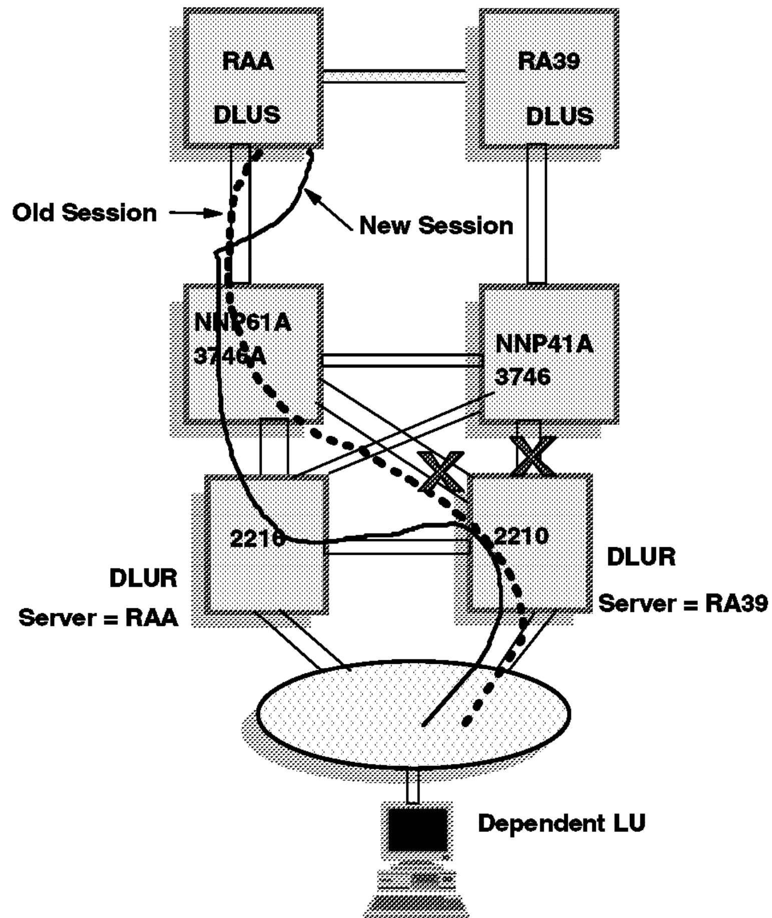

Figure 213. New Path for LU-LU RTP Pipe

From RA39, we displayed the RTP connection that the DLUR/S sessions were using, as in Figure 214 on page 212.

```
DISPLAY NET,ID=CNR00060,SCOPE=ALL
IST097I  DISPLAY  ACCEPTED
IST075I  NAME = CNR00060         , TYPE = PU_T2.1
IST1392I  DISCNTIM = 00010 DEFINED AT PU   FOR DISCONNECT
IST486I  STATUS= ACTIV--LX-, DESIRED STATE= ACTIV
IST1043I  CP NAME = CP2210T , CP NETID = USIBMRA , DYNAMIC LU = YES
IST1589I  XNETALS = YES
IST933I  LOGMODE=***NA***, COS=SNASVCMG
IST1476I  TCID X'38273F0C000000D1' - REMOTE TCID X'0000000031A0D580'
IST1481I  DESTINATION CP USIBMRA.CP2210T   - NCE X'8280           '
IST1587I  ORIGIN NCE X'D000000000000000'
IST1477I  ALLOWED DATA FLOW RATE =    399 KBITS/SEC
IST1516I  INITIAL DATA FLOW RATE =    399 KBITS/SEC
IST1511I  MAXIMUM NETWORK LAYER PACKET SIZE =   2048 BYTES
IST1478I  NUMBER OF UNACKNOWLEDGED BUFFERS =     0
IST1479I  RTP CONNECTION STATE = CONNECTED            - MNPS = NO
IST1480I  RTP END TO END ROUTE - PHYSICAL PATH
IST1460I  TGN  CPNAME             TG TYPE      HPR
IST1461I   21  USIBMRA.NNP41A     APPN         RTP
IST1461I   21  USIBMRA.NN2216A    APPN         RTP
IST1461I   29  USIBMRA.CP2210T    APPN         RTP
IST231I  RTP      MAJOR NODE = ISTRTPMN
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I  STATE TRACE = OFF
IST355I  LOGICAL UNITS:
IST080I  CP2210T  ACT/S----Y
IST314I  END
```

Figure 214. DLUR/S Pipe New Path

The new path for the DLUR/S sessions is:

```
CP2210T ──TG29── NN2216A ──TG21── NNP41A ──TG21── RA39
```

Finally, we performed the connectivity displays on the 2210 as shown in Figure 215 on page 213.

```
APPN >list rtp connections
RTP PARTNER TABLE:
 Remote Partner Name  Remote Boundary Name  TG Number
======================================================
        USIBMRA.RA39          USIBMRA.RA39     -1
         USIBMRA.RAA                           -1
RTP CONNECTION TABLE:
  TCID             CP Name  ISR  APPC  Pathswitch    Alive    COS TPF   T
=========================================================================
319D2B20   USIBMRA.NN2216A    0    2         180 17   180    CPSVCMG
31A0D580      USIBMRA.RA39    0    2         180 18   180   SNASVCMG
31A0D0F8       USIBMRA.RAA    2    0         240 19   180   #CONNECT
*
APPN >list cp cp sessions
           CP Name      Type    Status  Connwinner ID    Conloser ID
========================================================================
    USIBMRA.NNP41A       NN   Inactive    00000000         00000000
   USIBMRA.NN2216A       NN     Active    B88DB4F7   17    B88DB4F5
    USIBMRA.NNP61A       NN   Inactive    00000000         00000000
*
APPN >list link
    Name   Port Name  Intf       Adj CP Name  Type       HPR      State
=========================================================================
   ST61A        TRN5    5     USIBMRA.NNP61A    NN   ENABLED   RESET_LS
   ST41A        TRN5    5     USIBMRA.NNP41A    NN   ENABLED   RESET_LS
  ST2216        TRN0    0    USIBMRA.NN2216A    NN    ACTIVE     ACT_LS
     @@1        TRN0    0    USIBMRA.PU05170    EN  INACTIVE     ACT_LS
*
APPN >list port
Intf          Name       DLC Type        HPR          State
=============================================================
  5           TRN5       IBMTRNET       TRUE       RESET_PORT
  0           TRN0       IBMTRNET       TRUE       ACT_PORT
```

Figure 215. Displays Issued on 2210 Side

Most of the CP-CP sessions had disappeared. The only ones left were to the 2216 **17** over an RTP connection. The two RTP pipes that concerned us were active: the DLUR/S pipe to RA39 **18** and the LU-LU pipe to RAA **19**.

Chapter 11. HPR and Branch Extender

In our final network configuration to exercise HPR and DLUR, we configured branch extender capability on both 2216 and 2210 routers. Branch extender is a recent extension to the APPN architecture that allows large APPN networks to be divided into more manageable pieces by isolating remote branch locations from the backbone topology. Please see Chapter 4, “Branch Extender” on page 39 for a description of the branch extender function.

11.1 2216 and 2210 Branch Extender Configuration

Branch extender capability, like DLUR, is not enabled by default on the APPN routers. The process for configuring branch extender support is similar to that for DLUR. First you enable it at the node level, then you define it on each port and link. Because of the rules about branch extender configuration (see Chapter 4, “Branch Extender” on page 39), some combinations are not possible and therefore some configuration questions are not asked under certain circumstances.

We began the process of branch extender configuration on the 2216 by invoking `talk 6`, as before. The configuration process is identical on both routers so we have shown only the 2216 definitions. Figure 216 shows the node definitions we entered. We actually updated the configuration from the previous tests, hence the occasional warnings about changing existing records. For this test we defined trn1 as Port 0 (upstream) and trn2 as Port 1 (downstream).

```
2216 APPN config>set node
Enable APPN (Y)es (N)o (Y)?

Control point name (Max 8 characters) (NN2216A)?
Enable branch extender (Y)es (N)o (N)? y  1
Permit search for unregistered LUs (Y)es (N)o (N)?   2
Route addition resistance(0-255) (128)?
XID ID number for subarea connection (5 hex digits) (00000)?
Use enhanced #BATCH COS (Y)es (N)o (Y)?
Use enhanced #BATCHSC COS (Y)es (N)o (N)?
Use enhanced #INTER COS (Y)es (N)o (N)?
Use enhanced #INTERSC COS (Y)es (N)o (Y)?
Write this record? (Y)? y
The record has been written.
```

Figure 216. Node Definition for BX

The first new answer here is that to the BX question **1**. This immediately results in the 2216 asking whether to allow a search for unregistered LUs from its backbone NN server **2**. Please refer to 3.3, “DLUR/S Design Considerations” on page 32 for an explanation of why this might be necessary.

Next, we defined the port to be used for the uplink (upstream connection to the backbone network). Figure 217 on page 216 shows the port definitions we used for the trn1 port.

```
2216 APPN config>add port
APPN Port
Link Type: (P)PP, (FR)AME RELAY, (E)THERNET, (T)OKEN RING,
(M)PC, (S)DLC, (X)25, (FD)DI, (D)LSw, (A)TM, (I)P  ( ) t
Interface number(Default 0): (0)?
Port name (Max 8 characters) (TR000)? trn1
WARNING!! You are changing an existing record.
Enable APPN on this port (Y)es (N)o (Y)?
Port Definition
        Service any node: (Y)es (N)o (Y)?
        High performance routing: (Y)es (N)o (Y)?
        Maximum BTU size (768-17745) (2048)?
        Maximum number of link stations (1-976) (512)?
        Percent of link stations reserved for incoming calls (0-100) (0)?
        Percent of link stations reserved for outgoing calls (0-100) (0)?
        Local SAP address (04-EC) (4)?
        Local HPR SAP address (04-EC) (C8)?
        Branch uplink: (Y)es (N)o (N)? y  1
Edit TG Characteristics: (Y)es (N)o (N)?
Edit LLC Characteristics: (Y)es (N)o (N)?
Edit HPR defaults: (Y)es (N)o (N)?
Write this record? (Y)? y
The record has been written.
```

Figure 217. BX Port Definition

We responded in the affirmative to the question **1** about using this port as a BX uplink. This sets the default value to be used by link stations on this port. Each link station can be defined individually as an upstream or downstream connection, provided the BX rules are obeyed.

```
2216T3 APPN config>add link
APPN Station
Port name for the link station (  ) trn1
Station name (Max 8 characters) (  ) t61a
WARNING!! You are changing an existing record.
        Activate link automatically (Y)es (N)o (Y)?
        MAC address of adjacent node (400037462144)?
        Adjacent node type: 0 = APPN network node,
        1 = APPN end node or Unknown node type,
        2 = LEN end node, 3 = PU 2.0 node (1)? 0    2
        High performance routing: (Y)es (N)o (Y)?
        Edit Dependent LU Server: (Y)es (N)o (N)?
        Allow CP-CP sessions on this link (Y)es (N)o (Y)?
        CP-CP session level security (Y)es (N)o (N)?
        Configure CP name of adjacent node: (Y)es (N)o (N)?
        Link to preferred network node server: (Y)es (N)o (N)?
Edit TG Characteristics: (Y)es (N)o (N)?
Edit LLC Characteristics: (Y)es (N)o (N)?
Edit HPR defaults: (Y)es (N)o (N)?
Write this record? (Y)? y
The record has been written.
```

Figure 218. BX Link Definitions

In this dialog, we specified **2** that the adjacent node was a network node. This immediately dictates that the link station is an upstream connection, so the 2216 did not ask whether it was upstream or downstream. The other link station, to NNP41A, was defined in a similar fashion.

Having redefined the 3746 links as branch extender uplinks, we updated the active configuration file and restarted the 2216 as shown in Figure 219 on page 217.

```
2216 APPN config>exit
2216 Config>write
Config Save: Using bank B and config number 3
2216 Config>
2216 Config>CNTL P
2216 *t 5
2216T3 APPN >restart
Are you sure you want to restart APPN (Y)es (N)o (N)? y
```

Figure 219. Save Configuration and Restart 2216

There is another way to activate the new configuration: in the `APPN config` prompt (after `talk 6` and `protocol appn`), type `activate`. This command compares the active configuration with the new one, and either restarts the 2216 (as above) or dynamically updates the parameters, depending on what has changed.

Aside from the BX support, our 2216 and 2210 configurations were the same as in the previous chapter. In particular, the 2216 was acting as a DLUR for downstream nodes, with RAA as its DLUS. The 2210 was configured as a DLUR node with RA39 as its DLUS.

11.1.1 HPR Configuration for CS/2

For our downstream (branch workstation) node, we configured a CS/2 node BREN1 as an end node with two HPR-capable links; one to the 2216 and one to the 2210. The CS/2 node could not be a network node in the BX environment, and could establish CP-CP sessions with only one of the two gateway routers. One of the restrictions with BX is that you cannot have parallel concurrent gateways for the same downstream node, but we hoped to demonstrate that this is not a problem in an HPR environment, at least for independent LUs.

The BX rules do not permit the presence of a downstream DLUR node, so we defined our CS/2 with dependent LUs and SSCP support requested over the upstream (to the routers) connections. Each connection to a 221X has two dependent LUs defined on it. Figure 220 on page 218 shows part of the node definition file we used.

```
DEFINE_LOGICAL_LINK  LINK_NAME(HOST0001)   3
                     ADJACENT_NODE_TYPE(LEARN)
                     DLC_NAME(IBMTRNET)
                     ADAPTER_NUMBER(0)
                     DESTINATION_ADDRESS(X'400022160100004')   4
                     ETHERNET_FORMAT(NO)
                     CP_CP_SESSION_SUPPORT(YES)   5
                     PU_NAME(BREN1   )
                     SOLICIT_SSCP_SESSION(YES)   6
                     NODE_ID(X'05D04444')   7
                     USE_PUNAME_AS_CPNAME(NO)
                     MAX_ACTIVATION_ATTEMPTS(USE_ADAPTER_DEFINITION)
                     AUTO_REACTIVATE(NO_RETRY)
                     ACTIVATE_AT_STARTUP(NO)
                     LIMITED_RESOURCE(USE_ADAPTER_DEFINITION)
                     LINK_STATION_ROLE(USE_ADAPTER_DEFINITION)
                     EFFECTIVE_CAPACITY(USE_ADAPTER_DEFINITION)
                     COST_PER_CONNECT_TIME(USE_ADAPTER_DEFINITION)
                     COST_PER_BYTE(USE_ADAPTER_DEFINITION)
                     SECURITY(USE_ADAPTER_DEFINITION)
                     PROPAGATION_DELAY(USE_ADAPTER_DEFINITION)
                     HPR_SUPPORT(USE_ADAPTER_DEFINITION)
                     HPR_LLERP_SUPPORT(USE_ADAPTER_DEFINITION)
                     HPR_MLTG_NUMBER(0)
                     BRANCH_EXTENDER_UPLINK(USE_ADAPTER_DEFINITION)
                     MAX_I_FIELD_SIZE(USE_ADAPTER_DEFINITION)
                     LIMITED_RESOURCE_TIMEOUT(USE_ADAPTER_DEFINITION)
                     INACTIVITY_TIMEOUT(USE_ADAPTER_DEFINITION)
                     USER_DEFINED_1(USE_ADAPTER_DEFINITION)
                     USER_DEFINED_2(USE_ADAPTER_DEFINITION)
                     USER_DEFINED_3(USE_ADAPTER_DEFINITION);
```

Figure 220 (Part 1 of 3). NDF File for CS/2 End Node

```
DEFINE_LOGICAL_LINK  LINK_NAME(HOST0002)  8
                     ADJACENT_NODE_TYPE(LEARN)
                     DLC_NAME(IBMTRNET)
                     ADAPTER_NUMBER(0)
                     DESTINATION_ADDRESS(X'40002210A00004')  9
                     ETHERNET_FORMAT(NO)
                     CP_CP_SESSION_SUPPORT(YES)  10
                     PU_NAME(MPU00001)
                     SOLICIT_SSCP_SESSION(YES)  11
                     USE_PUNAME_AS_CPNAME(NO)
                     MAX_ACTIVATION_ATTEMPTS(USE_ADAPTER_DEFINITION)
                     AUTO_REACTIVATE(NO_RETRY)
                     ACTIVATE_AT_STARTUP(YES)
                     LIMITED_RESOURCE(USE_ADAPTER_DEFINITION)
                     LINK_STATION_ROLE(USE_ADAPTER_DEFINITION)
                     EFFECTIVE_CAPACITY(USE_ADAPTER_DEFINITION)
                     COST_PER_CONNECT_TIME(USE_ADAPTER_DEFINITION)
                     COST_PER_BYTE(USE_ADAPTER_DEFINITION)
                     SECURITY(USE_ADAPTER_DEFINITION)
                     PROPAGATION_DELAY(USE_ADAPTER_DEFINITION)
                     HPR_SUPPORT(USE_ADAPTER_DEFINITION)
                     HPR_LLERP_SUPPORT(USE_ADAPTER_DEFINITION)
                     HPR_MLTG_NUMBER(0)
                     BRANCH_EXTENDER_UPLINK(USE_ADAPTER_DEFINITION)
                     MAX_I_FIELD_SIZE(USE_ADAPTER_DEFINITION)
                     LIMITED_RESOURCE_TIMEOUT(USE_ADAPTER_DEFINITION)
                     INACTIVITY_TIMEOUT(USE_ADAPTER_DEFINITION)
                     USER_DEFINED_1(USE_ADAPTER_DEFINITION)
                     USER_DEFINED_2(USE_ADAPTER_DEFINITION)
                     USER_DEFINED_3(USE_ADAPTER_DEFINITION);
```

Figure 220 (Part 2 of 3). NDF File for CS/2 End Node

```
DEFINE_LUA  LU_NAME(@LUA0001)
            HOST_LINK_NAME(HOST0001)
            NAU_ADDRESS(2);

DEFINE_LUA  LU_NAME(@LUA0002)
            HOST_LINK_NAME(HOST0001)
            NAU_ADDRESS(3);

DEFINE_LUA  LU_NAME(@LUA0003)
            HOST_LINK_NAME(HOST0002)
            NAU_ADDRESS(4);

DEFINE_LUA  LU_NAME(@LUA0004)
            HOST_LINK_NAME(HOST0002)
            NAU_ADDRESS(5);
```

Figure 220 (Part 3 of 3). NDF File for CS/2 End Node

Link station HOST0001 **3** is connected to the 2216's downstream port **4**, and supports both CP-CP sessions **5** and SSCP sessions **6**. We have an APPN node with dependent LUs, so we require both. The PU that this link represents is given the IDNUM of 04444 **7**, because we need to distinguish it from the other link, which will inherit the node's IDNUM. If the IDNUMs on the two links are the same, the IBM-supplied ISTEXCCS exit on each DLU server will create the same LU names from the same local addresses.

Link station HOST0002 **8** is connected to the 2210′s downstream port **9**, and also supports both CP-CP sessions **10** and SSCP sessions **11**. The CP-CP sessions, of course, will only be established over one link at a time.

11.2 Example of Branch Extender with HPR

Figure 221 shows the network configuration that we set up to check out the branch extender function. The CS/2 node BREN1 has dependent LUs on both its upstream connections to the two routers. It is HPR-capable, and can establish CP-CP sessions with whichever BX it chooses. Thus if it loses its NN server it can transfer its allegiance to the other BX, and can maintain its independent LU sessions using HPR path switching (provided, of course, that the path switch timers have been correctly set). The dependent LU sessions can survive the loss of any link or node upstream from the BX, but they cannot survive the loss of the BX they are using as a DLUR node because that BX must be the RTP endpoint.

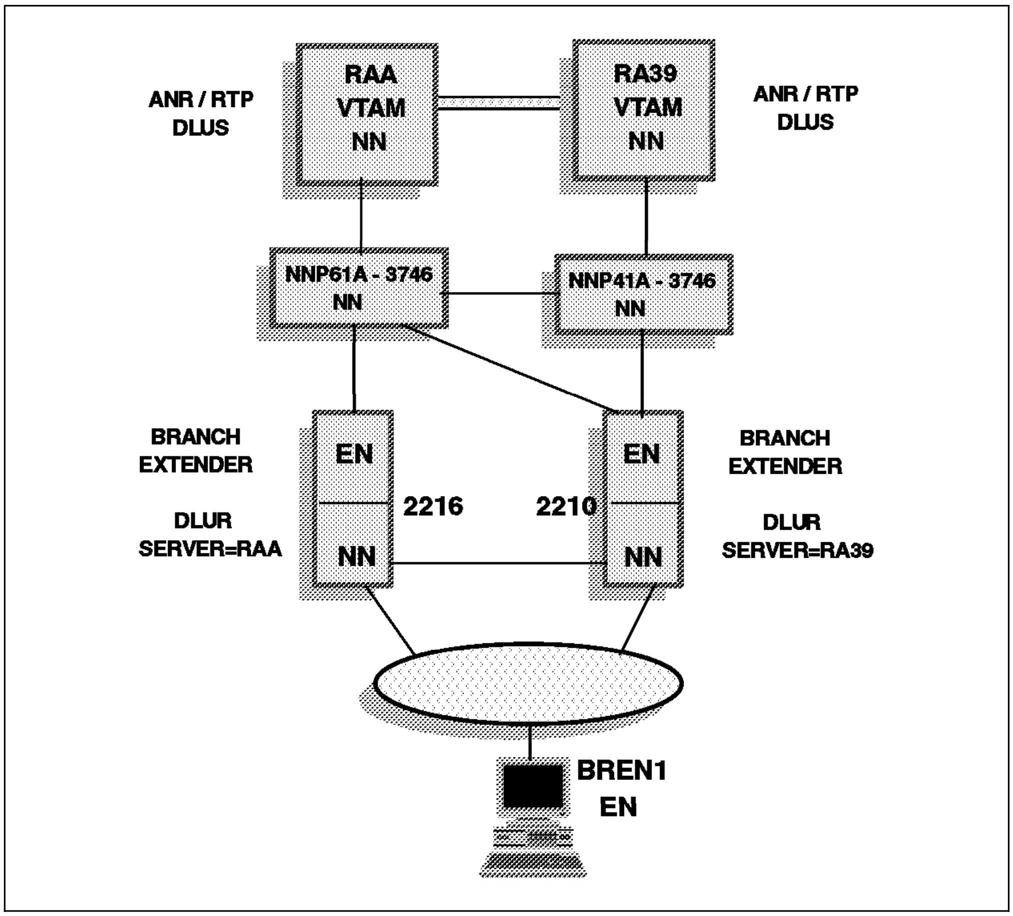

Figure 221. BX Configuration

After we started the CS/2 node, we displayed the active connections using the Subsystem Management facility as in Figure 222 on page 221.

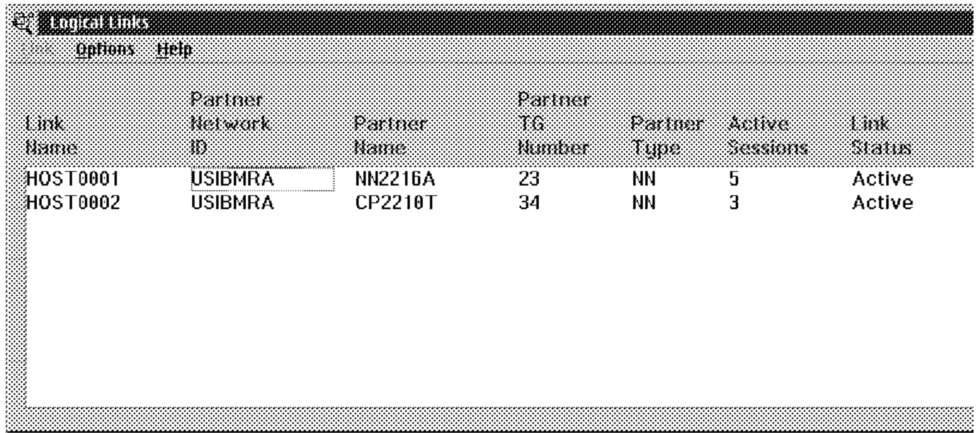

Figure 222. Logical Links to BX Nodes

We defined two dependent LUs on each link, to use the DLUR support in the 2210 and 2216. We did not use these in the test, but the display shows their SSCP-PU and SSCP-LU sessions (total three sessions on each link). The two extra sessions on HOST0001, of course, are the CP-CP sessions to the 2216, which the CS/2 node has chosen as its network node server.

The display of LU 6.2 sessions (Figure 223) shows that the extra sessions on HOST0001 are indeed the CP-CP sessions.

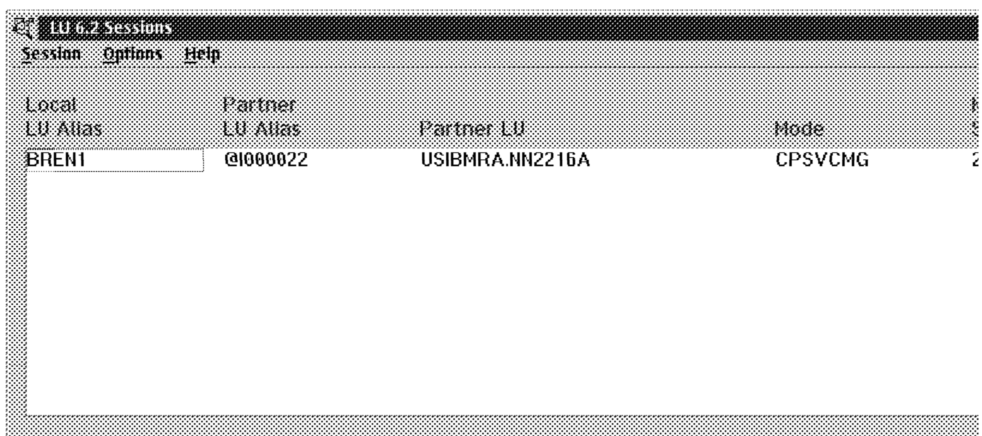

Figure 223. CP-CP Sessions to 2216 BX

Figure 224 on page 222 then showed us the network view as seen from the 2216.

```
2216T3 APPN >list link
    Name   Port Name  Intf       Adj CP Name  Type       HPR      State
==========================================================================
    T61A        TRN1    0     USIBMRA.NNP61A    NN    ACTIVE     ACT_LS
    T41A        TRN1    0     USIBMRA.NNP41A    NN   ENABLED   RESET_LS
     @@2        TRN2    1      USIBMRA.BREN1    EN    ACTIVE     ACT_LS
*
2216T3 APPN >list cp cp
           CP Name      Type    Status  Connwinner ID    Conloser ID
=========================================================================
    USIBMRA.NNP61A       NN     Active    34E957A1        34E957A3  1
    USIBMRA.NNP41A       NN   Inactive    00000000        00000000
     USIBMRA.BREN1       EN     Active    34E9586C        34E95867  2
*
2216T3 APPN >list rtp
RTP PARTNER TABLE:
 Remote Partner Name  Remote Boundary Name  TG Number
======================================================
         USIBMRA.RAA           USIBMRA.RAA     -1
RTP CONNECTION TABLE:
  TCID             CP Name  ISR  APPC  Pathswitch    Alive    COS TPF   T
==========================================================================
1E14570    USIBMRA.NNP61A    0    0          0        90   CPSVCMG
1E14E80    USIBMRA.NNP61A    0    2        180        90   CPSVCMG
1E149F8    USIBMRA.NNP61A    0    0          0       180    RSETUP
1E2EBC8       USIBMRA.RAA    0    2        180        90  SNASVCMG  3
*
2216T3 APPN >list appc
LU Name            Mode  Type  FSM
======================================
USIBMRA.NNP61A    CPSVCMG  Pri  ACT
USIBMRA.NNP61A    CPSVCMG  Sec  ACT
USIBMRA.RAA       CPSVRMGR Pri  ACT
USIBMRA.BREN1     CPSVCMG  Sec  ACT
USIBMRA.BREN1     CPSVCMG  Pri  ACT
USIBMRA.RAA       CPSVRMGR Sec  ACT
*
2216T3 APPN >list isr
  Adjacent CP Name  TG Number   ISR Sessions
===============================================
     USIBMRA.BREN1      23             0  4
```

Figure 224. Displays on 2216 BX

This display shows:

- The 2216 has two active links, both with HPR enabled. The connection to the 3746 NNP41A was down at this time because we did not use it in this test.
- The 2216 has CP-CP sessions with NNP61A **1** (as its network node server) and with BREN1 **2** (as a served end node). It appears as an EN upstream and as an NN downstream. If the NNP41A link was active, it would be usable for sessions and for RTP connections (including Route Setup pipes), but there would be no CP-CP sessions because NN2216A is only allowed one pair of these into the upstream network. NN2216A appears as an end node to the backbone network.

  We did not define in BREN1 what the preferred NN server should be, so BREN1 chose the first one it found.
- NN2216A has set up the DLUR/S sessions, on their RTP pipe **3**, to RAA. The Route Setup pipe to NNP61A was used for these.

- Although both 2216 and CS/2 support Control Flows over RTP, the CP-CP sessions between them did not use an RTP connection. At the time we did not trace the XID exchange to find out why this was so. Probably it was the CS/2 node that denied supporting Control Flows; we used more than one release of CS/2 in our tests whereas the 2216 MAS was always at the same level.
- There are, as yet, no ISR sessions **4**.

We performed similar displays from the 2210 console, as shown in Figure 225.

```
APPN >list link
    Name   Port Name  Intf       Adj CP Name  Type       HPR      State
==========================================================================
   ST61A        TRN5    5     USIBMRA.NNP61A    NN    ACTIVE     ACT_LS 5
   ST41A        TRN5    5     USIBMRA.NNP41A    NN    ACTIVE     ACT_LS 5
    @@13        TRN0    0      USIBMRA.BREN1    EN    ACTIVE     ACT_LS 5
*
APPN >list cp cp
           CP Name      Type    Status  Connwinner ID    Conloser ID
=========================================================================
     USIBMRA.BREN1       EN   Inactive    00000000        00000000
    USIBMRA.NNP41A       NN 6  Active    B8930D60        B8930D62
    USIBMRA.NNP61A       NN   Inactive    00000000        00000000
*
APPN >list rtp
RTP PARTNER TABLE:
 Remote Partner Name  Remote Boundary Name  TG Number
======================================================
        USIBMRA.RA39          USIBMRA.RA39     -1
RTP CONNECTION TABLE:
  TCID             CP Name  ISR  APPC  Pathswitch    Alive    COS TPF   T
==========================================================================
31ADDA28    USIBMRA.NNP41A    0    0           0        0    CPSVCMG
31ADF558    USIBMRA.NNP41A    0    2         180      180    CPSVCMG
31AEE100    USIBMRA.NNP41A    0    0           0      180     RSETUP
31AFBD50    USIBMRA.NNP61A    0    0           0      180     RSETUP
31AFC660      USIBMRA.RA39    0    2         180      180   SNASVCMG   7
*
APPN >list appc
LU Name            Mode  Type  FSM
======================================
USIBMRA.NNP41A    CPSVCMG  Pri  ACT
USIBMRA.NNP41A    CPSVCMG  Sec  ACT
USIBMRA.RA39      CPSVRMGR Pri  ACT
USIBMRA.RA39      CPSVRMGR Sec  ACT
*
APPN >list isr
  Adjacent CP Name  TG Number   ISR Sessions
===============================================
     USIBMRA.BREN1      34             0  8
```

Figure 225. Network Display from 2210 BX

On the 2210:

- There are three active connections **5**; one to each 3746 and one to BREN1.
- The only active CP-CP sessions **6** are with NNP41A. BREN1 chose NN2216A to be its server, and CP2210T has chosen NNP41A.
- The 2210 has its own DLUR/S connection, over an RTP pipe **7**, to RA39. This connection serves two dependent LUs on BREN1.

- There are no ISR sessions **8** on the link to BREN1 yet.

We also displayed the connections as seen from the 3746s. Figure 226 shows those on NNP41A.

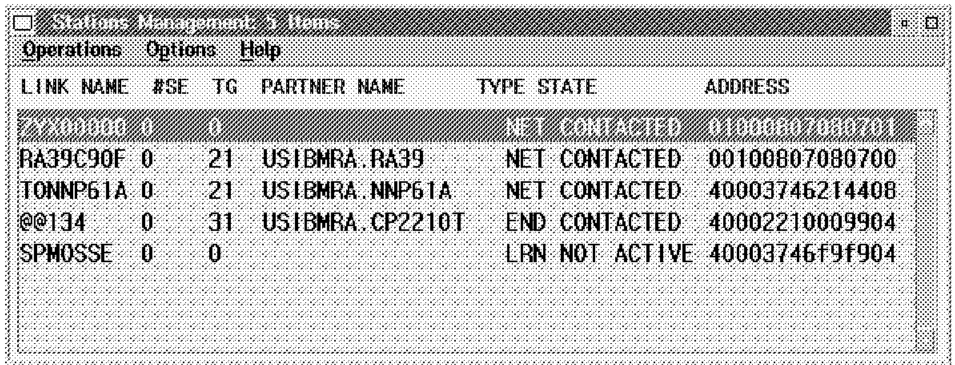

Figure 226. Active Connections on NNP41A

Note that NNP41A sees the 2210 (CP2210T) as an end node. Similarly, NNP61A is connected to both routers (Figure 227) and sees both as end nodes.

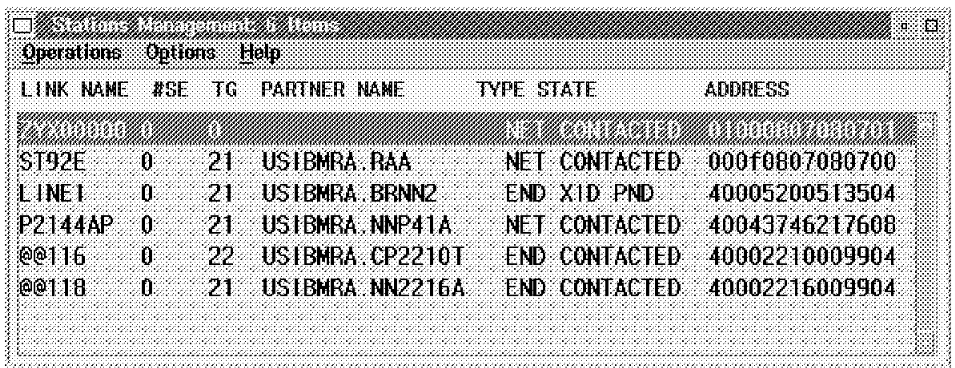

Figure 227. Active Connections on NNP61A

Both connections from CP2210T to the 3746s are known to the APPN network (the local topology databases in this case), even though only one of them at a time can carry CP-CP sessions.

On RAA, the DLUR/S sessions were carried on RTP pipe CNR008C7, which took the route:

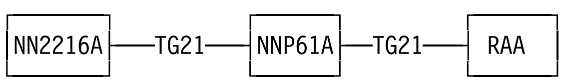

On RA39, the coresponding DLUR/S pipe took the following path:

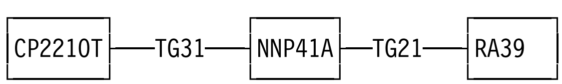

11.2.1 Session from Independent LU across BX

Since all RTP connections relating to DLUR/S (the LU 6.2 session pair and the dependent LU sessions themselves) must terminate at the BX nodes, we can learn nothing about HPR through a branch extender from a study of dependent LU sessions. Such sessions cannot cross a BX boundary as APPN/HPR sessions. Therefore, we established independent LU sessions, using APING, between our CS/2 end node and VTAM RA39. A simple invocation of APING results in two sessions, and two RTP connections, with the partner. The CNOS session runs on the RTP pipe with APPN COS SNASVCMG, and the other session runs with APPN COS #INTER.

We displayed the details of the #INTER session (Figure 228) to see that it was using HPR as its logical link.

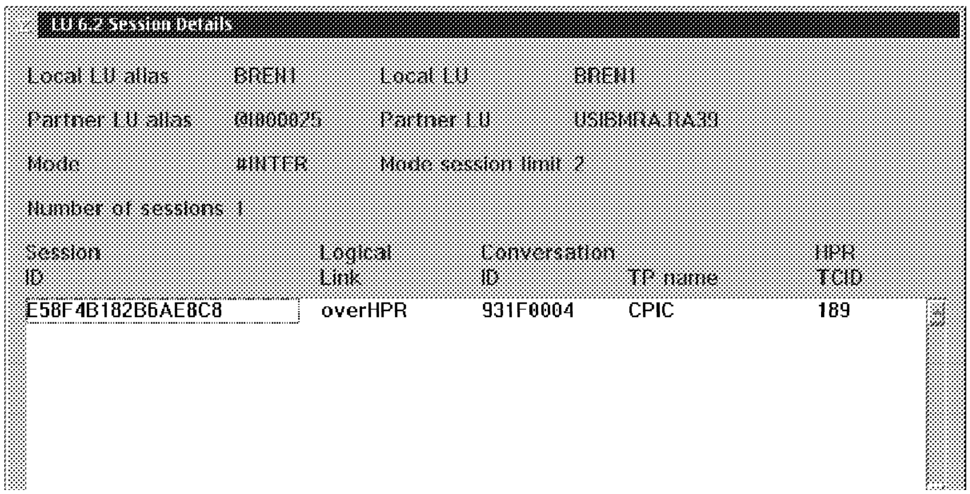

Figure 228. LU 6.2 Session Details (APING).

We then displayed the HPR connection summary, as in Figure 229.

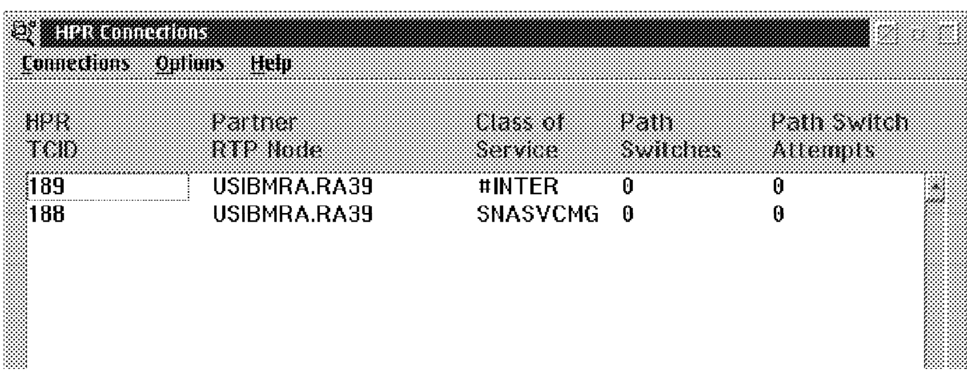

Figure 229. HPR Connection (APING)

The two expected RTP pipes to RA39 are there. We asked for the details of the #INTER pipe to see its route, as shown in Figure 230 on page 226.

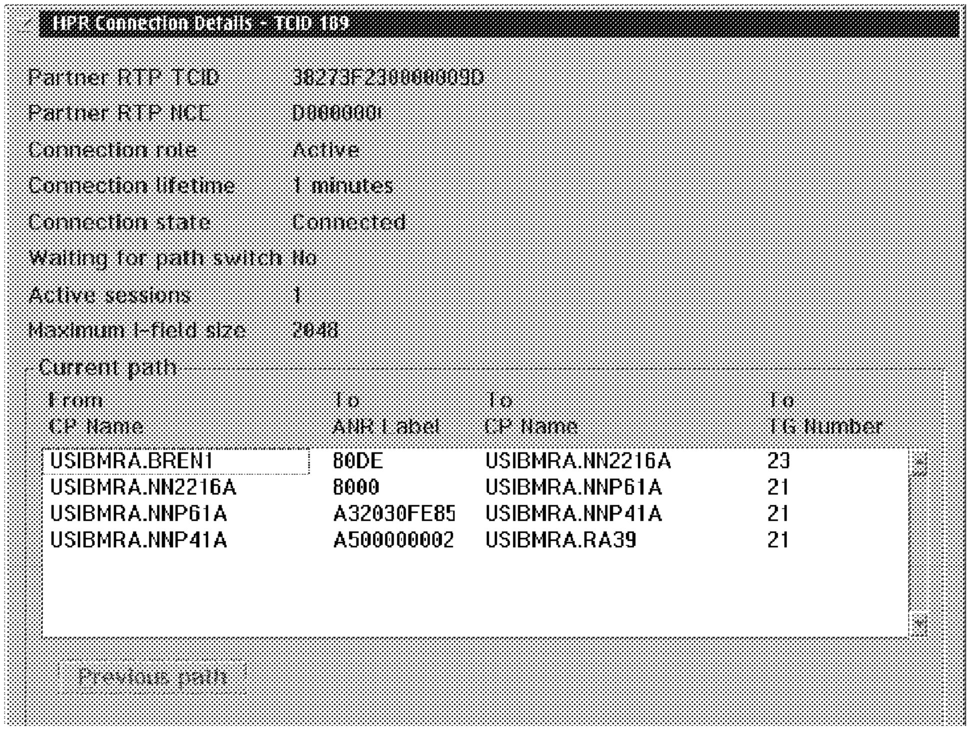

Figure 230. HPR Connection (APING)

This display is interesting on two counts:

1. The BX acts as an isolator between the downstream and upstream topologies, thus the usual APPN route selection mechanisms cause the session path to be calculated in two pieces, by NNS(PLU) in each topology. In this case, one piece was calculated by NN2216A (from BREN1 to NN2216A) and the other piece was calculated by NNP61A (from NN2216A to RA39). However, an HPR connection cannot be split in this way as NN2216A must act as an ANR node and should not have to be aware of the two routes glued together within it. Therefore, the BX (NN2216A) inserts CV85 subvectors in the RSCV to tell BREN1, the CP of the PLU, what the real path is. This is the same process that a border node uses if an HPR connection is to flow across it. Therefore, the RTP endpoint sends a Route Setup request with the full session route in it, and the full route is known at the endpoint as shown in Figure 230.

2. The route chosen, if you look at Figure 221 on page 220, is longer than it needs to be. The shortest route is clearly via CP2210T then NNP41A to RA39. However, the link between BREN1 and CP2210T *cannot* be used for any sessions in this configuration. NN2216A cannot inform NNP61A of the connection between BREN1 and CP2210T at session setup time. Indeed, NN2216A can only tell NNP61A that BREN1 is an LU on NN2216A itself. The route in the backbone network calculated by NNP61A therefore ends in NN2216A, and NN2216A can only extend the route to BREN1.

Displaying the #INTER RTP connection (CNR00077) from RA39 confirms the session path, which is:

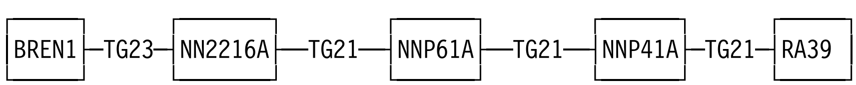

11.2.2 Path Switch with BX on the Path

We then broke the downlink connection between NN2216A and BREN1. This connection was carrying the two CP-CP sessions and the RTP connections for the APING sessions. Any dependent LU sessions would have had no chance and would have been terminated. To see what happened to the others, we displayed the network connections from the 2216 as shown in Figure 231.

```
2216T3 APPN >list link
    Name   Port Name  Intf       Adj CP Name  Type       HPR      State
========================================================================
    T61A        TRN1    0     USIBMRA.NNP61A    NN    ACTIVE     ACT_LS
    T41A        TRN1    0     USIBMRA.NNP41A    NN   ENABLED   RESET_LS
*
2216T3 APPN >list cp cp
           CP Name      Type    Status  Connwinner ID    Conloser ID
=======================================================================
    USIBMRA.NNP61A       NN     Active    34E957A1        34E957A3
    USIBMRA.NNP41A       NN   Inactive    00000000        00000000
     USIBMRA.BREN1       EN   Inactive    00000000        00000000
```

Figure 231. Display on 2216 after Link Failure

As expected, only the link to NNP61A and its CP-CP sessions remain. A similar display from the 2210 (Figure 232) is more interesting.

```
list cp cp
           CP Name      Type    Status  Connwinner ID    Conloser ID
=======================================================================
     USIBMRA.BREN1       EN     Active    B8930F86    1   B8930F82
    USIBMRA.NNP41A       NN     Active    B8930D60        B8930D62
    USIBMRA.NNP61A       NN   Inactive    00000000        00000000
APPN >list appc
LU Name            Mode  Type  FSM
=====================================
USIBMRA.NNP41A    CPSVCMG  Pri  ACT
USIBMRA.NNP41A    CPSVCMG  Sec  ACT
USIBMRA.RA39      CPSVRMGR Pri  ACT
USIBMRA.RA39      CPSVRMGR Sec  ACT
USIBMRA.BREN1     CPSVCMG  Sec  ACT  1
USIBMRA.BREN1     CPSVCMG  Pri  ACT  1
```

Figure 232. Displays on 2210 after Link Failure

The display reveals **1** that the CS/2 end node has changed its allegiance by establishing CP-CP sessions with CP2210T. If it has CP-CP connectivity into the network, it can initiate a path switch, so we displayed the RTP connection details again from BREN1. Figure 233 on page 228 shows what we saw.

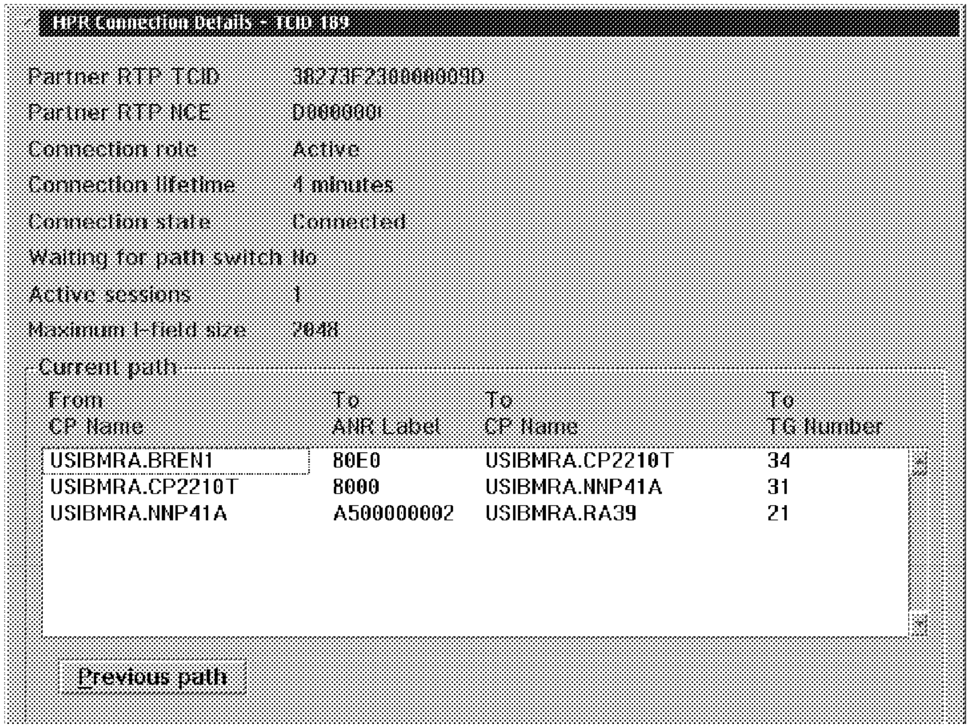

Figure 233. HPR Connection Details (New Path)

RTP path switch has occurred, and the new route goes via the new BX gateway (as it must) and then NNP41A. To confirm this, we displayed the same RTP pipe from RA39, as in Figure 234.

```
DISPLAY NET,ID=CNR00077,SCOPE=ALL
IST097I  DISPLAY  ACCEPTED
IST075I  NAME = CNR00077         , TYPE = PU_T2.1
IST1392I  DISCNTIM = 00010 DEFINED AT PU   FOR DISCONNECT
IST486I  STATUS= ACTIV--LX-, DESIRED STATE= ACTIV
IST1043I  CP NAME = BREN1   , CP NETID = USIBMRA , DYNAMIC LU = YES
IST1589I  XNETALS = YES
IST933I  LOGMODE=***NA***, COS=#INTER
IST1476I  TCID X'38273F230000009D' - REMOTE TCID X'0000000000000189'
IST1481I  DESTINATION CP USIBMRA.BREN1     - NCE X'80              '
IST1587I  ORIGIN NCE X'D000000000000000'
IST1477I  ALLOWED DATA FLOW RATE =   3085 KBITS/SEC
IST1516I  INITIAL DATA FLOW RATE =   1597 KBITS/SEC
IST1511I  MAXIMUM NETWORK LAYER PACKET SIZE =   2048 BYTES
IST1478I  NUMBER OF UNACKNOWLEDGED BUFFERS =     0
IST1479I  RTP CONNECTION STATE = CONNECTED             - MNPS = NO
ST1480I  RTP END TO END ROUTE - PHYSICAL PATH
IST1460I  TGN  CPNAME            TG TYPE      HPR
IST1461I   21  USIBMRA.NNP41A     APPN         RTP
IST1461I   31  USIBMRA.CP2210T    APPN         RTP
IST1461I   34  USIBMRA.BREN1      APPN         RTP
IST231I  RTP      MAJOR NODE = ISTRTPMN
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I  STATE TRACE = OFF
IST355I  LOGICAL UNITS:
IST080I  BREN1    ACT/S----Y
IST314I  END
```

Figure 234. New Path after BX Switch

The new path is shown below:

```
┌───────┐       ┌───────┐       ┌───────┐       ┌──────┐
│BREN1  ├─TG34──┤CP2210T├─TG31──┤NNP41A ├─TG21──┤RA39  │
└───────┘       └───────┘       └───────┘       └──────┘
```

In fact, the newly switched path is shorter than the original path. This is simply because the newly acquired NN (BX) server of BREN1 happens to be nearer the RTP partner than the original server.

This path switch worked well because the CP-CP sessions between BREN1 and its NN servers did not flow over RTP pipes, and were therefore terminated immediately after the link failure. The 2216 and 2210, in fact, will terminate such sessions at once even if they *do* flow over RTP connections. If there is no valid alternative path for CP-CP sessions the 221X nodes will not wait for the path switch timer to expire, thus allowing a timely RTP path switch for LU-LU sessions just as we saw.

Appendix A. Adaptive Rate-Based Flow and Congestion Control

The ARB mechanism used by HPR for both flow control (keeping the network running smoothly) and congestion control (keeping the receiving node running smoothly) is described in Chapter 9 of *Inside APPN - the Essential Guide to the Next-Generation SNA*. To assist in the reader's understanding, we offer a shortened description here together with some examples of traces taken during our tests, that will help to illustrate the concepts.

A.1 Introduction

The adaptive rate-based (ARB) congestion and flow control algorithm allows RTP connections to make more efficient use of network resources such as links and buffers. Input traffic (offered load sent by an RTP connection endpoint) entering the network is regulated by the ARB algorithm based on the conditions in the network and the partner RTP endpoint. When the network or partner RTP endpoint approaches congestion (in other words, there is increasing delay and throughput fails to keep pace with incoming traffic), input traffic is reduced. When the capacity of the network or partner RTP endpoint increases, input traffic is increased.

Figure 235 shows the relationship between network throughput and offered load for a given path.

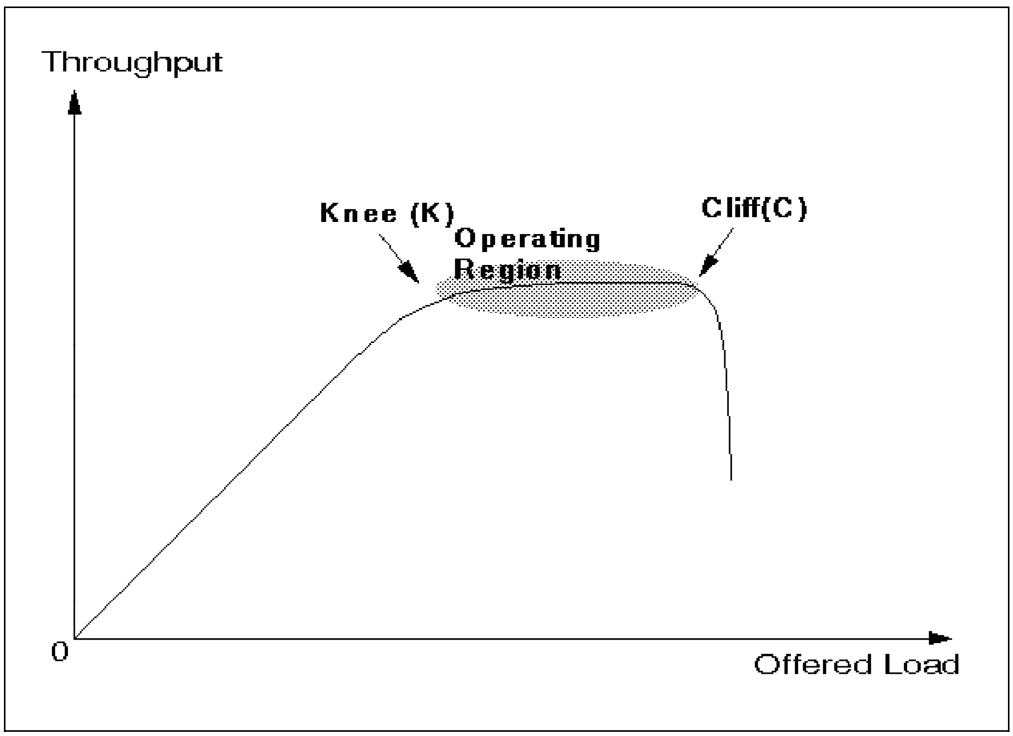

Figure 235. ARB Operating Region

The knee (point K) is the point beyond which the path starts to become congested (in other words, link transmission queues develop along the path resulting in higher network delays). Beyond point K, an increase in the send rate does not result in an increase of throughput. ARB detects this precongestion condition (saturation) and adjusts the sending rate accordingly, thus preventing the network from operating beyond the cliff (point C). The cliff is the point beyond which congestion results in packet loss and larger queueing delays at the links along the path. The network throughput decreases sharply when the offered load increases beyond point C. The ARB algorithm is designed to operate between points K and C (the operating region).

The ARB algorithm has the following properties:

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

The ARB algorithm employs a closed-loop, distributed, control mechanism based on information exchanged between two partner RTP connection endpoints. Figure 236 shows the relationship between an ARB sender and an ARB receiver over an RTP connection.

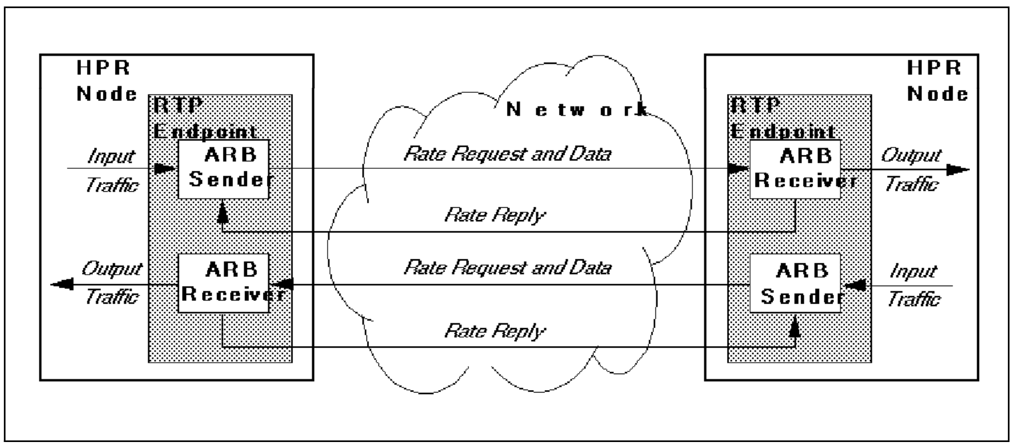

Figure 236. ARB Mechanism Overview

The ARB algorithm is implemented by the endpoints of an RTP connection. At each RTP endpoint there are two components, an ARB sender and an ARB receiver. The ARB procedures performed by the sender at one end are the same as those performed by the sender at the other end, similarly for the receivers. Note that intermediate nodes have no awareness of the ARB protocol and so do not participate in it. This is because ANR nodes have no awareness of the RTP pipes. This is in marked contrast with subarea VR flow control, which relies on feedback from intermediate nodes as well as endpoint nodes.

The data being regulated by the ARB algorithm always flows from a sender (ARB sender) to a receiver (ARB receiver). The sender continually queries the receiver, by sending it a rate request along with the data, in order to obtain information about the state of the network and the state of the node containing the receiver. The receiver responds by sending back a rate reply. The sender adjusts its send rate based on the information in the rate reply. The sender may reduce its send rate to forestall congestion or increase it to take advantage of the available network capacity.

The fixed characteristics of the path (the speed of the slowest link along the path and the total transmission time over the entire path) are used to calculate two very important parameters used in the ARB algorithm: the range begin time and the range end time. The range begin time (which we refer to as K time) is the amount of delay that will cause the network to reach the knee point (point K in Figure 235 on page 231). The range end time (which we call C time) is the amount of delay that will cause the network to reach the cliff (point C in Figure 235 on page 231).

The K time is the time taken to transmit 8000 bits over the slowest link in the RTP path between sender and receiver. The C time is the time taken to transmit 80000, 120000 or 160000 bits over the same link. Which of these three values is chosen depends on the number of hops and the number of slow-speed links in the path. This is so that longer paths can compete fairly with shorter paths for the same traffic.

These path characteristics are communicated by using an ARB setup message that is sent when the RTP connection is established (when the connection setup is done) and whenever the path changes because a path switch has occurred. Note that either RTP endpoint may send an ARB setup message as a result of a path switch.

The rate request, rate reply, and setup messages are transmitted in the ARB optional segment in the transport header of the NLP.

A.2 ARB Algorithm Overview

The sender, at regular intervals approximating the round trip delay on RTP connection, sends a rate request segment which is always piggybacked on an NLP containing data. This segment includes the sender′s measurement interval which is equal to the time that has elapsed since the last rate request segment was sent. Upon receipt of the rate request, the receiver determines whether any delay has occurred in the network. It does this by calculating the difference between the sender′s measurement interval and the receiver′s measurement interval. The receiver′s measurement interval is the time that has elapsed since the last rate request segment was received.

The receiver also takes into account previous delays remembered from earlier rate request messages. Based on this network delay information, the receiver will recommend appropriate actions to be taken by the sender. These actions are communicated in a rate reply segment that enables the sender to adjust its send rate appropriately. The rate reply segment may either be sent stand-alone or be carried along with data. The receiver, in addition to deriving its recommendations based on network delays, can also tell the sender to reduce its sending rate based on conditions within the receiver node (for example, buffer shortage). Figure 237 on page 234 illustrates these principles.

```
┌──────┐                                                    ┌────────┐
│Sender│                                                    │Receiver│
└──────┘                                                    └────────┘
 ▲    ─┼──Rate request:Sender_measurement_interval: t0──────────►  ▲
 │t1                                                               │
 ▼    ─┼──Rate─request:Sender_measurement_interval: t1──┐          │ t2
                                                        │          │
                                                        └────────► ▼
      ◄────Rate─reply:─Rate_Adjustment_Action: ─Normal──────────
                                                -Restraint
                                                -Slowdown1
                                                -Slowdown2
                                                -Critical
```

Figure 237. Rate Adjustment Overview

When the receiver gets the second rate request message it compares t1 (the interval between the sending of consecutive rate requests) with t2 (the interval between receipt of the same two requests). Depending on the relationship between t1 and t2, the receiver sends a rate reply recommending an action to the sender.

A.2.1 Receiver Actions

The receiver recommends action to the sender based on two sources of information:

1. Itself.

   The receiver ARB state is based on conditions at the ARB receiver node such as buffer availability. The possible values for this state are Normal, Restraint, Slowdown1, Slowdown2 and Critical. The algorithm for determining this value is dependent on product implementation, but a typical example might be related to the usage of the receiving buffer pool, as follows:

   - 0% < usage < 75% - Normal
   - 75% < usage < 80% - Restraint
   - 80% < usage < 85% - Slowdown1
   - 85% < usage < 90% - Slowdown2
   - 90% < usage < 100% - Critical

2. The network.

   Network delays are detected by measuring the change in measurement (rate request) intervals. The *delay change* value is the difference between the sender′s measurement interval (t1 in the diagram above) and the receiver′s measurement interval (t2 above). It equates to the difference in network delay experienced by the current rate request message versus that experienced by the previous message. A positive delay change means that the current rate request message took longer to traverse the network than the previous one, and therefore network queues are increasing. A negative delay change means that the current rate request took less time to cross the network than the previous one, and therefore the network queueing is decreasing. If the delay change is large enough, it will cause the ARB sender′s sending rate to be lowered.

   The receiver keeps a running total of the cumulative delay change. If the network is running smoothly, this running total will hover around zero. After corrective action is taken, the running total is reset to zero to allow for the new operating conditions.

The actions taken by the ARB algorithm depend on the value of the running total in the delay changes (known as the delay change sum):

- If the delay change sum is zero or negative, there is no problem and the indication returned to the sender is Normal.
- If the delay change sum is positive, but less than the K time, the network is experiencing some delay but has not yet reached the knee point. Therefore, the recommended action returned is still Normal.
- If the delay change sum is between the K time and the C time, the network is operating at its optimum level. The indication returned is Restraint, meaning that neither an increase nor a decrease in sending rate is required.
- If the delay change sum is greater than the C time but less than four times the C time, the network appears to have “fallen over the cliff” and the potential for excessive congestion exists. If this state is reached twice in succession, and the delay change is still positive, then the receiver sends Slowdown1 to the sender. If the state is temporary (only attained once) then the indication is still Restraint. When Slowdown1 is sent the delay change sum is reset to zero.
- If the delay change sum is greater than four times the C time, Restraint is sent. This is because the likelihood of reaching this state suddenly is extremely low; the actions outlined above should prevent it. This condition is most probably due to internal delays in one of the RTP endpoints, so the receiver assumes the network is functioning normally and returns Restraint.

A.2.2 Sender Actions

The adaptation of the sending rate depends on more than just the information in the rate reply:

- The indication (normal, restraint and so on) received from the receiver.
- The current operating mode of the sender. This mode (green, yellow or red) is a reflection of the previous history of the connection and is reviewed whenever a rate reply message is received. The sender starts the RTP connection in green mode.
- Other factors such as lost data, timeouts and an idle connection.

The action taken by the sender on receipt of the rate reply indicators is as follows:

- Normal.

  If the mode is green, the send rate is increased by an amount determined by the link capacity (not a percentage of the current send rate). If the mode is yellow or red, the send rate is left alone but the mode is changed (to green or yellow respectively).

- Restraint.

  The send rate is left unchanged, but the mode is changed from red to yellow or yellow to green, as appropriate.

- Slowdown1.

The send rate is reduced by 12½%, unless the lowest link speed is 128 kbps or more (in which case the send rate is reduced by 25%). The mode is set to yellow.

- Slowdown2.

  The send rate is reduced by 25% and the mode is set to yellow.

- Critical.

  The send rate is reduced by 50%, but never to less than 1 kbps. The mode is set to red.

In addition, other factors may affect the sending rate:

- If data has been lost, or the sender times out waiting for an acknowledgement to sent data, the send rate is reduced by 50% (but never to below 1 kbps) unless the operating mode was red. If the mode was red, corrective action is assumed to be in place already. The mode is set to red in any case upon receipt of a lost data indication.
- If the connection has been idle for the interval specified in the alive timer, the sending rate is reduced by 12½% (but never lower than the initial rate) and the mode is set to green.

Table 1 summarizes this algorithm. The first column represents the event that takes place, and the other columns show the result depending on what the current operating mode was. The numbers in the results columns indicate to which state the sender is switched: 1 means green, 2 means yellow, 3 means red. The words in brackets denote the effect on the sending rate of each event.

Table 1. Rate Adjustment Logic

	GREEN 1	**YELLOW 2**	**RED 3**
rcv, Rate_reply, normal	1 (increase)	1 (no change)	2 (no change)
rcv, Rate_reply, restraint	1 (no change)	1 (no change)	2 (no change)
rcv, Rate_reply, slowdown1	2 (decrease small)	2 (decrease small)	2 (decrease small)
rcv, Rate_reply, Slowdown2	2 (decrease medium)	2 (decrease medium)	2 (decrease medium)
rcv, Rate_reply, Critical	3 (decrease large)	3 (decrease large)	3 (decrease large)
short_req_timer_expires	3 (decrease large)	3 (decrease large)	3 (decrease large)
rcv, data_loss_indication	3 (decrease large)	3 (decrease large)	3 (no change)
send, status_exchange	1 (resync)	2 (resync)	3 (resync)
idle_state_indication	1 (decrease small)	1 (decrease small)	1 (decrease small)

A.3 RTP Connection Fairness

When multiple RTP connections having the same transmission priority are sharing a link, the ARB algorithm works in such a way so they all get an equal share of the link's bandwidth. Thus an individual RTP connection is prevented from hogging more than its fair share of bandwidth. Fairness occurs naturally because of the way the send rate is incremented and decremented. The basic idea is that all RTP connections increment their send rate using the same increment values. RTP connections with higher send rates will increment by the same absolute amount as those with lower send rates. However, decrementing the send rate is done by percentage. This means that RTP connections with higher send rates will decrement faster than those with lower send rates. This eventually leads to all the RTP connections stabilizing at the same send rate.

A.4 ARB Flows

The following sample flows show how the ARB segments are carried over an RTP connection. Figure 238 summarizes the flows that we discuss.

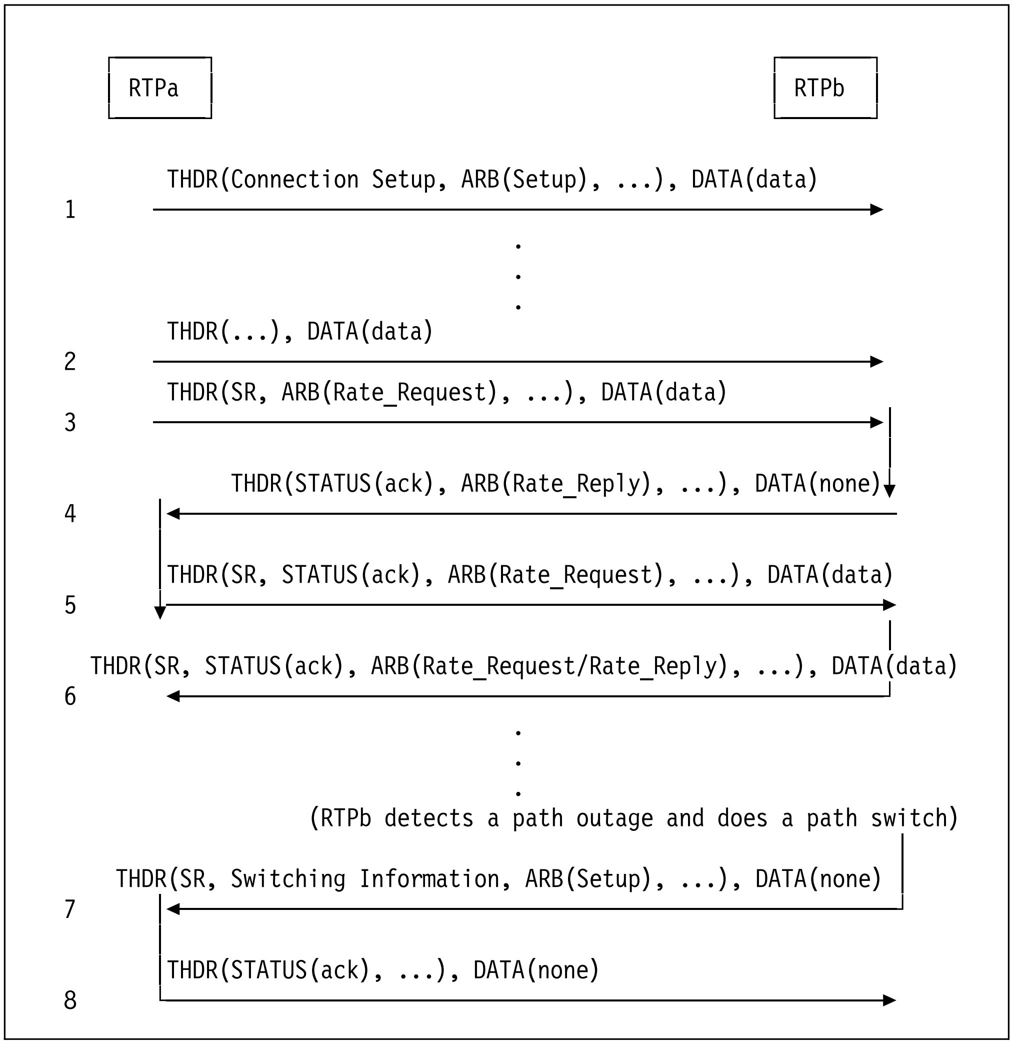

Figure 238. ARB Segments Flowing on RTP Connection

Notes:

1. RTPa establishes a connection with RTPb by sending a connection setup segment. The ARB setup segment is included (along with the switching information segment - not shown) so that ARB parameters can be initialized.
2. RTPa sends data to RTPb. No ARB rate request segment is included at this time.

3. The smoothed round-trip time interval has elapsed so RTPa sends an ARB rate request segment along with the data. Status is always requested (SR) on rate requests.
4. RTPb processes the received ARB rate request and returns a rate reply. RTPb has no data to send so it sends the rate reply stand-alone. It is not necessary to request status (SR) on a rate reply.
5. The smoothed round-trip time interval has elapsed again so RTPa sends an ARB rate request segment along with the data.
6. RTPb processes the received ARB rate request and returns a rate reply. This time RTPb has data to send. Also, the smoothed round-trip time interval has elapsed so a rate request needs to be sent. So RTPb sends a rate reply, rate request, and data all in one message.
7. RTPb has performed a path switch and sends the new path information (in the switching information segment) and the accompanying ARB parameters (in the ARB setup segment). Both RTPb and RTPa reinitialize all the ARB parameters and restart the ARB algorithm from the beginning.

A.5 Trace Examples

The following examples illustrate the previous discussion.

A.5.1 ARB (Setup)

Figure 239 on page 239 shows the ARB status of a node, followed by the sending of an NLP containing an ARB setup segment.

```
Time stamp: 11:02:32.59
DLC type: HPR
TCID: 0x0000001D
||TCID = 0x0000001d
||Allowed send rate (Kbps) = 1600  1
||Actual send rate (Kbps) = 0
||Accumulation of queueing time = 0
||Current queueing time = 0
||Round trip delay = 0
||Rate increment step = 128
||Minimum rate increment = 32
||Maximum rate increment = 128
||ARB mode = Green 2
||Rate increment count = 0
||Rate decrement count = 0
||RTT event counts before measurement interval adjustment = 0
||ARB waits = 0
||Burst size = 8192
||Burst time = 40
||Remaining burst size = 8192
||Remaining burst time = 40
||Next burst window = 0x0044e726
||Bytes transmitted this burst = 0
||Measured round trip time = 0
||Minimum round trip time = 100000
||Status request time = 0x00000000
||Measurement interval = 0x000000c8
||Current measurement interval = 0x00000000
||Next measurement interval = 0x0044e7c5
||Last measurement request sent = 0x00000000
||Last measurement request received = 0x00000000
||Transmitted data = 0
||Outstanding rate request = No
||Partner requests a reate reply = No
||ARB reply flags = 0x0000
||Maximum send rate (forward path) = 16000
||Maximum send rate (return path) = 16000
||Maximum queueing time (forward path) = 10
||Maximum queueing time (return path) = 10
||Accumulated transmit time (forward path) = 75
||Accumulated transmit time (return path) = 75
||Green threshold = 1
||Ring 0 timestamp = 0
```

Figure 239 (Part 1 of 3). ARB Setup Flow

```
Line:   9189 Send RTP Header
Time stamp: 11:02:32.59
DLC type: HPR
TCID: 0x0000001D
 ||RTP Header
 ||   Switching mode = ANR
 ||   Transmission priority = Network
 ||   Packet is time sensitive = Yes
 ||   Minor congestion exists = No
 ||   Significant congestion exists = No
 ||   ANR Routing = 0x8014d0201025ff
 ||   TCID assigned by sender = Yes
 ||   TCID = 0x000000000000001d
 ||   Connection Setup segment present = Yes
 ||   Start of message = Yes
 ||   End of message = Yes
 ||   Receiver must reply with a status segment = Yes
 ||   Receiver must reply ASAP = Yes
 ||   Sender will retransmit this packet = Yes
 ||   Last message on this connection = No
 ||   CQF type = Originator
 ||   Optional segments present = Yes
 ||   Data offset = 0x00bc
 ||   Data length = 0x0000008c
 ||   Byte sequence number = 0x00000000
 ||   Connection Qualifier/Source Identifier
 ||      Network Address HPR control vector
 ||       (0x0002) Network address is for point to point connection = Yes
 ||         Network Identifier HPR control vector
 ||          (0x0002) Network identifier = USIBMRA
 ||         Node Identifier (CP Name) HPR control vector
 ||          (0x0002) Node identifier = CM5HPRNN
 ||         NCE Identifier control vector
 ||          (0x0002) NCE identifier = 80
 ||         NCE Instance Identifier control vector
 ||          (0x0002) NCE instance identifer = 0x34db2207
 ||   Optional segments:
 ||      Connection Setup segment
 ||         Version = 1.1
 ||         Target resource ID present = Yes
 ||         ARB is used for flow/congestion control = Yes
 ||         Connection is reliable = Yes
 ||         Topic Identifier control vector
 ||          (0x0002) Topic is user defined = No
 ||          (0x0003) Topic identifier = RSETUP
 ||         Network Identifier HPR control vector
 ||          (0x0002) Network identifier = USIBMRA
 ||         Node Identifier (CP Name) HPR control vector
 ||          (0x0002) Node identifier = NNP61A
 ||         NCE Instance Identifier control vector
 ||          (0x0002) NCE instance identifer = 0x00000f5a
```

Figure 239 (Part 2 of 3). ARB Setup Flow

```
Switching Information segment
   HPR Switching Information control vector
    (0x0002) REFIFOing allowed = No
             Origin is mobile = No
             Locate search is required = No
             Limited resource links exist along path = No
             NCE is used for all LUs = Yes
    (0x0004) Maximum packet size on return path = 2058
    (0x0008) Path switch time (milliseconds) = 0
    (0x000C) RTP liveness timer (seconds) = 61
      ANR Path control vector
       (0x0003) ANR label represents a subarea network route = No
       (0x0004) ANR label = a320302482
   HPR Return Route TG Descriptor control vector
    (0x0002) Boundary function performed = No
    (0x0003) TG entry count = 1
      TG descriptor control vector
         TG Identifier TG Descriptor subfield
          (0x0002) TG number = 25
          (0x0004) Partner name = USIBMRA.CM5HPRNN
          (0x0014) Link connection network = No
                   Additional configuration information = No
                   HPR = Supported
                   TG type = Boundary Function based or APPN
                   Intersubnet link = No
                   Extended border node = Not supported
                   RTP Tower = Supported
Adaptive Rate-Based segment
   Message type = Setup 3
   Rate adjustment action = Normal
   Time to transmit a 1000 byte data packet (microseconds) = 1000
   Time to transmit a 10000 byte data packet (microseconds) = 10000
   Link capacity of the slowest link (Kbps) = 16000   4
   Time to transmit 1200 bits (microseconds) = 75
```

Figure 239 (Part 3 of 3). ARB Setup Flow

Notes:

1. If the minimum link capacity **4** is 16000 kbps, the initial allowed send rate is 10% of this (1600 kbps **1**). The ARB mechanism will adjust this value until the operating region (between knee and cliff) is reached.
2. **3** indicates the ARB setup segment.
3. **2** shows that the current ARB operating mode is green.

A.5.2 ARB (Request)

Figure 240 on page 242 is part of a CS/2 trace showing an ARB request. The smoothed round-trip time interval has elapsed so the ARB request is sent.

```
Line:    612 Receive RTP Header
Time stamp: 11:01:25.60
DLC type: HPR
TCID: 0x00000016
||RTP Header
||   Switching mode = ANR
||   Transmission priority = Network
||   Packet is time sensitive = Yes
||   Minor congestion exists = No
||   Significant congestion exists = No
||   ANR Routing = 0x80ff
||   TCID assigned by sender = No
||   TCID = 0x0000000000000016
||   Connection Setup segment present = No
||   Start of message = Yes
||   End of message = Yes
||   Receiver must reply with a status segment = Yes
||   Receiver must reply ASAP = Yes
||   Sender will retransmit this packet = Yes
||   Last message on this connection = No
||   CQF type = Not present
||   Optional segments present = Yes
||   Data offset = 0x0020
||   Data length = 0x00000089
||   Byte sequence number = 0x0001ef6b
||   Optional segments:
||      Adaptive Rate-Based segment
||         Message type = Rate request
||         Rate adjustment action = Normal
||   1     Time elapsed since last rate request sent (microseconds) =0000
```

Figure 240. ARB Rate Request

Note **1** the last measurement interval is sent to the receiver. This is the first rate request on an RTP connection.

A.5.3 ARB (Reply)

Figure 241 on page 243 is the reply to the rate request.

```
Line:    973 Send RTP Header
Time stamp: 11:01:25.62
DLC type: HPR
TCID: 0x00000016
||RTP Header
||   Switching mode = ANR
||   Transmission priority = Network
||   Packet is time sensitive = Yes
||   Minor congestion exists = No
||   Significant congestion exists = No
||   ANR Routing = 0x8014d0201025ff
||   TCID assigned by sender = No
||   TCID = 0x000000000b46c648
||   Connection Setup segment present = No
||   Start of message = No
||   End of message = No
||   Receiver must reply with a status segment = No
||   Receiver must reply ASAP = No
||   Sender will retransmit this packet = Yes
||   Last message on this connection = No
||   CQF type = Not present
||   Optional segments present = Yes
||   Data offset = 0x0034
||   Data length = 0x00000000
||   Byte sequence number = 0x000001c8
||   Optional segments:
||      Status segment
||         Packets have been lost = No
||         Connection is idle = No
||         Status report number = 0x0001
||         Status acknowledgment number = 0x0005
||         Received sequence number = 0x0001eff5
||         Delivered sequence number = 0x00000000
||      Adaptive Rate-Based segment
||         Message type = Rate reply
||     1   Rate adjustment action = Normal
```

Figure 241. ARB Rate Reply

The adjustment action recommended **1** is Normal. The sender is authorized to increase its sending rate.

A.5.4 ARB (Request/Reply)

In Figure 242 on page 244, the rate request and rate reply are combined into a single segment.

```
Line:  21957 Send RTP Header
Time stamp: 11:04:57.73
DLC type: HPR
TCID: 0x00000012
||RTP Header
||   Switching mode = ANR
||   Transmission priority = Network
||   Packet is time sensitive = Yes
||   Minor congestion exists = No
||   Significant congestion exists = No
||   ANR Routing = 0x8014d0201025ff
||   TCID assigned by sender = No
||   TCID = 0x000000000b448bf0
||   Connection Setup segment present = No
||   Start of message = Yes
||   End of message = Yes
||   Receiver must reply with a status segment = Yes
||   Receiver must reply ASAP = Yes
||   Sender will retransmit this packet = Yes
||   Last message on this connection = No
||   CQF type = Not present
||   Optional segments present = Yes
||   Data offset = 0x0034
||   Data length = 0x00000012
||   Byte sequence number = 0x000000ab
||   Optional segments:
||      Status segment
||         Packets have been lost = No
||         Connection is idle = No
||         Status report number = 0x0001
||         Status acknowledgment number = 0x0001
||         Received sequence number = 0x000001d9
||         Delivered sequence number = 0x00000000
||      Adaptive Rate-Based segment
||         Message type = Rate request/Rate reply
||         Rate adjustment action = Normal
||         Time elapsed since last rate request sent (microseconds) = 267264
```

Figure 242. ARB Rate Request/Reply

A.5.5 Slowdown 1 Example

In Figure 243 on page 245, the rate reply indicates Slowdown1 so the ARB mode of the sender is switched to yellow.

```
||      Adaptive Rate-Based segment
||         Message type = Rate request
||         Rate adjustment action = Normal
||         Time elapsed since last rate request sent (microseconds) = 69


||      Adaptive Rate-Based segment
||         Message type = Rate reply
||         Rate adjustment action = Slowdown1
||         Receiver's receive rate (Kbps) = 0


Line:   5285 ARB Segment
Time stamp: 11:02:02.62
DLC type: HPR
TCID: 0x00000017
||TCID = 0x00000017
||Allowed send rate (Kbps) = 32
||Actual send rate (Kbps) = 0
||Accumulation of queueing time = 0
||Current queueing time = -17
||Round trip delay = 70
||Rate increment step = 16
||Minimum rate increment = 32
||Maximum rate increment = 128
||ARB mode = Yellow
```

Figure 243. ARB Slowdown1 Message

Appendix B. A Complete Scenario

In this section we show a detailed illustration of an HPR and DLUR scenario, with traces and displays. The purpose is to give you an example with which to compare your own environment when a problem occurs and problem determination is called for. Figure 244 shows the network setup which we used.

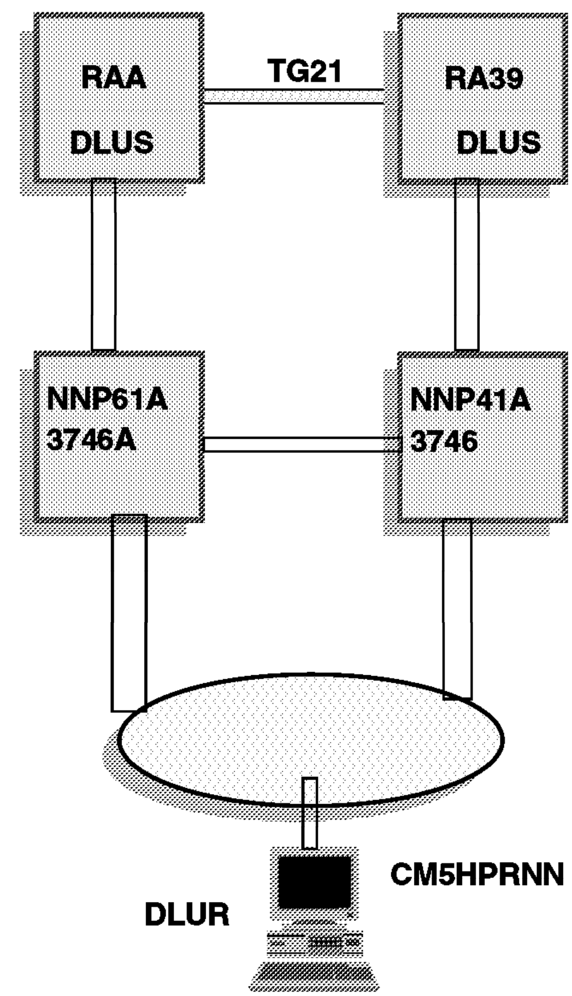

Figure 244. Test Configuration

We ran this test with the following equipment:

- A CS/2 PC with HPR and DLUR support, configured as a network node
- Two connections to 3746 NNs from the CS/2 node
- One connection between each 3746 and a host
- Two hosts, both NNs (and therefore DLUSs) linked via MPC, and both RTP-capable

The traces and displays were taken on the CS/2 node, mainly because formatting a trace is so easy on this platform. Not all the activity was traced; we confined ourselves to the main flows as this redbook is quite long enough already.

B.1 Network Startup

In this section we see the CP-CP session establishment, the Route Setup and the establishment of the SSCP control sessions (LU activation).

B.1.1 XID Exchange, LINK0001

As a result of link activation, XIDs are exchanged between CM5HPRNN (the CS/2 node) and first NNP61A then NNP41A. What is new for HPR is control vector CV61 which gives the HPR capabilities of each node. CV61 has two subfields:

- Subfield X′80′ is the LAN subfield and contains the LLC SAP that is to be used by the adjacent node when sending NLPs.
- Subfield X′81′ is used with control flows over RTP and provides the information that would otherwise be exchanged via Route Setup. If CP-CP sessions are to be established over an RTP connection, then there is no other possible source of information than the XID exchange.

```
Line:     58   Recv XID
Time stamp: 08:26:55.54
DLC type: IBMTRNET
Adapter number: 00
Destination address: 40003746214408
ALS ID: 4008279326F58625
XID:
|| (0x0000) Format = 3
||          Node type = 2
|| (0x0001) Total length = 111
|| (0x0002) Node ID = 0x05d66918
|| (0x0008) Init-self = Can receive
||          Stand-alone BIND = Can receive
||          BIND segment generation = Can generate
||          BIND segment receipt = Can receive
||          FID type = FID 2
|| (0x0009) ACTPU suppression = ACTPU not requested
||          APPN network node = Yes
||          CP-CP sessions requested = Yes
||          CP-CP sessions supported = Yes
||          Exchange state = Negotiation proceeding
||          Secondary initiated non-activation exchange = Supported
||          CP name change = Supported
|| (0x000A) Adaptive BIND pacing sender = Supported
||          Adaptive BIND pacing receiver = Supported
||          Quiesce TG requested = No
||          PU Capabilities control vector = Not supported
||          Sender is an APPN peripheral border node = No
||          Adaptive BIND pacing = Supported for all LUs, unless overridden
|| (0x000F) Parallel TGs = Supported
||          DLUR XID sender prefers ACTPU over CP-SVR pipe = Yes
||          DLUS-served LU registration = Supported
||          Downstream Branch Extender = Not supported
||          Upstream Branch Extender = Not supported
|| (0x0010) TG number = 26
|| (0x0011) DLC type = SDLC
|| (0x0012) DLC data length = 11
|| (0x0013) ABM = Supported
||          Link station role = Primary
||          Short hold mode status = Not reconnection
||          Short hold mode capability = Not supported
||
```

```
||         Transmit/receive capability = Secondary
|| (0x0014) ABM nonactivation XID exchange initiator = No
|| (0x0015) Maximum receive BTU length = 2058
|| (0x0017) SDLC CR profile = SNA link profile
|| (0x0018) SIM/RIM options = Not supported
|| (0x001B) Maximum I-frames before ACK = 6
|| (0x001D) Network name control vector:
||           (0x0002) Network name type = CP
||           (0x0003) Name = RA.NNP61A
|| (0x002E) Product set ID control vector:
||                f5f7f0f0 f9f0f0f2 f6                <.........       > <5700
```

CV X′61′ HPR capabilities CV

```
|| (0x0047) HPR Capabilities control vector
||           (0x0002) Error recovery = Available if requested by partner
||                    RTP Tower = Supported
||                    Control Flows over RTP Tower = Supported
||           (0x0005) ANR label = a320304c82
```

Subfield X′81′

```
|| (0x0051)    Control Flows over RTP Tower subfield
||              (0x0002) Maximum send packet size = 0x080a
||              (0x0004) Path switch time = 60000
||              (0x0008) Sender is mobile = No
||                       Multilink transmission groups = Not supported
||              (0x0009) Control Point NCE instance identifier = 0x00000ea6
||              (0x000D) Route Setup NCE instance identifier = 0x00000ea6
||              (0x0012) Control Point NCE identifier = d0201025
||              (0x0017) Route Setup NCE identifier = d0201025
```

Subfield X′80′

```
|| (0x006C)    Token Ring/Ethernet 802.2 LLC subfield
||              (0x0002) LLC SAP = 08
```

```
Line:     67   Send XID
Time stamp: 08:26:55.54
DLC type: IBMTRNET
Adapter number: 00
Destination address: 40003746214408
ALS ID: 4008279326F58625
XID:
|| (0x0000) Format = 3
||          Node type = 2
|| (0x0001) Total length = 132
|| (0x0002) Node ID = 0x05d00000
|| (0x0008) Init-self = Cannot receive
||          Stand-alone BIND = Can receive
||          BIND segment generation = Can generate
||          BIND segment receipt = Can receive
||          FID type = FID 2
```

```
(0x0009) ACTPU suppression = ACTPU not requested
         APPN network node = Yes
         CP-CP sessions requested = Yes
         CP-CP sessions supported = Yes
         Exchange state = Negotiation proceeding
         Secondary initiated non-activation exchange = Supported
         CP name change = Supported
(0x000A) Adaptive BIND pacing sender = Supported
         Adaptive BIND pacing receiver = Supported
         Quiesce TG requested = No
         PU Capabilities control vector = Not supported
         Sender is an APPN peripheral border node = No
         Adaptive BIND pacing = Supported for all LUs, unless overridden
(0x000F) Parallel TGs = Supported
         DLUR XID sender prefers ACTPU over CP-SVR pipe = No
         DLUS-served LU registration = Supported
         Downstream Branch Extender = Not supported
         Upstream Branch Extender = Not supported
(0x0010) TG number = 26
(0x0011) DLC type = SDLC
(0x0012) DLC data length = 11
(0x0013) ABM = Supported
         Link station role = Secondary
         Short hold mode status = Not reconnection
         Short hold mode capability = Not supported
         Transmit/receive capability = Secondary
(0x0014) ABM nonactivation XID exchange initiator = No
(0x0015) Maximum receive BTU length = 2058
(0x0017) SDLC CR profile = SNA link profile
(0x0018) SIM/RIM options = Not supported
(0x001B) Maximum I-frames before ACK = 4
(0x001D) Network name control vector:
          (0x0002) Network name type = CP
          (0x0003) Name = USIBMRA.CM5HPRNN
(0x0030) Network name control vector:
          (0x0002) Network name type = LS
          (0x0003) Name = LINK0001
(0x003B) Product set ID control vector:
            Hex dump:
               10280011 11040e02 f5f6f2f2 f8f7f8f0   <.(..............> <....
               f0f4f1f0 16110313 0011f0f0 f0f00000   <................> <0410
               00000000 00000000 0000                <..........      > <....
```

CV X'61'

```
(0x0065) HPR Capabilities control vector
          (0x0002) Error recovery = None
                   RTP Tower = Supported
                   Control Flows over RTP Tower = Supported
          (0x0005) ANR label = 801c
```

Subfield X'80'

```
(0x006C)    Token Ring/Ethernet 802.2 LLC subfield
             (0x0002) LLC SAP = 08
```

Subfield X'81'

```
(0x006F)    Control Flows over RTP Tower subfield
             (0x0002) Maximum send packet size = 0x080a
             (0x0004) Path switch time = 60000
             (0x0008) Sender is mobile = No
                      Multilink transmission groups = Not supported
             (0x0009) Control Point NCE instance identifier = 0x34e1a64d
             (0x000D) Route Setup NCE instance identifier = 0x34e1a64d
             (0x0012) Control Point NCE identifier = 80
             (0x0014) Route Setup NCE identifier = 80
```

B.1.2 CP-CP Sessions with NNP61A

The XIDs have been exchanged on LINK0001. The next thing that happens is the RTP connection setup exchange for the CP-CP session pipe (not shown). Then the CP-CP sessions start. We have shown only one of the two sessions, the CONWINNER session of CM5HPRNN.

```
Line:     78 Send MU
Time stamp: 08:26:55.67
DLC type: HPR
TH: FID5, Exp, OIS, SNF=0x0003, S SA2586f52ddc291f93
RH: RQ, SC, FI, OIC, RQD1
BIND rq
 (0x0001) Type = Negotiable
 (0x0002) FM profile = 19
 (0x0003) TS profile = 7
          FM usage - primary:
 (0x0004)    Chaining use = Multiple-RU chains allowed
             Request control mode = Immediate request mode
             Chain response protocol = Definite or exception response
             Two-phase commit = Not supported
             Compression = Will not be used
             Send end bracket = Will not send
          FM usage - secondary:
 (0x0005)    Chaining use = Multiple-RU chains allowed
             Request control mode = Immediate request mode
             Chain response protocol = Definite or exception response
             Two-phase commit = Not supported
             Compression = Will not be used
             Send end bracket = Will not send
          FM Usage - common:
 (0x0006)    Whole BIUs required = No
             FM header usage = Allowed
             Brackets are used = Yes
             Bracket reset state = INB
             Bracket termination rule = Conditional
             Alternate code set allowed = No
             BIND queueing allowed = No
 (0x0007)    Normal-flow send/receive mode = Full-duplex
             Recovery responsibility = Symmetric
             Contention winner = Primary
             Alternate code set = ASCII-7
             Control vectors included = Yes
             Half-duplex flip-flop primary reset state = Send
          TS usage:
```

```
(0x0008)    Secondary to primary pacing stages = One
            Secondary send window size = 1
(0x0009)    Adaptive pacing = Supported
            Secondary receive window size = 0
(0x000A)    Secondary maximum send RU size = 512
(0x000B)    Primary maximum send RU size = 512
(0x000C)    Primary to secondary pacing stages = One
            Primary send window size = 0
(0x000D)    Primary receive window size = 1
         PS profile:
(0x000E)    LU type = 6
(0x000F)    LU-6 level = 2
(0x0016)    Extended security sense codes = Supported
(0x0017)    Conversation-level security = Accepted
            Session security level = Enhanced
            Password substitution = Supported
            Already-verified indicator = Not accepted
(0x0018)    Synchronization level supported = Confirm
            Session reinitiation responsibility = Operator controlled
            Parallel sessions supported = Yes
            CNOS supported = Yes
(0x0019)    Limited resource = No
            Length-checked compression options = No Compression
         Cryptography options:
(0x001A)    Private cryptography support = No
            Session-level cryptography support = No
(0x001C) Primary LU name = USIBMRA.CM5HPRNN
         User Data:
            Structured user data:
(0x002F)       Mode name = CPSVCMG
(0x0038)       Session instance identifier = 0x01f586252cf58625
(0x0042)       Network-qualified PLU name = USIBMRA.CM5HPRNN
(0x0054)       Random data = 0x003a118a205da7c964
(0x0060) Secondary LU name = USIBMRA.NNP61A
(0x006E) Fully qualified PCID control vector:
          (0x0002) PCID = 0xc96f2a091256a273
          (0x000B) Network qualified CP name = USIBMRA.CM5HPRNN
(0x0089) COS/TPF control vector:
          (0x0002) Transmission priority = Network
          (0x0004) COS name = CPSVCMG

Line:     91 Recv MU
Time stamp: 08:26:55.79
DLC type: HPR
TH: FID5, Exp, OIS, SNF=0x0003, R SA2586f52ddc291f93
RH: +RSP, SC, FI, RQD1
BIND +rsp
(0x0001) Type = Negotiable
(0x0002) FM profile = 19
(0x0003) TS profile = 7
         FM usage - primary:
(0x0004)    Chaining use = Multiple-RU chains allowed
            Request control mode = Immediate request mode
            Chain response protocol = Definite or exception response
            Two-phase commit = Not supported
            Compression = Will not be used
            Send end bracket = Will not send
         FM usage - secondary:
(0x0005)    Chaining use = Multiple-RU chains allowed
```

```
            Request control mode = Immediate request mode
            Chain response protocol = Definite or exception response
            Two-phase commit = Not supported
            Compression = Will not be used
            Send end bracket = Will not send
         FM Usage - common:
(0x0006)    Whole BIUs required = No
            FM header usage = Allowed
            Brackets are used = Yes
            Bracket reset state = INB
            Bracket termination rule = Conditional
            Alternate code set allowed = No
            BIND queueing allowed = No
(0x0007)    Normal-flow send/receive mode = Full-duplex
            Recovery responsibility = Symmetric
            Contention winner = Primary
            Alternate code set = ASCII-7
            Control vectors included = Yes
            Half-duplex flip-flop primary reset state = Send
         TS usage:
(0x0008)    Secondary to primary pacing stages = One
            Secondary send window size = 0
(0x0009)    Adaptive pacing = Supported
            Secondary receive window size = 0
(0x000A)    Secondary maximum send RU size = 512
(0x000B)    Primary maximum send RU size = 512
(0x000C)    Primary to secondary pacing stages = One
            Primary send window size = 0
(0x000D)    Primary receive window size = 0
         PS profile:
(0x000E)    LU type = 6
(0x000F)    LU-6 level = 2
(0x0016)    Extended security sense codes = Supported
(0x0017)    Conversation-level security = Accepted
            Session security level = Enhanced
            Password substitution = Supported
            Already-verified indicator = Not accepted
(0x0018)    Synchronization level supported = Confirm
            Session reinitiation responsibility = Operator controlled
            Parallel sessions supported = Yes
            CNOS supported = Yes
(0x0019)    Limited resource = No
            Length-checked compression options = No Compression
         Cryptography options:
(0x001A)    Private cryptography support = No
            Session-level cryptography support = No
         User Data:
            Structured user data:
(0x001F)       Mode name = CPSVCMG
(0x0028)       Session instance identifier = 0x02
(0x002B)       Network-qualified SLU name = USIBMRA.NNP61A
(0x003B)       Random data = 0x001404a75f914bd50a
(0x0047) Fully qualified PCID control vector:
          (0x0002) PCID = 0xc96f2a091256a273
          (0x000B) Network qualified CP name = USIBMRA.CM5HPRNN
(0x0062) Session Address control vector
          (0x0002) Session address assignor = Sender
                   Session address = 0x05a23bdd00201015
```

B.1.3 XID Exchange with NNP41A

When the connection to NNP41A is activated, the same process happens that occurred on the link to NNP61A. Here we show only the last two flows on the XID exchange.

```
Line:   2128   Recv XID
Time stamp: 08:27:19.23
DLC type: IBMTRNET
Adapter number: 00
Destination address: 40043746217608
ALS ID: 501227936DF58625
XID:
|| (0x0000) Format = 3
||          Node type = 2
|| (0x0001) Total length = 111
|| (0x0002) Node ID = 0x05d0671d
|| (0x0008) Init-self = Can receive
||          Stand-alone BIND = Can receive
||          BIND segment generation = Can generate
||          BIND segment receipt = Can receive
||          FID type = FID 2
|| (0x0009) ACTPU suppression = ACTPU not requested
||          APPN network node = Yes
||          CP-CP sessions requested = Yes
||          CP-CP sessions supported = Yes
||          Exchange state = Negotiation proceeding
||          Secondary initiated non-activation exchange = Supported
||          CP name change = Supported
|| (0x000A) Adaptive BIND pacing sender = Supported
||          Adaptive BIND pacing receiver = Supported
||          Quiesce TG requested = No
||          PU Capabilities control vector = Not supported
||          Sender is an APPN peripheral border node = No
||          Adaptive BIND pacing = Supported for all LUs, unless overridden
|| (0x000F) Parallel TGs = Supported
||          DLUR XID sender prefers ACTPU over CP-SVR pipe = Yes
||          DLUS-served LU registration = Supported
||          Downstream Branch Extender = Not supported
||          Upstream Branch Extender = Not supported
|| (0x0010) TG number = 33
|| (0x0011) DLC type = SDLC
|| (0x0012) DLC data length = 11
|| (0x0013) ABM = Supported
||          Link station role = Primary
||          Short hold mode status = Not reconnection
||          Short hold mode capability = Not supported
||          Transmit/receive capability = Secondary
|| (0x0014) ABM nonactivation XID exchange initiator = No
|| (0x0015) Maximum receive BTU length = 2058
|| (0x0017) SDLC CR profile = SNA link profile
|| (0x0018) SIM/RIM options = Not supported
|| (0x001B) Maximum I-frames before ACK = 6
|| (0x001D) Network name control vector:
||           (0x0002) Network name type = CP
||           (0x0003) Name = USIBMRA.NNP41A
|| (0x002E) Product set ID control vector:
||             Hex dump:
||                10170016 11011300 11f3f7f4 f6f9f0f0   <................> <....
||                f5f7f0f0 f9f1f8f0 f5                  <.........       > <5700
```

CV X′61′, subfields X′80′ and X′81′

```
|| (0x0047) HPR Capabilities control vector
||           (0x0002) Error recovery = Available if requested by partner
||                    RTP Tower = Supported
||                    Control Flows over RTP Tower = Supported
||           (0x0005) ANR label = a400304481
|| (0x0051)    Control Flows over RTP Tower subfield
||              (0x0002) Maximum send packet size = 0x080a
||              (0x0004) Path switch time = 60000
||              (0x0008) Sender is mobile = No
||                       Multilink transmission groups = Not supported
||              (0x0009) Control Point NCE instance identifier = 0x00001a5e
||              (0x000D) Route Setup NCE instance identifier = 0x00001a5e
||              (0x0012) Control Point NCE identifier = d0201025
||              (0x0017) Route Setup NCE identifier = d0201025
|| (0x006C)    Token Ring/Ethernet 802.2 LLC subfield
||              (0x0002) LLC SAP = 08

Line:   2137   Send XID
Time stamp: 08:27:19.23
DLC type: IBMTRNET
Adapter number: 00
Destination address: 40043746217608
ALS ID: 501227936DF58625
XID:
|| (0x0000) Format = 3
||          Node type = 2
|| (0x0001) Total length = 132
|| (0x0002) Node ID = 0x05d00000
|| (0x0008) Init-self = Cannot receive
||          Stand-alone BIND = Can receive
||          BIND segment generation = Can generate
||          BIND segment receipt = Can receive
||          FID type = FID 2
|| (0x0009) ACTPU suppression = ACTPU not requested
||          APPN network node = Yes
||          CP-CP sessions requested = Yes
||          CP-CP sessions supported = Yes
||          Exchange state = Negotiation proceeding
||          Secondary initiated non-activation exchange = Supported
||          CP name change = Supported
|| (0x000A) Adaptive BIND pacing sender = Supported
||          Adaptive BIND pacing receiver = Supported
||          Quiesce TG requested = No
||          PU Capabilities control vector = Not supported
||          Sender is an APPN peripheral border node = No
||          Adaptive BIND pacing = Supported for all LUs, unless overridden
|| (0x000F) Parallel TGs = Supported
||          DLUR XID sender prefers ACTPU over CP-SVR pipe = No
||          DLUS-served LU registration = Supported
||          Downstream Branch Extender = Not supported
||          Upstream Branch Extender = Not supported
|| (0x0010) TG number = 33 || (0x0011) DLC type = SDLC
|| (0x0012) DLC data length = 11
|| (0x0013) ABM = Supported
||          Link station role = Secondary
```

```
          Short hold mode status = Not reconnection
          Short hold mode capability = Not supported
          Transmit/receive capability = Secondary
(0x0014) ABM nonactivation XID exchange initiator = No
(0x0015) Maximum receive BTU length = 2058
(0x0017) SDLC CR profile = SNA link profile
(0x0018) SIM/RIM options = Not supported
(0x001B) Maximum I-frames before ACK = 4
(0x001D) Network name control vector:
          (0x0002) Network name type = CP
          (0x0003) Name = USIBMRA.CM5HPRNN
(0x0030) Network name control vector:
          (0x0002) Network name type = LS
          (0x0003) Name = LINK0002
(0x003B) Product set ID control vector:
            Hex dump:
               10280011 11040e02 f5f6f2f2 f8f7f8f0   <.(..............> <....
               f0f4f1f0 16110313 0011f0f0 f0f00000   <................> <0410
               00000000 00000000 0000                <..........      > <....
```

CV X′61′, subfields X′80′ and X′81′

```
(0x0065) HPR Capabilities control vector
          (0x0002) Error recovery = None
                   RTP Tower = Supported
                   Control Flows over RTP Tower = Supported
          (0x0005) ANR label = 801e
(0x006C)    Token Ring/Ethernet 802.2 LLC subfield
             (0x0002) LLC SAP = 08
(0x006F)    Control Flows over RTP Tower subfield
             (0x0002) Maximum send packet size = 0x080a
             (0x0004) Path switch time = 60000
             (0x0008) Sender is mobile = No
                      Multilink transmission groups = Not supported
             (0x0009) Control Point NCE instance identifier = 0x34e1a64d
             (0x000D) Route Setup NCE instance identifier = 0x34e1a64d
             (0x0012) Control Point NCE identifier = 80
             (0x0014) Route Setup NCE identifier = 80
```

B.1.4 CP-CP Sessions with NNP41A

Again, NNP41A and CM5HPRNN set up the CP-CP RTP pipe. We show here the activation of CM5HPRNN′s CONWINNER session with NNP41A.

```
Line:   2159 Send MU
Time stamp: 08:27:19.36
DLC type: HPR
TH: FID5, Exp, OIS, SNF=0x0004, S SA2586f575ac371f93
RH: RQ, SC, FI, OIC, RQD1
BIND rq
 (0x0001) Type = Negotiable
 (0x0002) FM profile = 19
 (0x0003) TS profile = 7
          FM usage - primary:
 (0x0004)    Chaining use = Multiple-RU chains allowed
             Request control mode = Immediate request mode
             Chain response protocol = Definite or exception response
             Two-phase commit = Not supported
```

```
            Compression = Will not be used
            Send end bracket = Will not send
         FM usage - secondary:
(0x0005)    Chaining use = Multiple-RU chains allowed
            Request control mode = Immediate request mode
            Chain response protocol = Definite or exception response
            Two-phase commit = Not supported
            Compression = Will not be used
            Send end bracket = Will not send
         FM Usage - common:
(0x0006)    Whole BIUs required = No
            FM header usage = Allowed
            Brackets are used = Yes
            Bracket reset state = INB
            Bracket termination rule = Conditional
            Alternate code set allowed = No
            BIND queueing allowed = No
(0x0007)    Normal-flow send/receive mode = Full-duplex
            Recovery responsibility = Symmetric
            Contention winner = Primary
            Alternate code set = ASCII-7
            Control vectors included = Yes
            Half-duplex flip-flop primary reset state = Send
         TS usage:
(0x0008)    Secondary to primary pacing stages = One
            Secondary send window size = 1
(0x0009)    Adaptive pacing = Supported
            Secondary receive window size = 0
(0x000A)    Secondary maximum send RU size = 512
(0x000B)    Primary maximum send RU size = 512
(0x000C)    Primary to secondary pacing stages = One
            Primary send window size = 0
(0x000D)    Primary receive window size = 1
         PS profile:
(0x000E)    LU type = 6
(0x000F)    LU-6 level = 2
(0x0016)    Extended security sense codes = Supported
(0x0017)    Conversation-level security = Accepted
            Session security level = Enhanced
            Password substitution = Supported
            Already-verified indicator = Not accepted
(0x0018)    Synchronization level supported = Confirm
            Session reinitiation responsibility = Operator controlled
            Parallel sessions supported = Yes
            CNOS supported = Yes
(0x0019)    Limited resource = No
            Length-checked compression options = No Compression
         Cryptography options:
(0x001A)    Private cryptography support = No
            Session-level cryptography support = No
(0x001C) Primary LU name = USIBMRA.CM5HPRNN
         User Data:
            Structured user data:
(0x002F)       Mode name = CPSVCMG
(0x0038)       Session instance identifier = 0x01f5862574f58625
(0x0042)       Network-qualified PLU name = USIBMRA.CM5HPRNN
(0x0054)       Random data = 0x00e4d2076a41ba4f90
(0x0060) Secondary LU name = USIBMRA.NNP41A
(0x006E) Fully qualified PCID control vector:
```

```
||           (0x0002) PCID = 0xc96f2a091356a273
||           (0x000B) Network qualified CP name = USIBMRA.CM5HPRNN
|| (0x0089) COS/TPF control vector:
||           (0x0002) Transmission priority = Network
||           (0x0004) COS name = CPSVCMG

Line:   2172 Recv MU
Time stamp: 08:27:19.49
DLC type: HPR
||TH: FID5, Exp, OIS, SNF=0x0004, R SA2586f575ac371f93
||RH: +RSP, SC, FI, RQD1
||BIND +rsp
|| (0x0001) Type = Negotiable
|| (0x0002) FM profile = 19
|| (0x0003) TS profile = 7
||          FM usage - primary:
|| (0x0004)    Chaining use = Multiple-RU chains allowed
||             Request control mode = Immediate request mode
||             Chain response protocol = Definite or exception response
||             Two-phase commit = Not supported
||             Compression = Will not be used
||             Send end bracket = Will not send
||          FM usage - secondary:
|| (0x0005)    Chaining use = Multiple-RU chains allowed
||             Request control mode = Immediate request mode
||             Chain response protocol = Definite or exception response
||             Two-phase commit = Not supported
||             Compression = Will not be used
||             Send end bracket = Will not send
||          FM Usage - common:
|| (0x0006)    Whole BIUs required = No
||             FM header usage = Allowed
||             Brackets are used = Yes
||             Bracket reset state = INB
||             Bracket termination rule = Conditional
||             Alternate code set allowed = No
||             BIND queueing allowed = No
|| (0x0007)    Normal-flow send/receive mode = Full-duplex
||             Recovery responsibility = Symmetric
||             Contention winner = Primary
||             Alternate code set = ASCII-7
||             Control vectors included = Yes
||             Half-duplex flip-flop primary reset state = Send
||          TS usage:
|| (0x0008)    Secondary to primary pacing stages = One
||             Secondary send window size = 0
|| (0x0009)    Adaptive pacing = Supported
||             Secondary receive window size = 0
|| (0x000A)    Secondary maximum send RU size = 512
|| (0x000B)    Primary maximum send RU size = 512
|| (0x000C)    Primary to secondary pacing stages = One
||             Primary send window size = 0
|| (0x000D)    Primary receive window size = 0
||          PS profile:
|| (0x000E)    LU type = 6
|| (0x000F)    LU-6 level = 2
|| (0x0016)    Extended security sense codes = Supported
|| (0x0017)    Conversation-level security = Accepted
||             Session security level = Enhanced
```

```
||             Password substitution = Supported
||             Already-verified indicator = Not accepted
|| (0x0018)    Synchronization level supported = Confirm
||             Session reinitiation responsibility = Operator controlled
||             Parallel sessions supported = Yes
||             CNOS supported = Yes
|| (0x0019)    Limited resource = No
||             Length-checked compression options = No Compression
||          Cryptography options:
|| (0x001A)    Private cryptography support = No
||             Session-level cryptography support = No
||          User Data:
||             Structured user data:
|| (0x001F)       Mode name = CPSVCMG
|| (0x0028)       Session instance identifier = 0x02
|| (0x002B)       Network-qualified SLU name = USIBMRA.NNP41A
|| (0x003B)       Random data = 0x00543d5b0ab92ce266
|| (0x0047) Fully qualified PCID control vector:
||           (0x0002) PCID = 0xc96f2a091356a273
||           (0x000B) Network qualified CP name = USIBMRA.CM5HPRNN
|| (0x0062) Session Address control vector
||           (0x0002) Session address assignor = Sender
||                    Session address = 0x058bb89800201015
```

B.1.5 Route Setup

The next thing that happens is the activation of the DLUR/S pipe, because the DLUR/S logical link is activated at startup time by CM5HPRNN. We have not shown the Locate flows to find where the DLUS is, nor have we shown the establishment of the Route Setup HPR pipe when the BIND is ready to flow. Here we depict the Route Setup request that flows over the long-lived pipe to NNP61A on its way to RAA (the DLUS).

```
Line:   3288 Send MU
Time stamp: 08:27:32.14
DLC type: HPR
TCID: 0000007A
Transmission priority: Network
||TH: FID2, Exp, OIS, LFSID=0x00000, SNF=0x0000
||RH: RQ, NC, FI, OIC, RQN
||Route Setup
||   Type = Request
||   Route Setup triggered by path switch = No
||   Destination hop index = 2  1
```

1 The destination hop index contains the index (integer) into the RSCV (CV 2B) for the destination node.

```
||   Locate search is required = No
||   Destination is mobile = No
||   NCE is used for all LUs/BFs = No
```

2 The Path Switch time indicates the maximum time (in milliseconds) that the destination requires for a path switch.

```
Path switch time (milliseconds) = 0 2
Network name control vector:
 (0x0002) Network name type = LU
 (0x0003) Name = USIBMRA.RAA
Route selection control vector:
 (0x0002) Maximum hop count = 2
 (0x0003) Current hop count = 1
   TG descriptor control vector
      TG Identifier TG Descriptor subfield
       (0x0002) TG number = 26
       (0x0004) Partner name =USIBMRA.NNP61A
       (0x0012) Link connection network = No
                Additional configuration information = No
                HPR = Supported
                TG type = Boundary Function based or APPN
                Intersubnet link = No
                Extended border node = Not supported
                RTP Tower = Supported
   TG descriptor control vector
      TG Identifier TG Descriptor subfield
       (0x0002) TG number = 21
       (0x0004) Partner name = RAA
       (0x0007) Link connection network = No
                Additional configuration information = No
                HPR = Supported
                TG type = Boundary Function based or APPN
                Intersubnet link = No
                Extended border node = Not supported
                RTP Tower = Supported
Fully qualified PCID control vector:
 (0x0002) PCID = 0xe887a74907bc569a
 (0x000B) Network qualified CP name = USIBMRA.CM5HPRNN
COS/TPF control vector:
 (0x0002) Transmission priority = Network
 (0x0004) COS name = SNASVCMG
```

CV X′80′

```
Route Information control vector
 (0x0002) Route direction = Forward
          REFIFOing required = No
 (0x0004) Maximum packet size = 2058
 (0x0008) Accumulated transmission time (microseconds) = 75
 (0x000C) Minimum link capacity (Kbits/second) = 16000
 (0x0010) Limited resource liveness timer (seconds) = 0
```

CV X′67′, contains the description of the path using a series of ANR label entries.

```
             ANR Path control vector
    (0x0003) ANR label represents a subarea network route = No
    (0x0004) ANR label = 801c
```

Route Setup reply

```
Line:   3300 Recv MU
Time stamp: 08:27:32.33
DLC type: HPR
TCID: 0000007A
Transmission priority: Network
||TH: FID2, Exp, OIS, LFSID=0x00000, SNF=0x0000
||RH: RQ, NC, FI, OIC, RQN
||Route setup
||   Type = Reply
||   Destination hop index = 2
||   Locate search is required = Yes
||   Destination is mobile = No
||   NCE is used for all LUs/BFs = No
||   Path switch time (milliseconds) = 480000
||   Network name control vector:
||    (0x0002) Network name type = CP
||    (0x0003) Name = USIBMRA.RAA
||   NCE Instance Identifier control vector
||    (0x0002) NCE instance identifer = 0xaff10bb2
||   NCE Identifier control vector
||    (0x0002) NCE identifier = d000000000000000
||   Route selection control vector:
||    (0x0002) Maximum hop count = 2
||    (0x0003) Current hop count = 0
||      TG descriptor control vector
||         TG Identifier TG Descriptor subfield
||          (0x0002) TG number = 21
||          (0x0004) Partner name = USIBMRA.NNP61A
||          (0x0012) Link connection network = No
||                   Additional configuration information = No
||                   HPR = Supported
||                   TG type = Boundary Function based or APPN
||                   Intersubnet link = No
||                   Extended border node = Supported
||                   RTP Tower = Supported
||          (0x0013) Subarea number = 0x8000000a  SA: 10
||      TG descriptor control vector
||         TG Identifier TG Descriptor subfield
||          (0x0002) TG number = 26
||          (0x0004) Partner name = CM5HPRNN
||          (0x000C) Link connection network = No
||                   Additional configuration information = No
||                   HPR = Supported
||                   TG type = Boundary Function based or APPN
||                   Intersubnet link = No
||                   Extended border node = Not supported
||                   RTP Tower = Supported
||   COS/TPF control vector:
||    (0x0002) Transmission priority = Low
||    (0x0004) COS name = SNASVCMG
||   Fully qualified PCID control vector:
||    (0x0002) PCID = 0xe887a74907bc569a
||    (0x000B) Network qualified CP name = USIBMRA.CM5HPRNN
||   Route Information control vector
||    (0x0002) Route direction = Forward
```

```
||             REFIFOing required = No
||    (0x0004) Maximum packet size = 2058
||    (0x0008) Accumulated transmission time (microseconds) = 83
||    (0x000C) Minimum link capacity (Kbits/second) = 16000
||    (0x0010) Limited resource liveness timer (seconds) = 0
||      ANR Path control vector
||       (0x0003) ANR label represents a subarea network route = No
||       (0x0004) ANR label = 801c
||       (0x0007) ANR label represents a subarea network route = No
||       (0x0008) ANR label = a400000001
||   Route Information control vector
||    (0x0002) Route direction = Reverse
||             REFIFOing required = No
||    (0x0004) Maximum packet size = 2058
||    (0x0008) Accumulated transmission time (microseconds) = 108
||    (0x000C) Minimum link capacity (Kbits/second) = 15974
||    (0x0010) Limited resource liveness timer (seconds) = 0
||      ANR Path control vector
||       (0x0003) ANR label represents a subarea network route = No
||       (0x0004) ANR label = 803000c300000000
||       (0x000D) ANR label represents a subarea network route = No
||       (0x000E) ANR label = a320304c82
```

The path for this DLUR/S session is the following:

```
┌──────────┐          ┌───────┐       ┌─────┐
│ CM5HPRNN ├───TG26───┤NNP61A ├─TG21──┤RAA  │
└──────────┘          └───────┘       └─────┘
```

Figure 245. DLUR/S Path

B.1.6 BIND for DLUR/S Session

Once the path for the RTP connection has been determined, that connection is initialized with connection setup flows. Next, the BIND can flow to start the session. Here we show the BIND for CM5HPRNN's CONWINNER DLUR/S session.

```
Line:   3316 Send MU
Time stamp: 08:27:32.33
DLC type: HPR
||TH: FID5, Exp, OIS, SNF=0x0005, S SA2586f5a87c3f1f93
||RH: RQ, SC, FI, OIC, RQD1
||BIND rq
|| (0x0001) Type = Negotiable
|| (0x0002) FM profile = 19
|| (0x0003) TS profile = 7
||          FM usage - primary:
|| (0x0004)    Chaining use = Multiple-RU chains allowed
||             Request control mode = Immediate request mode
||             Chain response protocol = Definite or exception response
||             Two-phase commit = Not supported
||             Compression = Will not be used
||             Send end bracket = Will not send
||          FM usage - secondary:
|| (0x0005)    Chaining use = Multiple-RU chains allowed
```

```
            Request control mode = Immediate request mode
            Chain response protocol = Definite or exception response
            Two-phase commit = Not supported
            Compression = Will not be used
            Send end bracket = Will not send
         FM Usage - common:
(0x0006)    Whole BIUs required = No
            FM header usage = Allowed
            Brackets are used = Yes
            Bracket reset state = INB
            Bracket termination rule = Conditional
            Alternate code set allowed = No
            BIND queueing allowed = No
(0x0007)    Normal-flow send/receive mode = Full-duplex
            Recovery responsibility = Symmetric
            Contention winner = Primary
            Alternate code set = ASCII-7
            Control vectors included = Yes
            Half-duplex flip-flop primary reset state = Send
         TS usage:
(0x0008)    Secondary to primary pacing stages = One
            Secondary send window size = 1
(0x0009)    Adaptive pacing = Supported
            Secondary receive window size = 0
(0x000A)    Secondary maximum send RU size = 512
(0x000B)    Primary maximum send RU size = 512
(0x000C)    Primary to secondary pacing stages = One
            Primary send window size = 0
(0x000D)    Primary receive window size = 1
         PS profile:
(0x000E)    LU type = 6
(0x000F)    LU-6 level = 2
(0x0016)    Extended security sense codes = Supported
(0x0017)    Conversation-level security = Accepted
            Session security level = Enhanced
            Password substitution = Supported
            Already-verified indicator = Not accepted
(0x0018)    Synchronization level supported = Confirm
            Session reinitiation responsibility = Operator controlled
            Parallel sessions supported = Yes
            CNOS supported = Yes
(0x0019)    Limited resource = No
            Length-checked compression options = No Compression
         Cryptography options:
(0x001A)    Private cryptography support = No
            Session-level cryptography support = No
(0x001C) Primary LU name = USIBMRA.CM5HPRNN
         User Data:
            Structured user data:
(0x002F)       Mode name = CPSVRMGR
(0x0039)       Session instance identifier = 0x01f58625a7f58625
(0x0043)       Network-qualified PLU name = USIBMRA.CM5HPRNN
(0x0055)       Random data = 0x008c1a706f72678a21
(0x0061) Secondary LU name = USIBMRA.RAA
(0x006C) Fully qualified PCID control vector:
          (0x0002) PCID = 0xc96f2a091856a273
          (0x000B) Network qualified CP name = USIBMRA.CM5HPRNN
(0x0087) COS/TPF control vector:
          (0x0002) Transmission priority = Network
```

```
          (0x0004) COS name = SNASVCMG
 (0x0093) Route selection control vector:
          (0x0002) Maximum hop count = 2
          (0x0003) Current hop count = 2
 (0x0097)   TG descriptor control vector
 (0x0099)      TG Identifier TG Descriptor subfield
                (0x0002) TG number = 26
                (0x0004) Partner name = USIBMRA.NNP61A
                (0x0012) Link connection network = No
                         Additional configuration information = No
                         HPR = Supported
                         TG type = Boundary Function based or APPN
                         Intersubnet link = No
                         Extended border node = Not supported
                         RTP Tower = Supported
 (0x00AC)   TG descriptor control vector
 (0x00AE)      TG Identifier TG Descriptor subfield
                (0x0002) TG number = 21
                (0x0004) Partner name = RAA
                (0x0007) Link connection network = No
                         Additional configuration information = No
                         HPR = Supported
                         TG type = Boundary Function based or APPN
                         Intersubnet link = No
                         Extended border node = Not supported
                         RTP Tower = Supported

Line:   3331 Recv MU
Time stamp: 08:27:32.42
DLC type: HPR
TH: FID5, Exp, OIS, SNF=0x0005, R SA2586f5a87c3f1f93
RH: +RSP, SC, FI, RQD1
BIND +rsp
 (0x0001) Type = Negotiable
 (0x0002) FM profile = 19
 (0x0003) TS profile = 7
          FM usage - primary:
 (0x0004)    Chaining use = Multiple-RU chains allowed
             Request control mode = Immediate request mode
             Chain response protocol = Definite or exception response
             Two-phase commit = Not supported
             Compression = Will not be used
             Send end bracket = Will not send
          FM usage - secondary:
 (0x0005)    Chaining use = Multiple-RU chains allowed
             Request control mode = Immediate request mode
             Chain response protocol = Definite or exception response
             Two-phase commit = Not supported
             Compression = Will not be used
             Send end bracket = Will not send
          FM Usage - common:
 (0x0006)    Whole BIUs required = No
             FM header usage = Allowed
             Brackets are used = Yes
             Bracket reset state = INB
             Bracket termination rule = Conditional
             Alternate code set allowed = No
             BIND queueing allowed = No
 (0x0007)    Normal-flow send/receive mode = Half-duplex flip-flop
```

```
               Recovery responsibility = Symmetric
               Contention winner = Primary
               Alternate code set = ASCII-7
               Control vectors included = Yes
               Half-duplex flip-flop primary reset state = Send
            TS usage:
(0x0008)       Secondary to primary pacing stages = One
               Secondary send window size = 0
(0x0009)       Adaptive pacing = Supported
               Secondary receive window size = 0
(0x000A)       Secondary maximum send RU size = 512
(0x000B)       Primary maximum send RU size = 512
(0x000C)       Primary to secondary pacing stages = One
               Primary send window size = 0
(0x000D)       Primary receive window size = 0
            PS profile:
(0x000E)       LU type = 6
(0x000F)       LU-6 level = 2
(0x0016)       Extended security sense codes = Not supported
(0x0017)       Conversation-level security = Not accepted
               Session security level = Basic
               Password substitution = Not supported
               Already-verified indicator = Not accepted
(0x0018)       Synchronization level supported = Confirm
               Session reinitiation responsibility = Operator controlled
               Parallel sessions supported = Yes
               CNOS supported = Yes
(0x0019)       Limited resource = No
               Length-checked compression options = No Compression
            Cryptography options:
(0x001A)       Private cryptography support = No
               Session-level cryptography support = No
            User Data:
               Structured user data:
(0x001F)          Mode name = CPSVRMGR
(0x0029)          Session instance identifier = 0x026f2a091856a273
(0x0033)          Network-qualified SLU name = USIBMRA.RAA
(0x0041) Fully qualified PCID control vector:
             (0x0002) PCID = 0xc96f2a091856a273
             (0x000B) Network qualified CP name = USIBMRA.CM5HPRNN
(0x005C) Session Address control vector
             (0x0002) Session address assignor = Sender
                      Session address = 0x00000000f2000004
(0x0066) Route selection control vector:
             (0x0002) Maximum hop count = 2
             (0x0003) Current hop count = 2
(0x006A)       TG descriptor control vector
(0x006C)          TG Identifier TG Descriptor subfield
                    (0x0002) TG number = 26
                    (0x0004) Partner name = USIBMRA.NNP61A
                    (0x0012) Link connection network = No
                             Additional configuration information = No
                             HPR = Supported
                             TG type = Boundary Function based or APPN
                             Intersubnet link = No
                             Extended border node = Not supported
                             RTP Tower = Supported
(0x007F)       TG descriptor control vector
(0x0081)          TG Identifier TG Descriptor subfield
```

```
                    (0x0002) TG number = 21
                    (0x0004) Partner name = RAA
                    (0x0007) Link connection network = No
                             Additional configuration information = No
                             HPR = Supported
                             TG type = Boundary Function based or APPN
                             Intersubnet link = No
                             Extended border node = Not supported
                             RTP Tower = Supported
```

B.1.7 Dependent LU Activation

When the DLUR/S pipe is ready, the activation of the dependent resources on CM5HPRNN can commence. All the flows on the SSCP-PU and SSCP-LU sessions comprise FID-2 PIUs (the whole PIUs) encapsulated in GDS X′1500′ variables on the LU 6.2 flows. The transaction program that is attached to process these is X′22F0F0F6′.

Here we show the first flow on the DLUR/S pipe, which starts with a REQACTPU request from the DLUR.

```
Line:   3343 Send MU
Time stamp: 08:27:32.42
DLC type: HPR
TH: FID5, OIS, SNF=0x0001, R SA00000000f2000004
RH: RQ, FMD, FI, OIC, RQE1, PI, CEBI
FMH-5
   Command code = Attach
   User ID already verified = No
   Password is substituted = No
   PIP present = No
   Conversation type = Basic
   Synchronization level = None
   Transaction program name = ?006 (APPN Receive_Encap_Msg)
   Logical unit of work identifier:
      LU name = USIBMRA.CM5HPRNN
      Instance number = 0xafff44a61c0d
      Sequence number = 0x0001
   Conversation correlator = 0x58451f93a4f58625
FID2 Encapsulation Variable
   TH: FID2, OIS, LFSID=0x00000, SNF=0x0000
   RH: RQ, FMD, FI, OIC, RQD1
   REQACTPU rq
      Format = Internal PU
   DLUR/S Capabilities control vector
    (0x0002) Dependent LU support level = 0x01
    (0x0003) Node type = Network Node
    (0x0004) DLUR/S node type = DLUR
    (0x0005) Flow reduction sequence number = 0x00000000
    (0x0009) RECEIVE_TDU_TP = Not supported
   Network name control vector:
    (0x0002) Network name type = PU
    (0x0003) Name = MPU00001
   XID Image FID2 Encapsulation control vector
    (0x0002) XID I-field image:
             Hex dump:
                020605da a61a                     <......          > <....
   TG descriptor control vector
      DLC Signaling Type TG Descriptor subfield
```

```
||       (0x0002) DLC type = INTPU
||      DLC Signaling Information TG Descriptor subfield
||       (0x0002) Block number = 005d
||                ID number = 000aa61a
||   Fully qualified PCID control vector:
||    (0x0002) PCID = 0xc96f2a091456a273
||    (0x000B) Network qualified CP name = USIBMRA.CM5HPRNN


Line:   3386 Recv MU
Time stamp: 08:27:32.64
DLC type: HPR
||TH: FID5, OIS, SNF=0x0001, R SA2586f5adc0481f93
||RH: RQ, FMD, FI, OIC, RQE1, PI, CEBI
||FMH-5
||   Command code = Attach
||   User ID already verified = No
||   Password is substituted = No
||   PIP present = No
||   Conversation type = Basic
||   Synchronization level = None
||   Transaction program name = ?006 (APPN Receive_Encap_Msg)
||FID2 Encapsulation Variable
||   TH: FID2, OIS, LFSID=0x00000, SNF=0x0000
||   RH: +RSP, FMD, FI, RQD1
||   REQACTPU +rsp
||   Fully qualified PCID control vector:
||    (0x0002) PCID = 0xc96f2a091456a273
||    (0x000B) Network qualified CP name = USIBMRA.CM5HPRNN
||   DLUR/S Capabilities control vector
||    (0x0002) Dependent LU support level = 0x01
||    (0x0003) Node type = Network Node
||    (0x0004) DLUR/S node type = DLUS
||    (0x0005) Flow reduction sequence number = 0x00000000
||    (0x0009) RECEIVE_TDU_TP = Supported
```

Once the REQACTPU, and the ACTPU (not shown) have completed, ACTLU can proceed.

```
Line:   3428 Recv MU
Time stamp: 08:27:32.70
DLC type: HPR
||TH: FID5, OIS, SNF=0x0003, R SA2586f5adc0481f93
||RH: RQ, FMD, FI, OIC, RQE1, PI, BB, CEBI
||FMH-5
||   Command code = Attach
||   User ID already verified = No
||   Password is substituted = No
||   PIP present = No
||   Conversation type = Basic
||   Synchronization level = None
||   Transaction program name = ?006 (APPN Receive_Encap_Msg)
||FID2 Encapsulation Variable
||   TH: FID2, Exp, OIS, LFSID=0x00002, SNF=0x19a9
||   RH: RQ, SC, FI, OIC, RQD1
```

```
      ACTLU rq
         Enhanced address management = Supported
         LU address = Dynamic
         Activation type = ERP
         FM profile = 0
         TS profile = 1
         Network name control vector:
          (0x0002) Network name type = LU
          (0x0003) Name = USIBMRA.WAA61A02
      Assign LU Characteristics control vector
       (0x0002) Reserved session resources = 0x0100
       (0x0004) Route extension (REX) stage pacing = Adaptive allowed
                Pacing window size = 0x00
       (0x0005) Subarea stage pacing = Adaptive allowed
                Pacing window size = 0x07
       (0x0006) Maximum LU-LU sessions = Locally defined default value
      Fully qualified PCID control vector:
       (0x0002) PCID = 0xc96f2a091456a273
       (0x000B) Network qualified CP name = USIBMRA.CM5HPRNN

Line:   3441 Send MU
Time stamp: 08:27:32.71
DLC type: HPR
   TH: FID5, OIS, SNF=0x0003, R SA00000000f2000004
   RH: RQ, FMD, FI, OIC, RQE1, BB, CEBI
   FMH-5
      Command code = Attach
      User ID already verified = No
      Password is substituted = No
      PIP present = No
      Conversation type = Basic
      Synchronization level = None
      Transaction program name = ?006 (APPN Receive_Encap_Msg)
      Logical unit of work identifier:
         LU name = USIBMRA.CM5HPRNN
         Instance number = 0xafff44a61c0d
         Sequence number = 0x0001
      Conversation correlator = 0x58451f93b9f58625
   FID2 Encapsulation Variable
      TH: FID2, Exp, OIS, LFSID=0x00002, SNF=0x19a9
      RH: +RSP, SC, FI, RQD1
      ACTLU +rsp
         Activation type = Cold
         FM profile = 0
         TS profile = 1
         SSCP-LU session capabilities control vector:
          (0x0001) Maximum RU size = No limit
          (0x0002) Unsolicited character-coded requests = Not supported
                   Unsolicited field-formatted requests = Not supported
         LU-LU session services capabilities control vector:
          (0x0002) PLU capability = Inhibited
                   SLU capability = Disabled
```

CV X′0C′. Even if the ACTLU response is positive, the LU may be declared disabled until an application (in this case PComm) is ready to use it. When the application has been connected to the LU and is ready, a NOTIFY RU is sent on the SSCP-LU session indicating a change of status to Enabled.

```
||      (0x0003) LU-LU session limit = 256
||      (0x0005) LU-LU session count = 0
||      (0x0007) SESSST capability = Sent if SLU
||               XRF session activation CV supported on BIND = No
||               Peripheral node extended BIND receive support = Yes
||               Network-qualified name receive support = No
||               Subarea node extended BIND support = Yes
||   Fully qualified PCID control vector:
||    (0x0002) PCID = 0xc96f2a091456a273
||    (0x000B) Network qualified CP name = USIBMRA.CM5HPRNN
```

This process is repeated for each LU defined in CM5HPRNN.
We have not shown the other LUs in this trace.

B.1.8 Route Setup to RA39

Since RA39 has also been defined as being a DLU server, its own Route Setup flows for its DLUR/S pipe.

```
Line:   3756 Send MU
Time stamp: 08:27:41.07
DLC type: HPR
TCID: 0000007C
Transmission priority: Network
||TH: FID2, Exp, OIS, LFSID=0x00000, SNF=0x0000
||RH: RQ, NC, FI, OIC, RQN
||Route setup
||   Type = Request
||   Route Setup triggered by path switch = No
||   Destination hop index = 2
||   Locate search is required = No
||   Destination is mobile = No
||   NCE is used for all LUs/BFs = No
||   Path switch time (milliseconds) = 0
||   Network name control vector:
||    (0x0002) Network name type = LU
||    (0x0003) Name = USIBMRA.RA39
||   Route selection control vector:
||    (0x0002) Maximum hop count = 2
||    (0x0003) Current hop count = 1
||      TG descriptor control vector
||         TG Identifier TG Descriptor subfield
||          (0x0002) TG number = 33
||          (0x0004) Partner name = USIBMRA.NNP41A
||          (0x0012) Link connection network = No
||                   Additional configuration information = No
||                   HPR = Supported
||                   TG type = Boundary Function based or APPN
```

```
                    Intersubnet link = No
                    Extended border node = Not supported
                    RTP Tower = Supported
      TG descriptor control vector
         TG Identifier TG Descriptor subfield
          (0x0002) TG number = 21
          (0x0004) Partner name = RA39
          (0x0008) Link connection network = No
                   Additional configuration information = No
                   HPR = Supported
                   TG type = Boundary Function based or APPN
                   Intersubnet link = No
                   Extended border node = Not supported
                   RTP Tower = Supported
   Fully qualified PCID control vector:
    (0x0002) PCID = 0xe887a74908bc569a
    (0x000B) Network qualified CP name = USIBMRA.CM5HPRNN
   COS/TPF control vector:
    (0x0002) Transmission priority = Network
    (0x0004) COS name = SNASVCMG
   Route Information control vector
    (0x0002) Route direction = Forward
             REFIFOing required = No
    (0x0004) Maximum packet size = 2058
    (0x0008) Accumulated transmission time (microseconds) = 75
    (0x000C) Minimum link capacity (Kbits/second) = 16000
    (0x0010) Limited resource liveness timer (seconds) = 0
      ANR Path control vector
       (0x0003) ANR label represents a subarea network route = No
       (0x0004) ANR label = 801e

Line:   3768 Recv MU
Time stamp: 08:27:41.23
DLC type: HPR
TCID: 0000007C
Transmission priority: Network
TH: FID2, Exp, OIS, LFSID=0x00000, SNF=0x0000
RH: RQ, NC, FI, OIC, RQN
Route setup
   Type = Reply
   Destination hop index = 2
   Locate search is required = Yes
   Destination is mobile = No
   NCE is used for all LUs/BFs = No
   Path switch time (milliseconds) = 480000
   Network name control vector:
    (0x0002) Network name type = CP
    (0x0003) Name = USIBMRA.RA39
   NCE Instance Identifier control vector
    (0x0002) NCE instance identifer = 0xaff8273e
   NCE Identifier control vector
    (0x0002) NCE identifier = d000000000000000
   Route selection control vector:
    (0x0002) Maximum hop count = 2
    (0x0003) Current hop count = 0
      TG descriptor control vector
         TG Identifier TG Descriptor subfield
          (0x0002) TG number = 21
          (0x0004) Partner name = USIBMRA.NNP41A
```

```
          (0x0012) Link connection network = No
                   Additional configuration information = No
                   HPR = Supported
                   TG type = Boundary Function based or APPN
                   Intersubnet link = No
                   Extended border node = Supported
                   RTP Tower = Supported
          (0x0013) Subarea number = 0x80000027
      TG descriptor control vector
         TG Identifier TG Descriptor subfield
          (0x0002) TG number = 33
          (0x0004) Partner name = CM5HPRNN
          (0x000C) Link connection network = No
                   Additional configuration information = No
                   HPR = Supported
                   TG type = Boundary Function based or APPN
                   Intersubnet link = No
                   Extended border node = Not supported
                   RTP Tower = Supported
   COS/TPF control vector:
    (0x0002) Transmission priority = Low
    (0x0004) COS name = SNASVCMG
   Fully qualified PCID control vector:
    (0x0002) PCID = 0xe887a74908bc569a
    (0x000B) Network qualified CP name = USIBMRA.CM5HPRNN
   Route Information control vector
    (0x0002) Route direction = Forward
             REFIFOing required = No
    (0x0004) Maximum packet size = 2058
    (0x0008) Accumulated transmission time (microseconds) = 83
    (0x000C) Minimum link capacity (Kbits/second) = 16000
    (0x0010) Limited resource liveness timer (seconds) = 0
      ANR Path control vector
       (0x0003) ANR label represents a subarea network route = No
       (0x0004) ANR label = 801e
       (0x0007) ANR label represents a subarea network route = No
       (0x0008) ANR label = a500000002
   Route Information control vector
    (0x0002) Route direction = Reverse
             REFIFOing required = No
    (0x0004) Maximum packet size = 2058
    (0x0008) Accumulated transmission time (microseconds) = 108
    (0x000C) Minimum link capacity (Kbits/second) = 15974
    (0x0010) Limited resource liveness timer (seconds) = 0
      ANR Path control vector
       (0x0003) ANR label represents a subarea network route = No
       (0x0004) ANR label = 800900b100000000
       (0x000D) ANR label represents a subarea network route = No
       (0x000E) ANR label = a400304481
```

Now the CPSVRMGR BIND can flow.

```
Line:   3784 Send MU
Time stamp: 08:27:41.24
DLC type: HPR
  TH: FID5, Exp, OIS, SNF=0x0006, S SA2586f5ee146e1f93
  RH: RQ, SC, FI, OIC, RQD1
  BIND rq
   (0x0001) Type = Negotiable
   (0x0002) FM profile = 19
   (0x0003) TS profile = 7
            FM usage - primary:
   (0x0004)    Chaining use = Multiple-RU chains allowed
               Request control mode = Immediate request mode
               Chain response protocol = Definite or exception response
               Two-phase commit = Not supported
               Compression = Will not be used
               Send end bracket = Will not send
            FM usage - secondary:
   (0x0005)    Chaining use = Multiple-RU chains allowed
               Request control mode = Immediate request mode
               Chain response protocol = Definite or exception response
               Two-phase commit = Not supported
               Compression = Will not be used
               Send end bracket = Will not send
            FM Usage - common:
   (0x0006)    Whole BIUs required = No
               FM header usage = Allowed
               Brackets are used = Yes
               Bracket reset state = INB
               Bracket termination rule = Conditional
               Alternate code set allowed = No
               BIND queueing allowed = No
   (0x0007)    Normal-flow send/receive mode = Full-duplex
               Recovery responsibility = Symmetric
               Contention winner = Primary
               Alternate code set = ASCII-7
               Control vectors included = Yes
               Half-duplex flip-flop primary reset state = Send
            TS usage:
   (0x0008)    Secondary to primary pacing stages = One
               Secondary send window size = 1
   (0x0009)    Adaptive pacing = Supported
               Secondary receive window size = 0
   (0x000A)    Secondary maximum send RU size = 512
   (0x000B)    Primary maximum send RU size = 512
   (0x000C)    Primary to secondary pacing stages = One
               Primary send window size = 0
   (0x000D)    Primary receive window size = 1
            PS profile:
   (0x000E)    LU type = 6
   (0x000F)    LU-6 level = 2
   (0x0016)    Extended security sense codes = Supported
   (0x0017)    Conversation-level security = Accepted
               Session security level = Enhanced
               Password substitution = Supported
               Already-verified indicator = Not accepted
   (0x0018)    Synchronization level supported = Confirm
               Session reinitiation responsibility = Operator controlled
               Parallel sessions supported = Yes
               CNOS supported = Yes
```

```
(0x0019)    Limited resource = No
            Length-checked compression options = No Compression
         Cryptography options:
(0x001A)    Private cryptography support = No
            Session-level cryptography support = No
(0x001C) Primary LU name = USIBMRA.CM5HPRNN
         User Data:
            Structured user data:
(0x002F)       Mode name = CPSVRMGR
(0x0039)       Session instance identifier = 0x01f58625edf58625
(0x0043)       Network-qualified PLU name = USIBMRA.CM5HPRNN
(0x0055)       Random data = 0x00a4534473e58f9289
(0x0061) Secondary LU name = USIBMRA.RA39
(0x006D) Fully qualified PCID control vector:
          (0x0002) PCID = 0xc96f2a091d56a273
          (0x000B) Network qualified CP name = USIBMRA.CM5HPRNN
(0x0088) COS/TPF control vector:
          (0x0002) Transmission priority = Network
          (0x0004) COS name = SNASVCMG
(0x0094) Route selection control vector:
          (0x0002) Maximum hop count = 2
          (0x0003) Current hop count = 2
(0x0098)    TG descriptor control vector
(0x009A)       TG Identifier TG Descriptor subfield
                (0x0002) TG number = 33
                (0x0004) Partner name = USIBMRA.NNP41A
                (0x0012) Link connection network = No
                         Additional configuration information = No
                         HPR = Supported
                         TG type = Boundary Function based or APPN
                         Intersubnet link = No
                         Extended border node = Not supported
                         RTP Tower = Supported
(0x00AD)    TG descriptor control vector
(0x00AF)       TG Identifier TG Descriptor subfield
                (0x0002) TG number = 21
                (0x0004) Partner name = RA39
                (0x0008) Link connection network = No
                         Additional configuration information = No
                         HPR = Supported
                         TG type = Boundary Function based or APPN
                         Intersubnet link = No
                         Extended border node = Not supported
                         RTP Tower = Supported

Line:   3799 Recv MU
Time stamp: 08:27:41.29
DLC type: HPR
TH: FID5, Exp, OIS, SNF=0x0006, R SA2586f5ee146e1f93
RH: +RSP, SC, FI, RQD1
BIND +rsp
(0x0001) Type = Negotiable
(0x0002) FM profile = 19
(0x0003) TS profile = 7
         FM usage - primary:
(0x0004)    Chaining use = Multiple-RU chains allowed
            Request control mode = Immediate request mode
            Chain response protocol = Definite or exception response
            Two-phase commit = Not supported
```

```
                Compression = Will not be used
                Send end bracket = Will not send
             FM usage - secondary:
   (0x0005)     Chaining use = Multiple-RU chains allowed
                Request control mode = Immediate request mode
                Chain response protocol = Definite or exception response
                Two-phase commit = Not supported
                Compression = Will not be used
                Send end bracket = Will not send
             FM Usage - common:
   (0x0006)     Whole BIUs required = No
                FM header usage = Allowed
                Brackets are used = Yes
                Bracket reset state = INB
                Bracket termination rule = Conditional
                Alternate code set allowed = No
                BIND queueing allowed = No
   (0x0007)     Normal-flow send/receive mode = Half-duplex flip-flop
                Recovery responsibility = Symmetric
                Contention winner = Primary
                Alternate code set = ASCII-7
                Control vectors included = Yes
                Half-duplex flip-flop primary reset state = Send
             TS usage:
   (0x0008)     Secondary to primary pacing stages = One
                Secondary send window size = 0
   (0x0009)     Adaptive pacing = Supported
                Secondary receive window size = 0
   (0x000A)     Secondary maximum send RU size = 512
   (0x000B)     Primary maximum send RU size = 512
   (0x000C)     Primary to secondary pacing stages = One
                Primary send window size = 0
   (0x000D)     Primary receive window size = 0
             PS profile:
   (0x000E)     LU type = 6
   (0x000F)     LU-6 level = 2
   (0x0016)     Extended security sense codes = Not supported
   (0x0017)     Conversation-level security = Not accepted
                Session security level = Basic
                Password substitution = Not supported
                Already-verified indicator = Not accepted
   (0x0018)     Synchronization level supported = Confirm
                Session reinitiation responsibility = Operator controlled
                Parallel sessions supported = Yes
                CNOS supported = Yes
   (0x0019)     Limited resource = No
                Length-checked compression options = No Compression
             Cryptography options:
   (0x001A)     Private cryptography support = No
                Session-level cryptography support = No
             User Data:
                Structured user data:
   (0x001F)        Mode name = CPSVRMGR
   (0x0029)        Session instance identifier = 0x026f2a091d56a273
   (0x0033)        Network-qualified SLU name = USIBMRA.RA39
   (0x0042) Fully qualified PCID control vector:
              (0x0002) PCID = 0xc96f2a091d56a273
              (0x000B) Network qualified CP name = USIBMRA.CM5HPRNN
   (0x005D) Session Address control vector
```

```
||          (0x0002) Session address assignor = Sender
||                   Session address = 0x000000000f000004
|| (0x0067) Route selection control vector:
||          (0x0002) Maximum hop count = 2
||          (0x0003) Current hop count = 2
|| (0x006B)    TG descriptor control vector
|| (0x006D)       TG Identifier TG Descriptor subfield
||                 (0x0002) TG number = 33
||                 (0x0004) Partner name = USIBMRA.NNP41A
||                 (0x0012) Link connection network = No
||                          Additional configuration information = No
||                          HPR = Supported
||                          TG type = Boundary Function based or APPN
||                          Intersubnet link = No
||                          Extended border node = Not supported
||                          RTP Tower = Supported
|| (0x0080)    TG descriptor control vector
|| (0x0082)       TG Identifier TG Descriptor subfield
||                 (0x0002) TG number = 21
||                 (0x0004) Partner name = RA39
||                 (0x0008) Link connection network = No
||                          Additional configuration information = No
||                          HPR = Supported
||                          TG type = Boundary Function based or APPN
||                          Intersubnet link = No
||                          Extended border node = Not supported
||                          RTP Tower = Supported
```

The path for the DLUR/S session to RA39 is:

```
┌─────────┐            ┌───────┐            ┌──────┐
│CM5HPRNN ├───TG33─────┤NNP41A ├───TG21─────┤RA39  │
└─────────┘            └───────┘            └──────┘
```

Figure 246. DLUR/S Path for RA39

REQACTPU is requested, in order to activate the DLUR PU.

```
Line:   3811 Send MU
Time stamp: 08:27:41.29
DLC type: HPR
||TH: FID5, OIS, SNF=0x0001, R SA000000000f000004
||RH: RQ, FMD, FI, OIC, RQE1, PI, CEBI
||FMH-5
||   Command code = Attach
||   User ID already verified = No
||   Password is substituted = No
||   PIP present = No
||   Conversation type = Basic
```

```
   Synchronization level = None
   Transaction program name = ?006 (APPN Receive_Encap_Msg)
   Logical unit of work identifier:
      LU name = USIBMRA.CM5HPRNN
      Instance number = 0xafff44aea03a
      Sequence number = 0x0001
   Conversation correlator = 0xd02b1f93eaf58625
FID2 Encapsulation Variable
   TH: FID2, OIS, LFSID=0x00000, SNF=0x0000
   RH: RQ, FMD, FI, OIC, RQD1
   REQACTPU rq
      Format = Internal PU
   DLUR/S Capabilities control vector
    (0x0002) Dependent LU support level = 0x01
    (0x0003) Node type = Network Node
    (0x0004) DLUR/S node type = DLUR
    (0x0005) Flow reduction sequence number = 0x00000000
    (0x0009) RECEIVE_TDU_TP = Not supported
   Network name control vector:
    (0x0002) Network name type = PU
    (0x0003) Name = MPU00002
   XID Image FID2 Encapsulation control vector
    (0x0002) XID I-field image:
             Hex dump:
                020605da a41a                        <......         > <....
   TG descriptor control vector
      DLC Signaling Type TG Descriptor subfield
       (0x0002) DLC type = INTPU
      DLC Signaling Information TG Descriptor subfield
       (0x0002) Block number = 005d
                ID number = 000aa41a
   Fully qualified PCID control vector:
    (0x0002) PCID = 0xc96f2a091956a273
    (0x000B) Network qualified CP name = USIBMRA.CM5HPRNN

Line:   3881 Recv MU
Time stamp: 08:27:41.54
DLC type: HPR
TH: FID5, OIS, SNF=0x0001, R SA2586f5f3fc711f93
RH: RQ, FMD, FI, OIC, RQE1, PI, CEBI
FMH-5
   Command code = Attach
   User ID already verified = No
   Password is substituted = No
   PIP present = No
   Conversation type = Basic
   Synchronization level = None
   Transaction program name = ?006 (APPN Receive_Encap_Msg)
FID2 Encapsulation Variable
   TH: FID2, OIS, LFSID=0x00000, SNF=0x0000
   RH: +RSP, FMD, FI, RQD1
   REQACTPU +rsp
   Fully qualified PCID control vector:
    (0x0002) PCID = 0xc96f2a091956a273
    (0x000B) Network qualified CP name = USIBMRA.CM5HPRNN
   DLUR/S Capabilities control vector
    (0x0002) Dependent LU support level = 0x01
    (0x0003) Node type = Network Node
    (0x0004) DLUR/S node type = DLUS
```

```
||    (0x0005) Flow reduction sequence number = 0x00000000
||    (0x0009) RECEIVE_TDU_TP = Supported


Line:   3893 Recv MU
Time stamp: 08:27:41.55
DLC type: HPR
||TH: FID5, OIS, SNF=0x0002, R SA2586f5f3fc711f93
||RH: RQ, FMD, FI, OIC, RQE1, RLWI, PI, BB, CEBI
||FMH-5
||   Command code = Attach
||   User ID already verified = No
||   Password is substituted = No
||   PIP present = No
||   Conversation type = Basic
||   Synchronization level = None
||   Transaction program name = ?006 (APPN Receive_Encap_Msg)
||FID2 Encapsulation Variable
||   TH: FID2, Exp, OIS, LFSID=0x00000, SNF=0x0205
||   RH: RQ, SC, FI, OIC, RQD1
||   ACTPU rq
||      Format = 0
||      Activation type = ERP
||      FM profile = 0
||      TS profile = 1
||      SSCP ID = 0x050000000027
||      Network name control vector:
||       (0x0002) Network name type = PU
||       (0x0003) Name = USIBMRA.MPU0002
||   Extended SDLC Secondary Station control vector
||    (0x0002) Node type = T2
||             Continue link-level contact during auto network shutdown = No
||             Secondary station polling type = SNRM
||             Secondary station modem test = Not supported
||             Data mode = Half-duplex
||    (0x0004) Maximum unanswered frames sent = 0x01
||    (0x0005) Maximum consecutive BTUs sent = 0x01
||    (0x0006) Immediate retry on error = No
||    (0x0009) Maximum sendable BTU size = 0x0109
||    (0x000B) Total transmissions = 0x0000
||    (0x000D) Total error retries = 0x0000
||    (0x000F) Average bytes expected when polled = 0x0000
||    (0x0011) Link connection segments = 0x01
||    (0x0012) Link segment 1 local modem address = 0x00
||             Link segment 2 local modem address = 0x00
||   TG descriptor control vector
||      DLC Signaling Type TG Descriptor subfield
||       (0x0002) DLC type = INTPU
||      DLC Signaling Information TG Descriptor subfield
||       (0x0002) Block number = 005d
||                ID number = 000aa41a
||   Fully qualified PCID control vector:
||    (0x0002) PCID = 0xc96f2a091956a273
||    (0x000B) Network qualified CP name = USIBMRA.CM5HPRNN

Line:   3908 Send MU
Time stamp: 08:27:41.56
DLC type: HPR
||TH: FID5, OIS, SNF=0x0002, R SA000000000f000004
```

```
RH: RQ, FMD, FI, OIC, RQE1, PI, BB, CEBI
FMH-5
   Command code = Attach
   User ID already verified = No
   Password is substituted = No
   PIP present = No
   Conversation type = Basic
   Synchronization level = None
   Transaction program name = ?006 (APPN Receive_Encap_Msg)
   Logical unit of work identifier:
      LU name = USIBMRA.CM5HPRNN
      Instance number = 0xafff44aea03a
      Sequence number = 0x0001
   Conversation correlator = 0xd02b1f93fcf58625
FID2 Encapsulation Variable
   TH: FID2, Exp, OIS, LFSID=0x00000, SNF=0x0205
   RH: +RSP, SC, FI, RQD1
   ACTPU +rsp
      Format = 1
      Activation type = ERP
      Contents ID =
      PU FMD-RU-usage control vector:
       (0x0001) Adjacent PU load capability = Cannot load
                FMD request capability = Can receive
   XID Image FID2 Encapsulation control vector
    (0x0002) XID I-field image:
             Hex dump:
                020605da a41a                        <......          > <....
   Fully qualified PCID control vector:
    (0x0002) PCID = 0xc96f2a091956a273
    (0x000B) Network qualified CP name = USIBMRA.CM5HPRNN
```

The VTAM Welcome message (USSMSG10) is received.

```
Line:   3998 Recv MU
Time stamp: 08:28:14.14
DLC type: HPR
TH: FID5, OIS, SNF=0x0006, R SA2586f5f3fc711f93
RH: RQ, FMD, FI, FIC, RQE1, BB
FMH-5
   Command code = Attach
   User ID already verified = No
   Password is substituted = No
   PIP present = No
   Conversation type = Basic
   Synchronization level = None
   Transaction program name = ?006 (APPN Receive_Encap_Msg)
FID2 Encapsulation Variable
   TH: FID2, OIS, LFSID=0x00002, SNF=0x003b
   RH: RQ, FMD, OIC, RQD1
   User data - remainder of RU follows:
      Hex dump:
         4015d4e2 c7f1f040 e2d5c115 4015c9d5   <@......@....@...> < .MSG10 SNA
```

B.1.9 Communications Server/2 Displays

At this point in the test, we used the CS/2 administration facilities to display the network connectivity. First, we displayed the logical links as shown in Figure 247.

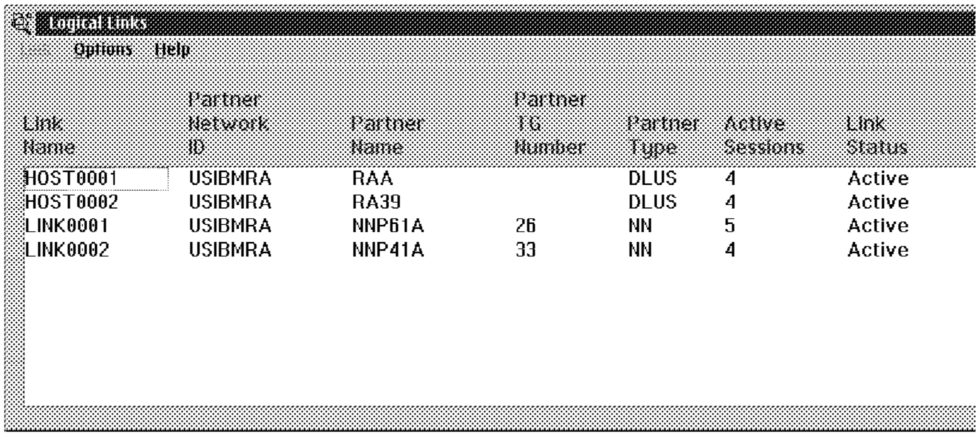

Logical Links

Options Help

Link Name	Partner Network ID	Partner Name	Partner TG Number	Partner Type	Active Sessions	Link Status
HOST0001	USIBMRA	RAA		DLUS	4	Active
HOST0002	USIBMRA	RA39		DLUS	4	Active
LINK0001	USIBMRA	NNP61A	26	NN	5	Active
LINK0002	USIBMRA	NNP41A	33	NN	4	Active

Figure 247. Logical Links Display

The links used for the CP-CP sessions and the DLUR/S sessions are all active.

At the same time, we displayed the active RTP connections, as in Figure 248.

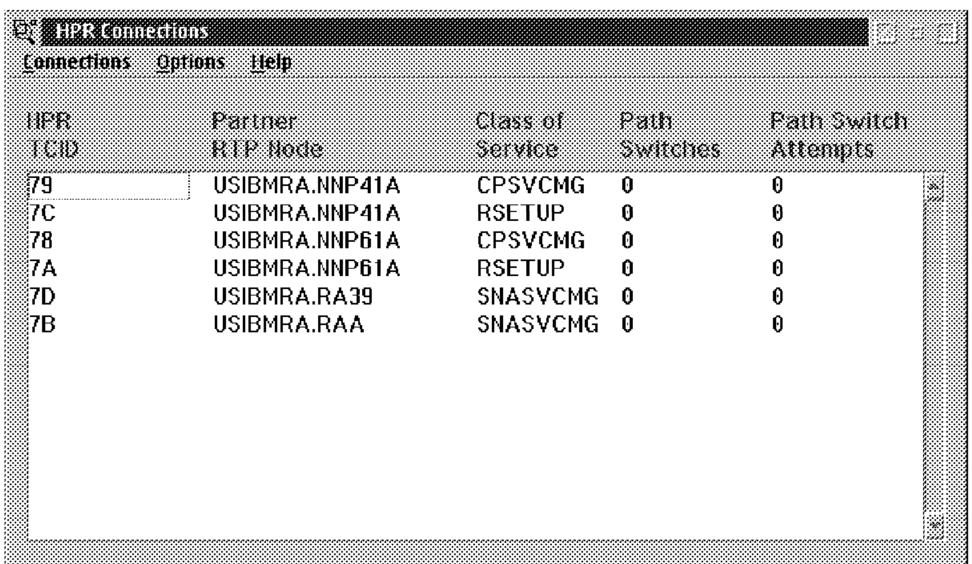

HPR Connections

Connections Options Help

HPR TCID	Partner RTP Node	Class of Service	Path Switches	Path Switch Attempts
79	USIBMRA.NNP41A	CPSVCMG	0	0
7C	USIBMRA.NNP41A	RSETUP	0	0
78	USIBMRA.NNP61A	CPSVCMG	0	0
7A	USIBMRA.NNP61A	RSETUP	0	0
7D	USIBMRA.RA39	SNASVCMG	0	0
7B	USIBMRA.RAA	SNASVCMG	0	0

Figure 248. HPR Connections

All the path switch counters are at zero.

Next, we looked at individual RTP connections from the previous display. Figure 249 on page 280 shows the one identified by TCID 79. This turns out to be a CP-CP session pipe to NNP41A.

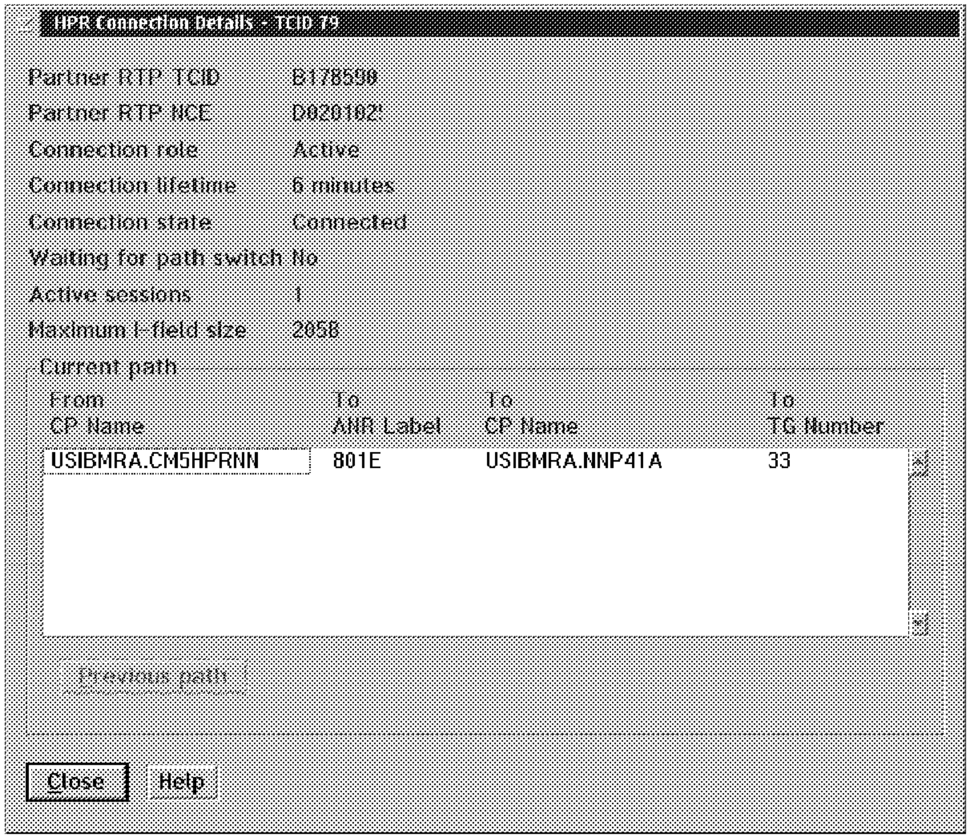

Figure 249. TCID 79

TCID 7C (Figure 250 on page 281) is a Route Setup pipe to NNP41A.

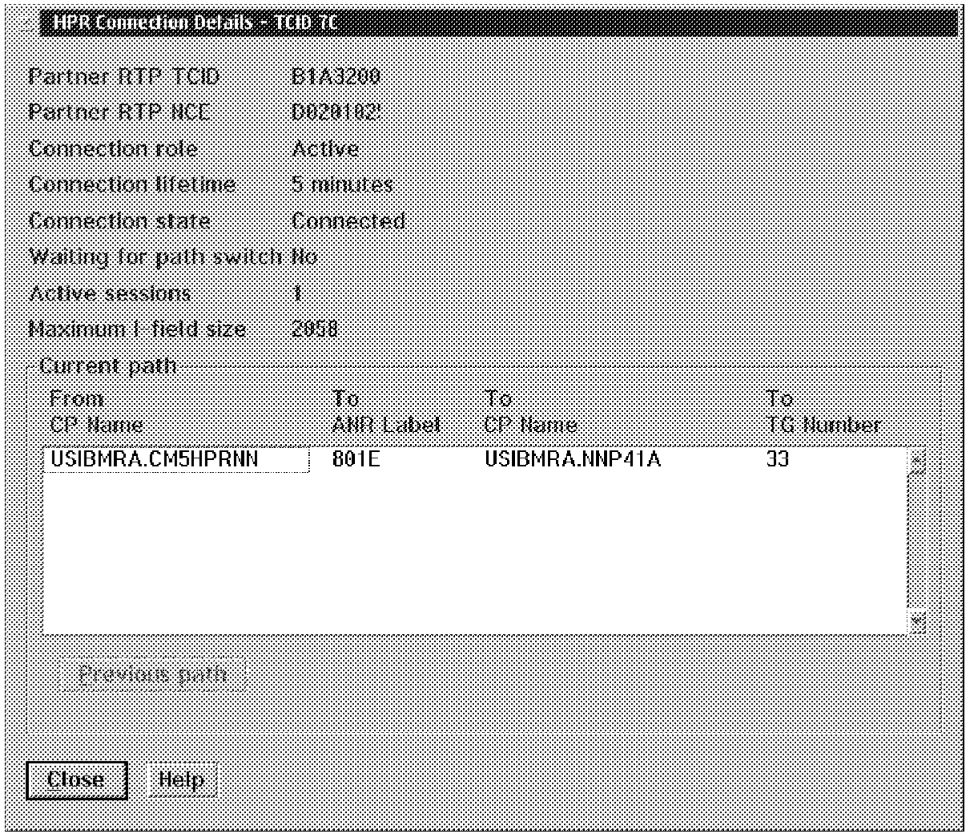

Figure 250. TCID 7C

TCID 7D (Figure 251 on page 282) is the DLUR/S pipe to RA39, which takes the following route:

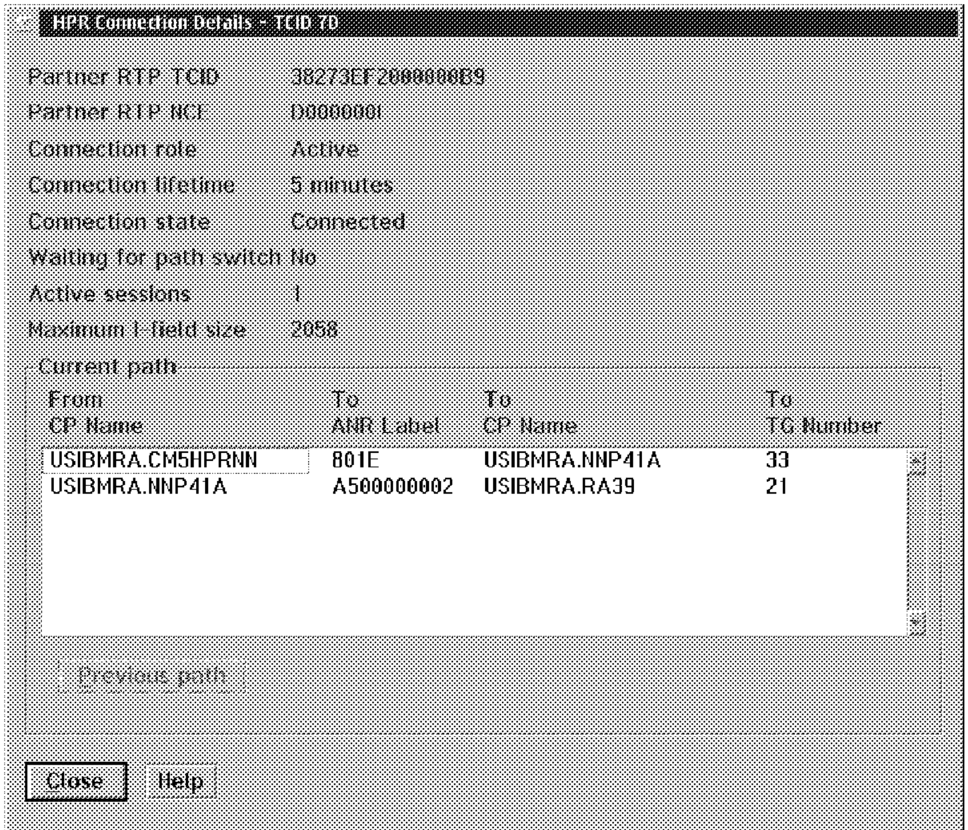

Figure 251. TCID 7D

Similarly, TCID 7B (not shown) is the DLUR/S pipe to RAA via the route:

B.2 Logon to NetView on RAA

Our dependent LU is now in session with the RA39 SSCP, with the VTAM USS10 message displayed on the screen. From there, we log on to RAAAN, which is NetView on VTAM RAA.

```
Line:    822 Send MU
Time stamp: 08:37:39.91
DLC type: HPR
||TH: FID5, OIS, SNF=0x0008, R SA000000000f000004
||RH: RQ, FMD, FI, OIC, RQE1, BB, CEBI
||FMH-5
||   Command code = Attach
||   User ID already verified = No
||   Password is substituted = No
||   PIP present = No
||   Conversation type = Basic
||   Synchronization level = None
||   Transaction program name = ?006 (APPN Receive_Encap_Msg)
||   Logical unit of work identifier:
```

```
||      LU name = USIBMRA.CM5HPRNN
||      Instance number = 0xafff44aea03a
||      Sequence number = 0x0001
||   Conversation correlator = 0xd02b1f93aaf78625
||FID2 Encapsulation Variable
||   TH: FID2, OIS, LFSID=0x00002, SNF=0x0002
||   RH: RQ, FMD, OIC, RQD1
||   User data - remainder of RU follows:
||      Hex dump:
||         93968796 95408197 97938984 4d998181 logon applid(raaan)
||         81955d
||   Fully qualified PCID control vector:
||    (0x0002) PCID = 0xc96f2a091956a273
||    (0x000B) Network qualified CP name = USIBMRA.CM5HPRNN

Line:    832 Recv MU
Time stamp: 08:37:39.94
DLC type: HPR
||TH: FID5, OIS, SNF=0x0009, R SA2586f5f3fc711f93
||RH: RQ, FMD, FI, OIC, RQE1, BB, CEBI
||FMH-5
||   Command code = Attach
||   User ID already verified = No
||   Password is substituted = No
||   PIP present = No
||   Conversation type = Basic
||   Synchronization level = None
||   Transaction program name = ?006 (APPN Receive_Encap_Msg)
||FID2 Encapsulation Variable
||   TH: FID2, OIS, LFSID=0x00002, SNF=0x0002
||   RH: +RSP, FMD, RQD1
||   Fully qualified PCID control vector:
||    (0x0002) PCID = 0xc96f2a091956a273
||    (0x000B) Network qualified CP name = USIBMRA.CM5HPRNN
```

RAA has received the request and performs a Route Setup to establish the HPR path on which this new session will be set up. The Route Setup message flows on the existing long-lived pipe from NNP61A.

```
Line:    839 Recv MU
Time stamp: 08:37:40.03
DLC type: HPR
TCID: 0000007A
Transmission priority: Network
||TH: FID2, Exp, OIS, LFSID=0x00000, SNF=0x0000
||RH: RQ, NC, FI, OIC, RQN
||Route setup
||   Type = Request
||   Route Setup triggered by path switch = No
||   Destination hop index = 2
||   Locate search is required = No
||   Destination is mobile = No
||   NCE is used for all LUs/BFs = No
||   Path switch time (milliseconds) = 0
||   Network name control vector:
||    (0x0002) Network name type = LU
||    (0x0003) Name = USIBMRA.LU41A1
```

```
Route selection control vector:
 (0x0002) Maximum hop count = 2
 (0x0003) Current hop count = 2
   TG descriptor control vector
      TG Identifier TG Descriptor subfield
        (0x0002) TG number = 21
        (0x0004) Partner name = USIBMRA.NNP61A
        (0x0012) Link connection network = No
                 Additional configuration information = No
                 HPR = Supported
                 TG type = Boundary Function based or APPN
                 Intersubnet link = No
                 Extended border node = Supported
                 RTP Tower = Supported
        (0x0013) Subarea number = 0x8000000a
   TG descriptor control vector
      TG Identifier TG Descriptor subfield
        (0x0002) TG number = 26
        (0x0004) Partner name = CM5HPRNN
        (0x000C) Link connection network = No
                 Additional configuration information = No
                 HPR = Supported
                 TG type = Boundary Function based or APPN
                 Intersubnet link = No
                 Extended border node = Not supported
                 RTP Tower = Supported
COS/TPF control vector:
 (0x0002) Transmission priority = Medium
 (0x0004) COS name = #CONNECT
Fully qualified PCID control vector:
 (0x0002) PCID = 0xf7ff6164529f6dff
 (0x000B) Network qualified CP name = USIBMRA.RAA
Route Information control vector
 (0x0002) Route direction = Forward
          REFIFOing required = No
 (0x0004) Maximum packet size = 2058
 (0x0008) Accumulated transmission time (microseconds) = 108
 (0x000C) Minimum link capacity (Kbits/second) = 15974
 (0x0010) Limited resource liveness timer (seconds) = 0
   ANR Path control vector
    (0x0003) ANR label represents a subarea network route = No
    (0x0004) ANR label = 803000c300000000
    (0x000D) ANR label represents a subarea network route = No
    (0x000E) ANR label = a320304c82
```

The path is as follows:

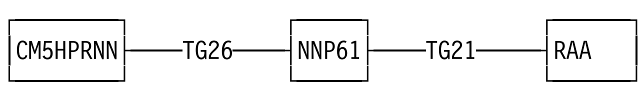

Figure 252. DLU Session Path

Now the Route Setup reply.

```
Line:    852 Send MU
Time stamp: 08:37:40.03
DLC type: HPR
TCID: 0000007A
Transmission priority: Network
||TH: FID2, Exp, OIS, LFSID=0x00000, SNF=0x0000
||RH: RQ, NC, FI, OIC, RQN
||Route setup
||   Type = Reply
||   Destination hop index = 2
||   Locate search is required = No
||   Destination is mobile = No
||   NCE is used for all LUs/BFs = Yes
||   Path switch time (milliseconds) = 240000
||   COS/TPF control vector:
||    (0x0002) Transmission priority = Medium
||    (0x0004) COS name = #CONNECT
||   Fully qualified PCID control vector:
||    (0x0002) PCID = 0xf7ff6164529f6dff
||    (0x000B) Network qualified CP name = USIBMRA.RAA
||   Route Information control vector
||    (0x0002) Route direction = Forward
||             REFIFOing required = No
||    (0x0004) Maximum packet size = 2058
||    (0x0008) Accumulated transmission time (microseconds) = 108
||    (0x000C) Minimum link capacity (Kbits/second) = 15974
||    (0x0010) Limited resource liveness timer (seconds) = 0
||      ANR Path control vector
||       (0x0003) ANR label represents a subarea network route = No
||       (0x0004) ANR label = 803000c300000000
||       (0x000D) ANR label represents a subarea network route = No
||       (0x000E) ANR label = a320304c82
||   Network name control vector:
||    (0x0002) Network name type = CP
||    (0x0003) Name = USIBMRA.CM5HPRNN
||   NCE Identifier control vector
||    (0x0002) NCE identifier = 80
||   NCE Instance Identifier control vector
||    (0x0002) NCE instance identifier = 0x34e1a64d
||   Route selection control vector:
||    (0x0002) Maximum hop count = 1
||    (0x0003) Current hop count = 0
||      TG descriptor control vector
||         TG Identifier TG Descriptor subfield
||          (0x0002) TG number = 26
||          (0x0004) Partner name = USIBMRA.NNP61A
||          (0x0012) Link connection network = No
||                   Additional configuration information = No
||                   HPR = Supported
||                   TG type = Boundary Function based or APPN
||                   Intersubnet link = No
||                   Extended border node = Not supported
||                   RTP Tower = Supported
||   Route Information control vector
||    (0x0002) Route direction = Reverse
```

```
||             REFIFOing required = No
||    (0x0004) Maximum packet size = 2058
||    (0x0008) Accumulated transmission time (microseconds) = 75
||    (0x000C) Minimum link capacity (Kbits/second) = 16000
||    (0x0010) Limited resource liveness timer (seconds) = 0
||      ANR Path control vector
||       (0x0003) ANR label represents a subarea network route = No
||       (0x0004) ANR label = 801c
```

A BIND is received from RAAAN. This BIND is not encapsulated in the LU 6.2 pipe.

```
Line:    866 Recv MU
Time stamp: 08:37:40.22
DLC type: HPR
||TH: FID5, Exp, OIS, SNF=0x19b9, S SA000000000f00000a
||RH: RQ, SC, FI, OIC, RQD1
||BIND rq
|| (0x0001) Type = Non-negotiable
|| (0x0002) FM profile = 3
|| (0x0003) TS profile = 3
||          FM usage - primary:
|| (0x0004)    Chaining use = Multiple-RU chains allowed
||             Request control mode = Immediate request mode
||             Chain response protocol = Definite or exception response
||             Two-phase commit = Not supported
||             Compression = Will not be used
||             Send end bracket = May send
||          FM usage - secondary:
|| (0x0005)    Chaining use = Multiple-RU chains allowed
||             Request control mode = Immediate request mode
||             Chain response protocol = Exception response
||             Two-phase commit = Not supported
||             Compression = Will not be used
||             Send end bracket = Will not send
||          FM Usage - common:
|| (0x0006)    Whole BIUs required = No
||             FM header usage = Not allowed
||             Brackets are used = Yes
||             Bracket reset state = BETB
||             Bracket termination rule = Conditional
||             Alternate code set allowed = No
||             BIND queueing allowed = No
|| (0x0007)    Normal-flow send/receive mode = Half-duplex flip-flop
||             Recovery responsibility = Contention loser
||             Contention winner = Secondary
||             Alternate code set = ASCII-7
||             Control vectors included = Yes
||             Half-duplex flip-flop primary reset state = Receive
||          TS usage:
|| (0x0008)    Secondary to primary pacing stages = One
||             Secondary send window size = 0
|| (0x0009)    Adaptive pacing = Supported
||             Secondary receive window size = 0
|| (0x000A)    Secondary maximum send RU size = 1024
|| (0x000B)    Primary maximum send RU size = 3840
|| (0x000C)    Primary to secondary pacing stages = One
```

```
              Primary send window size = 0
(0x000D)      Primary receive window size = 0
           PS profile:
(0x000E)      LU type = 2
(0x000F)      Query = Supported
(0x0014)      Default screen size = 0x0
(0x0016)      Alternate screen size = 0x0
(0x0018)      Screen size = Not specified
              Length-checked compression options = No Compression
           Cryptography options:
(0x001A)      Private cryptography support = No
              Session-level cryptography support = No
(0x001C) Primary LU name = USIBMRA.RAAAN
           User Data:
(0x002C) Secondary LU name = USIBMRA.LU41A1
(0x003A) Fully qualified PCID control vector:
           (0x0002) PCID = 0xf70794547c2c2221
           (0x000B) Network qualified CP name = USIBMRA.RA39
(0x0051) Network name control vector:
           (0x0002) Network name type = CP
           (0x0003) Name = USIBMRA.RAA
(0x005F) Route selection control vector:
           (0x0002) Maximum hop count = 2
           (0x0003) Current hop count = 2
(0x0063)      TG descriptor control vector
(0x0065)         TG Identifier TG Descriptor subfield
                   (0x0002) TG number = 21
                   (0x0004) Partner name = USIBMRA.NNP61A
                   (0x0012) Link connection network = No
                            Additional configuration information = No
                            HPR = Supported
                            TG type = Boundary Function based or APPN
                            Intersubnet link = No
                            Extended border node = Supported
                            RTP Tower = Supported
                   (0x0013) Subarea number = 0x8000000a
(0x007C)      TG descriptor control vector
(0x007E)         TG Identifier TG Descriptor subfield
                   (0x0002) TG number = 26
                   (0x0004) Partner name = CM5HPRNN
                   (0x000C) Link connection network = No
                            Additional configuration information = No
                            HPR = Supported
                            TG type = Boundary Function based or APPN
                            Intersubnet link = No
                            Extended border node = Not supported
                            RTP Tower = Supported
(0x008B) COS/TPF control vector:
           (0x0002) Transmission priority = Medium
           (0x0004) COS name = #CONNECT
(0x0097) Mode control vector:
           (0x0003) Mode name = D4C32XX3
```

The LU acknowledges this BIND with a BIND response.

```
Line:    880 Send MU
Time stamp: 08:37:40.28
DLC type: HPR
||TH: FID5, Exp, OIS, SNF=0x19b9, R SA000000000f00000a
||RH: +RSP, SC, FI, RQD1
||BIND +rsp
|| (0x0001) Type = Non-negotiable
|| (0x0002) FM profile = 3
|| (0x0003) TS profile = 3
||          FM usage - primary:
|| (0x0004)    Chaining use = Multiple-RU chains allowed
||             Request control mode = Immediate request mode
||             Chain response protocol = Definite or exception response
||             Two-phase commit = Not supported
||             Compression = Will not be used
||             Send end bracket = May send
||          FM usage - secondary:
|| (0x0005)    Chaining use = Multiple-RU chains allowed
||             Request control mode = Immediate request mode
||             Chain response protocol = Exception response
||             Two-phase commit = Not supported
||             Compression = Will not be used
||             Send end bracket = Will not send
||          FM Usage - common:
|| (0x0006)    Whole BIUs required = No
||             FM header usage = Not allowed
||             Brackets are used = Yes
||             Bracket reset state = BETB
||             Bracket termination rule = Conditional
||             Alternate code set allowed = No
||             BIND queueing allowed = No
|| (0x0007)    Normal-flow send/receive mode = Half-duplex flip-flop
||             Recovery responsibility = Contention loser
||             Contention winner = Secondary
||             Alternate code set = ASCII-7
||             Control vectors included = Yes
||             Half-duplex flip-flop primary reset state = Receive
||          TS usage:
|| (0x0008)    Secondary to primary pacing stages = One
||             Secondary send window size = 0
|| (0x0009)    Adaptive pacing = Not supported
||             Secondary receive window size = 0
|| (0x000A)    Secondary maximum send RU size = 1024
|| (0x000B)    Primary maximum send RU size = 3840
|| (0x000C)    Primary to secondary pacing stages = One
||             Primary send window size = 0
|| (0x000D)    Primary receive window size = 0
||          PS profile:
|| (0x000E)    LU type = 2
|| (0x000F)    Query = Supported
|| (0x0014)    Default screen size = 0x0
|| (0x0016)    Alternate screen size = 0x0
|| (0x0018)    Screen size = Not specified
||             Length-checked compression options = No Compression
||          Cryptography options:
|| (0x001A)    Private cryptography support = No
||             Session-level cryptography support = No
||          User Data:
|| (0x001F) Fully qualified PCID control vector:
```

```
            (0x0002) PCID = 0xf70794547c2c2221
            (0x000B) Network qualified CP name = USIBMRA.RA39
 (0x0036) Session Address control vector
            (0x0002) Session address assignor = Sender
                     Session address = 0x2586f7ae448d1f93
 (0x0040) Route selection control vector:
            (0x0002) Maximum hop count = 2
            (0x0003) Current hop count = 2
 (0x0044)    TG descriptor control vector
 (0x0046)       TG Identifier TG Descriptor subfield
                  (0x0002) TG number = 21
                  (0x0004) Partner name = USIBMRA.NNP61A
                  (0x0012) Link connection network = No
                           Additional configuration information = No
                           HPR = Supported
                           TG type = Boundary Function based or APPN
                           Intersubnet link = No
                           Extended border node = Supported
                           RTP Tower = Supported
                  (0x0013) Subarea number = 0x8000000a
 (0x005D)    TG descriptor control vector
 (0x005F)       TG Identifier TG Descriptor subfield
                  (0x0002) TG number = 26
                  (0x0004) Partner name = CM5HPRNN
                  (0x000C) Link connection network = No
                           Additional configuration information = No
                           HPR = Supported
                           TG type = Boundary Function based or APPN
                           Intersubnet link = No
                           Extended border node = Not supported
                           RTP Tower = Supported
```

Then the SESSST (session started) RU is sent to the SSCP. This flows on the SSCP-LU session in the DLUR/S pipe.

```
Line:    890 Send MU
Time stamp: 08:37:40.36
DLC type: HPR
TH: FID5, OIS, SNF=0x0009, R SA000000000f000004
RH: RQ, FMD, FI, OIC, RQE1, BB, CEBI
FMH-5
   Command code = Attach
   User ID already verified = No
   Password is substituted = No
   PIP present = No
   Conversation type = Basic
   Synchronization level = None
   Transaction program name = ?006 (APPN Receive_Encap_Msg)
   Logical unit of work identifier:
      LU name = USIBMRA.CM5HPRNN
      Instance number = 0xafff44aea03a
      Sequence number = 0x0001
   Conversation correlator = 0xd02b1f93aff78625
FID2 Encapsulation Variable
   TH: FID2, OIS, LFSID=0x00002, SNF=0x0000
   RH: RQ, FMD, FI, OIC, RQN
   SESSST rq
```

```
||      Format = 3
||      Unknown session key = 0x2a
||   Fully qualified PCID control vector:
||    (0x0002) PCID = 0xc96f2a091956a273
||    (0x000B) Network qualified CP name = USIBMRA.CM5HPRNN


Line:    920 Send MU
Time stamp: 08:37:40.42
DLC type: HPR
||TH: FID5, OIS, SNF=0x000a, R SA000000000f000004
||RH: RQ, FMD, FI, OIC, RQE1, BB, CEBI
||FMH-5
||   Command code = Attach
||   User ID already verified = No
||   Password is substituted = No
||   PIP present = No
||   Conversation type = Basic
||   Synchronization level = None
||   Transaction program name = ?006 (APPN Receive_Encap_Msg)
||   Logical unit of work identifier:
||      LU name = USIBMRA.CM5HPRNN
||      Instance number = 0xafff44aea03a
||      Sequence number = 0x0001
||   Conversation correlator = 0xd02b1f93b2f78625
||FID2 Encapsulation Variable
||   TH: FID2, OIS, LFSID=0x00002, SNF=0x003c
||   RH: -RSP, FMD, SDI, RQD1
||   Sense data: 0x081b0000
||   Fully qualified PCID control vector:
||    (0x0002) PCID = 0xc96f2a091956a273
||    (0x000B) Network qualified CP name = USIBMRA.CM5HPRNN
```

This is a negative response with sense code 081B0000, meaning “receiver in transmit mode”. This is quite normal, and results from VTAM sending the USS 0 message (not shown in the trace) while the LU is doing something else.

Next, the LU receives an UNBIND because of NetView issuing a CLSDST PASS. Note the BIND forthcoming indication. We will eventually be logged on to RAAAN007, not RAAAN.

```
Line:    929 Recv MU
Time stamp: 08:37:40.56
DLC type: HPR
||TH: FID5, Exp, OIS, SNF=0x19bc, R SA2586f7ae448d1f93
||RH: RQ, SC, FI, OIC, RQD1
||UNBIND rq
||   Type = BIND forthcoming
||   Sense data = 00000000
||   Fully qualified PCID control vector:
||    (0x0002) PCID = 0xf70794547c2c2221
||    (0x000B) Network qualified CP name = USIBMRA.RA39
Line:    934 Send MU
Time stamp: 08:37:40.57
DLC type: HPR
```

```
||TH: FID5, Exp, OIS, SNF=0x19bc, R SA000000000f00000a
||RH: +RSP, SC, FI, RQD1
||UNBIND +rsp
Line:    938 Send MU
Time stamp: 08:37:40.58
DLC type: HPR
||TH: FID5, OIS, SNF=0x000b, R SA000000000f000004
||RH: RQ, FMD, FI, OIC, RQE1, BB, CEBI
||FMH-5
||   Command code = Attach
||   User ID already verified = No
||   Password is substituted = No
||   PIP present = No
||   Conversation type = Basic
||   Synchronization level = None
||   Transaction program name = ?006 (APPN Receive_Encap_Msg)
||   Logical unit of work identifier:
||      LU name = USIBMRA.CM5HPRNN
||      Instance number = 0xafff44aea03a
||      Sequence number = 0x0001
||   Conversation correlator = 0xd02b1f93b3f78625
||FID2 Encapsulation Variable
||   TH: FID2, OIS, LFSID=0x00002, SNF=0x0000
||   RH: RQ, FMD, FI, OIC, RQN
||   SESSEND rq
||      Format = 3
||   Fully qualified PCID control vector:
||    (0x0002) PCID = 0xc96f2a091956a273
||    (0x000B) Network qualified CP name = USIBMRA.CM5HPRNN
```

Next, a Route Setup is received for the new session with RAAAN007. The previous RTP connection was closed when the last (and only) session on it ended, so we need a new one. Note that this would not happen with TSO which does the CLSDST PASS before the BIND is sent.

```
Line:    951 Recv MU
Time stamp: 08:37:40.77
DLC type: HPR
TCID: 0000007A
Transmission priority: Network
||TH: FID2, Exp, OIS, LFSID=0x00000, SNF=0x0000
||RH: RQ, NC, FI, OIC, RQN
||Route setup
||   Type = Request
||   Route Setup triggered by path switch = No
||   Destination hop index = 2
||   Locate search is required = No
||   Destination is mobile = No
||   NCE is used for all LUs/BFs = No
||   Path switch time (milliseconds) = 0
||   Network name control vector:
||    (0x0002) Network name type = LU
||    (0x0003) Name = USIBMRA.LU41A1
||   Route selection control vector:
||    (0x0002) Maximum hop count = 2
||    (0x0003) Current hop count = 2
```

```
      TG descriptor control vector
         TG Identifier TG Descriptor subfield
          (0x0002) TG number = 21
          (0x0004) Partner name = USIBMRA.NNP61A
          (0x0012) Link connection network = No
                   Additional configuration information = No
                   HPR = Supported
                   TG type = Boundary Function based or APPN
                   Intersubnet link = No
                   Extended border node = Supported
                   RTP Tower = Supported
          (0x0013) Subarea number = 0x8000000a
      TG descriptor control vector
         TG Identifier TG Descriptor subfield
          (0x0002) TG number = 26
          (0x0004) Partner name = CM5HPRNN
          (0x000C) Link connection network = No
                   Additional configuration information = No
                   HPR = Supported
                   TG type = Boundary Function based or APPN
                   Intersubnet link = No
                   Extended border node = Not supported
                   RTP Tower = Supported
   COS/TPF control vector:
    (0x0002) Transmission priority = Medium
    (0x0004) COS name = #CONNECT
   Fully qualified PCID control vector:
    (0x0002) PCID = 0xf7ff6164529f6e02
    (0x000B) Network qualified CP name = USIBMRA.RAA
   Route Information control vector
    (0x0002) Route direction = Forward
             REFIFOing required = No
    (0x0004) Maximum packet size = 2058
    (0x0008) Accumulated transmission time (microseconds) = 108
    (0x000C) Minimum link capacity (Kbits/second) = 15974
    (0x0010) Limited resource liveness timer (seconds) = 0
      ANR Path control vector
       (0x0003) ANR label represents a subarea network route = No
       (0x0004) ANR label = 803000c300000000
       (0x000D) ANR label represents a subarea network route = No
       (0x000E) ANR label = a320304c82

Line:    964 Send MU
Time stamp: 08:37:40.77
DLC type: HPR
TCID: 0000007A
Transmission priority: Network
TH: FID2, Exp, OIS, LFSID=0x00000, SNF=0x0000
RH: RQ, NC, FI, OIC, RQN
Route setup
   Type = Reply
   Destination hop index = 2
   Locate search is required = No
   Destination is mobile = No
   NCE is used for all LUs/BFs = Yes
   Path switch time (milliseconds) = 240000
   COS/TPF control vector:
    (0x0002) Transmission priority = Medium
    (0x0004) COS name = #CONNECT
```

```
    Fully qualified PCID control vector:
     (0x0002) PCID = 0xf7ff6164529f6e02
     (0x000B) Network qualified CP name = USIBMRA.RAA
    Route Information control vector
     (0x0002) Route direction = Forward
              REFIFOing required = No
     (0x0004) Maximum packet size = 2058
     (0x0008) Accumulated transmission time (microseconds) = 108
     (0x000C) Minimum link capacity (Kbits/second) = 15974
     (0x0010) Limited resource liveness timer (seconds) = 0
       ANR Path control vector
        (0x0003) ANR label represents a subarea network route = No
        (0x0004) ANR label = 803000c300000000
        (0x000D) ANR label represents a subarea network route = No
        (0x000E) ANR label = a320304c82
    Network name control vector:
     (0x0002) Network name type = CP
     (0x0003) Name = USIBMRA.CM5HPRNN
    NCE Identifier control vector
     (0x0002) NCE identifier = 80
    NCE Instance Identifier control vector
     (0x0002) NCE instance identifer = 0x34e1a64d
    Route selection control vector:
     (0x0002) Maximum hop count = 1
     (0x0003) Current hop count = 0
       TG descriptor control vector
          TG Identifier TG Descriptor subfield
           (0x0002) TG number = 26
           (0x0004) Partner name = USIBMRA.NNP61A
           (0x0012) Link connection network = No
                    Additional configuration information = No
                    HPR = Supported
                    TG type = Boundary Function based or APPN
                    Intersubnet link = No
                    Extended border node = Not supported
                    RTP Tower = Supported
    Route Information control vector
     (0x0002) Route direction = Reverse
              REFIFOing required = No
     (0x0004) Maximum packet size = 2058
     (0x0008) Accumulated transmission time (microseconds) = 75
     (0x000C) Minimum link capacity (Kbits/second) = 16000
     (0x0010) Limited resource liveness timer (seconds) = 0
       ANR Path control vector
        (0x0003) ANR label represents a subarea network route = No
        (0x0004) ANR label = 801c
```

Again a BIND, this time for the RAAAN007 session.

```
Line:    978 Recv MU
Time stamp: 08:37:40.93
DLC type: HPR
TH: FID5, Exp, OIS, SNF=0x19bd, S SA000000001000000a
RH: RQ, SC, FI, OIC, RQD1
BIND rq
 (0x0001) Type = Non-negotiable
 (0x0002) FM profile = 3
```

```
(0x0003) TS profile = 3
         FM usage - primary:
(0x0004)    Chaining use = Multiple-RU chains allowed
            Request control mode = Immediate request mode
            Chain response protocol = Definite or exception response
            Two-phase commit = Not supported
            Compression = Will not be used
            Send end bracket = May send
         FM usage - secondary:
(0x0005)    Chaining use = Multiple-RU chains allowed
            Request control mode = Immediate request mode
            Chain response protocol = Exception response
            Two-phase commit = Not supported
            Compression = Will not be used
            Send end bracket = Will not send
         FM Usage - common:
(0x0006)    Whole BIUs required = No
            FM header usage = Not allowed
            Brackets are used = Yes
            Bracket reset state = BETB
            Bracket termination rule = Conditional
            Alternate code set allowed = No
            BIND queueing allowed = No
(0x0007)    Normal-flow send/receive mode = Half-duplex flip-flop
            Recovery responsibility = Contention loser
            Contention winner = Secondary
            Alternate code set = ASCII-7
            Control vectors included = Yes
            Half-duplex flip-flop primary reset state = Receive
         TS usage:
(0x0008)    Secondary to primary pacing stages = One
            Secondary send window size = 0
(0x0009)    Adaptive pacing = Supported
            Secondary receive window size = 0
(0x000A)    Secondary maximum send RU size = 1024
(0x000B)    Primary maximum send RU size = 3840
(0x000C)    Primary to secondary pacing stages = One
            Primary send window size = 0
(0x000D)    Primary receive window size = 0
         PS profile:
(0x000E)    LU type = 2
(0x000F)    Query = Supported
(0x0014)    Default screen size = 0x0
(0x0016)    Alternate screen size = 0x0
(0x0018)    Screen size = Not specified
            Length-checked compression options = No Compression
         Cryptography options:
(0x001A)    Private cryptography support = No
            Session-level cryptography support = No
(0x001C) Primary LU name = USIBMRA.RAAAN007
         User Data:
(0x002F) Secondary LU name = USIBMRA.LU41A1
(0x003D) Fully qualified PCID control vector:
          (0x0002) PCID = 0xf7ff6164529f6e01
          (0x000B) Network qualified CP name = USIBMRA.RAA
(0x0053) Route selection control vector:
          (0x0002) Maximum hop count = 2
          (0x0003) Current hop count = 2
(0x0057)    TG descriptor control vector
```

```
(0x0059)       TG Identifier TG Descriptor subfield
                (0x0002) TG number = 21
                (0x0004) Partner name = USIBMRA.NNP61A
                (0x0012) Link connection network = No
                         Additional configuration information = No
                         HPR = Supported
                         TG type = Boundary Function based or APPN
                         Intersubnet link = No
                         Extended border node = Supported
                         RTP Tower = Supported
                (0x0013) Subarea number = 0x8000000a
(0x0070)    TG descriptor control vector
(0x0072)       TG Identifier TG Descriptor subfield
                (0x0002) TG number = 26
                (0x0004) Partner name = CM5HPRNN
                (0x000C) Link connection network = No
                         Additional configuration information = No
                         HPR = Supported
                         TG type = Boundary Function based or APPN
                         Intersubnet link = No
                         Extended border node = Not supported
                         RTP Tower = Supported
(0x007F) COS/TPF control vector:
          (0x0002) Transmission priority = Medium
          (0x0004) COS name = #CONNECT
(0x008B) Mode control vector:
          (0x0003) Mode name = D4C32XX3
         Unknown control vector key = 0x5f
         Hex dump:
            5f1600f7 0794547c 2c22210c e4e2c9c2   <_.....T|,"!.....>  <[..7.m.
            d4d9c14b d9c1f3f9                     <...K....         > <MRA.RA3
Line:    992 Send MU
Time stamp: 08:37:40.94
DLC type: HPR
TH: FID5, Exp, OIS, SNF=0x19bd, R SA000000001000000a
RH: +RSP, SC, FI, RQD1
BIND +rsp
(0x0001) Type = Non-negotiable
(0x0002) FM profile = 3
(0x0003) TS profile = 3
         FM usage - primary:
(0x0004)    Chaining use = Multiple-RU chains allowed
            Request control mode = Immediate request mode
            Chain response protocol = Definite or exception response
            Two-phase commit = Not supported
            Compression = Will not be used
            Send end bracket = May send
         FM usage - secondary:
(0x0005)    Chaining use = Multiple-RU chains allowed
            Request control mode = Immediate request mode
            Chain response protocol = Exception response
            Two-phase commit = Not supported
            Compression = Will not be used
            Send end bracket = Will not send
         FM Usage - common:
(0x0006)    Whole BIUs required = No
            FM header usage = Not allowed
            Brackets are used = Yes
            Bracket reset state = BETB
```

```
                Bracket termination rule = Conditional
                Alternate code set allowed = No
                BIND queueing allowed = No
(0x0007)        Normal-flow send/receive mode = Half-duplex flip-flop
                Recovery responsibility = Contention loser
                Contention winner = Secondary
                Alternate code set = ASCII-7
                Control vectors included = Yes
                Half-duplex flip-flop primary reset state = Receive
             TS usage:
(0x0008)        Secondary to primary pacing stages = One
                Secondary send window size = 0
(0x0009)        Adaptive pacing = Not supported
                Secondary receive window size = 0
(0x000A)        Secondary maximum send RU size = 1024
(0x000B)        Primary maximum send RU size = 3840
(0x000C)        Primary to secondary pacing stages = One
                Primary send window size = 0
(0x000D)        Primary receive window size = 0
             PS profile:
(0x000E)        LU type = 2
(0x000F)        Query = Supported
(0x0014)        Default screen size = 0x0
(0x0016)        Alternate screen size = 0x0
(0x0018)        Screen size = Not specified
                Length-checked compression options = No Compression
             Cryptography options:
(0x001A)        Private cryptography support = No
                Session-level cryptography support = No
             User Data:
(0x001F) Fully qualified PCID control vector:
              (0x0002) PCID = 0xf7ff6164529f6e01
              (0x000B) Network qualified CP name = USIBMRA.RAA
(0x0035) Session Address control vector
              (0x0002) Session address assignor = Sender
                       Session address = 0x2586f7b5448d1f93
(0x003F) Route selection control vector:
              (0x0002) Maximum hop count = 2
              (0x0003) Current hop count = 2
(0x0043)        TG descriptor control vector
(0x0045)           TG Identifier TG Descriptor subfield
                    (0x0002) TG number = 21
                    (0x0004) Partner name = USIBMRA.NNP61A
                    (0x0012) Link connection network = No
                             Additional configuration information = No
                             HPR = Supported
                             TG type = Boundary Function based or APPN
                             Intersubnet link = No
                             Extended border node = Supported
                             RTP Tower = Supported
                    (0x0013) Subarea number = 0x8000000a
(0x005C)        TG descriptor control vector
(0x005E)           TG Identifier TG Descriptor subfield
                    (0x0002) TG number = 26
                    (0x0004) Partner name = CM5HPRNN
                    (0x000C) Link connection network = No
                             Additional configuration information = No
                             HPR = Supported
                             TG type = Boundary Function based or APPN
```

```
||                                Intersubnet link = No
||                                Extended border node = Not supported
||                                RTP Tower = Supported
```

Another SESSST RU to the owning SSCP (RA39).

```
Line:   1002 Send MU
Time stamp: 08:37:41.00
DLC type: HPR
||TH: FID5, OIS, SNF=0x000c, R SA000000000f000004
||RH: RQ, FMD, FI, OIC, RQE1, BB, CEBI
||FMH-5
||   Command code = Attach
||   User ID already verified = No
||   Password is substituted = No
||   PIP present = No
||   Conversation type = Basic
||   Synchronization level = None
||   Transaction program name = ?006 (APPN Receive_Encap_Msg)
||   Logical unit of work identifier:
||      LU name = USIBMRA.CM5HPRNN
||      Instance number = 0xafff44aea03a
||      Sequence number = 0x0001
||   Conversation correlator = 0xd02b1f93b6f78625
||FID2 Encapsulation Variable
||   TH: FID2, OIS, LFSID=0x00002, SNF=0x0000
||   RH: RQ, FMD, FI, OIC, RQN
||   SESSST rq
||      Format = 3
||      Unknown session key = 0x2a
||   Fully qualified PCID control vector:
||    (0x0002) PCID = 0xc96f2a091956a273
||    (0x000B) Network qualified CP name = USIBMRA.CM5HPRNN
```

Now the session begins with SDT (start data traffic), followed by some 3270 data.

```
Line:   1018 Recv MU
Time stamp: 08:37:41.03
DLC type: HPR
||TH: FID5, Exp, OIS, SNF=0x19be, R SA2586f7b5448d1f93
||RH: RQ, SC, FI, OIC, RQD1
||SDT rq
Line:   1022 Send MU
Time stamp: 08:37:41.04
DLC type: HPR
||TH: FID5, Exp, OIS, SNF=0x19be, R SA000000001000000a
||RH: +RSP, SC, FI, RQD1
||SDT +rsp


Line:   1026 Recv MU
Time stamp: 08:37:41.06
DLC type: HPR
```

```
||TH: FID5, OIS, SNF=0x0001, R SA2586f7b5448d1f93
||RH: RQ, FMD, OIC, RQD1, BB, CD
||User data - remainder of RU follows:
||   Hex dump:
||      f3000501 ff02                          <......          > <3.....
Line:   1030 Send MU
Time stamp: 08:37:41.06
DLC type: HPR
||TH: FID5, OIS, SNF=0x0001, R SA000000001000000a
||RH: +RSP, FMD, RQD1

Line:   1033 Send MU
Time stamp: 08:37:41.11
DLC type: HPR
||TH: FID5, OIS, SNF=0x0001, R SA000000001000000a
||RH: RQ, FMD, OIC, RQE1, CD
||User data - remainder of RU follows:
||   Hex dump:
||      88001781 81010000 50001801 011b0400   <.......P.......> <h..aa...&.....

Line:   1054 Recv MU
Time stamp: 08:37:41.15
DLC type: HPR
||TH: FID5, OIS, SNF=0x0002, R SA2586f7b5448d1f93
||RH: RQ, FMD, OIC, RQD1, EB
||User data - remainder of RU follows:
||   Hex dump:
||      f5c61140 c31df0d5 d5404040 40d5d511   <...@.....@@@@...> <5F. C.ONN    N
```

B.2.1 Communications Server/2 Displays

A dependent LU session is now in progress, so we do some more displays on the CS/2 node to see the status of the network connections. First, we look at the HPR connections (Figure 253).

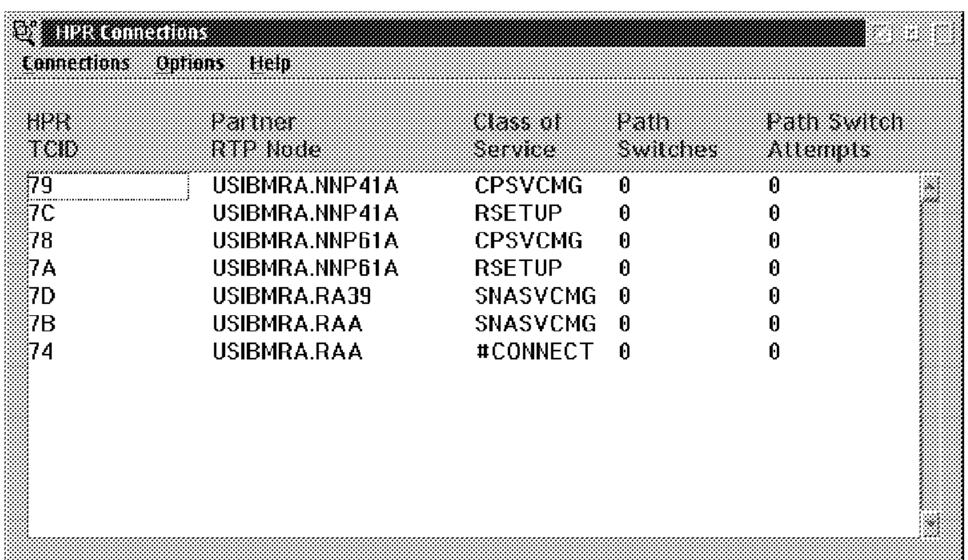

Figure 253. HPR Connections

A new HPR connection is now present with the TCID 74, the COS used being #CONNECT as we saw in the Route Setup flows. The details of this connection are shown in Figure 254 on page 299.

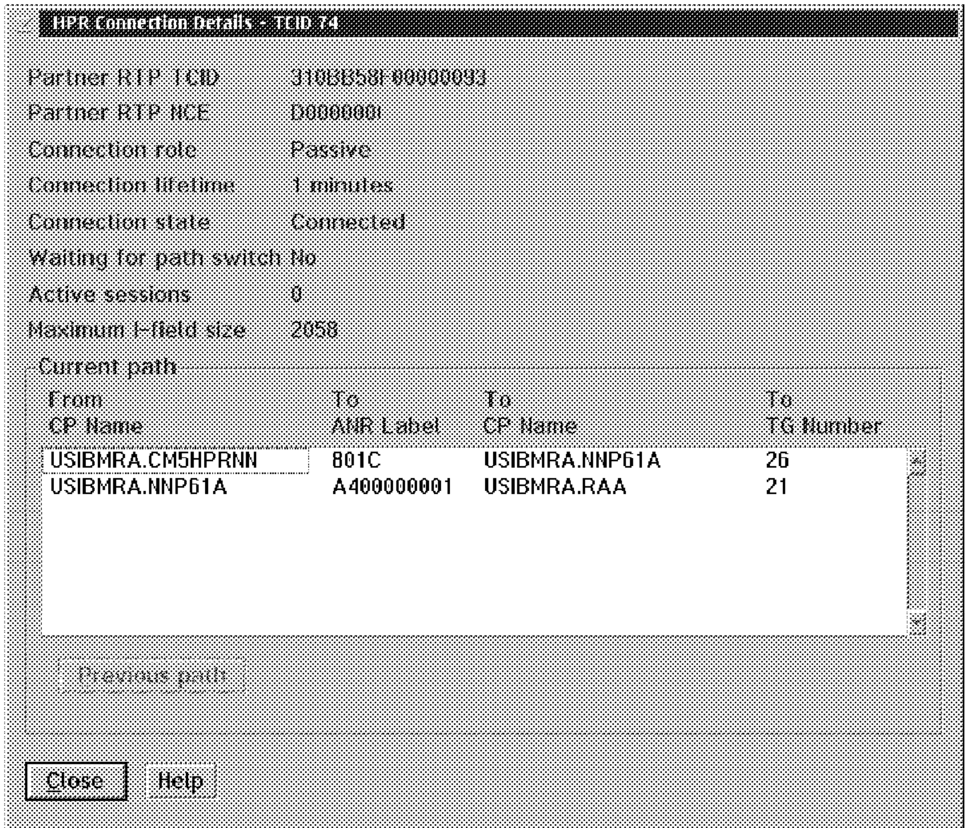

Figure 254. TCID 74

The path used for this connection is:

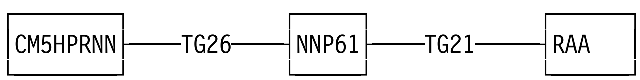

B.3 Path Switch

Now we break LINK0001, which is the connection to NNP61A being used for our dependent LU session (among others). Figure 255 on page 300 shows CS/2′s view of the connections now.

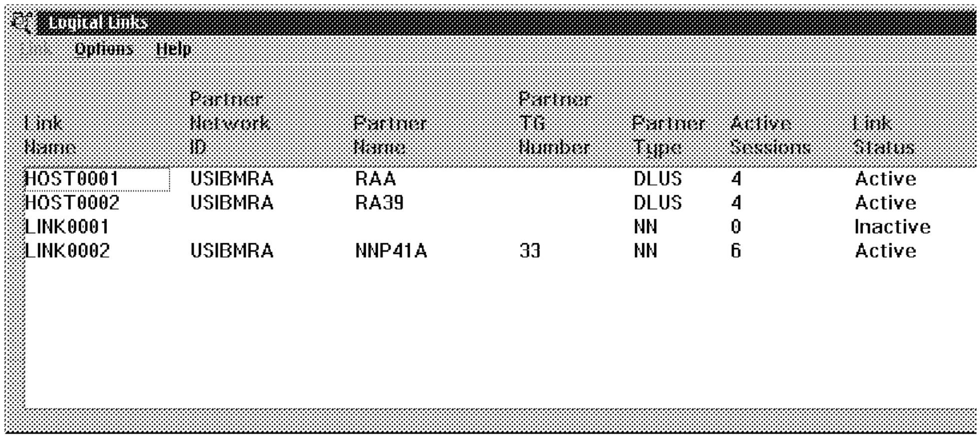

Link Name	Partner Network ID	Partner Name	Partner TG Number	Partner Type	Active Sessions	Link Status
HOST0001	USIBMRA	RAA		DLUS	4	Active
HOST0002	USIBMRA	RA39		DLUS	4	Active
LINK0001				NN	0	Inactive
LINK0002	USIBMRA	NNP41A	33	NN	6	Active

Figure 255. LINK0001 Inactivated

Clearly a path switch must occur, and the first node to notice the problem (CM5HPRNN, because an adjacent link failed) will initiate it. Before it can do the path switch it must calculate a new route (being a network node) and send a Route Setup message to obtain the ANR labels and ARB information for the new route. Because both OLU and DLU are on network nodes, CM5HPRNN can do this without sending any Locates into the network.

B.3.1 Route Setup

As the result of the LINK0001 inactivation, a Route Setup triggered by the path switch function is forwarded to RAA.

```
Line:    126 Send MU
Time stamp: 08:42:58.89
DLC type: HPR
TCID: 0000007C
Transmission priority: Network
||TH: FID2, Exp, OIS, LFSID=0x00000, SNF=0x0000
||RH: RQ, NC, FI, OIC, RQN
||Route setup
||   Type = Request
||   Route Setup triggered by path switch = Yes
||   Destination hop index = 3
||   Locate search is required = No
||   Destination is mobile = No
||   NCE is used for all LUs/BFs = No
||   Path switch time (milliseconds) = 0
||   NCE Identifier control vector
||    (0x0002) NCE identifier = d000000000000000
||   Route selection control vector:
||    (0x0002) Maximum hop count = 3
||    (0x0003) Current hop count = 1
||      TG descriptor control vector
||         TG Identifier TG Descriptor subfield
||          (0x0002) TG number = 33
||          (0x0004) Partner name = USIBMRA.NNP41A
||          (0x0012) Link connection network = No
||                   Additional configuration information = No
||                   HPR = Supported
||                   TG type = Boundary Function based or APPN
||                   Intersubnet link = No
||                   Extended border node = Not supported
```

```
                  RTP Tower = Supported
      TG descriptor control vector
         TG Identifier TG Descriptor subfield
          (0x0002) TG number = 21
          (0x0004) Partner name = RA39
          (0x0008) Link connection network = No
                   Additional configuration information = No
                   HPR = Supported
                   TG type = Boundary Function based or APPN
                   Intersubnet link = No
                   Extended border node = Not supported
                   RTP Tower = Supported
      TG descriptor control vector
         TG Identifier TG Descriptor subfield
          (0x0002) TG number = 21
          (0x0004) Partner name = RAA
          (0x0007) Link connection network = No
                   Additional configuration information = No
                   HPR = Supported
                   TG type = Boundary Function based or APPN
                   Intersubnet link = No
                   Extended border node = Supported
                   RTP Tower = Supported
          (0x0008) Subarea number = 0x80000027
   Fully qualified PCID control vector:
    (0x0002) PCID = 0xe887a74909bc569a
    (0x000B) Network qualified CP name = USIBMRA.CM5HPRNN
   COS/TPF control vector:
    (0x0002) Transmission priority = Medium
    (0x0004) COS name = #CONNECT
   Route Information control vector
    (0x0002) Route direction = Forward
             REFIFOing required = No
    (0x0004) Maximum packet size = 2058
    (0x0008) Accumulated transmission time (microseconds) = 75
    (0x000C) Minimum link capacity (Kbits/second) = 16000
    (0x0010) Limited resource liveness timer (seconds) = 0
      ANR Path control vector
       (0x0003) ANR label represents a subarea network route = No
       (0x0004) ANR label = 801e
```

The new path for this RTP pipe is:

```
CM5HPRNN ——TG33—— NNP41A ——TG21—— RA39 ——TG21—— RAA
```

Another path switch occurs for the DLUR/S session with RAA, which was also on the failing link. Here is the new Route Setup.

```
Line:    138 Send MU
Time stamp: 08:42:58.90
DLC type: HPR
TCID: 0000007C
Transmission priority: Network
TH: FID2, Exp, OIS, LFSID=0x00000, SNF=0x0000
RH: RQ, NC, FI, OIC, RQN
Route setup
   Type = Request
   Route Setup triggered by path switch = Yes
   Destination hop index = 3
   Locate search is required = No
   Destination is mobile = No
   NCE is used for all LUs/BFs = No
   Path switch time (milliseconds) = 0
   NCE Identifier control vector
    (0x0002) NCE identifier = d000000000000000
   Route selection control vector:
    (0x0002) Maximum hop count = 3
    (0x0003) Current hop count = 1
      TG descriptor control vector
         TG Identifier TG Descriptor subfield
          (0x0002) TG number = 33
          (0x0004) Partner name = USIBMRA.NNP41A
          (0x0012) Link connection network = No
                   Additional configuration information = No
                   HPR = Supported
                   TG type = Boundary Function based or APPN
                   Intersubnet link = No
                   Extended border node = Not supported
                   RTP Tower = Supported
      TG descriptor control vector
         TG Identifier TG Descriptor subfield
          (0x0002) TG number = 21
          (0x0004) Partner name = RA39
          (0x0008) Link connection network = No
                   Additional configuration information = No
                   HPR = Supported
                   TG type = Boundary Function based or APPN
                   Intersubnet link = No
                   Extended border node = Not supported
                   RTP Tower = Supported
      TG descriptor control vector
         TG Identifier TG Descriptor subfield
          (0x0002) TG number = 21
          (0x0004) Partner name = RAA
          (0x0007) Link connection network = No
                   Additional configuration information = No
                   HPR = Supported
                   TG type = Boundary Function based or APPN
                   Intersubnet link = No
                   Extended border node = Supported
                   RTP Tower = Supported
          (0x0008) Subarea number = 0x80000027
   Fully qualified PCID control vector:
    (0x0002) PCID = 0xe887a7490abc569a
    (0x000B) Network qualified CP name = USIBMRA.CM5HPRNN
   COS/TPF control vector:
    (0x0002) Transmission priority = Network
```

```
    (0x0004) COS Name = SNASVCMG
   Route Information control vector
    (0x0002) Route direction = Forward
             REFIFOing required = No
    (0x0004) Maximum packet size = 2058
    (0x0008) Accumulated transmission time (microseconds) = 75
    (0x000C) Minimum link capacity (Kbits/second) = 16000
    (0x0010) Limited resource liveness timer (seconds) = 0
      ANR Path control vector
       (0x0003) ANR label represents a subarea network route = No
       (0x0004) ANR label = 801e
```

The Route Setup replies are now received by CS/2.

```
Line:    150 Recv MU
Time stamp: 08:42:59.13
DLC type: HPR
TCID: 0000007C
Transmission priority: Network
  TH: FID2, Exp, OIS, LFSID=0x00000, SNF=0x0000
  RH: RQ, NC, FI, OIC, RQN
  Route setup
   Type = Reply
   Destination hop index = 3
   Locate search is required = Yes
   Destination is mobile = No
   NCE is used for all LUs/BFs = No
   Path switch time (milliseconds) = 240000
   Network name control vector:
    (0x0002) Network name type = CP
    (0x0003) Name = USIBMRA.RAA
   NCE Identifier control vector
    (0x0002) NCE identifier = d000000000000000
   Route selection control vector:
    (0x0002) Maximum hop count = 3
    (0x0003) Current hop count = 0
      TG descriptor control vector
         TG Identifier TG Descriptor subfield
          (0x0002) TG number = 21
          (0x0004) Partner name = USIBMRA.RA39
          (0x0010) Link connection network = No
                   Additional configuration information = No
                   HPR = Supported
                   TG type = Boundary Function based or APPN
                   Intersubnet link = No
                   Extended border node = Supported
                   RTP Tower = Supported
          (0x0011) Subarea number = 0x8000000a
      TG descriptor control vector
         TG Identifier TG Descriptor subfield
          (0x0002) TG number = 21
          (0x0004) Partner name = NNP41A
          (0x000A) Link connection network = No
                   Additional configuration information = No
                   HPR = Supported
                   TG type = Boundary Function based or APPN
                   Intersubnet link = No
```

```
                  Extended border node = Supported
                  RTP Tower = Supported
         (0x000B) Subarea number = 0x80000027
     TG descriptor control vector
        TG Identifier TG Descriptor subfield
         (0x0002) TG number = 33
         (0x0004) Partner name = CM5HPRNN
         (0x000C) Link connection network = No
                  Additional configuration information = No
                  HPR = Supported
                  TG type = Boundary Function based or APPN
                  Intersubnet link = No
                  Extended border node = Not supported
                  RTP Tower = Supported
  COS/TPF control vector:
   (0x0002) Transmission priority = Medium
   (0x0004) COS name = #CONNECT
  Fully qualified PCID control vector:
   (0x0002) PCID = 0xe887a74909bc569a
   (0x000B) Network qualified CP name = USIBMRA.CM5HPRNN
  Route Information control vector
   (0x0002) Route direction = Forward
            REFIFOing required = No
   (0x0004) Maximum packet size = 2058
   (0x0008) Accumulated transmission time (microseconds) = 119
   (0x000C) Minimum link capacity (Kbits/second) = 16000
   (0x0010) Limited resource liveness timer (seconds) = 0
     ANR Path control vector
      (0x0003) ANR label represents a subarea network route = No
      (0x0004) ANR label = 801e
      (0x0007) ANR label represents a subarea network route = No
      (0x0008) ANR label = a500000002
      (0x000E) ANR label represents a subarea network route = No
      (0x000F) ANR label = 8006009c00000000
  Route Information control vector
   (0x0002) Route direction = Reverse
            REFIFOing required = No
   (0x0004) Maximum packet size = 2058
   (0x0008) Accumulated transmission time (microseconds) = 144
   (0x000C) Minimum link capacity (Kbits/second) = 15974
   (0x0010) Limited resource liveness timer (seconds) = 0
     ANR Path control vector
      (0x0003) ANR label represents a subarea network route = No
      (0x0004) ANR label = 802a00bb00000000
      (0x000D) ANR label represents a subarea network route = No
      (0x000E) ANR label = 800900b100000000
      (0x0017) ANR label represents a subarea network route = No
      (0x0018) ANR label = a400304481

Line:    168 Recv MU
Time stamp: 08:42:59.14
DLC type: HPR
TCID: 0000007C
Transmission priority: Network
TH: FID2, Exp, OIS, LFSID=0x00000, SNF=0x0000
RH: RQ, NC, FI, OIC, RQN
Route setup
  Type = Reply
  Destination hop index = 3
```

```
Locate search is required = Yes
Destination is mobile = No
NCE is used for all LUs/BFs = No
Path switch time (milliseconds) = 480000
Network name control vector:
 (0x0002) Network name type = CP
 (0x0003) Name = USIBMRA.RAA
NCE Identifier control vector
 (0x0002) NCE identifier = d000000000000000
Route selection control vector:
 (0x0002) Maximum hop count = 3
 (0x0003) Current hop count = 0
   TG descriptor control vector
      TG Identifier TG Descriptor subfield
       (0x0002) TG number = 21
       (0x0004) Partner name = USIBMRA.RA39
       (0x0010) Link connection network = No
                Additional configuration information = No
                HPR = Supported
                TG type = Boundary Function based or APPN
                Intersubnet link = No
                Extended border node = Supported
                RTP Tower = Supported
       (0x0011) Subarea number = 0x8000000a
   TG descriptor control vector
      TG Identifier TG Descriptor subfield
       (0x0002) TG number = 21
       (0x0004) Partner name = NNP41A
       (0x000A) Link connection network = No
                Additional configuration information = No
                HPR = Supported
                TG type = Boundary Function based or APPN
                Intersubnet link = No
                Extended border node = Supported
                RTP Tower = Supported
       (0x000B) Subarea number = 0x80000027
   TG descriptor control vector
      TG Identifier TG Descriptor subfield
       (0x0002) TG number = 33
       (0x0004) Partner name = CM5HPRNN
       (0x000C) Link connection network = No
                Additional configuration information = No
                HPR = Supported
                TG type = Boundary Function based or APPN
                Intersubnet link = No
                Extended border node = Not supported
                RTP Tower = Supported
COS/TPF control vector:
 (0x0002) Transmission priority = Low
 (0x0004) COS name = SNASVCMG
Fully qualified PCID control vector:
 (0x0002) PCID = 0xe887a7490abc569a
 (0x000B) Network qualified CP name = USIBMRA.CM5HPRNN
Route Information control vector
 (0x0002) Route direction = Forward
          REFIFOing required = No
 (0x0004) Maximum packet size = 2058
 (0x0008) Accumulated transmission time (microseconds) = 119
 (0x000C) Minimum link capacity (Kbits/second) = 16000
```

```
||    (0x0010) Limited resource liveness timer (seconds) = 0
||      ANR Path control vector
||       (0x0003) ANR label represents a subarea network route = No
||       (0x0004) ANR label = 801e
||       (0x0007) ANR label represents a subarea network route = No
||       (0x0008) ANR label = a500000002
||       (0x000E) ANR label represents a subarea network route = No
||       (0x000F) ANR label = 8006009c00000000
||   Route Information control vector
||    (0x0002) Route direction = Reverse
||             REFIFOing required = No
||    (0x0004) Maximum packet size = 2058
||    (0x0008) Accumulated transmission time (microseconds) = 144
||    (0x000C) Minimum link capacity (Kbits/second) = 15974
||    (0x0010) Limited resource liveness timer (seconds) = 0
||      ANR Path control vector
||       (0x0003) ANR label represents a subarea network route = No
||       (0x0004) ANR label = 802a00bb00000000
||       (0x000D) ANR label represents a subarea network route = No
||       (0x000E) ANR label = 800900b100000000
||       (0x0017) ANR label represents a subarea network route = No
||       (0x0018) ANR label = a400304481
```

During this time we are still using the 3270 session. We press the Enter key and the application responds.

```
Line:    186 Send MU
Time stamp: 08:43:02.75
DLC type: HPR
||TH: FID5, OIS, SNF=0x0002, R SA000000001000000a
||RH: RQ, FMD, OIC, RQE1, BB, CD
||User data - remainder of RU follows:
||   Hex dump:
||      7dd17c11 d17c4040 40404040 404011d3   <}.|..|@@@@@@@@..> <'J@.J@

Line:    193 Recv MU
Time stamp: 08:43:02.79
DLC type: HPR
||TH: FID5, OIS, SNF=0x0003, R SA2586f7b5448d1f93
||RH: RQ, FMD, OIC, RQD1, EB
||User data - remainder of RU follows:
||   Hex dump:
||      f5c61140 c31df0d5 d5404040 40d5d511   <...@.....@@@@...> <5F. C.ONN    N

Line:    228 Send MU
Time stamp: 08:43:02.80
DLC type: HPR
||TH: FID5, OIS, SNF=0x0003, R SA000000001000000a
||RH: +RSP, FMD, RQD1
```

A TDU is now sent to inform other nodes that TG26 from CM5HPRNN to NNP61A is inoperative.

```
Line:    273 Send MU
Time stamp: 08:43:06.08
DLC type: HPR
||TH: FID5, OIS, SNF=0x0097, R SA058bb89800201015
||RH: RQ, FMD, FI, OIC, RQE1, BB, CEBI
||FMH-5
||   Command code = Attach
||   User ID already verified = No
||   Password is substituted = No
||   PIP present = No
||   Conversation type = Basic
||   Synchronization level = None
||   Transaction program name = ?004 (APPN Topology database update)
||   Logical unit of work identifier:
||      LU name = USIBMRA.CM5HPRNN
||      Instance number = 0xafff449a2dda
||      Sequence number = 0x0001
||   Conversation correlator = 0xd02b1f93abf88625
||Topology database update
||   Flow-Reduction Sequence Numbers TDU control vector
||    (0x0002) Current FRSN = 0x00000149
||    (0x0006) Last FRSN sent = 0x00000147
||   Node Descriptor TDU control vector
||    (0x0003) Network qualified CP name = USIBMRA.CM5HPRNN
||    (0x0014) CP name identifies a connection network = No
||   TG descriptor control vector
||      TG Identifier TG Descriptor subfield
||       (0x0002) TG number = 26
||       (0x0004) Partner name = USIBMRA.NNP61A
||       (0x0012) Link connection network = No
||                Additional configuration information = No
||                HPR = Not supported
||                TG type = Boundary Function based or APPN
||                Intersubnet link = No
||                Extended border node = Not supported
||                RTP Tower = Not supported
||   TG Characteristics control vector
||    (0x0002) Sequence number = 0x0000000e
||    (0x0006) Operational = No
||             TG will be garbage collected next cycle = No
||             Quiescing = No
||             CP-CP sessions supported = Supported but not active
||    (0x0007) Effective capacity = 0x85 (15.97 Mbps)
||    (0x000D) Connect cost = 0x00
||    (0x000E) Byte cost = 0x00
||    (0x0010) Security = Non-secure
||    (0x0011) Propagation delay = LAN 0x4c (384.00 us)
||    (0x0012) Modem class = 0x00
||    (0x0013) User defined 1 = 0x80
||    (0x0014) User defined 2 = 0x80
||    (0x0015) User defined 3 = 0x80
```

```
Line:    284 Send MU
Time stamp: 08:43:08.85
DLC type: HPR
||TH: FID5, OIS, SNF=0x0003, R SA000000001000000a
||RH: RQ, FMD, OIC, RQE1, BB, CD
||User data - remainder of RU follows:
||   Hex dump:
||      7dd17c11 d17c4040 40404040 404011d3   <}.|..|@@@@@@@@..> <´J@.J@
||      4c11d45c 40404040 40404040 11d56c40   <L..\@@@@@@@@..l@> <<.M*        .N
||      40404040 40404011 d67c4040 40404040   <@@@@@@@..|@@@@@@> <       .O@
||      4040                                  <@@              > <
```

B.3.2 Communications Server/2 Displays

Again we use the CS/2 Subsystem Management to display the new paths. We have already seen the logical links in Figure 255 on page 300. Figure 256 shows the HPR connections, which are all there except the CP-CP and Route Setup pipes to NNP61A.

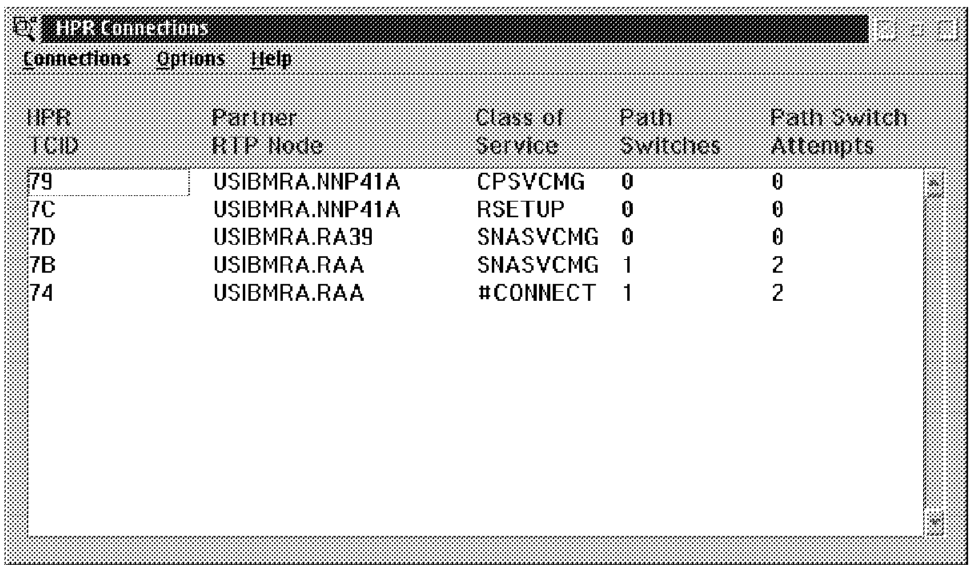

Figure 256. HPR Connection Summary

TCID 7B is the DLUR/S pipe to RAA, so we display its details in Figure 257 on page 309.

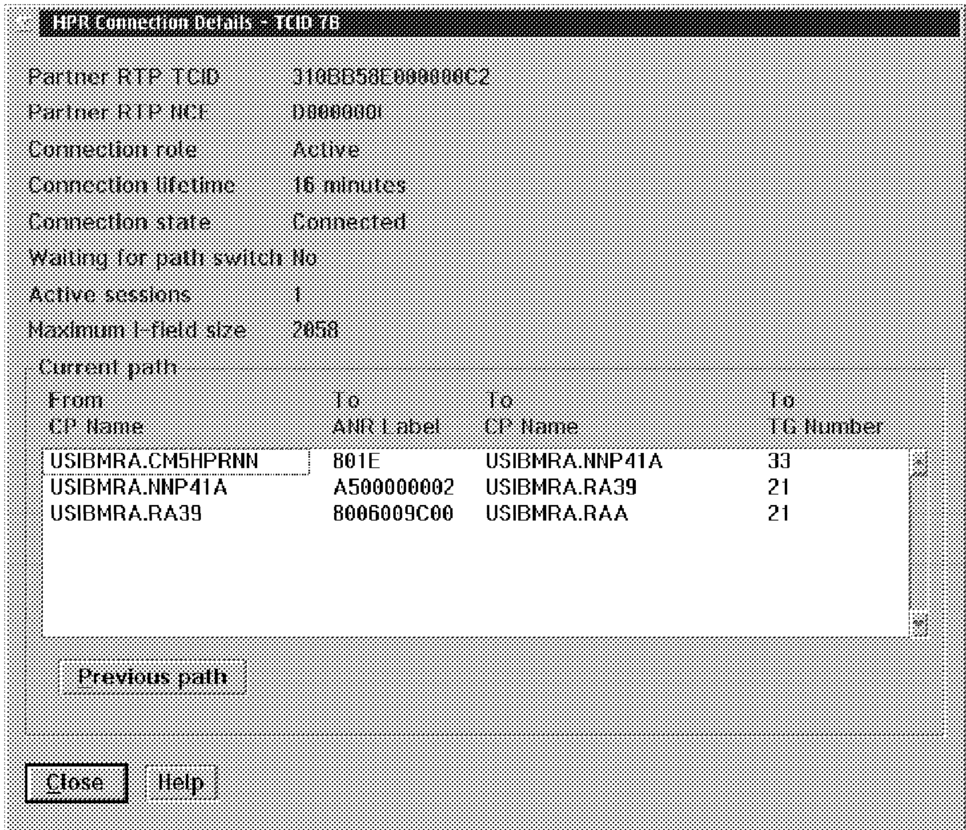

Figure 257. New Path for TCID 7B

The new path is:

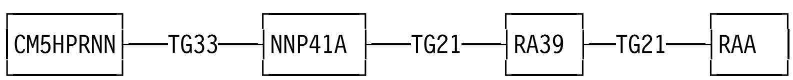

The Previous path button allows us to view the old route, so without hesitation we use it to see Figure 258 on page 310.

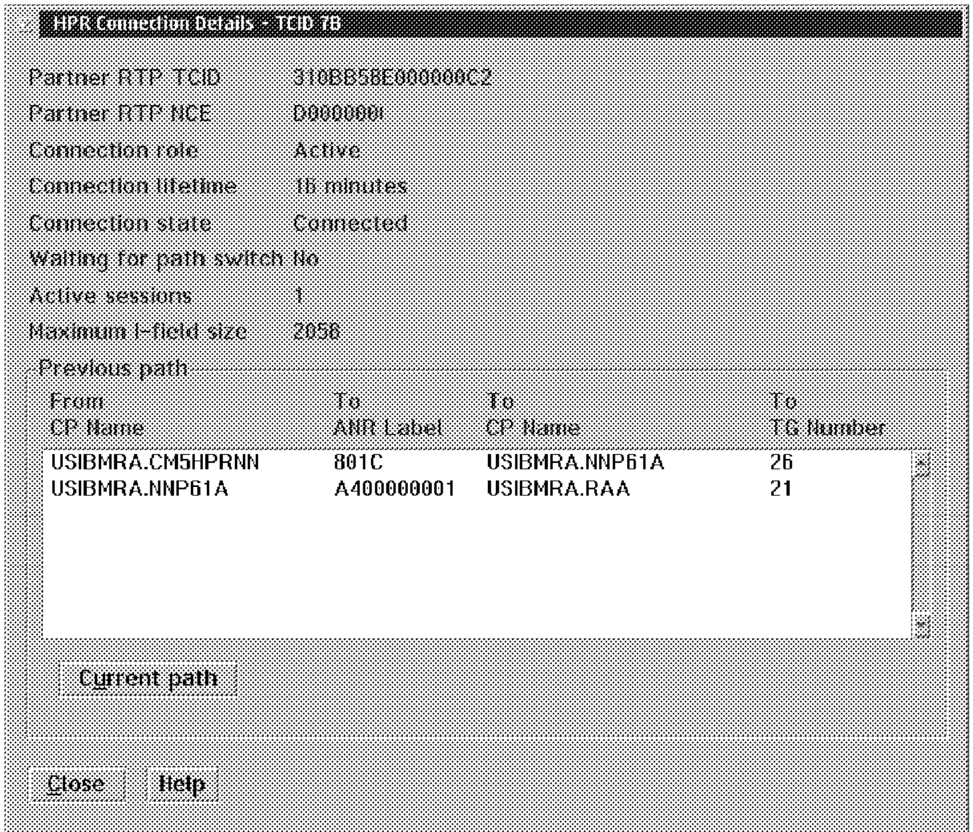

Figure 258. Old Path for TCID 7B

The previous path was:

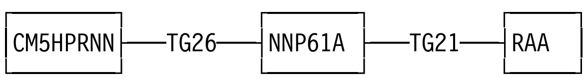

TCID 74 is the LU-LU session pipe, so we display its details in Figure 259 on page 311.

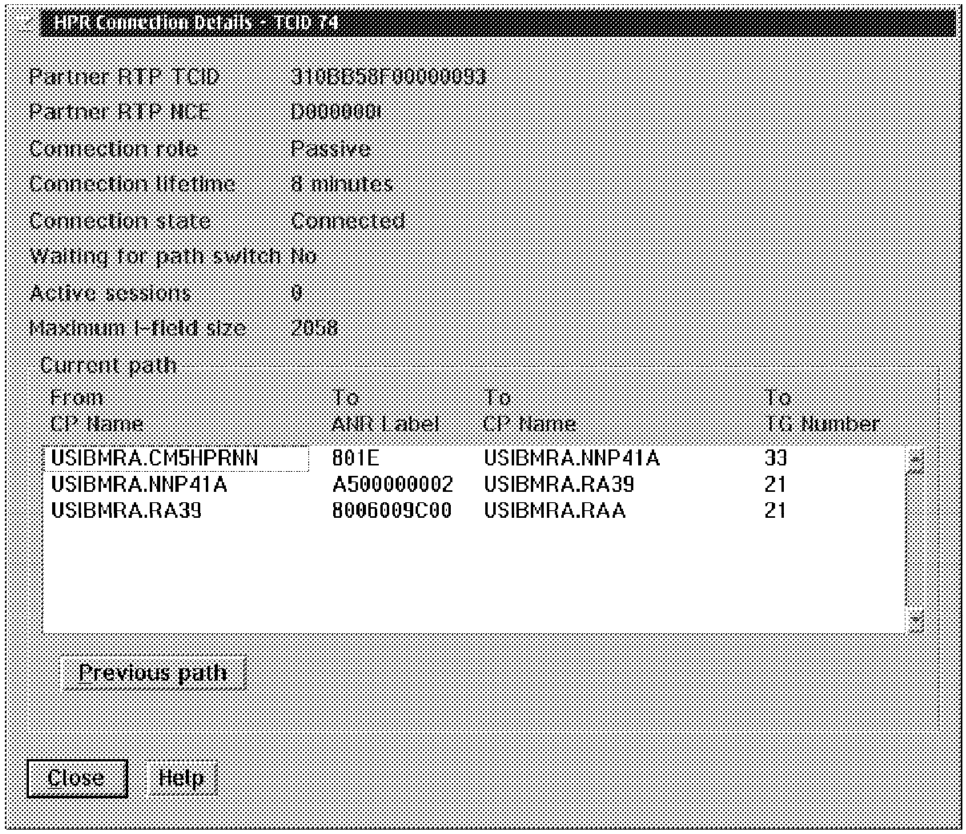

Figure 259. New Path for TCID 74

The LU-LU session takes the path:

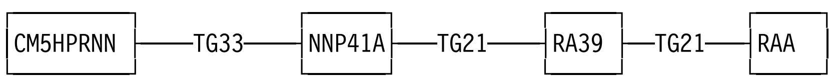

B.3.3 Summary

Figure 260 on page 312 summarizes the path switch described in the previous section. It is the same scenario as Figure 193 on page 196.

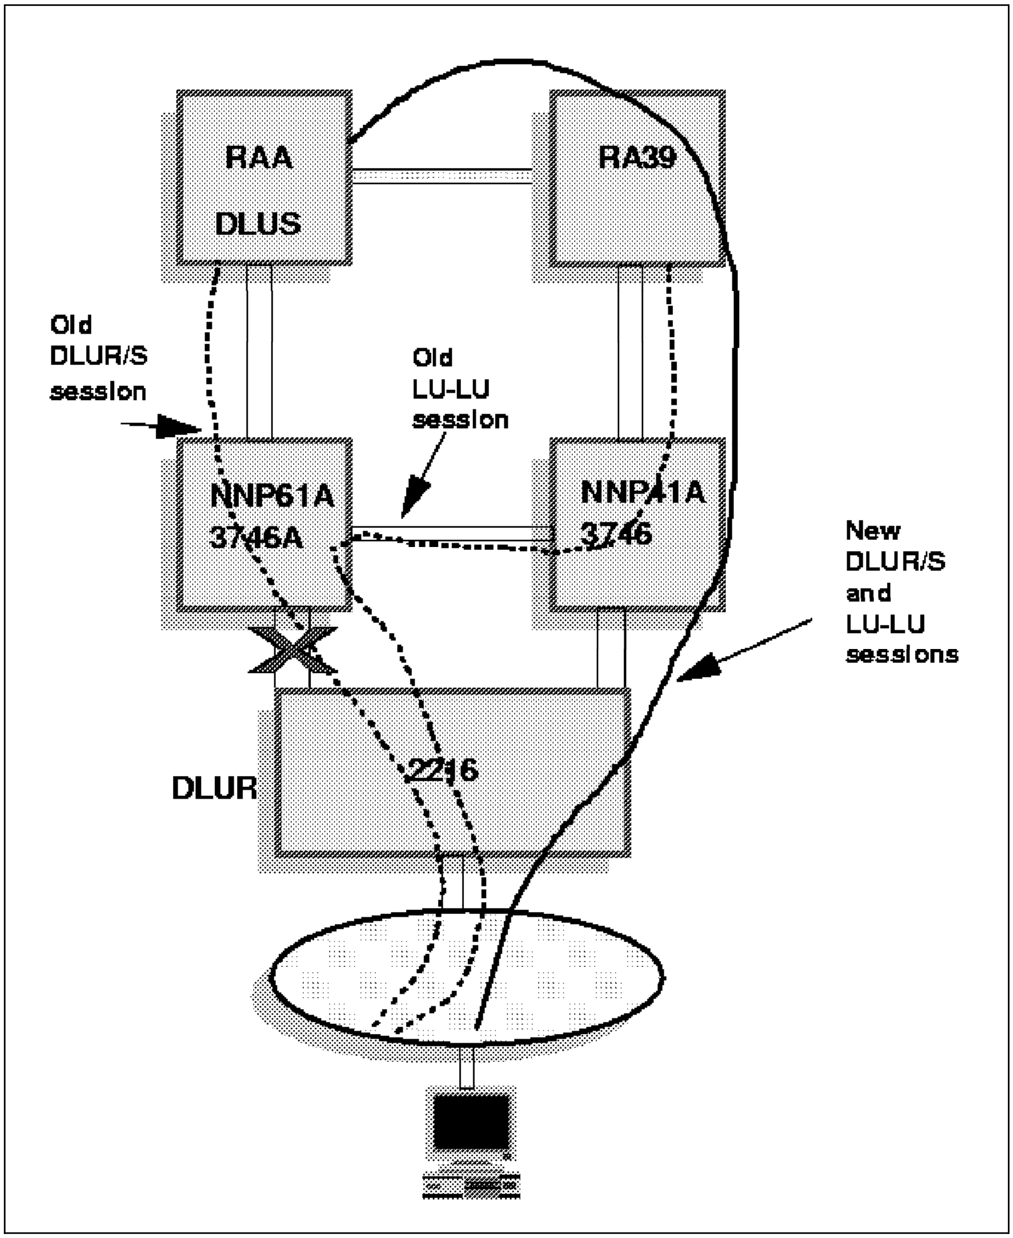

Figure 260. Test Summary

Appendix C. 3746 CCM Configuration Parameters

The following table summarizes the configuration parameters available for an ESCON connection on the 3746.

Table 2 (Page 1 of 2). ESCON Configuration Parameters

CCM Parameter	NCP/VTAM Equivalent		Purpose
	Keyword	Statement	
ESCON Port Configuration, see Figure 92 on page 121.			
Network	n/a	n/a	Selects support for APPN, IP, and/or subarea traffic.
Fiber status	n/a	n/a	Selects status of ESCON channel.
Port name	Line name	LINE	Port name.
NPA eligible	NPACOLL	LINE	Eligible for NPM performance data collection (requires NPM V2R3).
Attachment	n/a	n/a	Indicates whether the ESCON port is attached directly, via an ESCD, or via a chain of ESCDs.
ESCD number	n/a	n/a	Identifies the ESCON Director to which this port is attached.
ESCD model	n/a	n/a	Identifies which type of ESCON director is being used.
Control unit link address	n/a	n/a	Specifies the port at the ESCON director to which the optical fiber from the 3746 ESCON coupler is attached.
ESCON Host Links Configuration, see Figure 95 on page 124.			
Network	n/a	n/a	Selects support for APPN, IP, and/or subarea traffic.
Host mode	n/a	n/a	Mode of the host configuration with respect to partitioning. Basic, LPAR or EMIF.
Host name	n/a	n/a	Specifies the machine to which the ESCON channel adapter is attached and is used to identify IOCP output.
Partition name	n/a	n/a	Specifies the partition to which the CHPID is attached.
CHIPID	n/a	n/a	Specifies the Channel Path Identifier generated in the IOCP at the Host for the ESCON host link.
Partition number	n/a	n/a	Identifies the logical host within a physical host. Valid only when host mode EMIF or LPAR is selected.
Host link address (HLA)	n/a	n/a	Identifies the port in the ESCON Director to which the optical fiber from the host is attached.
Host Link Configuration - APPN Parameters, see Figure 96 on page 125.			
Accept any incoming call?	CALL	LINE	Defines whether link stations can be created dynamically on this host link. Not available for ESCON.
Automatic reactivation	n/a	n/a	Restart link after failure.
NPA eligible	NPACOLL	LINE	Eligible for NPM performance data collection.
Maximum received PIU size	TRANSFR*BFRS-18	LINE/BUILD	Maximum frame size 3746 is able to receive.
Maximum sent PIU size	MAXBFRU*IOBUF	HOST/PU	Maximum frame size 3746 is allowed to send. **Note:** May change after XID negotiation during link establishment.
Propagation delay	PDELAY	LINE	Used for APPN route calculation.
Security	SECURITY	LINE	Used for APPN route calculation.
Relative cost per byte	COSTBYTE	LINE	Used for APPN route calculation.
Relative cost per unit of time	COSTTIME	LINE	Used for APPN route calculation.
ESCON Station Configuration, see Figure 97 on page 126.			
Network	PUTYPE	LINE	Specifies APPN, IP or subarea (T5) connection
VTAM/TPF name	n/a	n/a	Link station name.
PU type	PUTYPE	LINE	Used to define the role of the adjacent node.
Unit address (UA)	ADDR	PU	Station number.
IPL through that station?	IPL	LINE	Enables NCP loading and dumping through that station.
On which CCU?	n/a	n/a	Indicates owner (APPN, CCU-A, CCU-B or the link station).

Table 2 (Page 2 of 2). ESCON Configuration Parameters			
CCM Parameter	**NCP/VTAM Equivalent**		**Purpose**
	Keyword	**Statement**	
Station Configuration - APPN Parameters, see Figure 98 on page 127.			
Activated at startup?	ISTATUS	PU	Link station activated at (re)start of 3746.
Automatic reactivation	n/a	n/a	Restart station after failure.
CP-CP session support?	CPCP	PU	Indication if CP-CP session initiation is required.
NPA eligible	NPACOLL	PU	Eligible for NPM performance data collection (future).
HPR support	LLERP / HPR	PU	Select HPR support as none or yes (ERP required).
Propagation delay	PDELAY	PU	Used for APPN route calculation.
Security	SECURITY	PU	Used for APPN route calculation.
Effective capacity	CAPACITY	PU	Link capacity to station (speed).
Relative cost per byte	COSTBYTE	PU	Used for APPN route calculation.
Relative cost per unit of time	COSTTIME	PU	Used for APPN route calculation.
Adjacent node identifier	IDBLK, IDNUM	PU	IDBLK, IDNUM of adjacent node. No meaning for ESCON links.
XID receipt supported?	XID	PU	No meaning for ESCON links. Relevant only when using dependent LU requester function.
Primary dependent LU server (DLUS)	n/a	n/a	No meaning for ESCON links.
Backup DLUS	n/a	n/a	No meaning for ESCON links.
ESCON Station DLC Parameters, see Figure 100 on page 128.			
Channel adapter slowdown timer (CASDL)	CASDL	PU	Maximum amount of time that the ESCON station can block inbound traffic due to slowdown before signaling that the station is inoperative.
Attention timer (TIMEOUT)	TIMEOUT	PU	Amount of time to wait for a response to an attention signal sent to the host before initiating channel disconnect.
Delay timer (delay)	DELAY	PU	Maximum amount of time to wait between the time data is available to the host and the time the attention signal is sent to a host node.
Total transmit threshold	SRT	PU	Number of transmissions associated with the station before informing the host that the threshold was reached.
Total retry threshold	SRT	PU	Number of retries associated with the station before informing the host that the threshold was reached.

The following table summarizes the configuration parameters available for a token-ring connection on the 3746.

Table 3 (Page 1 of 3). Token-Ring Configuration Parameters			
3746 Parameter	**NCP/VTAM Equivalent**		**Purpose**
	Keyword	**Statement**	
Token-Ring Port Configuration, see Figure 103 on page 131			
Network	n/a	n/a	Specifies APPN or IP connection.
APPN name	linename	LINE	Line name.
Local MAC address	LOCADD	LINE	Token-ring port address of 3746.
APPN local SAP	n/a (always 4)	n/a	DLC service access point (SAP) of 3746.
HPR local SAP	n/a	n/a	DLC service access point (SAP) of 3746 for HPR use
Token-Ring Port - APPN Parameters, see Figure 104 on page 132			
Accept any incoming call?	CALL	LINE	Defines whether link stations can be created dynamically on this host link.
Automatic reactivation	n/a	n/a	Restart port after failure.
NPA eligible	NPACOLL	LINE	Eligible for NPM performance data collection (requires NPM V2R3).
Maximum received PIU size	RCVBUFC	LINE	Maximum frame size 3746 is able to receive.

Table 3 (Page 2 of 3). Token-Ring Configuration Parameters

3746 Parameter	NCP/VTAM Equivalent		Purpose
	Keyword	Statement	
Maximum sent PIU size	MAXTSL	LINE	Maximum frame size 3746 is allowed to send.
HPR support	HPR, LLERP	PU	HPR support and ERP capability.
Propagation delay	PDELAY	LINE	Used for APPN route calculation.
Security	SECURITY	LINE	Used for APPN route calculation.
Relative cost per byte	COSTBYTE	LINE	Used for APPN route calculation.
Relative cost per unit of time	COSTTIME	LINE	Used for APPN route calculation.
Token-Ring Station - Station Parameters, see Figure 106 on page 134			
T1 reply timer	T1TIMER	BUILD	Value specifying time within which reply should be received.
T2 acknowledgement timer	T2TIMER	PU	Value specifying maximum time within which reply is returned.
Inactivity timer	TITIMER	PU	Value specifying maximum time within which data should be received.
Maximum transmitted frames	MAXOUT	PU	Maximum transmit window size.
Maximum received frames	MAXOUT or 127	PU	Maximum number of frames received before partner requires acknowledgment.
RNR limit	RNRLIMT	PU	Specifies how long a remote station is allowed to refuse data before being identified as inoperative.
Authorize infinite retries			Indicates that the retry process is infinite.
Retries per sequence	RETRIES(m,,)	PU	Number of retry attempts in a sequence after a transmission has failed. Total attempts within a sequence is m+1.
Retry sequence	RETRIES(,,n)	PU	The number of retry sequences. The total number of sequences is n+1.
Pause between retry sequences	RETRIES(,t,)	PU	The period between two retry sequences.
Token-Ring Station Configuration, see Figure 107 on page 135			
Name	PU name	PU	Name of token-ring station.
Remote MAC address	DIALNO	PATH	Destination token-ring address.
Remote SAP	DIALNO	PATH	DLC service access point (SAP) of station.
HPR remote SAP	n/a	n/a	DLC service access point (SAP) of station for HPR use.
Token-Ring Station - APPN Parameters, see Figure 108 on page 136			
Activated at startup?	ISTATUS	PU	Link station activated at (re)start of 3746.
Automatic reactivation	n/a	n/a	Restart station after failure.
CP-CP session support?	CPCP	PU	Indication if CP-CP session initiation is required.
NPA eligible	NPACOLL	PU	Eligible for NPM performance data collection.
HPR support	ERP / HPR	PU	HPR support, no or yes with ERP type required.
Propagation delay	PDELAY	PU	Used for APPN route calculation.
Security	SECURITY	PU	Used for APPN route calculation.
Effective capacity	CAPACITY	PU	Logical link capacity.
Relative cost per byte	COSTBYTE	PU	Used for APPN route calculation.
Relative cost per unit of time	COSTTIME	PU	Used for APPN route calculation.
Adjacent node identifier	n/a	n/a	IDBLK, IDNUM of adjacent node. For Dependent LU Requester function only.
XID receipt supported?	XID	PU	XID supported by remote node? Relevant only when using dependent LU requester function.
Primary dependent LU server (DLUS)	n/a	n/a	Fully qualified name of primary DLUS. Relevant only when using dependent LU requester function.
Backup DLUS	n/a	n/a	Fully qualified name of backup DLUS (if appropriate). Relevant only when using dependent LU requester function.
Token-Ring Station - DLC Parameters			
T1 reply timer	T1TIMER	PU	Value specifying time within which reply should be received.

Table 3 (Page 3 of 3). Token-Ring Configuration Parameters			
3746 Parameter	**NCP/VTAM Equivalent**		**Purpose**
	Keyword	**Statement**	
T2 acknowledgement timer	T2TIMER	PU	Value specifying maximum time within which reply is returned.
Inactivity timer	TITIMER	PU	Value specifying maximum time within data should be received.
Maximum transmitted frames	MAXOUT	PU	Maximum transmit window size.
Maximum received frames	127 or MAXOUT	PU	Maximum number of frames received before partner requires acknowledgement.
RNR limit	RNRLIMT	PU	Specifies how long a remote station is allowed to refuse data before being identified as inoperative.
Authorize infinite retries			Indicates that the retry process is infinite.
Retries per sequence	RETRIES(m,,)	LINE	Number of retry attempts in a sequence after a transmission has failed. Total attempts within a sequence is m+1.
Retry sequence	RETRIES(,,n)	PU	The number of retry sequences. The total number of sequences is n+1.
Pause between retry sequences	RETRIES(,t,)	PU	The period between two retry sequences.
Token-Ring Connection Network			
Network identifier	VNNAME	GROUP, LINE	network ID of virtual routing node (VRN).
CN Name	VNNAME	GROUP, LINE	Name of virtual routing node (VRN).

Appendix D. HPR Format Overview

There are various types of packets that can flow between HPR nodes. A packet on an HPR-capable link may contain an XID-3 I-frame, a FID-2 PIU, or a network layer packet (NLP). Figure 261 illustrates.

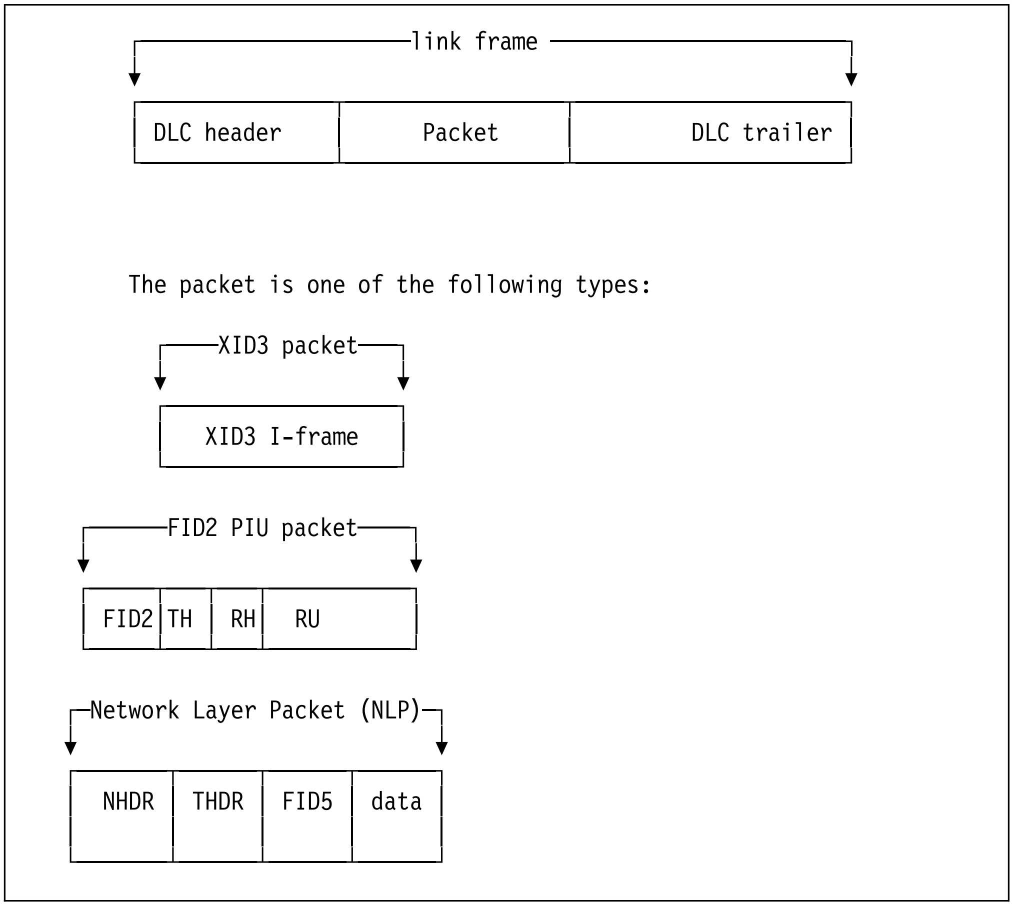

Figure 261. HPR Frame and Packet Format

Thus a frame on an HPR-capable link may contain:

- DLC header

  The term DLC used here can mean a variety of things, from the simple SDLC to a frame relay, X.25 or ATM network. What matters is that it *appears* to the higher SNA layers as a DLC, and is capable of distinguishing XIDs from information frames. In the simple DLCs, XID is represented by the control code X′AF′ or X′BF′. In the more complex cases, protocols such as QLLC or BAN are needed to transmit useful information between HPR-capable SNA nodes.

- XID-3 data

  HPR uses XID-3 (as defined in APPN) with the addition of a new control vector, CV61.

- FID-2 PIU

FID-2 PIUs contain a FID-2 transmission header where the first four bits indicate the FID type (B′0010′ indicates FID-2). These four bits are used to distinguish FID-2 PIUs from NLPs.

- Network Layer Packet (NLP)

  The first four bits of an NLP are either B′1100′ or B′1101′. The NLP consists of the following:

  – The network header (NHDR) contains ANR routing information, including the destination NCE that identifies the target of the NLP.

  – The transport header (THDR) contains RTP information such as ARB segments, switching segments, connection setup segments and so on. It contains the TCID that distinguishes between RTP pipes on the same node.

  – The FID-5 header is present only if user data is present. The FID-5 header contains the session address allocated by each RTP partner to distinguish between the sessions on a pipe. The FID-5 header is followed by the data, beginning with the request or response header.

- DLC trailer

  For the simpler DLCs, this contains the cyclic redundancy check field.

D.1 FID-2 PIU/NLP Usage on Links

FID-2 PIUs and NLPs can flow over the same link and be interleaved. Each FID-2 PIU flows in a DLC frame and each NLP flows in a DLC frame. There is no blocking of FID-2 PIUs and NLPs into a single DLC frame.

A receiver can distinguish a FID-2 PIU from an NLP by examining the first bits of the packet. For FID-2 PIUs, these bits are always set to B′0010′ indicating a format identifier of 2. The first four bits of an NLP will never be B′0010′.

Both FID-2 PIUs and NLPs use the same transmission priorities: network, high, medium, and low. Link outbound queues (one for each priority) are used to implement priority routing. FID-2 PIUs and NLPs are enqueued according to their priority and so any given priority queue can contain both FID-2 PIUs and NLPs. There is no additional priority implied based on the type of packet. FID-2 PIUs and NLPs are treated equally so the order they appear in the queue is the same as the arrival order.

FID-2 PIUs are assigned a priority based on the ISR routing tables in the intermediate node, as the priority is not carried in the header. An RTP connection has no tables in the ANR node; the priority field is in the NHDR.

Whether both FID-2 PIUs and NLPs actually flow over a given link depends on several factors:

- CP-CP sessions flow on NLPs or FID-2 PIUs depending on whether both sides support the Control Flows over RTP option.
- Route Setup messages flow as NLPs or FID-2 PIUs depending on whether both sides support the Control Flows option.
- If the link is a peripheral subarea connection then SSCP-PU, SSCP-LU, and the route extension parts of LU-LU sessions use FID-2 PIUs with dependent session addresses (OAF and DAF).

- All other LU-LU session traffic uses NLPs as long as the session is flowing over an RTP connection at this point. If not, the traffic uses FID-2 PIUs with LFSIDs. This applies to dependent as well as independent LU session traffic, except for the route extension portion from the subarea boundary function.

D.2 Formats

This section summarizes the new and changed data formats for HPR. For a detailed description of the formats, please refer to *SNA Formats*.

D.2.1 XID

HPR uses XID-3 exchange just as does base APPN. Additional information is exchanged between the partner nodes in a new control vector:

- HPR capabilities control vector (CV61)

  The presence of this CV indicates that it is desired that the link run HPR protocols. It is used in the XID exchange and can include:

  – IEEE 802.2 LLC subfield (X′80′)

  – Control Flows Over RTP Tower subfield (X′81′)

D.2.2 Network Layer Packet (NLP)

This is the basic format for data flowing on an HPR connection. It comprises:

- Network Layer Header (NHDR)

  The NHDR is constructed by the origin sender, processed by each intermediate node, and received and processed by the final destination receiver. It contains the session priority and the ANR labels for the route.

- RTP Transport Header (THDR)

  The THDR is created and processed only by the RTP endpoints. It includes the connection TCID, the byte sequence number, and a number of optional segments and control vectors:

  – Connection Setup (CS) Segment

  – Status Segment

  This segment is used to convey state information from one end of the logical connection to the other.

  – Client Out of Band Bits (COB) Segment

  – Connection Identifier Exchange (CIE) Segment

  This segment is sent to the partner in order to communicate a Transport Connection Identifier (TCID) that is to be used by the partner in all messages it sends on this RTP connection.

  – RTP Connection Fault Segment

  This segment is used to communicate errors to the partner RTP endpoint.

  – Switching Information (SI) Segment

  This segment is always present when the connection setup segment is present. It may also be present when it is necessary to convey new path information to the partner on a path switch.

- Adaptive Rate-Based (ARB) Segment
- RTP Control Vectors
  - Node Identifier Control Vector (X′00′)

    This CV identifies a node and always contains the non-qualified CP name of the node.
  - Network Identifier Control Vector (X′03′)

    This CV identifies a network.
  - Network Address Control Vector (X′05′)

    This CV is used in the Connection Qualifier field in the THDR.
  - NCE Identifier Control Vector (X′26′)
  - Topic Identifier Control Vector (X′28′)

    Identifies the intended listening application.
  - Network Connection Endpoint (NCE) Instance Identifier CV (X′39′)
  - HPR Switching Information CV (X′83′) Identifier

    This CV contains information about the path used by the RTP connection.
  - HPR Return Route TG Descriptor CV (X′85′)

    Please note that this CV85 is not the same as the subvector 85 used in CV46.
  - ANR Path CV (X′67′)

    This CV contains the description of a path using a series of ANR label entries.

- FID-5 Transmission Header

  This is used to identify the session on which data is flowing, and precedes the data (if present) in the NLP.
- Session Address Control Vector (CV62)

  This CV is used to convey the session address to be used by the primary LU when sending data to the secondary LU. It is assigned by the secondary LU.

D.2.3 FID-2 Route Setup

This is the format of the Route Setup when sent between nodes where one or both do not support the Control Flows Over RTP option. In this case the Route Setup is carried in a FID-2 PIU. If the nodes support Control Flows, the Route Setup is carried as a GDS variable in the data portion of an NLP.

D.2.4 New and Changed GDS Variables

The Route Setup GDS variable is new with HPR, and some of the other APPN variables have been changed:

- Route Setup GDS Variable X′12CE′

  This GDS variable may be included in either the DATA portion af an NLP or the RU portion of a FID-2 PIU.
- Cross-Domain Initiate (X′12C5′ GDS)

This GDS variable has been enhanced to allow an HPR-only route to be requested.

- Find Resource (X′12CA′) GDS Variable

  This GDS variable has been modified to carry the LU NCE.

- Found Resource (X′12CB′) GDS Variable

  This GDS variable has been modified to carry the LU NCE.

Appendix E. Special Notices

This publication is intended to help network systems programmers and network analysts to migrate existing subarea networks to combined subarea and APPN/HPR networks. The information in this publication is not intended as the specification of any programming interfaces that are used in this book. See the PUBLICATIONS section of the IBM Programming Announcements for eNetwork Communications Server for OS/390, ACF/NCP, Communications Server for OS/2 Warp, Multiprotocol Access Service and Multiprotocol Routing Services for more information about what publications are considered to be product documentation.

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

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, North Castle Drive, Armonk, NY 10504-1785.

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 information about non-IBM (″vendor″) products in this manual has been supplied by the vendor and IBM assumes no responsibility for its accuracy or completeness. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer′s ability to evaluate and integrate them into the customer′s operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

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.

Any performance data contained in this document was determined in a controlled environment, and therefore, the results that may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.

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

Appendix F. Related Publications

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

F.1 International Technical Support Organization Publications

For information on ordering these ITSO publications see "How to Get ITSO Redbooks" on page 327.

- *Inside APPN - The Essential Guide to the Next-Generation SNA*, SG24-3669-03
- *IBM Connectivity Guide*, SG24-4169-02
- *Subarea to APPN Migration: VTAM and APPN Implementation*, SG24-4656-01
- *VTAM V4R4 for MVS/ESA Implementation Guide*, SG24-2100
- *SNA in a Parallel Sysplex Environment*, SG24-2113
- *IBM 3746 Nways Controller APPN Implementation Guide*, SG24-2536-01
- *IBM 2216 Nways Multiaccess Connector Description and Configuration Scenarios - Volume I*, SG24-4957
- *IBM 2210 Nways Multiprotocol Router Description and Configuration Scenarios - Volume I*, SG24-4446
- *IBM 2210 Nways Multiprotocol Router and IBM 2216 Nways Multiaccess Connector Description and Configuration Scenarios - Volume II*, SG24-4956
- *IBM eNetwork Communications Server for AIX: Understanding and Migrating to Version 5: Part 1 - Configuration and New Features*, SG24-5215
- *2216 ESCON Solutions*, SG24-2137
- *3746, 2210, 2216 and 2220 Interconnectivity: Frame Relay and Related Functions*, SG24-2146
- *IBM eNetwork Communications Server for OS/2 Warp Version 5.0 Enhancements*, SG24-2147
- *IBM eNetwork Communications Server for Windows NT Version 6.0*, SG24-5232 (available 4Q98)

F.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

F.3 Other Publications

These publications are also relevant as further information sources:

- *SNA Formats*, GA27-3136-16
- *eNetwork Communications Server for OS/390 SNA Network Implementation Guide*, SC31-8563
- *eNetwork Communications Server for OS/390 SNA Resource Definition Reference*, SC31-8565
- *eNetwork Communications Server for OS/390 SNA Resource Definition Samples*, SC31-8566
- *eNetwork Communications Server for OS/390 SNA Operation*, SC31-8567
- *eNetwork Communications Server for OS/390 SNA Planning and Migration Guide*, SC31-8622
- *eNetwork Communications Server for OS/390 SNA Diagnosis*, LY43-0079 (available to IBM-licensed customers only)
- *eNetwork Communications Server for OS/390 SNA Customization*, LY43-0110 (available to IBM-licensed customers only)
- *NCP, SSP and EP Resource Definition Reference*, SC31-6224
- *Nways Multiprotocol Access Services Protocol Configuration and Monitoring Reference Volume 2*, SC30-3885-03

The latest online information on eNetwork Communications Server for OS/390 may be found on the Web at the following address:

`http://www.software.ibm.com/enetwork/commserver/about/csos390.html`

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:

- **Redbooks Web Site on the World Wide Web**

  http://w3.itso.ibm.com/

- **PUBORDER** — to order hardcopies in the United States
- **Tools Disks**

  To get LIST3820s of redbooks, type one of the following commands:

  ```
  TOOLCAT REDPRINT
  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
  ```

- **REDBOOKS Category on INEWS**
- **Online** — send orders to: USIB6FPL at IBMMAIL or DKIBMBSH at IBMMAIL

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 **USA International Access Code -408-256-5422 (Outside USA)** — ask for:

Index # 4421 Abstracts of new redbooks
Index # 4422 IBM redbooks
Index # 4420 Redbooks for last six months

- **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

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.

List of Abbreviations

ACF/NCP	Advanced Communication Function / Network Control Program
AIX	Advanced Interactive eXecutive
ANNC	APPN Node to Node Connection
ANR	Automatic Network Routing
APAR	Authorized Programming Analysis Report
API	Application Programming Interface
APPC	Advanced Program to Program Communication
APPN	Advanced Peer to Peer Networking
ARB	Adaptive Rate-Based
AS/400	Application System / 400
ATM	Asynchronous Transfer Mode
BF-TG	Boundary Function Transmission Group
BX	Branch Extender
CCM	Controller Configuration and Management
CDLC	Channel Data Link Control
CDRM	Cross Domain Resource Manager
CDRSC	Cross Domain ReSourCe
CDS	Central Directory Server
CMC	Communications Management Configuration
CNN	Composite Network Node
COS	Class Of Service
CP	Control Point
CRC	Cyclic Redundancy Check
CS OS/390	eNetwork Communications Server for OS/390
CS/AIX	Communications Server for AIX
CS/NT	Communications Server for Windows NT
CS/2	Communications Server/2
CSM	Communication Storage Manager
CTC	Channel To Channel
CV	Control Vector
DCAF	Distributed Console Access Facility
DIX	Digital, Intel, Xerox
DLC	Data Link Control
DLSw	Data Link Switching
DLU	Destination Logical Unit
DLUR/S	Dependent LU Requester / Server
DSME	Directory Services Management Exit
EBN	Extended Border Node
EGA	ESCON Generation Assistant
EMIF	ESCON Multiple Image Facility
EN	End Node
ER	Explicit Route
ESCD	ESCON Director
ESCON	Enterprise Systems CONnection
FDDI	Fiber Distributed Data Interchange
FID	Format IDentifier
FMH	Function Management Header
FQPCID	Fully Qualified Procedure Correlated IDentifier
FRFH	Frame Relay Frame Handler
GDLC	Generalized Data Link Control
GDS	Generalized Data Stream
GUI	Graphical User Interface
HPDT	High Peformance Data Transfer
HPR	High-Performance Routing
IANA	Internet Assigned Numbers Authority
IBM	International Business Machines Corporation
IC-TG	InterChange Transmission Group
ICN	InterChange Node
IEEE	Institute of Electrical and Electronic Engineers
IP	Internet Protocol
IPL	Initial Program Load

ISDN	Integrated Services Digital Network
ISR	Intermediate Session Routing
ITSO	International Technical Support Organization
LAN	Local Area Network
LDLC	Logical Data Link Control
LEN	Low Entry Networking
LFSID	Local-Form Session IDentifier
LLC	Logical Link Control
LPAR	Logical PARtition
LSA	Link Services Architecture
LU	Logical Unit
MAC	Media Access Control
MAS	Multiprotocol Access Services
MAE	MultiAccess Enclosure
MDH	Migration Data Host
MLTG	MultiLink Transmission Group
MOSS	Maintenance Operator SubSystem
MNPS	MultiNode Persistent Sessions
MPC	MultiPath Channel
MVS/ESA	Multiple Virtual Storage / Enterprise Systems Architecture
NAU	Network Accessible Unit
NCP	Network Control Program
NCE	Network Connection Endpoint
NDF	Node Definition File
NHDR	Network layer HeaDeR
NLP	Network Layer Packet
NMVT	Network Management Vector Transport
NN	Network Node
NNP	Network Node Processor
NNS	Network Node Server
NTRI	NCP Token-Ring Interface
OLU	Origin Logical Unit
OS/2	Operating System / 2
OS/390	Operating System / 390
OS/400	Operating System / 400
OSA	Open Systems Adapter
PBN	Peripheral Border Node
PC	Personal Computer
PIU	Path Information Unit
PLU	Primary Logical Unit
PPP	Point to Point Protocol
PS/2	Personal System/2
PU	Physical Unit
RJE	Remote Job Entry
RSCV	Route Selection Control Vector
RTP	Rapid Transport Protocol
RU	Request Unit
SAP	Service Access Point
SDLC	Synchronous Data Link Control
SLU	Secondary Logical Unit
SNA	Systems Network Architecture
SNI	SNA Network Interconnection
SSCP	System Services Control Point
SVC	Switched Virtual Channel
TCP	Transmission Control Protocol
TDB	Topology DataBase
TDU	Topology Database Update
TCID	Transport Connection IDentifier
TG	Transmission Group
TH	Transmission Header
THDR	Transport layer HeaDeR
TIC	Token-ring Interface Coupler
TP	Transmission Priority
TRLE	Transport Resource List Element
TSO	Time Sharing Option
UDP	User Datagram Protocol
UI	Unnumbered Information
USS	Unformatted System Services
VM/ESA	Virtual Machine / Enterprise Systems Architecture
VR	Virtual Route
VR-TG	Virtual Route-based Transmission Group
VTAM	Virtual Telecommunications Access Method
WAC	Wide Area Connector
XCA	External Communications Adapter

XCF Cross-system Coupling Facility

XID EXchange IDentifier

Index

Numerics

A

V

X

ITSO Redbook Evaluation

Subarea to APPN Migration: HPR and DLUR Implementation
SG24-5204-00

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 914 432 8264
- Send your comments in an Internet note to redbook@us.ibm.com

Which of the following best describes you?
__Customer __Business Partner __Solution Developer __IBM employee
__None of the above

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!)

____

____

____

____

____

Subarea to APPN Migration: HPR and DLUR Implementation SG24-5204-00

SG24-5204-00
Printed in the U.S.A.