VTAM V4R2 Early User Experiences

Document Number GG24-4250-00

June 1994

International Technical Support Organization
Raleigh Center

Take Note!

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

First Edition (June 1994)

This edition applies to Version 4, Release 2 of ACF/VTAM for MVS/ESA, Program Number 5695-117.

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

An ITSO Technical Bulletin Evaluation Form for reader′s feedback appears facing Chapter 1. If the form has been removed, comments may be addressed to:

IBM Corporation, International Technical Support Organization
Dept. 985 Building 657
P.O. Box 12195
Research Triangle Park, NC 27709-2195

Abstract

This document describes experiences gained in implementing VTAM for MVS/ESA Version 4 Release 2. Its purpose is to assist others in implementing the new functions in this product, by providing tested examples and definitions which have actually worked in a VTAM V4R2 environment.

Most of the new functions in VTAM V4R2 enhance the APPN support first implemented in V4R1, and the bulk of this document is devoted to describing these improvements. Several other important features are described and the remainder are summarized, but were not tested by us.

This document is intended for SNA network systems programmers and IBM technical support personnel, who plan to implement the functions in this release of VTAM. It assumes a working knowledge of SNA, APPN and previous releases of VTAM.

(224 pages)

Contents

Figures

Special Notices

This publication is intended to help network systems programmers to implement ACF/VTAM V4R2 for MVS/ESA. The information in this publication is not intended as the specification of any programming interfaces that are provided by ACF/VTAM. See the PUBLICATIONS section of the IBM Programming Announcement for ACF/VTAM V4R2 for more information about what publications are considered to be product documentation.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM′s product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM′s intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, 208 Harbor Drive, Stamford, CT 06904 USA.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer′s ability to evaluate and integrate them into the customer′s operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

The following terms, which are denoted by an asterisk (*) in this publication, are trademarks of the International Business Machines Corporation in the United States and/or other countries:

ACF/VTAM	Advanced Peer-to-Peer Networking
AnyNet	APPN
AS/400	CICS
CICS/ESA	ESCON
IBM	Micro Channel
MVS/ESA	NetView
OS/2	OS/400
PS/2	VM/ESA
VTAM	

The following terms, which are denoted by a double asterisk (**) in this publication, are trademarks of other companies:

IPX	Novell, Inc.

Preface

This document is intended to provide a technical reference for those planning to implement the new functions in VTAM V4R2. It contains working examples and definitions which have been tested in a VTAM V4R2 environment at the International Technical Support Organization, Raleigh.

The document is intended for SNA systems programmers and other technical support personnel, who are already familiar with VTAM and NCP customization and setup. An understanding of the APPN architecture is also assumed.

The document does not constitute a reference to all the functions and other technical parameters available in VTAM V4R2; it is intended to be read in conjunction with the product installation and reference manuals, which are listed in "Related Publications" on page xviii.

How This Document is Organized

The document is organized as follows:

- Chapter 1, "Introduction"

  This chapter describes the background, the scope and the objectives of this book. It also summarizes our experiences and impressions of VTAM V4R2.

- Chapter 2, "Raleigh Test Environment"

  This chapter documents the hardware and software environment in which we perform our tests.

- Chapter 3, "APPN Host-to-Host Connection"

  This chapter describes how a channel connection between hosts can be used as an APPN link.

- Chapter 4, "Connection Network"

  This chapter shows how any-to-any APPN connectivity can be achieved on a local area network.

- Chapter 5, "Border Node"

  This chapter describes how disparate APPN networks can be linked in such a way that topology updates are isolated, yet searches and sessions remain unrestricted.

- Chapter 6, "Virtual Route-Based Transmission Group"

  This chapter shows how VTAM V4R2 improves route calculation within a subarea network that carries APPN sessions.

- Chapter 7, "CNN Routing"

  This chapter describes a second method used by VTAM to optimize route calculation in a mixed subarea/APPN network.

- Chapter 8, "Dependent LU Requester/Server"

  This chapter shows how dependent LUs can participate fully in an APPN network, without the connectivity and routing restrictions imposed on them by the subarea implementation.

- Chapter 9, "Automatic Logon Enhancements"

This chapter describes how VTAM will now retry automatic logon requests under a wider variety of circumstances than previously.

- Chapter 10, “Compression”

  This chapter describes additions to VTAM′s existing data compression function, that improve performance and allow greater control.

- Chapter 11, “Search Reduction”

  This chapter shows how VTAM can eliminate unnecessary searches for nonexistent resources, thus reducing overhead.

- Chapter 12, “Generic Resources”

  This chapter describes how VTAM supports groups of applications running on a Systems Complex configuration.

Related Publications

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

- *VTAM V4R2 for MVS/ESA Migration Guide*, GC31-6491
- *VTAM V4R2 for MVS/ESA Release Guide*, GC31-6492
- *VTAM V4R2 for MVS/ESA Network Implementation Guide*, SC31-6494
- *VTAM V4R2 for MVS/ESA Resource Definition Reference*, SC31-6498
- *VTAM V4R2 for MVS/ESA Operation*, SC31-6495
- *VTAM V4R2 for MVS/ESA Messages and Codes*, SC31-6493
- *SNA Formats*, GA27-3136
- *VTAM V4R2 for MVS/ESA Customization*, LY43-0063 (available to IBM-licensed customers only)
- *SNA Network Product Formats*, LY43-0081 (available to IBM-licensed customers only)

International Technical Support Organization Publications

- *VTAM V4R1 for MVS/ESA Implementation Guide*, GG24-4011
- *VTAM V4R1 for MVS/ESA Planning Guide*, GG24-3941
- *APPN Architecture & Product Implementations Tutorial*, GG24-3669
- *VTAM V3R4.2 AnyNet Implementation*, GG24-4066

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

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

Acknowledgments

The advisor for this project was:

Michael Li
International Technical Support Organization, Raleigh Center

The authors of this document are:

Jerzy Buczak
IBM United Kingdom

Bernd Kampmann
IBM Germany

Aung Maung
IBM Australia

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

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

Fran Collins, International Technical Support Organization, Raleigh Center

Joost Fonville, International Technical Support Organization, Raleigh Center

Peter Lenhard, International Technical Support Organization, Raleigh Center

David Billing, Research Triangle Park

Roy Brabson, Research Triangle Park

Eric Chen, Research Triangle Park

Johnathan Harter, Research Triangle Park

Steve Hawley, Research Triangle Park

Art McDonald, Research Triangle Park

Lisa Uthe, Research Triangle Park

Nancy Gates, Networking Support Center, Gaithersburg

Thanks also to the following people whose support and assistance made this book possible:

Gray Heffner, International Technical Support Organization, Raleigh Center

Carla Sadtler, International Technical Support Organization, Raleigh Center

Shawn Walsh, International Technical Support Organization, Raleigh Center

Chapter 1. Introduction

In 1992, VTAM* V4R1 introduced APPN* to the mainframe environment. Before then, subarea and APPN networks could interwork (via low entry networking connections) but could not be said to be truly integrated. In VTAM V4R1, for the first time, the strengths of subarea (control and management) were combined with the strengths of APPN (dynamic, peer-to-peer communication). APPN sessions could now be routed over a subarea network, and dependent LU sessions could traverse an APPN network. The search and route selection mechanisms were brought together in VTAM′s interchange node implementation.

However, in this first release of APPN VTAM there were a few restrictions that made it of limited benefit to some customers′ implementations. Specifically:

- There was no APPN equivalent of SNI, thus subarea connections had to be used across all network boundaries.
- Channel-to-channel connection was not supported using APPN.
- The route calculation was not fully integrated between APPN and subarea portions of a network; thus in a complex network a less than optimum route was likely to be chosen.
- Dependent LUs had to be adjacent to subarea network.

VTAM V4R2 addresses all these points, as well as providing several important new features, for both APPN and non-APPN users.

1.1 Scope and Objectives

We recognize that, because of the factors mentioned above, many customers will migrate straight to VTAM V4R2 from VTAM V3, without implementing V4R1. We cannot therefore assume that all readers will be familiar with the APPN functions in V4R1. However, these functions are described in the companion pair of ITSO publications *VTAM V4R1 Planning Guide* and *VTAM V4R1 Implementation Guide*. We therefore assume that the reader has these volumes at hand, and is familiar — in outline at least — with their contents. We encourage the reader to refer to these volumes.

Our objective in this document is to describe the new features in VTAM V4R2, with examples based on our experience wherever possible. We have not tested all the new features: for some we did not have all the required resources available, and others are of relatively minor significance. All are described in the books referred to in “Related Publications” on page xviii; we include a summary of them in 1.2 below, indicating which are described in more detail in the body of the book.

One function which was present in VTAM V4R1 is no longer available in this release: the OSI remote programming interface.

1.2 Summary of New Features in VTAM V4R2

- **APPN Multiple Network Connectivity** provides the ability for APPN networks with different net IDs to be connected, giving similar facilities to SNI. It is described in Chapter 5, "Border Node" on page 33.
- **Dependent LU Server** gives dependent LUs the full APPN connectivity and routing facilities that independent LUs have enjoyed. It is described in Chapter 8, "Dependent LU Requester/Server" on page 127.
- **APPN Host-to-Host Connection** allows APPN connections directly between hosts across a channel. It is described in Chapter 3, "APPN Host-to-Host Connection" on page 11.
- **Connection Network** improves the efficiency of APPN links over a LAN. It is described in Chapter 4, "Connection Network" on page 21.
- **Virtual Route-Based Transmission Group** and **CNN Routing** are features that improve route selection in a mixed subarea/APPN network. They are described in Chapter 6, "Virtual Route-Based Transmission Group" on page 83 and Chapter 7, "CNN Routing" on page 101 respectively.
- **Search Reduction** prevents excessive searches for nonexistent resources. It is described in Chapter 11, "Search Reduction" on page 181.
- **Compression** improves the efficiency of the existing data compression support in previous VTAM releases. It is described in Chapter 10, "Compression" on page 171.
- **Autologon Enhancements** allow a wider range of session requests to be redriven automatically when certain events occur. They are described in Chapter 9, "Automatic Logon Enhancements" on page 167.
- The **Generic Resources** feature provides VTAM support for a SYSPLEX configuration. Although we were not able to test this, we describe it in Chapter 12, "Generic Resources" on page 191.
- The **ADJLIST** function for adjacent SSCP selection gives greater control over the routing of cross-network subarea sessions. You can now specify exactly - without writing an exit - which route is to be used to reach which resource. ADJLIST links a CDRSC definition to a list of adjacent SSCP names.
- **Expanded Addressing Pool** allows VTAM to handle more sessions than were allowed by the previous limit of 65535 element addresses per subarea. The potential for large numbers of parallel sessions in an APPN network, as well as for large numbers of LAN stations attached via a 3172, makes this feature desirable.
- **Expanded Dial Number** permits more than 32 digits to be passed to the dialling function when VTAM sets up a switched link, using the DLCADDR parameter rather than the DIALNO on the PATH statement. This will be particularly useful for NPSI, which often requires more than 32 digits to pass to an X.25 network; although the X.25 DTE address is restricted to 15 digits, NPSI needs additional information to be passed with the address. Expanded Dial Number is also required for the dial-out function used for dependent LU requester connection, as described in 8.6, "CP-SVR Pipe Initiated by VTAM" on page 157. As products (such as NPSI) migrate from the old method to the new, the use of *both* DIALNO and DLCADDR on the same definition is recommended.

- Enhancements to **APPC/VTAM** provide additional function for LU 6.2 application programs:
  - Full-duplex protocol support
  - RECEIVE IMMEDIATE to retrieve available data without waiting for the buffer to be filled
  - Improved processing for abnormal deallocation
  - A new command to obtain conversation status
- **Delayed Dial Disconnection** improves the flexibility of the existing support, whereby VTAM will drop a switched connection when the last session using that connection is terminated. A time-out value may now be specified, so that VTAM waits for that time before dropping the connection. This is to allow certain products to re-establish sessions immediately, without incurring the costs and overheads of redialling.
- **MAXSESS** currently limits the number of sessions that an independent LU on a leased NCP link may have. VTAM V4R2 extends the use of this parameter by allowing its definition on switched, CDRSC, DR and MODEL major nodes. Thus dynamically defined and switched independent LUs may be restricted in this manner.
- **Display Commands** have been improved in various ways such as the use of "wildcard" characters on resource names. One new command - DISPLAY RSCLIST - is shown in use in Figure 146 on page 106 and in Figure 229 on page 170.
- The **Message Flooding Table** may now be defined by the user, instead of being predefined in the VTAM code. The installation therefore has much flexibility in determining which messages are to be suppressed if they occur very frequently.
- **Additional serviceability aids** include more APPN function in the VTAM dump formatter; an extension to the API which allows FRR information to be kept across additional request types; and a MODIFY QUERY command (for use by network management applications only, not by operators) which exchanges tuning information with an NCP.
- Last but not least, the **AnyNet*** feature, previously available with VTAM V3R4.2, is available with VTAM V4R2, again as a separately orderable feature. It is improved over VTAM V3R4.2 to include support for any LU session over IP (not just APPC). The extensive function in AnyNet is beyond the scope of this book, and readers are referred to the available publications on that subject such as *VTAM V3R4.2 AnyNet Implementation*.

1.3 Conclusions

Our tests have been performed during the field trials of VTAM V4R2, and the code is not at the level which will be made generally available to customers. Nevertheless, we have found the product to be extremely robust and free from bugs. We have no hesitation in recommending its implementation by MVS/ESA* customers as soon as practicable.

A much bigger question is whether to implement APPN following the installation of VTAM V4R2 in a subarea network. This release of VTAM contains APPN functions on a par with subarea functions, yet there are still some network configurations for which APPN will bring minimal benefits. At present, the only external factor which mandates APPN is SYSPLEX.

In Table 1 on page 4 we have drawn up a list of factors which influence the applicability of APPN VTAM to an installation now that the V4R2 function is available. If your network contains mostly the features on the left of the table, APPN is likely to offer few benefits. If it embraces mostly the factors on the right side of the table, or if some of those factors are of major significance, then you should start planning for APPN now.

Table 1. Factors Affecting the Choice of APPN for VTAM

APPN VTAM Not Very Useful	**APPN VTAM Very Useful**
CMC with no data hosts (or single VTAM)	Many data hosts
Few VTAMs, many NCPs	Many VTAMs, few NCPs
No LEN connections	Many LEN connections
Network definitions are simple	Network definitions are complex, time consuming, and easy to get wrong
Network configuration is stable	Network changes are frequent
Subarea and APPN networks are geographically distinct	Subarea and APPN networks share locations and routing requirements
Many BSC 3270s	No BSC 3270s
Many channel-to-channel connections not supported by AHHC	No channel-to-channel connections, or all ESCON* / 3088 / VM
No existing APPN network	Moderate to large existing APPN network
No local area network	Extensive use of LANs
Centralized environment with dependent terminals	Large distributed client/server environment
Centralized administration	Decentralized administration
Load distribution not required	Load distribution between routes is desirable

A brief explanation of these factors is given here, but the reader is encouraged to study the available documentation (this book and those quoted in "Related Publications" on page xviii) for a full understanding.

- **Data hosts** own no NCPs, have no multilink TGs, and are therefore prime candidates for conversion into APPN end nodes, with a corresponding saving in definitions such as path tables. This is not possible, however, if a data host is acting as a backup CMC and needs subarea routes into the network.
- **The ratio of VTAMs to NCPs** is an important factor because an NCP cannot be an APPN node on its own. If one VTAM owns many NCPs, the whole forms a composite network node and no reduction in path tables is possible. If on the other hand the network contains many VTAMs each with one or two NCPs, there is potential for splitting it into a number of APPN nodes linked by FID2 connections which require no routing definitions.
- **LEN connections** form a barrier between subarea and APPN networks, since no form of search request can cross them. They require predefinition of any resources which are to be searched for across them, whether in VTAM or in the APPN node attached to VTAM. Converting them to APPN, if the attached products permit, will remove the need for these definitions.

- APPN requires fewer **network definitions** than subarea networking, therefore less effort is involved and the likelihood of making an error is reduced.
- **Network changes** generally require definition changes and therefore raise the possibility of loss of availability if the new definitions are in error. APPN lessens this risk.
- **Parallel subarea and APPN networks** often result in duplication of resources since they do not share links and the sessions take different paths. They also impact availability since it is very difficult to use subarea and APPN paths to back each other up. APPN VTAM, by integrating the two types of SNA, alleviates these problems.
- **BSC 3270** sessions cannot traverse an APPN link, so must remain wholly within a subarea network. This restriction does not apply to channel-attached non-SNA 3270s, which are more widespread.
- Certain types of **channel-to-channel** links are not supported by APPN. ESCON, 3088 or VM virtual channel-to-channel connections will work.
- A **large existing APPN network** without VTAM will not have the benefit of a central directory server, since only VTAM today provides this function. Adding a VTAM CDS node to such a network will improve performance by cutting down the number of broadcast searches. Products such as AS/400*, CM/2 V1R1, and 3174 (with LIC C5) will make use of CDS if one is provided.
- An extensive **local area network** typically comprises intelligent workstations with connectivity to multiple hosts, peers or servers. The use of an APPN connection network will ease the effort required to optimize these connections.
- A **client/server environment** can benefit greatly from APPN, not only because APPC communications are typically implemented using T2.1 connections. Other important client/server protocols based on TCP/IP, NETBIOS and IPX** can easily be transported across an APPN network using products such as AnyNet and the LAN to LAN WAN Program.
- If the **administration** of various parts of a network is performed by independent groups of people, much effort can be required to coordinate the definitions. Since APPN requires fewer definitions, this effort will be minimized in an APPN environment.
- APPN performs **load distribution** across a network automatically; routes of equal weight will be chosen at random, ensuring a balance between the possible paths to a destination. In a subarea network this function requires either the coding of an exit, or judicious design and careful use of mode table entries.

Chapter 2. Raleigh Test Environment

This chapter describes the test environment that we use at the ITSC in Raleigh to build the test scenarios described in the book. Figure 1 shows a diagram of the environment. We use the same configuration wherever possible in our tests, any required changes being described in the test scenarios. These changes might include:

- Changing a link from FID2 to FID4, or vice versa
- Re-assigning a channel adapter to a different VTAM
- Re-customizing a PS/2 as network node or end node

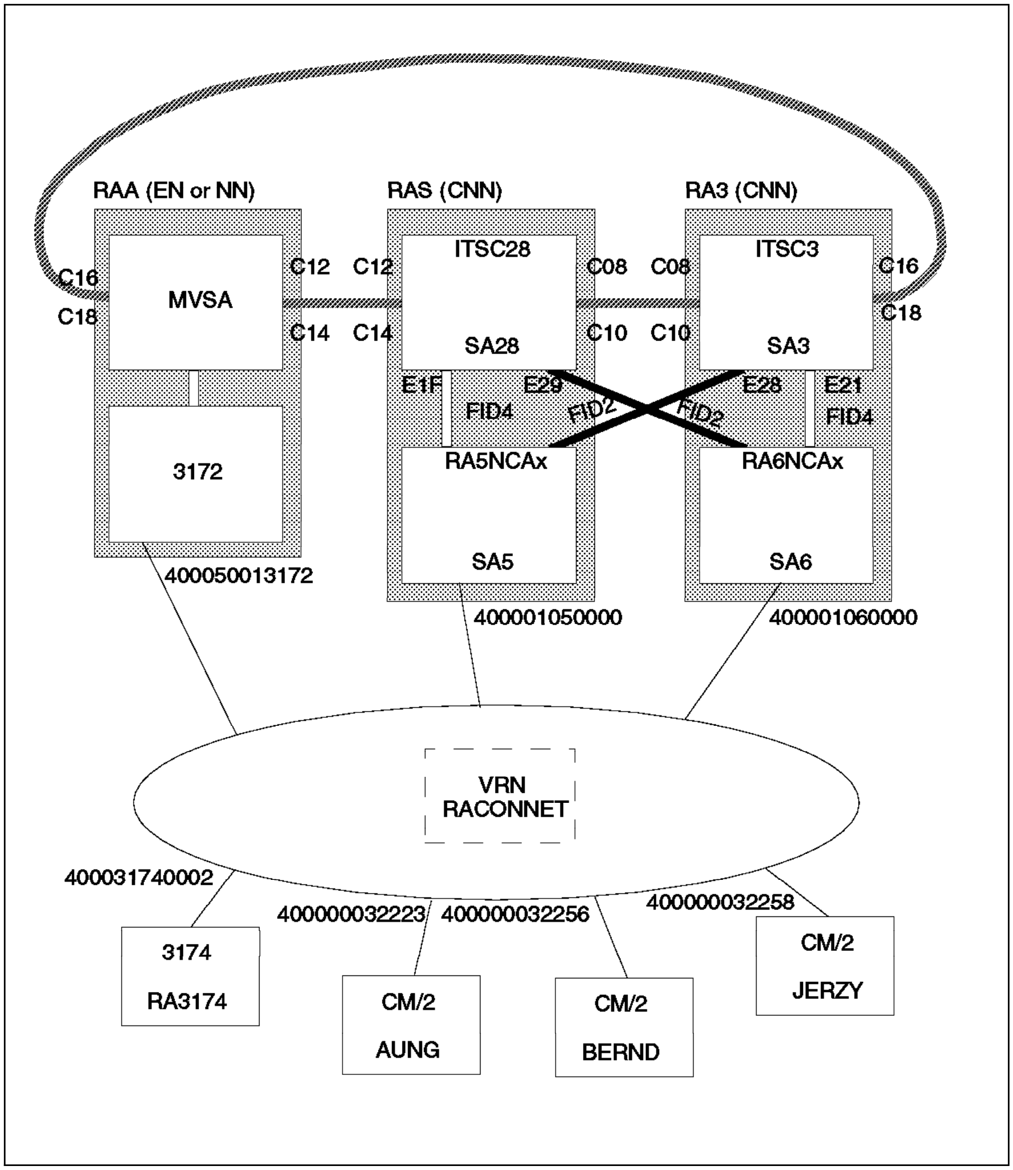

Figure 1. Diagram of the Test Environment

2.1 Test Systems

This section describes the hardware and software of the test environment. Due to the timing of this project, we are required to use pre-release software for VTAM, CM/2 and 3174. We advise that you check the product documentation to ensure that you have the correct level of software for the function that you wish to implement.

We use three MVS/ESA V4R3 systems in our tests as shown in Figure 1 on page 7. All are virtual systems running under VM/ESA* V1R2 which is running in an IBM* 9121-320 processor. All three systems run VTAM V4R2.

In addition to the three VTAM V4R2 systems, our environment is connected to RAK, which is a VTAM V4R1 system, for some of our scenarios. RAK provides a connection to IBM′s world-wide SNA network which we use to verify that LUs in our APPN network environment can set up sessions with LUs in the corporate SNA network across multiple SNI connections.

Attached to RAA, we use a 3172 Model 1 for LAN connection. The 3172 is running ICP V2R2.

Attached to RAS, we use one side of a 3745-410. It is running NCP V7R1. We use SSP V4R1 to generate the NCPs.

Attached to RA3, we use one side of a 3745-610. It is also running NCP V7R1.

The multi-path channel (MPC) links between the MVS systems are virtual channel-to-channel connections provided by VM.

The local terminals are also virtual. We access them by dialing the appropriate MVS system from the VM logon screen.

The PS/2s* used in the tests are PS/2 Model 80s, each with 16MB of RAM. Each PS/2 runs OS/2* 2.1. The PS/2s use standard 16/4 Micro Channel* Token-Ring Adapters. We use Communications Manager/2 1.1 with the latest enhancements as delivered to support the new APPN functions in VTAM V4R2.

2.2 Software Requirements for VTAM V4R2

This section defines the software level requirements for the VTAM V4R2 functions that we describe in the book. A full list of software requirements can be found in the "Programming Requirements" section of the VTAM V4R2 announcement.

- **NCP V7R1** is required where the NCP provides the boundary function support for **connection network** and **multiple network connectivity (border node connection)**.
- **MVS/ESA V4R3** is required for **APPN host-to-host channel**.
- **Dependent LU Server** support requires **Dependent LU Requester** software. At the time of writing, DLUR implementations are made available to us for **CM/2 1.1 (in VTAM-provided DLUR modules)** and for **3174 Licensed Internal Code Release 5**. The VTAM-provided DLUR in CM/2 is limited to VTAM V4R2 only. IBM has announced its intention to provide enhanced DLUR support that will ship as a standard part of a future release of CM/2.

- Formatting of VTAM **buffer traces** using ACF/TAP requires **SSP V3R8**.
- For **compatibility with border nodes** you need the **APPN function set 1013** implemented in the network nodes that will be connecting across the border nodes. At the time of writing this function is made available to us via refreshes to CM/2 1.1.

Chapter 3. APPN Host-to-Host Connection

In this chapter we describe how APPN protocols are now supported over a channel-to-channel connection between two VTAM hosts. This support is called APPN Host-to-Host Connection, or AHHC.

3.1 Benefits

1. For the first time, it is possible to make an APPN connection across a channel-to-channel link. This greatly increases the range of configurations that can be adapted to take advantage of the APPN dynamics. For example, many customers have data hosts which are connected to the network only via (real or virtual) channel-to-channel links. Converting these to APPN will save system definitions, in particular path tables, and the effort that goes into maintaining them.
2. The new APPN border node support (see Chapter 5, "Border Node" on page 33) is independent of the connection type, and therefore allows a cross-network connection over a channel-to-channel link. In a subarea network, SNI always requires an NCP to be present at the gateway between networks; the options for network design are thus much wider with APPN.
3. The ability to use multiple physical connections as a single TG improves reliability; it is possible to lose a number of channel-to-channel connections without breaking the sessions using the multipath channel. (VTAM V4R1 introduced this resilience for subarea connections; prior to that the loss of a channel-to-channel link would result in session outage).

3.2 How AHHC Works

In order to understand AHHC, it is important to distinguish between the physical connection used to join two processors together, and the channel protocol (DLC) used to communicate between the processors. The physical connection may be:

- An ESCON channel
- A piece of hardware such as the IBM 3088
- A software emulation that connects two virtual machines, such as VM′s virtual channel-to-channel connection
- Two pieces of hardware joined by a remote link, such as a pair of IBM 3172s

The DLC used by MVS/ESA can be either the traditional channel-to-channel (CTC) protocol, or the more recent multipath channel (MPC) protocol. In this chapter we use CTC and MPC to refer to *DLC protocols* rather than the physical connections.

AHHC will only work with the MPC protocol, rather than the more traditional CTC. MPC was first introduced with MVS/ESA V4R3, and supported for subarea connectivity by VTAM V4R1, as described in *VTAM V4R1 Implementation Guide*. MPC works with ESCON, with the 3088, and with the VM virtual channel-to-channel, which emulates a 3088. It does *not* work with older channel-to-channel hardware, nor with the 3172 in its channel extender configuration.

MPC uses multiple subchannels (channel/unit addresses) for communication, each subchannel being unidirectional. Each “read” subchannel on one processor must have a corresponding “write” subchannel on the other processor. Therefore, a minimum of two subchannels is required for an MPC link; more may be defined for greater availability.

AHHC is seen by APPN as a single-link transmission group (the APPN architecture does not currently support multilink TGs), but provides resilience against the loss of one or more subchannels as long as at least one path remains in each direction. Each subchannel carries an XID2 exchange to identify the individual connections, but the XID3 exchange is done only once for the whole MPC group. After the XID3 exchange, the same protocols flow over the MPC link as flow over a FID2 channel to an NCP.

3.3 Experiences with AHHC

The configuration used for the tests, and the results observed, are described in this section.

3.3.1 Description

Figure 2 shows the configuration used to exercise AHHC.

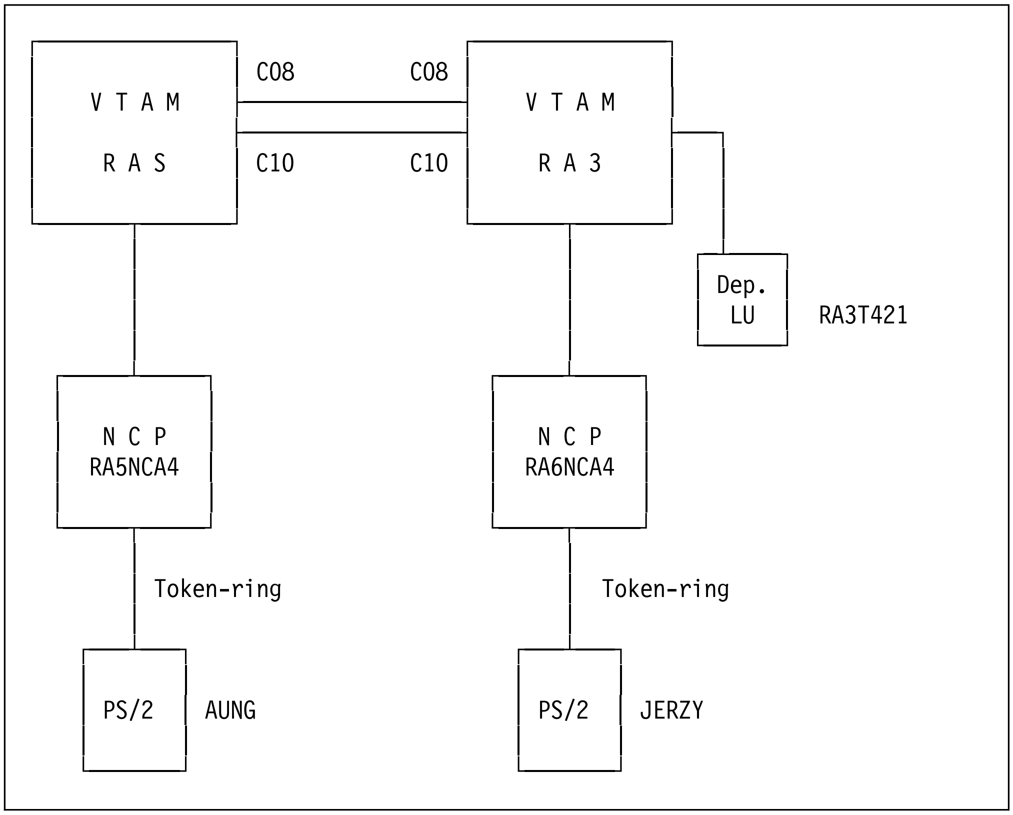

Figure 2. APPN Host-to-Host Connection

Both VTAMs are network nodes, and both PCs (attached via a token-ring LAN) are end nodes served by the VTAMs to which they are connected. There is no FID2 or FID4 connection between the VTAMs other than the AHHC link shown. The AHHC link comprises just two subchannels, one in each direction. RA3T421 is a dependent LU channel-attached to RA3.

3.3.2 Definitions

A “LOCAL” major node is used to define the AHHC connection, as shown in Figure 3 on page 14.

```
RASPUTIN VBUILD  TYPE=LOCAL
* APPN MPC LINK FROM RAS TO RA3
RASPC8A  PU   TRLE=RASKMPC,PUTYPE=2,XID=YES   1


RA3PUTIN VBUILD  TYPE=LOCAL
* APPN MPC LINK FROM RA3 TO RAS
RA3PCA8  PU   TRLE=RA3KMPC,PUTYPE=2,XID=YES   1
```

Figure 3. VTAM Definitions for AHHC in RAS and RA3

The major difference between this and a connection to an NCP is the use of the TRLE operand **1** on the PU instead of a channel address. TRLE points to the AHHC definitions that are contained in the TRL major nodes shown in Figure 4.

```
RASKAHHC VBUILD  TYPE=TRL
* APPN MPC LINK FROM RAS TO RA3
RASKMPC  TRLE   LNCTL=MPC,READ=(C08),WRITE=(C10)   1


RA3KAHHC VBUILD  TYPE=TRL
* APPN MPC LINK FROM RA3 TO RAS
RA3KMPC  TRLE   LNCTL=MPC,READ=(C10),WRITE=(C08)   1
```

Figure 4. TRL Major Nodes for RAS and RA3

Each TRLE (Transport Resource List Element) entry in this major node identifies an AHHC link; in our case we have just a single TRLE in each VTAM **1** which comprises one read and one write subchannel. Care is needed to ensure that the read subchannels on one side correspond to the write subchannels on the other.

3.3.3 Test Results

With the PCs already connected to their network node servers, we start by activating the TRL major nodes RASPUTIN and RA3PUTIN in RAS and RA3 respectively. Like PATH major nodes, TRL major nodes cannot be displayed or deactivated by name; to display their status we issue D NET,TRL as shown in Figure 5.

```
NCCF                      N E T V I E W    RA3AO BUCZAK   01/28/94 10:59:04
* RA3AO    D NET,TRL
  RA3AO    IST097I  DISPLAY  ACCEPTED
' RA3AO
IST350I  DISPLAY TYPE = TRL
IST1314I  TRLE = RA3KMPC   STATUS = ACTIV       CONTROL = MPC
IST314I  END
-------------------------------------------------------------------------------
```

Figure 5. Display of TRL Resources

The TRLEs themselves in fact do not become active until the activation of the link station which uses them; immediately after activating the TRL major node the TRLEs will display as NEVAC.

Activation of the link stations (the LOCAL major nodes) from each side results in the APPN connection and the CP-CP sessions becoming active, as shown in Figure 6 on page 15. RA3PUTIN has already been activated from RA3 at this stage, so the connections come up immediately.

```
NCCF                      N E T V I E W    RASAN BUCZAK   01/27/94 11:03:01
* RASAN    ACT RASPUTIN
C RASAN    VARY NET,ACT,ID=RASPUTIN
  RASAN    IST097I  VARY     ACCEPTED
  RASAN    IST093I  RASPUTIN ACTIVE
  RASAN    IST093I  RASPC8A  ACTIVE
Q RASAN P% IST1086I  APPN CONNECTION FOR USIBMRA.RA3       IS ACTIVE - TGN =
           21
Q RASAN P% IST1096I  CP-CP SESSIONS WITH USIBMRA.RA3       ACTIVATED
-------------------------------------------------------------------------------
```

Figure 6. AHHC Link Station Activation

Once the connection is active, a display of a TRLE can be seen as shown in Figure 7.

```
NCCF                      N E T V I E W    RA3AO BUCZAK   01/28/94 11:09:26
C RA3AO    DISPLAY NET,ID=RA3KMPC,SCOPE=ALL
  RA3AO    IST097I  DISPLAY  ACCEPTED
' RA3AO
IST075I  NAME = RA3KMPC          , TYPE = TRLE      1
IST486I  STATUS= ACTIV     , DESIRED STATE= ACTIV
IST087I  TYPE = LEASED             , CONTROL = MPC  2
IST1221I  WRITE  DEV = 0C08  STATUS = ACTIVE        3
IST1221I  READ   DEV = 0C10  STATUS = ACTIVE        3
IST1068I  PHYSICAL RESOURCE (PHYSRSC) = RA3PCA8     4
IST314I  END
-------------------------------------------------------------------------------
```

Figure 7. TRLE Display

The display shows that the resource is a TRLE **1**, that the link is a multipath channel **2**, the write and read subchannel addresses **3**, and the name of the link station using the AHHC connection **4**.

A display of the topology centered on RA3 shows that the AHHC connection is treated exactly as any other APPN connection. Figure 8 on page 16 illustrates this.

```
 NCCF                      N E T V I E W    RA3AO BUCZAK   01/27/94 11:05:54
* RA3AO    D NET,TOPO,ID=RA3,LIST=ALL
  RA3AO    IST097I  DISPLAY  ACCEPTED
' RA3AO
IST350I  DISPLAY TYPE = TOPOLOGY
IST1295I  CP NAME           NODETYPE ROUTERES CONGESTION  CP-CP WEIGHT
IST1296I  USIBMRA.RA3       NN       128      NONE        *NA*  *NA*
IST1297I                    ICN/MDH  CDSERVR  RSN
IST1298I                    YES      NO       66
IST1223I                    BN       NATIVE
IST1224I                    NO       YES
IST1299I  TRANSMISSION GROUPS ORIGINATING AT CP USIBMRA.RA3
IST1357I                                                  CPCP
IST1300I  DESTINATION CP    TGN      STATUS   TGTYPE      VALUE WEIGHT
IST1301I  USIBMRA.RAK       21   1   OPER     INTERM      YES   *NA*
IST1301I  USIBMRA.RAS       21   2   OPER     INTERM      YES   *NA*
IST1301I  USIBMRA.RAS       22       INOP     INTERM      YES   *NA*
IST1301I  USIBMRA.BERND     21       INOP     INTERM      YES   *NA*
IST1301I  USIBMRA.RAS       23       INOP     INTERM      YES   *NA*
IST1301I  USIBMRA.RAS       24       INOP     INTERM      YES   *NA*
IST1301I  USIBMRA.RAA       21       INOP     INTERM      YES   *NA*
IST1301I  USIBMRA.RACONNET  21       INOP     INTERM      NO    *NA*
IST1301I  USIBMRA.JERZY     21   3   OPER     ENDPT       YES   *NA*
IST314I  END
-------------------------------------------------------------------------------
```

Figure 8. Topology Display from RA3

The topology database shows several TGs used in previous tests, which have since become inoperative. The ones in operation are:

- A channel link from RA6NCA4 to RAK, a VTAM V4R1 system **1**
- The AHHC link to RAS **2**
- The token-ring LAN link to the served end node JERZY **3**

After activating the AHHC link, we set up a session between a local screen (RA3T421) attached to RA3 and TSO on RAS. We also set up a session between the two PC end nodes. Both these sessions use the AHHC link as shown in the session path displays below.

Figure 9 on page 17 shows the subarea configuration of one of the two CP-CP sessions between RAS and RA3, as seen from RAS.

```
 NLDM.CON                  SESSION CONFIGURATION DATA                PAGE    1
 -------------- PRIMARY ---------------+-------------- SECONDARY --------------
 NAME RAS      SA 0000001C  EL 0007   |   NAME RA3  1  SA 0000001C  EL 0067
 --------------------------------------+---------------------------------------
 DOMAIN RASAN      PCID USIBMRA.RAS.F627D1647EA6BDCA             DOMAIN RASAN
               +-------------+                 +-------------+
 RAS           |     CP      | ---         --- |     CP      | RA3
               +-------------+  APPN TP 03     +-------------+



                              APPNCOS CPSVCMG

                              LOGMODE CPSVCMG

                              SADJ CP RA3




SELECT PT, ST (PRI, SEC TRACE), RT (RESP TIME), P, AR, FC
CMD==>
```

Figure 9. CP Session to RA3 - Session Configuration from RAS

In the subarea part of the network, RA3 is seen as an adjacent control point residing on a link station owned by RAS, so it is shown as being in the same subarea (28) **1** as RAS.

The APPN view of the same session is shown in Figure 10.

```
 NLDM.AR              APPN SESSION ROUTE CONFIGURATION                PAGE   1
 -- PRIMARY ---+-- SECONDARY --+--------------- PCID ---------------+- DOMAIN -
 NAME RAS      | NAME RA3      | USIBMRA.RAS.F627D1647EA6BDCA       | RASAN
 --------------+---------------+------------------------------------+----------

  +---------+
  |   CP    |
  |RAS      |
  +----+----+
 TG021 |
  +----+----+
  |   CP    |
  |RA3      |
  +---------+






END OF DATA
SELECT PAR, SAR
CMD==>
```

Figure 10. CP Session between RAS and RA3 - APPN View

The LU 2 session between the dependent LU on RA3 and TSO on RAS is depicted in Figure 11 on page 18.

```
 NLDM.CON                  SESSION CONFIGURATION DATA                PAGE    1
 -------------- PRIMARY ---------------+-------------- SECONDARY --------------
 NAME RASAT01   SA 0000001C  EL 000A   |   NAME RA3T421   SA 0000001C  EL 006B
 --------------------------------------+---------------------------------------
 DOMAIN RASAN      PCID USIBMRA.RAS.F627D1647EA6BDCC         C-C DOMAIN RASAN
                +-------------+                   +-------------+
 RAS            |   CP/SSCP   | ---           --- |   CP/SSCP   | RAS
 ISTPUS28(0000) | SUBAREA PU  |  APPN TP 01       | SUBAREA PU  | ISTPUS28(0000)
                +------+------+                   +------+------+
                       |                                 |
                +------+------+                   +------+------+
 RASAT01 (000A) |     LU      |                   |     CUA     |
                +-------------+                   +------+------+
                               APPNCOS #CONNECT          |
                                                  +------+------+
                               LOGMODE M2BSCNQ    |     PU      | RASPC8A (0065)
                                                  +------+------+
                               SADJ CP RA3               |
                                                  +------+------+
                                                  |     ILU     | RA3T421 (006B)
                                                  +-------------+    1

SELECT PT, ST (PRI, SEC TRACE), RT (RESP TIME), P, AR, FC
CMD==>
```

Figure 11. LU 2 Session - Session Configuration from RAS

Note that the dependent LU RA3T421 is shown as an independent LU **1** on the link station to RA3, since RA3 is its adjacent CP. The APPN view of this session is identical to that shown in Figure 10 on page 17, since the same CPs own the session endpoints.

RAS's view of the LU 6.2 session between the PC end nodes is displayed in Figure 12 on page 19.

```
NLDM.CON                   SESSION CONFIGURATION DATA                PAGE    1
--------------- PRIMARY ----------------+--------------- SECONDARY ---------------
NAME JERZY     SA 0000001C  EL 0069   |   NAME AUNG      SA 00000005  EL 0098
----------------------------------------+----------------------------------------
DOMAIN RASAN  C-C  PCID USIBMRA.JERZY.D54B4E4B303E90D4           DOMAIN RASAN
               +-------------+                  +-------------+
RAS            |   CP/SSCP   | ---          ---  |             |
ISTPUS28(0000) | SUBAREA PU  |  APPN TP 02       | SUBAREA PU  | RA5NCA4 (0000)
               +------+------+  SUBA TP 00       +------+------+
                      |              VR 00              |
               +------+------+       ER 00       +------+------+
        (0064) |     CUA     |      RER 00       |    LINK     | J0005013
               +------+------+                   +------+------+
                      |       APPNCOS #INTER            |
               +------+------+                   +------+------+
RASPC8A (0065) |     PU      | LOGMODE #INTER    |     PU      | W32223  (001E)
               +------+------+ PADJ CP RA3       +------+------+
                      |        SADJ CP AUNG             |
               +------+------+                   +------+------+
JERZY   (0069) |     ILU     |                   |     ILU     | AUNG    (0098)
   1           +-------------+                   +-------------+    2

SELECT PT, ST (PRI, SEC TRACE), RT (RESP TIME), P, ER, VR, AR
CMD==>
```

Figure 12. LU 6.2 Session - Session Configuration from RAS

Here we can see that JERZY is shown as an independent LU **1** on the AHHC link station, since RA3 is the adjacent CP on the APPN route. AUNG is shown as an independent LU **2** attached to the NCP, since RA5NCA4 provides the subarea boundary function for the APPN link.

The APPN view of the LU 6.2 session is shown in Figure 13.

```
NLDM.AR              APPN SESSION ROUTE CONFIGURATION               PAGE   1
-- PRIMARY ---+-- SECONDARY --+--------------- PCID ---------------+- DOMAIN -
NAME JERZY    | NAME AUNG     | USIBMRA.JERZY.D54B4E4B303E90D4     | RASAN
--------------+---------------+------------------------------------+----------
                                      TG021 |
 +---------+                           +----+----+
 |   CP    |                           |   CP    |
 |JERZY    |                           |AUNG     |
 +----+----+                           +---------+
TG021 |
 +----+----+
 | CP(ICN) |
 |RA3      |
 +----+----+
TG021 |
 +----+----+
 |   CP    |
 |RAS      |
 +----+----+
      |

END OF DATA
SELECT PAR, SAR
CMD==>
```

Figure 13. LU 6.2 Session - APPN View

3.4 Hints and Tips

The implementation of AHHC is quite straightforward and we have encountered no difficulty in implementing it in several of our test scenarios. Points to note are:

- Remember to match the read and write subchannels on the two sides of the link.
- The TRL major node behaves rather like a path table. It is not possible to deactivate TRLEs, but one may add, delete or modify them by using the UPDATE= operand on the VARY ACT command.
- Remember the prerequisites:
  - MVS/ESA V4R3, or later
  - ESCON, 3088 or virtual CTC only

Chapter 4. Connection Network

This chapter describes new features that allow VTAM V4R2 to connect to a connection network. A connection network is represented by a single virtual routing node that exists on a shared access transport facility, such as a local area network (LAN). This virtual routing node is also referred to in the VTAM manuals as a virtual node.

The chapter contains sample definitions for a 3172 using an XCA (External Communications Adapter) major node and for a 3745 using NCP.

4.1 Benefits

A connection network enables direct communication between the APPN nodes connected to a shared transport facility such as a local area network.

Without a connection network the user has one of two choices:

1. Predefine a link between every pair of nodes on the LAN.
2. Accept that session routes can take a longer route than is strictly necessary, passing through the connections that *have* been defined.

A connection network enables optimum routing on a shared transport facility *without* the need to define every possible path.

4.2 How Connection Network Works

VTAM V4R2 implements the connection network function for token-ring, CSMA/CD and FDDI LANs. There is a known requirement for an extension of the architecture to support connection networks over X.25.

Figure 14 on page 22 shows an example of a token-ring-connected APPN network. It has a PS/2, a VTAM end node and a VTAM network node connected to a token-ring LAN. Without a connection network and without an explicitly defined link between the end nodes, the session data between the end nodes will need to be routed via the VTAM network node even though a direct logical connection between the two end nodes is possible using the LAN. This is shown in Figure 14 on page 22.

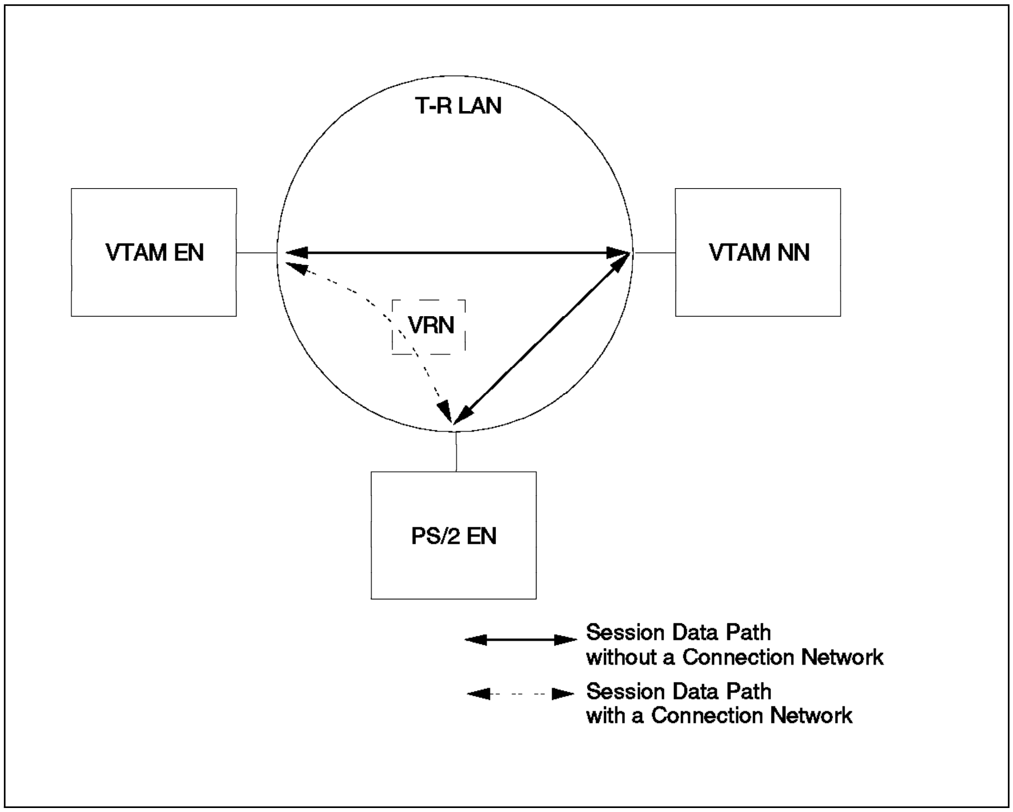

Figure 14. Sample Configuration

In a connection network, a “virtual routing node” (VRN) must be defined. The APPN nodes participating in the connection network have a TG to the virtual routing node. This means that once the optimal route is determined to be through the virtual routing node, the session data can flow between the two end nodes directly through the LAN. Figure 14 shows the direct LAN connection between the two end systems once the route is determined to be through the virtual routing node.

If the PS/2 wishes to establish a session with the VTAM EN, it sends a Locate search request to the VTAM NN (its network node server) for the destination LU in the VTAM EN. The VTAM NN performs route calculation as usual, but it knows that the VRN does not really exist, and that any route through it is really one hop instead of two. It therefore takes the weights of the two TGs in and out of the VRN, and uses their average instead of adding them. (The VRN itself has a node weight of zero, so its presence does not affect route calculation.) Having determined that such a route now has the lowest weight, the VTAM NN returns the RSCV to the PS/2 in the Locate reply.

The RSCV contains TG Descriptor control vectors describing TGs from the PS/2 EN to the virtual routing node and from the virtual routing node to the VTAM EN. The TG Descriptor control vector contains DLC signalling information (MAC and LSAP addresses) for the VTAM EN.

Using the information contained in the RSCV, the PS/2 EN can now establish a direct connection to the VTAM EN if one does not already exist. Once the connection is established, the session will be routed over it. This direct connection will not be reported to the topology. No CP-CP session is allowed over this direct connection. Once the session ends, after a suitable timeout period, the PS/2 EN will terminate the direct connection to VTAM EN.

VTAM V4R1, when configured as a network node, was able to calculate routes correctly over a connection network for other nodes to use. What is new in VTAM V4R2 is the ability to *use* a connection network. VTAM V4R2 can:

- Notify other nodes of its connection to the VRN.
- Adjust the route calculation to take account of the presence of a VRN.
- Dynamically define *dial-out* connections as well as dial-in.
- Use correctly the information presented in the RSCV, and establish a session across the connection network.

4.3 Experiences with Connection Network

This section describes the testing of the connection network function. We use a token-ring LAN environment for a connection network.

4.3.1 Description

In the test environment, we use three MVS VTAM V4R2 systems and a PS/2 Communications Manager/2 system as shown in Figure 15.

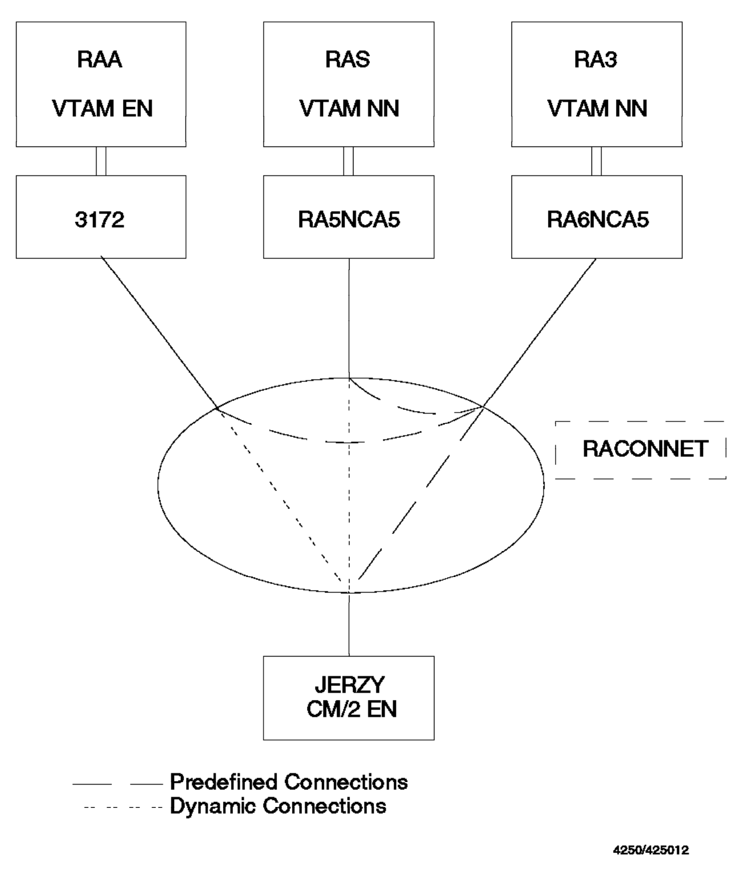

Figure 15. Test Environment for Connection Network

The first MVS system has a CP name of RAA (Raleigh System A) and is defined as an end node. It uses a 3172 to connect to the token-ring LAN.

The second MVS system has a CP name of RAS (Raleigh System S) and is defined as a network node. It uses a 3745 with a TIC (Token-Ring Interface Coupler) to connect to the LAN.

The third MVS system has a CP name of RA3 (Raleigh System 3) and is defined as a network node. It also uses a 3745 with a TIC to connect to the LAN.

The PS/2 system is running Communications Manager/2. It has a CP name of JERZY and is defined as an end node. RA3 is its network node server (NNS).

Figure 16 shows the APPN view of the network. Note that all four systems are connected to the same token-ring LAN. All CP-CP connections are established via the LAN.

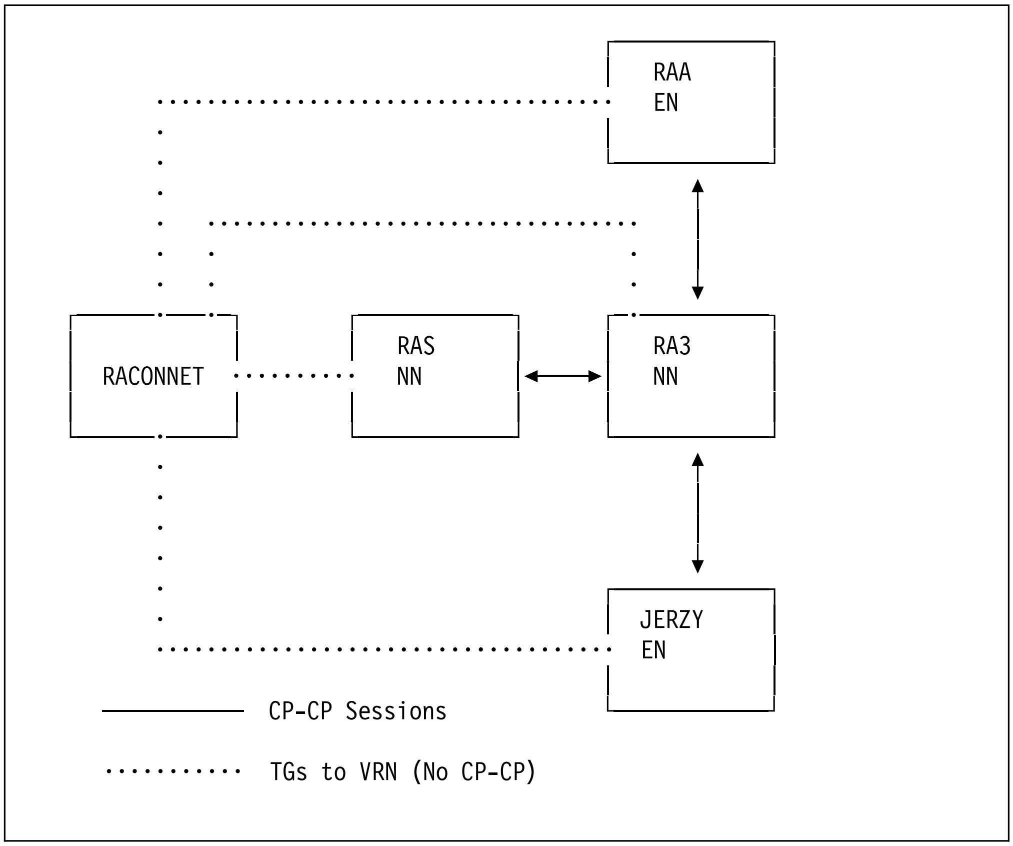

Figure 16. APPN View of the Test Environment

4.3.2 Definitions

This section contains the system definitions for the connection network test environment. The important definitions are VTAM key words **VNNAME** and **VNGROUP**, which define the virtual routing node name and the line group to be used for virtual routing node connections.

VTAM associates a connection network with a *physical* port, although the actual connections are made over *logical* lines. Therefore, on a shared NCP (or in the case of SSCP takeover) the VTAM which owns the physical port provides the connection network facilities.

For a connection network, VTAM needs the ability to define PUs dynamically for outgoing as well as incoming calls. It can do this for outgoing calls since it knows the destination MAC address from the session RSCV. Dynamically defined PUs on a connection network are given names of the form xxVnnnnn, where xx is a user-defined identifier (CN by default) and nnnnn is a five-digit number incremented by one every time a connection is made. Thus the first dynamic PU on a connection network is named CNV00001, as opposed to CN000001 which it would be named on a real connection.

Incoming PU names can also be predefined, or created via the configuration services exit (as happens in our tests).

4.3.2.1 RAA Definitions

In RAA, we define the XCA (External Communications Adapter) major node for the 3172, and a switched major node for the dial-out connection to the network node server (RAS). These definitions are shown in Figure 17 and Figure 18 on page 26 respectively. Note the VNGROUP **1** and VNNAME **2** parameters in the PORT definition for the token-ring adapter in the 3172. VNGROUP points to the real (logical) line group **3** for the 3172, and tells VTAM that the connection network is defined on this group of links. The DYNPU parameter is set to YES on the logical line group definition, although it defaults to YES if VNNAME and VNGROUP are coded.

```
RAA3172 VBUILD TYPE=XCA
**
RAA3172P PORT  ADAPNO=0,                                              * X
               CUADDR=EA2,                                            * X
               VNGROUP=RAAXCAG1, 1   VIRTUAL NODE GROUP               * X
               VNNAME=RACONNET,  2   VIRTUAL NODE NAME                * X
               MEDIUM=RING,                                           * X
               SAPADDR=4,                                             * X
               TIMER=60
**
RAAXCAG1 GROUP DIAL=YES,CALL=INOUT,DYNPU=YES 3
*
RAAXL1   LINE
RAAXP1   PU
RAAXL2   LINE
RAAXP2   PU
```

Figure 17. 3172 Definitions

```
        VBUILD TYPE=SWNET,MAXNO=20,MAXGRP=5
 RAA301   PU    ADDR=01,MAXDATA=265,MAXOUT=7,MAXPATH=5,PASSLIM=7,        X
                PUTYPE=2,                                                X
                CPNAME=RA3,                                              X
                NETID=USIBMRA,                                           X
                CONNTYPE=APPN,                                           X
                DYNLU=YES,                                               X
                CPCP=YES
 *
 RA3301P  PATH  DIALNO=0004400001060000,           IN NCP RA6NCA         X
                GRPNM=RAAXCAG1                      IN 3172 RAA3172
```

Figure 18. Switched PU Definition for Dial Out

4.3.2.2 RA3 Definitions

In RA3, we define the token-ring connection in the NCP major node, and a switched major node for the dial-out connection to the adjacent network node (RAS). These definitions can be seen in Figure 19 and Figure 20 on page 27 respectively. Note the VNNAME parameter **4** in the physical line definition for the TIC and VNGROUP **5** pointing to the logical group.

```
*----------------------------------------------------------------------
* NTRI PHYSICAL DEFINITIONS
*----------------------------------------------------------------------
RA6GTRP1 GROUP ECLTYPE=(PHYSICAL,ANY), ALLOW SA,PERIPHERAL AND IP      *
               ....
*
RA6LA88  LINE ADDRESS=(1088,FULL),                                     *
               LOCADD=400001060000,  LOCAL ADMINISTERED ADDRESS        *
               VNGROUP=RA6GTRL2, 5   LOGICAL GROUP FOR CONNECTION NET  *
               VNNAME=RACONNET, 4    CONNECTION NETWORK NAME           *
               ....
*
RA6PA881 PU ADDR=01,                 SNA NTRI ATTACHMENT ADDR          *
               ....
*
*----------------------------------------------------------------------
*PERIPHERAL LOGICAL GROUP  NTRI TIC 2 - CCUA                          *
*----------------------------------------------------------------------
RA6GTRL2 GROUP ECLTYPE=LOGICAL,                                        *
               ...
               ...
*
```

Figure 19. NCP Definitions

```
       VBUILD TYPE=SWNET,MAXNO=20,MAXGRP=5
*  DEFINITION FOR ITSC28
RA3U01   PU    ADDR=01,MAXDATA=265,MAXOUT=7,MAXPATH=5,PASSLIM=7,       X
               PUTYPE=2,                                               X
               CPNAME=RAS,                                             X
               NETID=USIBMRA,                                          X
               CONNTYPE=APPN,                                          X
               DYNLU=YES,                                              X
               CPCP=YES
*
RA3U01P  PATH  DIALNO=0004400001050000,           IN NCP RA5NCA.      X
               GRPNM=RA6GTRL2                      IN NCP RA6NCA.
```

Figure 20. Switched PU Definition for Dial Out

4.3.2.3 RAS Definitions

In RAS, we define the token-ring connection in the NCP major node in similar fashion to RA3. Figure 21 illustrates this.

```
*-----------------------------------------------------------------------
* NTRI PHYSICAL DEFINITIONS
*-----------------------------------------------------------------------
RA5GTRP1 GROUP ECLTYPE=(PHYSICAL,ANY), ALLOW SA,PERIPERAL AND IP      *
               ....
*
RA5LA88  LINE ADDRESS=(1088,FULL),                                    *
               LOCADD=400001050000,  LOCAL ADMINISTERED ADDRESS       *
               VNGROUP=RA5GTRL2,     LOGICAL GROUP FOR CONNECTION NET  *
               VNNAME=RACONNET,      CONNECTION NETWORK NAME          *
               ....
*
RA5PA881 PU ADDR=01,                 SNA NTRI ATTACHMENT ADDR         *
               ....
*
*-----------------------------------------------------------------------
*PERIPHERAL LOGICAL GROUP  NTRI TIC 2 - CCUA                         *
*-----------------------------------------------------------------------
RA5GTRL2 GROUP ECLTYPE=LOGICAL,                                       *
               ...
               DYNPU=YES,            CREATE PU DYNAMICALLY            *
               ...
*
```

Figure 21. NCP Definitions

4.3.2.4 JERZY Definitions

In the PS/2, using the CM/2 Setup facility, we define the node as an end node in the **SNA local node characteristics** profile. In the same profile, we define RA3 as its preferred network node server, providing the MAC address of the TIC in the 3745 connected to RA3.

We also define in the **DLC - Token-Ring and other LAN types** profile the virtual routing node name as a network-qualified name in the "connection network name" parameter.

The setup process produces a Network Definition File (NDF) containing all the configuration parameters. Figure 22 on page 28 shows the NDF listing for JERZY.

```
DEFINE_LOCAL_CP  FQ_CP_NAME(USIBMRA.JERZY    )
                  CP_ALIAS(JERZY   )
                  NAU_ADDRESS(INDEPENDENT_LU)
                  NODE_TYPE(EN)
                  NODE_ID(X'05D32258')
                  HOST_FP_SUPPORT(YES);

 DEFINE_CONNECTION_NETWORK  FQ_CN_NAME(USIBMRA.RACONNET )
                            ADAPTER_INFO( DLC_NAME(IBMTRNET)
                                          ADAPTER_NUMBER(0));

 DEFINE_LOGICAL_LINK  LINK_NAME(LINK0002)
                      ADJACENT_NODE_TYPE(NN)
                      PREFERRED_NN_SERVER(YES)
                      DLC_NAME(IBMTRNET)
                      ADAPTER_NUMBER(0)
                      DESTINATION_ADDRESS(X'400001060000')
                      CP_CP_SESSION_SUPPORT(YES)
                      ACTIVATE_AT_STARTUP(YES)
                      LIMITED_RESOURCE(NO)
                      ....

 DEFINE_DEFAULTS  IMPLICIT_INBOUND_PLU_SUPPORT(YES)
                  ....

 DEFINE_TP  TP_NAME(ATELLD)
            ....


 START_ATTACH_MANAGER;
```

Figure 22. NDF File for JERZY

4.3.3 Test Results

The aim of the test is to establish LU 6.2 sessions between JERZY and RAA and between JERZY and RAS, and to show that the session path uses the direct connections across the LAN rather than routing via RA3. The LU 6.2 sessions are established using an application called ATELL. The activities and results are outlined below.

In RAA, we activate the 3172. As shown in **1** in Figure 23, we get a console message confirming that the connection to the virtual routing node is established.

```
V NET,ACT,ID=RAA3172
IST097I  VARY     ACCEPTED
IST093I  RAA3172  ACTIVE
IST1168I  VIRTUAL NODE USIBMRA.RACONNET  CONNECTION ACTIVE  1
```

Figure 23. 3172 Activation

We then issue the VARY DIAL command to establish the link from RAA to RA3 as shown in Figure 24 on page 29. Note the CPCP=YES parameter **2** in the command. We get a console message confirming that the CP-CP session is established **3**.

```
V NET,DIAL,ID=RAA301,CPCP=YES 2
IST590I  CONNECTOUT ESTABLISHED FOR PU RAA301   ON LINE RAAXL1
IST1086I  APPN CONNECTION FOR USIBMRA.RA3       IS ACTIVE - TGN =  21
IST1096I  CP-CP SESSIONS WITH USIBMRA.RA3       ACTIVATED  3
```

Figure 24. Dial-Out Connection to RA3

During the activation of NCP from RAS and RA3, we see the messages shown below on the respective consoles.

```
            .
            .
            .
IST093I RA5PA881 ACTIVE
IST093I R55PA882 ACTIVE
IST1168I  VIRTUAL NODE USIBMRA.RACONNET  CONNECTION ACTIVE
            .
            .
            .
IST093I RA6PA881 ACTIVE
IST093I RA6PA882 ACTIVE
IST1168I  VIRTUAL NODE USIBMRA.RACONNET  CONNECTION ACTIVE
```

Figure 25. Activation of TIC

We now establish the link from RA3 to RAS. We check the console to ensure that the CP-CP session is established.

```
V NET,DIAL,ID=RA3U01,CPCP=YES
IST590I  CONNECTOUT ESTABLISHED FOR PU RA3U01   ON LINE J0006001
IST1086I  APPN CONNECTION FOR USIBMRA.RAS       IS ACTIVE - TGN =  21
IST1096I  CP-CP SESSIONS WITH USIBMRA.RAS       ACTIVATED
```

Figure 26. Dial-Out Connection to RA3

From the PS/2, we activate the link to RA3 by using the Subsystem Management Facility.

The display of the APPN topology in RA3 (see Figure 27 on page 30) now shows operative TGs to RAS and JERZY as well as to the virtual routing node (RACONNET).

```
* RA3AO    D NET,TOPO,LIST=ALL,ID=RA3
   RA3AO    IST097I  DISPLAY  ACCEPTED
 ′ RA3AO
 IST350I  DISPLAY TYPE = TOPOLOGY
 IST1295I  CP NAME           NODETYPE ROUTERES CONGESTION  CP-CP WEIGHT
 IST1296I  USIBMRA.RA3       NN       128      NONE        *NA*  *NA*
 IST1297I                    ICN/MDH  CDSERVR  RSN
 IST1298I                    YES      NO       68
 IST1223I                    BN       NATIVE
 IST1224I                    NO       YES
 IST1299I  TRANSMISSION GROUPS ORIGINATING AT CP USIBMRA.RA3
 IST1357I                                                  CPCP
 IST1300I  DESTINATION CP    TGN      STATUS   TGTYPE      VALUE WEIGHT
 IST1301I  USIBMRA.RAK       21       OPER     INTERM      YES   *NA*
 IST1301I  USIBMRA.RAS       21       OPER     INTERM      YES   *NA*
 IST1301I  USIBMRA.RACONNET  21       OPER     INTERM      NO    *NA*
 IST1301I  USIBMRA.JERZY     21       OPER     ENDPT       YES   *NA*
 IST314I  END
```

Figure 27. Topology Display on RA3

We can also display the topology centered on RACONNET, as the VRN appears in the topology database. Figure 28 demonstrates this.

```
 * RA3AO    D NET,TOPO,ID=RACONNET,LIST=ALL,APPNCOS=#INTER
   RA3AO    IST097I  DISPLAY  ACCEPTED
 ′ RA3AO
 IST350I  DISPLAY TYPE = TOPOLOGY
 IST1295I  CP NAME           NODETYPE ROUTERES CONGESTION  CP-CP WEIGHT
 IST1296I  USIBMRA.RACONNET  VN  1    128      NONE        *NA*  0  2
 IST1297I                    ICN/MDH  CDSERVR  RSN
 IST1298I                    NO       NO       0
 IST1223I                    BN       NATIVE
 IST1224I                    NO       YES
 IST1299I  TRANSMISSION GROUPS ORIGINATING AT CP USIBMRA.RACONNET
 IST1357I                                                  CPCP
 IST1300I  DESTINATION CP    TGN      STATUS   TGTYPE      VALUE WEIGHT
 IST1301I  USIBMRA.RA3       21       OPER     INTERM      NO    240  3
 IST1301I  USIBMRA.RAS       21       OPER     INTERM      NO    240
 IST314I  END
```

Figure 28. Topology Display on RA3

Note that:

- RACONNET is described as "VN" **1** in the display.
- The VRN has a node weight of zero **2**.
- The TGs from the VRN have non-zero weights **3** to allow the same degree of flexibility in route calculation across a connection network as across a real link. As explained above, the exit and entry TG weights will be averaged when calculating routes.

We are now ready to establish the LU 6.2 sessions. We invoke ATELL in JERZY to send a message from JERZY to ATELL in RAA. This results in an LU 6.2 session between ATELL in JERZY and ATELL in RAA. The ATELL test is completed successfully, and we can see the following message on the console for RAA indicating that a direct connection between JERZY and RAA has been established.

```
IST590I  CONNECTIN  ESTABLISHED FOR PU W32258   ON LINE RAAXL9
IST1086I  APPN CONNECTION FOR USIBMRA.JERZY     IS ACTIVE - TGN =   1
```

Figure 29. Console Messages Showing Direct Connection from JERZY to RAA

While the ATELL LU 6.2 session is active, we use the facilities provided by CM/2 Subsystem Management to verify in JERZY that a new logical connection (LLC connection) is established between JERZY and RAA and that the ATELL LU 6.2 session is using that connection.

Similar results are achieved by using ATELL from JERZY to RAS. A direct connection is seen on the console as shown below.

```
IST590I  CONNECTIN  ESTABLISHED FOR PU W32258   ON LINE J000500F
IST1086I  APPN CONNECTION FOR USIBMRA.JERZY     IS ACTIVE - TGN =    1
```

Figure 30. Console Messages Showing Direct Connection from JERZY to RAS

4.4 Hints and Tips

When dial-out connections are used, PUs in the receiving (dialed) system can be created by using a predefined switched major node or by using dynamic PU generation, either by the VTAM exit or by the built-in VTAM functions.

We find that the standard configuration services exit supplied with VTAM (ISTEXCCS) does not work for VTAM-to-VTAM connections, as it relies on a predefined range of Node ID values that does not include the values used by VTAM. Therefore, we either use a switched major node definition in the dialed system or we turn off the exit using the MODIFY EXIT command in the dialed system. Turning off the exit results in VTAM using its internal dynamic PU generation function.

If you find that the PU names generated by VTAM′s dynamic PU generation function do not meet your needs, you can define a switched major node or modify the supplied exit to meet your naming standards.

Chapter 5. Border Node

The address space of SNA allows more than 1 000 000 000 000 000 000 000 000 network accessible units (NAUs). That is enough for the foreseeable future. For administrative purposes, this huge amount is divided into separate networks (also called subnetworks) with unique net IDs. Their uniqueness is made certain by central registration of net IDs by IBM. Each (sub)network still has control over the names of more than 10 to the power of 12 NAUs.

Architecture and implementations provide several methods of crossing network borders from one net ID to another net ID. They are listed below. Only the last two methods in the list (border node) are covered in this chapter. The others are mentioned here to put border node into context. The methods are:

SNI
SNA Network Interconnect. The network border goes through a gateway NCP, which is controlled by a gateway SSCP. The SNI gateway is an end point for routes, but it is transparent to the search process and LU-LU sessions. SNI is very powerful, but rather complex. SNI is widely used.

NNNC
Non-Native Network Connection. It connects nodes with different net IDs using an LEN level connection. It is limited to PLU-initiated sessions (such as LU 6.2). It implies the disadvantages of LEN, in particular the requirement for predefinition of destination LUs. NNNC is of lesser importance.

APPN EN
APPN End Node. It connects an APPN end node to an APPN network node with a different net ID. This method is standard in *all* APPN end node implementations making it very easy to connect one node to different networks.

APPN PBN
APPN Peripheral Border Node. It connects two APPN network nodes with different net IDs, one of which is started as a border node and pretends to be an end node. A peripheral border node connection can be set up between a VTAM border node and any network node. Peripheral border node connections are also supported by the peripheral border node implementation of the AS/400. Some restrictions apply to peripheral border node connections. They are explained later in this chapter.

APPN EBN
APPN Extended Border Node. It connects two APPN network nodes with different net IDs. Both nodes are started as extended border nodes and are aware of the role of the partner. The function is implemented in VTAM V4R2. It is an easy way to connect different APPN networks.

This chapter concentrates on the border node functions in VTAM. Several scenarios show connections between VTAMs working as extended border nodes. Connections between a VTAM border node and a normal APPN network node are covered, too. This latter connection is a peripheral subnetwork connection. Apart from showing our test configurations, we document definitions, explain commands and messages, and write about our observations during the test.

5.1 Benefits

1. The border node function allows two VTAM network nodes with different net IDs to connect. This connection provides directory services for the search process and LU-LU connectivity.
2. The border node function allows two VTAM network nodes to isolate their topologies from one another, no matter if they have the same or different net IDs. This results in reduction of network overhead.
3. The border node function allows a VTAM network node to connect to any network node with a different net ID, provided this network node implements the APPN base function set 1013. The topology of one network is isolated from another, but directory services and LU-LU session connectivity spans the subnetworks via Border Node.
4. The COS definitions on both sides are isolated. This further reduces the requirement for coordinated definitions.
5. The extended border node function between two VTAM border nodes is activated with minimal definitions. The border nodes use dynamic processes to exchange their capabilities and information on reachable networks.
6. Cross network connections can be made through channel-to-channel links.

5.2 How the Border Node Works

A concise introduction to VTAM′s border node function is given in *VTAM Network Implementation Guide*, in a section named “Multiple Network Connectivity.” A detailed description of the APPN border node architecture can be found in *APPN Architecture and Product Implementations Tutorial*, which also has the merit of covering the considerations for some APPN products concerning the interoperability with VTAM as a border node.

Basically, an APPN border node is a network node with additional functions. The main functions among them are:

- Isolation of topology
- Avoidance of loops in search path and session path
- Dynamic discovery of net IDs
- Compatibility with “normal” network nodes

Of course, the normal functions of a network node are present in the border node, as well. They are for example:

- Intermediate session routing for LU-LU sessions
- Propagating searches for directory services

5.2.1.1 Isolation of Topology

During XID exchange, two adjacent extended border nodes inform each other of their capabilities and agree not to propagate the network topology between one another. They establish CP-CP sessions, and exchange information that they are border nodes. In short, they are aware of each other′s existence and capabilities.

The extended border node not only connects APPN networks with different net IDs. It also is capable of dividing a network with one net ID into separate topology subnetworks or clusters, still using the same net ID. This may be done to reduce the number of topology updates flowing on the CP-CP sessions, or to overcome a limitation in the size of the topology database present in small network nodes, such as the 3174.

Isolation of topology has a drawback: route calculation is done piece-wise. The optimum route is found in each topology network. The result may or may not be the optimum route end-to-end. This drawback does not matter when there is a limited number of borders. For large networks, it can be overcome by controlling the search process.

5.2.1.2 Avoidance of Loops in Search Path and Session Path

A broadcast search in APPN is a parallel process: all adjacent nodes are searched concurrently. The process stops only when all nodes have been searched. This is done to get the fastest possible response.

Across network borders the search process is serialized. This reduces the network overhead. A border node is able to shield other nodes from searches of LUs, which are not in their domain.

As there may be parallel border nodes it is necessary to prevent a search from crossing the network border several times. This is achieved by additional information in control vectors.

5.2.1.3 Dynamic Discovery of Net IDs

You do not have to predefine the location of resources or adjacent networks. Border nodes remember paths to specific resources and net IDs via XID exchanges and the results of previous locate searches. This is adequate for small or middle sized networks.

If you choose, you can have more control over the locations searched for a resource in a specific net ID. This is advisable in large networks for performance reasons.

5.2.1.4 Compatibility with Peripheral Border Nodes

The AS/400 implements a function, which can connect two network nodes in different net IDs. This function is called peripheral border node. It can connect two and only two networks with different net IDs. A peripheral border node portrays itself as a network node to the nodes in its native network.

The peripheral border node portrays itself as an end node to a network node in a non-native network.

Two such AS/400 peripheral border nodes can be connected across a network border. By negotiation they decide which will act as an end node.

If a peripheral border node connects to a VTAM V4R2 extended border node, the connection will be a peripheral border node connection managed by the extended border node. The peripheral border node will act as a "normal" network node for the connection, in the sense described in 5.2.1.5, "Compatibility with "Normal" Network Nodes" on page 36. VTAM will always take the role of the end node. This connection is shown in the left-hand side of Figure 31 on page 36. This connection type can only be used at the ends of a multi-network search path.

5.2.1.5 Compatibility with "Normal" Network Nodes

A border node can connect across a network boundary to a network node which does not have the border node function implemented or activated (the "normal" network node). During XID exchange, the border node and an adjacent network node inform each other that they are network nodes. If they have different net IDs, the border node in the next XID during the exchange sequence sets the node type to end node. (The connection becomes a peripheral subnetwork connection.) They establish CP-CP sessions. They do not exchange topology information. In short, the border node pretends to be an end node, and exploits the fact that all APPN network nodes accept connections to end nodes in other net IDs.

To properly route session BIND traffic over a peripheral border node connection the "normal" network node needs to understand the new information in control vectors. This is APPN architecture function set 1013 which is being implemented in all up-to-date network nodes. Details are given in *APPN Architecture and Product Implementations Tutorial*.

If a VTAM V4R2 border node connects to a network node, which has this function, the border node assumes the role of an end node, and the connection falls back to become a peripheral border connection. This kind of connection is shown in the right-hand side of Figure 31.

As you see, an extended border node can also have peripheral border connections. A peripheral border connection has less function than an inter-cluster link between extended border nodes. It can only cross one network border on its own. If there are extended border nodes on the session path, more than one border can be crossed. But the peripheral border node connection will have to be at the peripheral borders, that is towards the end points of the session.

We encourage you to read on, and study the examples for these functions.

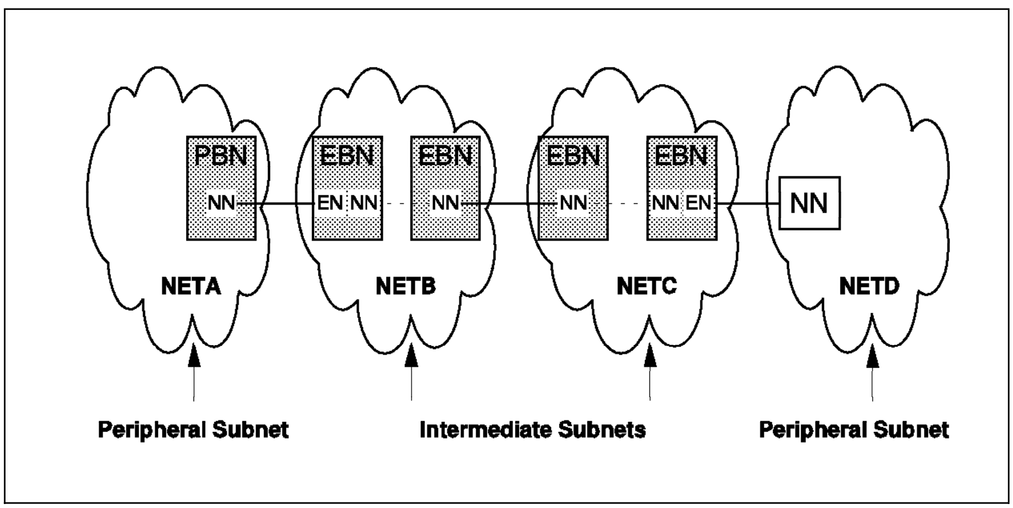

Figure 31. Border Nodes

5.3 Connecting Two Networks through VTAM Border Nodes

This section describes our tests of the simplest border node scenario. The configuration is symmetrical and uses defaults, wherever possible. The APPN nodes involved are VTAM with NCP. You can use the definitions as an example to connect two APPN networks with different net IDs.

5.3.1 Description

In the scenario shown in Figure 32, two composite network nodes are connected through an APPN connection across a token-ring LAN. Both composite network nodes consist of one VTAM and one NCP. The two composite network nodes are started as border nodes in networks with different net IDs. As both adjacent VTAMs are started as border nodes they assume the role of extended border nodes.

VTAM RAS in net ID USIBMRAT owns RA5NCA6, and VTAM RA3 in net ID USIBMRA owns RA6NCA4. Local terminals are used to set up sessions to TSO on the other side of the network border. The terminals are called RAST423 on VTAM RAS, and RA3T422 on VTAM RA3.

Additional APPN connections are from RA5NCA6 to a network node on a PS/2 (CP name AUNG) and from VTAM RA3 to an interchange node (CP name RAK), which in turn connects through several SNI gateways to the worldwide IBM network.

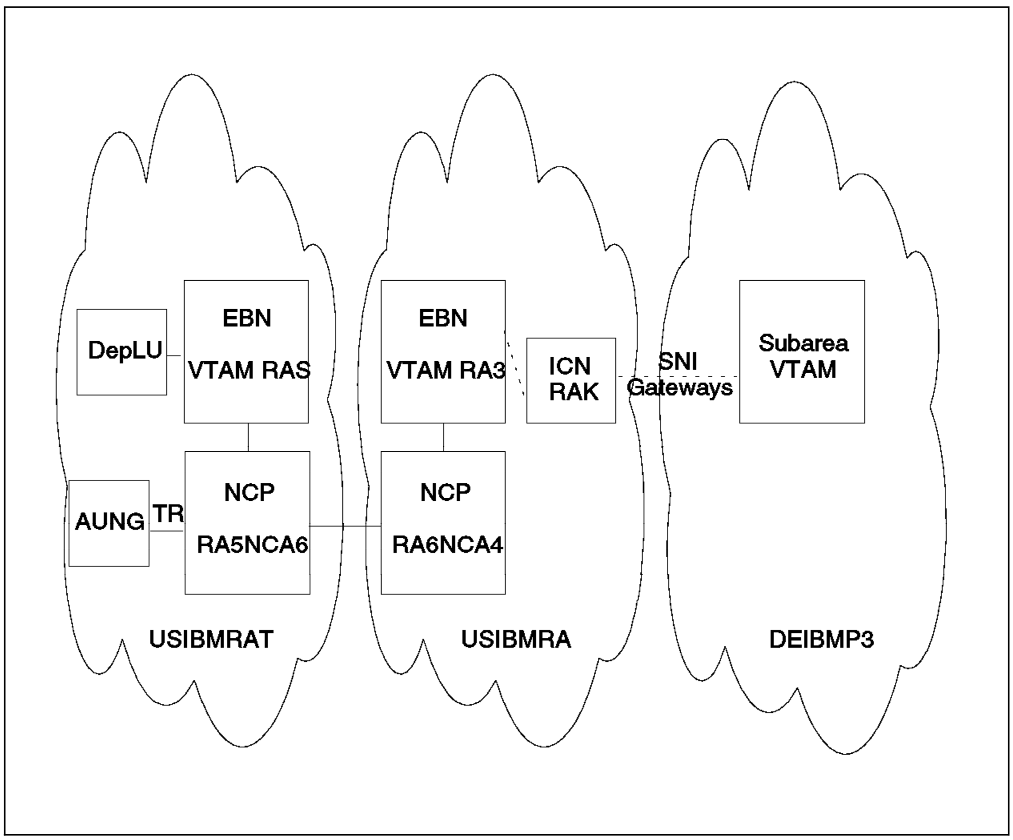

Figure 32. Connecting Two Net IDs through VTAM Border Nodes

5.3.2 Definitions

As the configuration in this scenario is symmetrical, the definitions are shown for just one VTAM and one NCP. A complete set of NCP definitions is shown in Appendix A.1, “Base NCPs” on page 193.

In order to activate the border node functions in VTAM, the VTAM start parameters need to be changed. This is one of the (few) changes, that cannot be done dynamically. It requires a restart of VTAM. In Figure 33 we show only the start parameters relating to the border node function. In this scenario we use the extended border node function; that means, the two VTAMs on *both* sides of the network border are started as border nodes.

```
DYNLU=YES,
NODETYPE=NN,
CPCP=YES,
BN=YES,NETID=USIBMRAT,                          BORDER NODE 1
XNETALS=YES,
```

Figure 33. VTAM Start Parameters Used for the Simple Border Node Scenario

APPN exploits dynamic discovery of resources. In order to allow VTAM to dynamically define CDRSCs for these resources we must specify CDRDYN=YES on the host CDRM coded in a CDRM major node. Since we are not using SSCP-SSCP sessions only this single CDRM definition is necessary. In our tests, we found it useful not to specify a net ID in the CDRM major node, making it possible to change the net ID in the start options.

XNETALS=YES becomes the default when specifying BN=YES. XNETALS=NO would prevent sessions from being set up across an APPN connection with a node with a different net ID.

The changes in the NCPs are minimal and relate to setting up our test scenario. Changing the NETID requires an NCP generation. As shown in Figure 34, we also specify DYNPU in order to work without a switched major node definition on the dialed-to side.

```
         BUILD BFRS=(240),                  NCP BUFFER SIZE         *
               ...
               NETID=USIBMRAT,    REQUIRED                          *
               ...
RA5GTRL2 GROUP ECLTYPE=LOGICAL,                                     *
               ...                                                  *
               DYNPU=YES,           CREATES SWITCH PU DYNAMICALLY   *
               ....
```

Figure 34. NCP Changes for Border Node

The dialing side requires a switched major node, though. The complete major node is shown in Figure 35 on page 39. The PU name will be referred to in the DIAL operator command.

```
       VBUILD TYPE=SWNET,MAXNO=20,MAXGRP=5
*  DEFINITION FOR ITSC28
RA3U01   PU    ADDR=01,MAXDATA=265,MAXOUT=7,MAXPATH=5,PASSLIM=7,       X
               PUTYPE=2,                                               X
               CPNAME=RAS,                                             X
               NETID=USIBMRAT,                                         X
               CONNTYPE=APPN,                                          X
               DYNLU=YES
*
RA3U01P  PATH  DIALNO=0004400001050000,           IN NCP RA5NCA.      X
               GRPNM=RA6GTRL2                      IN NCP RA6NCA.
```

Figure 35. Switched Major Node in VTAM RA3

5.3.3 Test Results

Before connecting the NCPs we display the topology on VTAM RAS. As you can see in Figure 36 **1**, VTAM RAS is defined as a border node but is not acting as one, yet **2**. Only after the first border node connection is established will RAS indicate that it is a border node in the APPN topology.

```
 D NET,TOPO,ID=RAS,LIST=ALL
 IST097I DISPLAY ACCEPTED
 IST350I DISPLAY TYPE = TOPOLOGY 722
 IST1295I CP NAME          NODETYPE ROUTERES CONGESTION  CP-CP WEIGHT
 IST1296I USIBMRAT.RAS     NN       128      NONE        *NA*  *NA*
 IST1297I                  ICN/MDH  CDSERVR  RSN
 IST1298I                  YES      NO       0
 IST1223I                  BN       NATIVE
 IST1224I                  NO  2    YES
 IST1299I TRANSMISSION GROUPS ORIGINATING AT CP USIBMRAT.RAS
 IST1308I RESOURCE WAS NOT FOUND IN THE TOPOLOGY DATABASE
 IST314I END
D NET,VTAMOPTS,OPTION=BN
IST097I DISPLAY ACCEPTED
IST1188I ACF/VTAM V4R2 STARTED AT 13:31:27 ON 01/27/94 731
IST1349I COMPONENT ID IS 5695-11701-201
IST1348I VTAM STARTED AS INTERCHANGE NODE
IST1189I BN       = YES  1
IST314I END
```

Figure 36. Topology of VTAM RAS before Dialing

The switched major node shown in Figure 35 is activated. The connection between the two networks is established by using the V NET,DIAL operator command. The command and the resulting messages are shown in Figure 37 on page 40. The ID used (RA3U01) is the PU in the switched major node. CPCP=YES is added, as switched lines do not allow CP-CP sessions by default.

```
v net,dial,id=ra3u01,cpcp=yes
IST097I VARY ACCEPTED
IST590I  CONNECTOUT ESTABLISHED FOR PU RA3U01   ON LINE J0006001
IST1086I  APPN CONNECTION FOR USIBMRAT.RAS      IS ACTIVE - TGN =  21
IST241I  VARY DIAL  COMMAND COMPLETE FOR RA3U01
IST1096I  CP-CP SESSIONS WITH USIBMRAT.RAS      ACTIVATED
```

Figure 37. Establishing APPN Connection between NCPs

After CP-CP sessions are established, the CPs exchange CP capabilities, as we can verify in a trace. The CPs identify themselves as border nodes, which do not accept regular topology updates. Minimal topology updates are done, though. This is reflected in Figure 38, where the adjacent border node is included in the count of NNs and BNs (marked **3**). Also, the link between the border nodes is included, with a TGTYPE of INTERCLUST (marked **4**). This identifies the link as going to a different topology subnetwork (topology cluster).

```
D NET,TOPO
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = TOPOLOGY 795
IST1306I LAST CHECKPOINT   ADJ  NN   EN   SERVED EN CDSERVR ICN  BN
IST1307I NONE              1    2 3  0    0         0       1    2 3
IST314I END
D NET,TOPO,ID=RAS,LIST=ALL
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = TOPOLOGY 799
IST1295I CP NAME           NODETYPE ROUTERES CONGESTION  CP-CP WEIGHT
IST1296I USIBMRAT.RAS      NN       128      NONE        *NA*  *NA*
IST1297I                   ICN/MDH  CDSERVR  RSN
IST1298I                   YES      YES      4
IST1223I                   BN       NATIVE
IST1224I                   YES      YES
IST1299I TRANSMISSION GROUPS ORIGINATING AT CP USIBMRAT.RAS
IST1357I                                                 CPCP
IST1300I DESTINATION CP    TGN      STATUS   TGTYPE      VALUE WEIGHT
IST1301I USIBMRA.RA3       21       OPER     INTERCLUST  YES   *NA*
IST314I END                                      4
```

Figure 38. Topology of VTAM RAS after Dialing

We establish cross-network LU-LU sessions in both directions. We use local terminals to log on to TSO in the other network. On VTAM RAS, we capture the NLDM view of one of the LU-LU sessions. This is shown in Figure 39 and Figure 40 on page 41. The trace shows that the search flows between border nodes are operating normally. Even though the border node function isolates topologies, it provides full connectivity for directory services and LU-LU sessions.

From VTAM RA3 we also activate a link to VTAM RAK, our “link to the rest of the world.” The connection is reflected in the topology as a normal transmission group between network nodes (INTERM). This is shown in Figure 41 on page 42, marked **5**. We verify that an LU on VTAM RAS can reach an application across an inter-cluster connection and several SNI gateways.

```
NLDM.AR               APPN SESSION ROUTE CONFIGURATION                PAGE   1
-- PRIMARY ---+-- SECONDARY --+--------------- PCID ---------------+- DOMAIN -
NAME RA3AT02  | NAME RAST423  | USIBMRA.RA3.F6FF41645F727A73       | RASAN
--------------+---------------+------------------------------------+----------

 +---------+
 |   CP    |
 |RA3      | NETWORK: USIBMRA
 +----+----+
TG021 |
 +----+----+
 | CP(ICN) |
 |RAS      | NETWORK: USIBMRAT
 +---------+
```

Figure 39. How NLDM on VTAM RAS Sees the Route

```
NLDM.CON                  SESSION CONFIGURATION DATA                 PAGE    1
-------------- PRIMARY ---------------+-------------- SECONDARY --------------
NAME RA3AT02   SA 00000005  EL 009B   |   NAME RAST423   SA 0000001C  EL 0058
--------------------------------------+---------------------------------------
DOMAIN RASAN  C-C  PCID USIBMRA.RA3.F6FF41645F727A73        NNNA DOMAIN RASAN
               +-------------+                 +-------------+
               |             | ---         --- |   CP/SSCP   | RAS
RA5NCA6 (0000) | SUBAREA PU  |  APPN TP 01     | SUBAREA PU  | ISTPUS28(0000)
               +------+------+  SUBA TP 00     +------+------+
                      |              VR 00            |
               +------+------+       ER 00     +------+------+
J0005013       |    LINK     |      RER 00     |     CUA     | 0423
               +------+------+                 +------+------+
                      |      APPNCOS #CONNECT         |
               +------+------+                 +------+------+
CN000001(001E) |     PU      | LOGMODE M2BSCNQ |     LU      | RAST423 (0058)
               +------+------+ PADJ CP RA3     +-------------+
                      |
               +------+------+
RA3AT02 (009B) |     ILU     |
               +-------------+
```

Figure 40. How NLDM on VTAM RAS Sees the Session

In this context, the VTAM start option BNDYN is worth a consideration. Its default is LIMITED. This value is not adequate when completely relying on dynamic network discovery. The LU-LU session establishment can fail as shown in Figure 42 on page 42. After changing BNDYN to FULL, the session is established successfully. We recommend including BNDYN=FULL in the start parameters for all border nodes when used in small or medium sized APPN networks. In these types of networks, the advantages of BNDYN=FULL for plug-and-play flexibility outweigh the extra searches that occur when VTAM is permitted to search all possible locations for a resource. This start option can be changed on-line by operator command.

In large networks or where search performance is critical, BNDYN=LIMITED will be preferred, in combination with adjacent cluster definitions as shown in 5.8.3.7, “Using Adjacent Cluster Definitions in VTAM” on page 77.

```
IST1295I CP NAME           NODETYPE ROUTERES CONGESTION  CP-CP WEIGHT
IST1296I USIBMRA.RA3       NN       128      NONE        *NA*  *NA*
IST1297I                   ICN/MDH  CDSERVR  RSN
IST1298I                   YES      YES      70
IST1223I                   BN       NATIVE
IST1224I                   YES      YES
IST1299I TRANSMISSION GROUPS ORIGINATING AT CP USIBMRA.RA3
IST1357I                                                 CPCP
IST1300I DESTINATION CP    TGN      STATUS   TGTYPE      VALUE WEIGHT
IST1301I USIBMRA.RAK       21       OPER     INTERM 5    YES   *NA*
IST1301I USIBMRAT.RAS      21       OPER     INTERCLUST  YES   *NA*
```

Figure 41. Topology Display in VTAM RA3

```
>IST663I  INIT SELF  REQUEST              FAILED  , SENSE=087D0001
>IST664I  REAL  OLU=USIBMRA.RA3T420     REAL  DLU=USIBMRAG.RAAAT
>IST889I  SID = F6FF416460405859
>IST894I  ADJSSCPS TRIED  FAILURE SENSE   ADJSSCPS TRIED  FAILURE SENSE
>IST895I     ISTAPNCP        08400007
>IST314I  END
F NET03,VTAMOPTS,BNDYN=FULL
IST097I MODIFY ACCEPTED
IST223I MODIFY COMMAND COMPLETED
```

Figure 42. Effect of BNDYN on the Search Process

We compare VTAM′s view of the topology to that of CM/2. As CM/2 does not implement the border node function, it cannot handle any information it receives about border nodes. Consequently, the topology information on VTAM RAS is different from that of CM/2 AUNG. As can be seen in Figure 43 **6**, VTAM also counts its adjacent border node. As Figure 45 on page 43 **7** shows, CM/2 does not.

The reason for this discrepancy is that VTAM RAS, as a border node, can see the partner border node on the other side of the network border. VTAM′s view is shown in Figure 44 on page 43. CM/2, on the other hand, is a normal network node. It can only see the network nodes inside the network (topology network), as is shown in Figure 46 on page 43.

```
* RASAN    D NET,TOPO
  RASAN    IST097I  DISPLAY  ACCEPTED
IST350I  DISPLAY TYPE = TOPOLOGY
IST1306I  LAST CHECKPOINT   ADJ  NN   EN   SERVED EN CDSERVR ICN  BN
IST1307I  NONE              2    3 6  1    0         0       1    2
IST314I  END
```

Figure 43. Topology Display by VTAM RAS in USIBMRAT

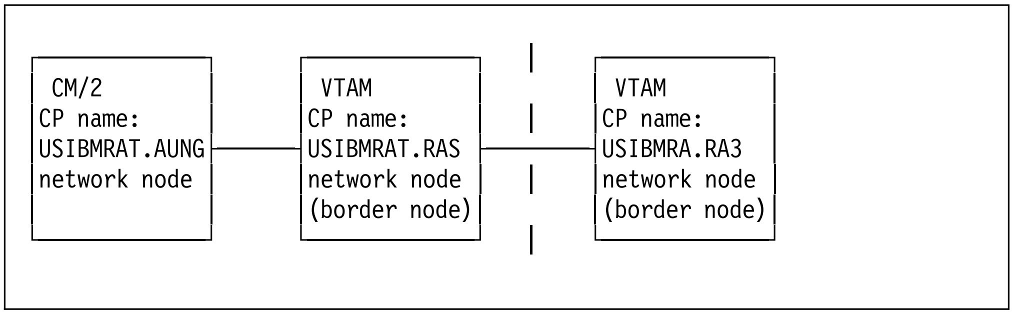

Figure 44. The View from VTAM RAS in USIBMRAT

```
*************************************
*        Topology Information        *
*************************************
Number of network nodes                    2 7

1>Network node CP name                     USIBMRAT.RAS
  Route additional resistance              128
  Congested?                               No
  Quiescing?                               No
  ISR depleted?                            No
  Number of TGs                            1


2>Network node CP name                     USIBMRAT.AUNG
  Route additional resistance              128
  Congested?                               No
  Quiescing?                               No
  ISR depleted?                            No
  Number of TGs                            2
```

Figure 45. Topology Display by CM/2 AUNG in USIBMRAT (Extract)

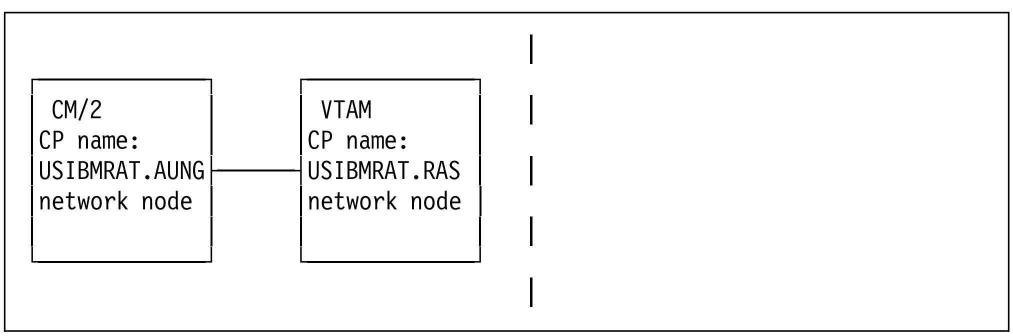

Figure 46. The View from CM/2 AUNG in USIBMRAT

5.4 Peripheral Subnet Boundaries

This section describes two test scenarios where VTAM provides inter-subnetwork boundaries using peripheral border node connections. Such a subnetwork boundary is called a "peripheral subnet boundary." We provide two example configurations and describe the cross-network session setup capabilities of each configuration.

5.4.1 Description

In the test environment, we use three MVS VTAM V4R2 systems and two PS/2 Communications Manager/2 systems. Figure 47 on page 45 shows the schematic diagram of the test environment. The system environment contains three networks.

- RAA, a VTAM V4R2 network node with the net ID of USIBMRAG, has an MPC connection to RAS. The MPC connections use AHHC as described in Chapter 3, "APPN Host-to-Host Connection" on page 11.

  It also has a 3172, which provides the LAN connection.

- BERND, a PS/2 running Communications Manager/2, is connected to RAA via the LAN and the 3172. It is defined as an end node and has the net ID of USIBMRAG.
- RAS is a VTAM V4R2 composite network node. It has the net ID of USIBMRA. It connects to RAA via an AHHC connection. It connects to RA3 via the token-ring LAN. Connectivity to the token-ring LAN is provided by a dedicated NCP.
- RA3 is also a VTAM V4R2 composite network node. It has the net ID of USIBMRAY. It connects to RAS via the token-ring LAN. Connectivity to the token-ring LAN is provided by a dedicated NCP.
- AUNG, a PS/2 running Communications Manager/2, is connected to RA3 via the LAN. It is defined as an end node and has the net ID of USIBMRAY.

We will use VTAM V4R2 on both sides of the peripheral subnet boundary in order to analyze the peripheral border node function. This is done for testing purposes only. Two adjacent VTAM V4R2 nodes can communicate via an extended subnet boundary and we encourage that whenever possible you use extended subnet boundaries.

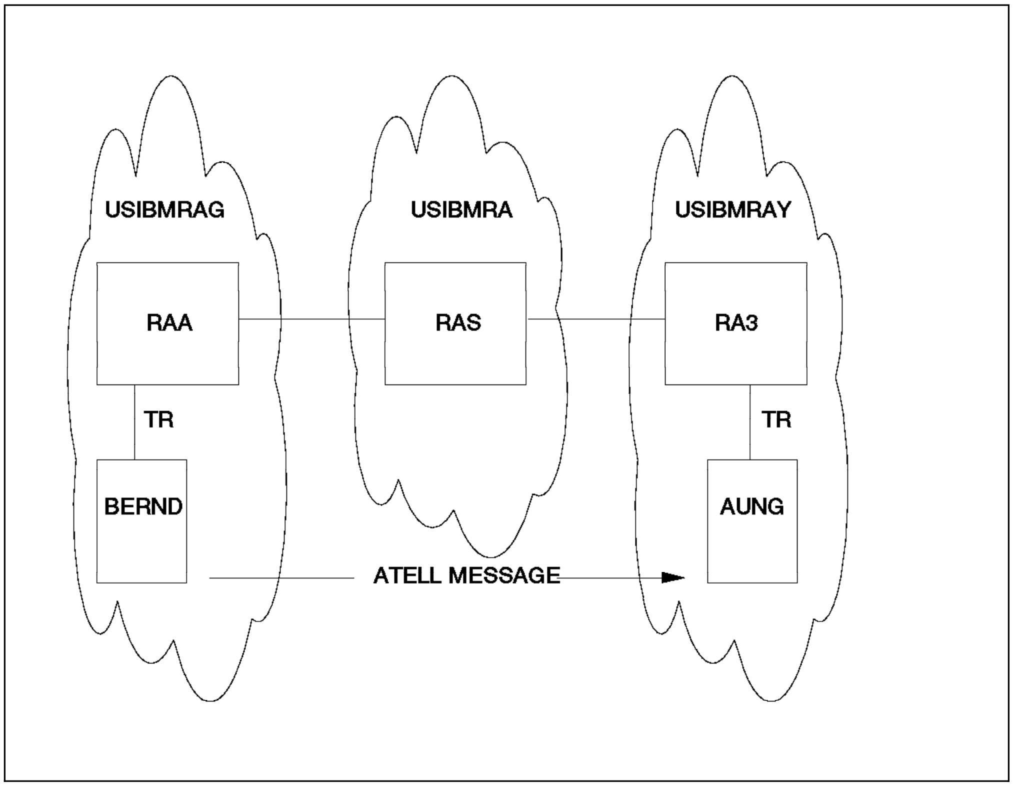

Figure 47. Schematic Diagram of the Test Environment

Figure 48 shows the inter-subnetwork links and peripheral subnet boundaries in the test configuration. We create the peripheral subnet boundaries by turning on the border node function on only one side of the inter-subnetwork link. So in the first test RAA and RA3 will be border nodes, and in the second test RAS will be a border node. This in described in detail in 5.4.3, "Test Results" on page 47.

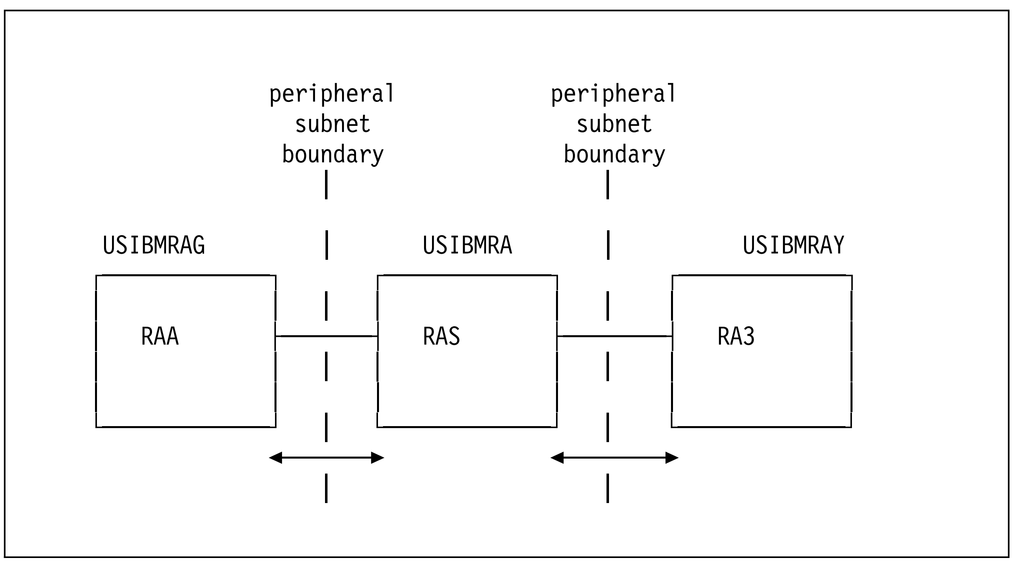

Figure 48. Inter-Cluster Links in the Test Environment

5.4.2 Definitions

This section contains the system definitions of the test environment for the peripheral subnet boundaries.

5.4.2.1 RAA Definitions

In RAA, we define the XCA (External Communications Adapter) major node for the 3172 and the AHHC connections to RAS and RA3. Figure 49, Figure 50 and Figure 51 show the definitions.

```
RAA3172 VBUILD TYPE=XCA
**
RAA3172P PORT  ADAPNO=0,                                                   * X
               CUADDR=EA2,                                                 * X
               MEDIUM=RING,                                                * X
               SAPADDR=4,                                                  * X
               TIMER=60
**
RAAXCAG1 GROUP DIAL=YES,CALL=INOUT,DYNPU=YES
*
RAAXL1   LINE
RAAXP1   PU
```

Figure 49. 3172 Definitions

```
RAAAHHC  VBUILD  TYPE=TRL
RAAMPCS  TRLE   LNCTL=MPC,READ=(C12),WRITE=(C14)
RAAMPC3  TRLE   LNCTL=MPC,READ=(C16),WRITE=(C18)
```

Figure 50. TRL Major Node in RAA

```
RAALCA   VBUILD  TYPE=LOCAL
RAAMPCSP PU   TRLE=RAAMPCS,PUTYPE=2,XID=YES
RAAMPC3P PU   TRLE=RAAMPC3,PUTYPE=2,XID=YES
```

Figure 51. LOCAL Major Node in RAA

5.4.2.2 RAS Definitions

In RAS, we define the AHHC connection to RAA and a switched major node for the dial-out connection from RA3.

```
RASAHHCA VBUILD  TYPE=TRL
RASMPCA  TRLE   LNCTL=MPC,READ=(C14),WRITE=(C12)
```

Figure 52. TRL Major Node in RAS

```
RASLCAA  VBUILD  TYPE=LOCAL
RASMPCAP PU   TRLE=RASMPCA,PUTYPE=2,XID=YES
```

Figure 53. LOCAL Major Node in RAS

```
         VBUILD TYPE=SWNET,MAXNO=20,MAXGRP=5
RASU01   PU     ADDR=01,MAXDATA=265,MAXOUT=7,MAXPATH=5,PASSLIM=7,      X
                PUTYPE=2,                                              X
                CPNAME=RA3,                                            X
                NETID=USIBMRAY,                                        X
                CONNTYPE=APPN,                                         X
                DYNLU=YES,                                             X
                CPCP=YES
*
RASU01P  PATH  DIALNO=0004400001060000,           IN NCP RA6NCA.      X
                GRPNM=RA5GTRL2                     IN NCP RA5NCA.
```

Figure 54. Definition for Switched Connection from RAS to RA3

5.4.2.3 RA3 Definitions

In RA3, we define the switched major node for the connection to RAS.

5.4.2.4 AUNG and BERND Definitions

We define each as an end node with the net ID of the respective network node server. We also define the connection to the preferred network node server by providing the appropriate MAC address.

5.4.3 Test Results

The aim of the test is to establish LU 6.2 sessions and LU 2 sessions across the two peripheral subnet boundaries. We use ATELL to establish LU 6.2 sessions and log on to TSO to establish LU 2 sessions.

We divide the tests into two configurations. The first configuration has the two nodes at either end providing the border node functions. The second configuration has the node in the middle providing the border node functions. Figure 55 on page 48 shows the two configurations.

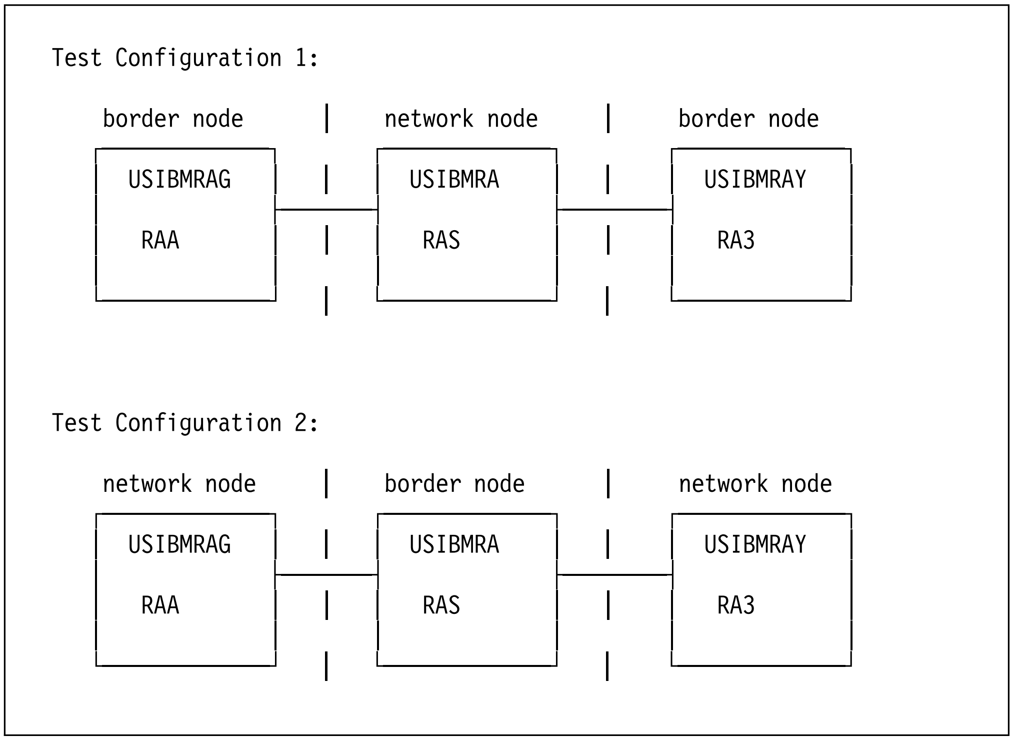

Figure 55. Two Test Configurations for Peripheral Subnet Boundaries

5.4.3.1 Initialization — Test Configuration 1

We start the VTAMs by issuing:

```
S NETA0,,,(LIST=A0,BN=YES,BNDYN=FULL,NETID=USIBMRAG)
S NET28,,,(LIST=S0,BN=NO,NETID=USIBMRA)
S NET03,,,(LIST=30,BN=YES,BNDYN=FULL,NETID=USIBMRAY)
```

Note the start option **BNDYN=FULL** which means that entries are dynamically added to the adjacent cluster table in VTAM even if the net ID of the entry added does not match the net ID to which the table applies. This means that VTAM will search adjacent subnetworks even if the adjacent subnetworks do not have the same net ID as the resource that VTAM is searching for.

The AHHC connections are established by activating the appropriate major nodes in each VTAM. The token-ring connections are established using the **VARY DIAL** command as shown below.

```
V NET,DIAL,ID=RASU01,CPCP=YES
```

Once the VTAM systems are active, the links from the PS/2s, AUNG and BERND, are activated.

5.4.3.2 Test Results — Test Configuration 1

Note that the inter-subnetwork links are seen by RAA and RA3, which are border nodes, as having the TGTYPE of INTERCLUST. Please see **1** in Figure 56 on page 49 and **2** in Figure 58 on page 49 for the results of the DISPLAY TOPO command. In contrast, the same subnetwork links are seen by RAS, which is not a border node, as links to end nodes. They are shown in the topology display for RAS as having the TGTYPE of ENDPT. Please see **3** and **4** in Figure 57 on page 49 for examples.

```
* RAAAN    D NET,TOPO,ID=RAA,LIST=ALL
  RAAAN    IST097I  DISPLAY  ACCEPTED
' RAAAN
IST350I  DISPLAY TYPE = TOPOLOGY
IST1295I  CP NAME           NODETYPE ROUTERES CONGESTION  CP-CP WEIGHT

IST1296I  USIBMRAG.RAA      NN       128      NONE        *NA*  *NA*
IST1297I                    ICN/MDH  CDSERVR  RSN
IST1298I                    YES      YES      4
IST1223I                    BN       NATIVE
IST1224I                    YES      YES
IST1299I  TRANSMISSION GROUPS ORIGINATING AT CP USIBMRAG.RAA
IST1357I                                                  CPCP
IST1300I  DESTINATION CP    TGN      STATUS   TGTYPE      VALUE WEIGHT
IST1301I  USIBMRA.RAS       21       OPER     INTERCLUST  YES   *NA*  1
IST1301I  USIBMRAG.BERND    21       OPER     ENDPT       YES   *NA*
IST314I  END
```

Figure 56. Topology in RAA

```
* RASAN    D NET,TOPO,ID=RAS,LIST=ALL
  RASAN    IST097I  DISPLAY  ACCEPTED
' RASAN
IST350I  DISPLAY TYPE = TOPOLOGY
IST1295I  CP NAME           NODETYPE ROUTERES CONGESTION  CP-CP WEIGHT
IST1296I  USIBMRA.RAS       NN       128      NONE        *NA*  *NA*
IST1297I                    ICN/MDH  CDSERVR  RSN
IST1298I                    YES      NO       109770
IST1223I                    BN       NATIVE
IST1224I                    NO       YES
IST1299I  TRANSMISSION GROUPS ORIGINATING AT CP USIBMRA.RAS
IST1357I                                                  CPCP
IST1300I  DESTINATION CP    TGN      STATUS   TGTYPE      VALUE WEIGHT
IST1301I  USIBMRAG.RAA      21       OPER     ENDPT       YES   *NA*  3
IST1301I  USIBMRAY.RA3      21       OPER     ENDPT       YES   *NA*  4
IST314I  END
```

Figure 57. Topology in RAS

```
* RA3AO    D NET,TOPO,ID=RA3,LIST=ALL
  RA3AO    IST097I  DISPLAY  ACCEPTED
' RA3AO
IST350I  DISPLAY TYPE = TOPOLOGY
IST1295I  CP NAME           NODETYPE ROUTERES CONGESTION  CP-CP WEIGHT
IST1296I  USIBMRAY.RA3      NN       128      NONE        *NA*  *NA*
IST1297I                    ICN/MDH  CDSERVR  RSN
IST1298I                    YES      YES      10
IST1223I                    BN       NATIVE
IST1224I                    YES      YES
IST1299I  TRANSMISSION GROUPS ORIGINATING AT CP USIBMRAY.RA3
IST1357I                                                  CPCP
IST1300I  DESTINATION CP    TGN      STATUS   TGTYPE      VALUE WEIGHT
IST1301I  USIBMRA.RAS       21       OPER     INTERCLUST  YES   *NA*  2
IST1301I  USIBMRAY.AUNG     21       OPER     ENDPT       YES   *NA*
IST314I  END
```

Figure 58. Topology in RA3

Now we invoke ATELL in USIBMRAG.BERND to send a message to USIBMRAY.AUNG. It fails because the resource cannot be located.

We try to log on to TSO in RA3. It fails for the same reason. We try both ATELL and TSO logon in the reverse direction — that is, going from USIBMRAY to USIBMRAG. They too fail.

We are successful in both ATELL and TSO logon across only one peripheral subnet boundary.

In summary, the test results show that, in the configuration tested where the middle node is not a border node, it is not possible to establish sessions across two peripheral subnet boundaries.

5.4.3.3 Initialization — Test Configuration 2

We start the VTAMs by issuing:

```
S NETA0,,,(LIST=A0,BN=NO,NETID=USIBMRAG)
S NET28,,,(LIST=S0,BN=YES,BNDYN=FULL,NETID=USIBMRA)
S NET03,,,(LIST=30,BN=NO,NETID=USIBMRAY)
```

Note that in this configuration, only the middle node (RAS) is defined as a border node. The required connections for the test are started as described previously.

5.4.3.4 Test Results — Test Configuration 2

This time, the results of the topology displays are reversed. The border node RAS sees the inter-subnetwork links as having the TGTYPE of INTERCLUST. Please see **1** and **2** in Figure 60 on page 51. The same subnetwork links are seen by RAA and RA3, which are not border nodes, as links to end nodes. They are shown in the topology displays as having the TGTYPE of ENDPT. Please see **3** in Figure 59 and **4** in Figure 61 on page 51.

```
* RAAAN    D NET,TOPO,ID=RAA,LIST=ALL
  RAAAN    IST097I  DISPLAY  ACCEPTED
' RAAAN
IST350I  DISPLAY TYPE = TOPOLOGY
IST1295I  CP NAME          NODETYPE ROUTERES CONGESTION  CP-CP WEIGHT
IST1296I  USIBMRAG.RAA     NN       128      NONE        *NA*  *NA*
IST1297I                   ICN/MDH  CDSERVR  RSN
IST1298I                   YES      NO       0
IST1223I                   BN       NATIVE
IST1224I                   NO       YES
IST1299I  TRANSMISSION GROUPS ORIGINATING AT CP USIBMRAG.RAA
IST1357I                                                 CPCP
IST1300I  DESTINATION CP    TGN      STATUS   TGTYPE      VALUE WEIGHT
IST1301I  USIBMRA.RAS       21       OPER     ENDPT       YES   *NA*  3
IST1301I  USIBMRAG.BERND    21       OPER     ENDPT       YES   *NA*
IST314I  END
```

Figure 59. Topology in RAA

```
* RASAN    D NET,TOPO,ID=RAS,LIST=ALL
  RASAN    IST097I  DISPLAY  ACCEPTED
' RASAN
IST350I  DISPLAY TYPE = TOPOLOGY
IST1295I  CP NAME           NODETYPE ROUTERES CONGESTION  CP-CP WEIGHT
IST1296I  USIBMRA.RAS       NN       128      NONE        *NA*  *NA*
IST1297I                    ICN/MDH  CDSERVR  RSN
IST1298I                    YES      YES      8
IST1223I                    BN       NATIVE
IST1224I                    YES      YES
IST1299I  TRANSMISSION GROUPS ORIGINATING AT CP USIBMRA.RAS
IST1357I                                                  CPCP
IST1300I  DESTINATION CP    TGN      STATUS   TGTYPE      VALUE WEIGHT
IST1301I  USIBMRAG.RAA      21       OPER     INTERCLUST  YES   *NA*  1
IST1301I  USIBMRAY.RA3      21       OPER     INTERCLUST  YES   *NA*  2
IST314I  END
```

Figure 60. Topology in RAS

```
* RA3AO    D NET,TOPO,ID=RA3,LIST=ALL
  RA3AO    IST097I  DISPLAY  ACCEPTED
' RA3AO
IST350I  DISPLAY TYPE = TOPOLOGY
IST1295I  CP NAME           NODETYPE ROUTERES CONGESTION  CP-CP WEIGHT
IST1296I  USIBMRAY.RA3      NN       128      NONE        *NA*  *NA*
IST1297I                    ICN/MDH  CDSERVR  RSN
IST1298I                    YES      NO       0
IST1223I                    BN       NATIVE
IST1224I                    NO       YES
IST1299I  TRANSMISSION GROUPS ORIGINATING AT CP USIBMRAY.RA3
IST1357I                                                  CPCP
IST1300I  DESTINATION CP    TGN      STATUS   TGTYPE      VALUE WEIGHT
IST1301I  USIBMRA.RAS       21       OPER     ENDPT       YES   *NA*  4
IST1301I  USIBMRAY.AUNG     21       OPER     ENDPT       YES   *NA*
IST314I  END
```

Figure 61. Topology in RA3

We now invoke ATELL in USIBMRAG.BERND to send a message to USIBMRAY.AUNG. It succeeds.

We try to log on to TSO in RA3. It also succeeds. We try both ATELL and TSO logon in the reverse direction — that is, going from USIBMRAY to USIBMRAG. They too succeed.

We are successful in both ATELL and TSO logon across both peripheral subnet boundaries in this configuration.

In summary, the test results show that, in the configuration tested where the middle node is a border node, it is possible to establish sessions across two peripheral subnet boundaries. This is because RAS, which is an extended border node, is able to use extended border node capabilities to establish sessions across two peripheral subnet boundaries.

5.5 Inter-Network Channel Links

This section gives examples in which APPN networks with different net IDs are connected through channel links. Three different net IDs are used. Crossing network borders through a channel link between VTAM and NCP can also be achieved in a subarea network with SNI. Crossing network borders through a channel link between two VTAMs is a new function, which *cannot* be done with SNI.

5.5.1 Description

As shown in Figure 62, in this scenario two composite network nodes and one VTAM network node are connected through APPN channel links. They are started as border nodes in networks with three different net IDs. The composite network nodes consist of one VTAM and one NCP each. VTAM RAS in net ID USIBMRAT owns NCP RA5NCA6, VTAM RA3 in net ID USIBMRA owns NCP RA6NCA4. VTAM RAA in net ID USIBMRAG is a pure network node without NCP. Local terminals are used to set up sessions to TSO across two network borders. The terminals are called RAAT420 on VTAM RAA, and RA3T420 on VTAM RA3.

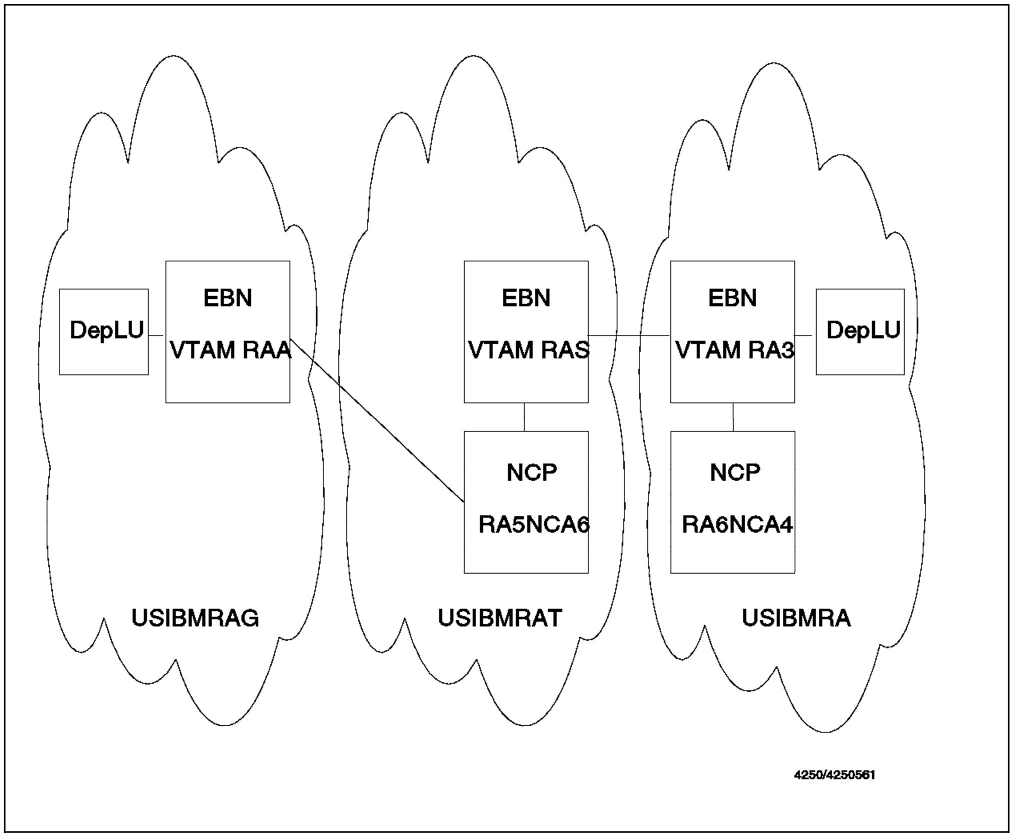

Figure 62. Connecting Three Net IDs through Channel Links

5.5.2 Definitions

All three VTAMs are started as border nodes. The relevant start parameters for VTAM RAA are shown in Figure 63.

The APPNCOS start parameter is worth a consideration. By setting APPNCOS, you have defined a single APPN COS that is to be used when “all else fails” in terms of figuring out

- Which APPN COS to use based on the received subarea COS/mode
- Which native APPN COS to use based on the received nonnative APPN COS (for border node users)

This should prevent any sessions from failing due to mismatch in COS values, but it will force all these “mismatched” sessions to the same COS (and possibly/likely the same route). That means interactive data may get mixed with batch data, secure with non-secure, and so on. Many classes of service are mapped using the APPNCOS value. As your network gets bigger, you may find that unacceptable. APPNCOS is useful to avoid COS mismatch situations, but in the long run you are probably better off finding out why the COS mismatch occurred in the first place and fixing that anomaly.

```
NETID=USIBMRAG,                             BORDER NODE 3           X
HOSTPU=ISTPUSAO,SSCPNAME=RAA,                                       X
APPNCOS=#CONNECT,                                                   X
NODETYPE=NN,                                                        X
DYNLU=YES,                                                          X
CPCP=YES,                                                           X
BN=YES,BNDYN=FULL,                           BORDER NODE 3          X
XNETALS=YES,                                                        X
```

Figure 63. VTAM Start Parameters Relevant for Border Node

The channel link from VTAM RAA to NCP RA5NCA6 is defined in the usual way for a FID2 channel. The definitions are shown in Figure 64 and Figure 65 on page 54. The parameter NETID need *not* be coded. The net IDs are exchanged dynamically.

```
RAARAT  VBUILD TYPE=LOCAL                 VTAM A TO NCP 5
RAALE28  PU    CUADDR=E28,                                          X
               MAXBFRU=34,                                          X
               PUTYPE=2,                                            X
               XID=YES
```

Figure 64. Cross Network APPN Channel Link (VTAM Side)

```
RA5GCHA1 GROUP LNCTL=CA,                                                *
               ISTATUS=ACTIVE
*
RA5LCHA4 LINE ADDRESS=P01,         CA PHYSICAL POSITION 1               *
               CA=TYPE6,                                                *
               CASDL=120,          INTERNAL BEFORE CHANNEL SLOWDOWN     *
               DELAY=0.0,          CHAN ATTNDELAY                       *
               NCPCA=ACTIVE,       NATIVE SUBCHANNEL (NSC) ACTIVE       *
               TIMEOUT=120         INTERVAL BEFORE CHANNEL DISCONTACT
RA5PCHA4 PU PUTYPE=2,              APPN    CONNECTION                   *
               XID=YES
```

Figure 65. Cross Network APPN Channel Link (NCP Side)

In SNI, a channel link is possible from a VTAM to an NCP in another net ID. It requires a substantial number of parameters in the NCP and some in VTAM. The APPN inter-network channel link is much easier to use.

The definitions for the APPN channel link between VTAM RAS and VTAM RA3 are *exactly* the same as for the AHHC scenario within one net ID. NETID is not coded in the TRL major nodes, nor in the LOCAL major nodes. To look at the definitions for the APPN channel connection you can go back to Figure 3 and Figure 4 on page 14.

5.5.3 Test Results

VTAMs and NCPs are started. The APPN channel connections are activated in the usual way. CP-CP sessions are established. LU-LU sessions are established from local terminals on VTAM RAA and RA3 to the respective TSOs on the other side. Figure 66 on page 55 shows all the sessions in VTAM RA3, among them the LU-LU session from terminal RA3T420.

```
D NET,SESSIONS,LIST=ALL
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = SESSIONS 229
IST873I        PLU                SLU              SID        STATUS
IST874I USIBMRA.RA3A0004  USIBMRA.RA3T421   F6FF416460405865 ACTIV
IST874I USIBMRA.RA3AT01   USIBMRAG.RAAT420  F6FF41646040585E ACTIV
IST874I USIBMRAG.RAAAT01  USIBMRA.RA3T420   F7FF61826D4D2F4E ACTIV  1
IST874I USIBMRAT.RAS      USIBMRA.RA3       F627D1F27F7C2602 ACTIV/CP
IST874I USIBMRA.RA3       USIBMRAT.RAS      F6FF416460405856 ACTIV/CP
IST874I USIBMRA.DSIAMLUT  USIBMRA.ISTPDCLU  F6FF416460405850 ACTIV
IST874I USIBMRA.DSIAMLUT  USIBMRA.ISTPDCLU  F6FF41646040584F ACTIV
IST924I -------------------------------------------------------------
IST878I NUMBER OF PENDING SESSIONS =         0
IST924I -------------------------------------------------------------
IST878I NUMBER OF ACTIVE  SESSIONS =         7
IST1162I    LU-LU                  =         5
IST1162I    CP-CP CONWINNER        =         1
IST1162I    CP-CP CONLOSER         =         1
IST924I -------------------------------------------------------------
IST878I NUMBER OF QUEUED  SESSIONS =         0
IST924I -------------------------------------------------------------
IST878I NUMBER OF TOTAL   SESSIONS =         7
IST924I -------------------------------------------------------------
IST1161I ACTIVE SSCP SESSIONS
IST1162I    SSCP-LU                =        35
IST1162I    SSCP-PU                =         7
IST1162I    SSCP-SSCP              =         0
IST314I END
```

Figure 66. Sessions Seen from VTAM RA3

VTAM provides the information to follow the route of a session through the networks. You can follow it from node to node. If you have NetView* with the APPNTAM feature installed, this information is combined and graphically displayed on one screen. Using VTAM displays, Figure 66 gives the starting point. It provides the session identifier (SID), marked with **1**. This same SID is used in the following displays, even in other networks. The D NET,SESSION command in Figure 67 is one way of obtaining the adjacent link station (ALS), marked **2**. The display of the ALS in Figure 68 on page 56 shows us the adjacent CP, the next node on the session path, which is RAS, marked **3**. RAS is “across the border” in net ID USIBMRAT.

```
                                     1
* RA3A0    D NET,SESSION,SID=F7FF61826D4D2F4E
  RA3A0    IST097I  DISPLAY  ACCEPTED
' RA3A0
IST350I  DISPLAY TYPE = SESSIONS
IST879I  PLU/OLU REAL = USIBMRAG.RAAAT01  ALIAS = USIBMRA.RAAAT01
IST879I  SLU/DLU REAL = USIBMRA.RA3T420   ALIAS = USIBMRAG.RA3T420
IST880I  SETUP STATUS = ACTIV
IST875I  ADJSSCP TOWARDS PLU = ISTAPNCP
IST875I  ALSNAME TOWARDS PLU = RA3PCA8  2
IST933I  LOGMODE=M2BSCNQ , COS=*BLANK*
IST875I  APPNCOS TOWARDS PLU = #CONNECT
IST314I  END
```

Figure 67. VTAM RA3 Points to Adjacent Link Station

```
* RA3AO    D NET,ID=RA3PCA8,E  2
  RA3AO    IST097I  DISPLAY  ACCEPTED
' RA3AO
IST075I  NAME = RA3PCA8          , TYPE = PU_T2.1
IST486I  STATUS= ACTIV--L--, DESIRED STATE= ACTIV
IST1043I  CP NAME = RAS  3  , CP NETID = USIBMRAT, DYNAMIC LU = YES
IST1105I  RESOURCE STATUS TGN CP-CP TG CHARACTERISTICS
IST1106I  RA3PCA8  AC/R    21 YES   982D0000000000000000017100808080
IST136I  LOCAL    SNA MAJOR NODE = RA3PUTIN
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST1314I  TRLE = RA3KMPC   STATUS = ACTIV       CONTROL = MPC
IST355I  LOGICAL UNITS:
IST080I  RAAAT01  ACT/S----Y RAS      ACT/S----Y RAAT420  ACT/S----Y
IST314I  END
```

Figure 68. Adjacent Link Station Points to Next Hop VTAM RAS

In VTAM RAS we display the same SID. See Figure 69. This again points to the adjacent link station, marked **4**. From there, Figure 70 **5** points to RAA, which is, at last, the CP name of the node in which the PLU resides. This way we are able to follow the session from end to end.

```
                         1
D NET,SESSION,SID=F7FF61826D4D2F4E
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = SESSIONS 412
IST879I PLU/OLU REAL = USIBMRAG.RAAAT01  ALIAS = USIBMRAT.RAAAT01
IST879I SLU/DLU REAL = USIBMRA.RA3T420   ALIAS = USIBMRAG.RA3T420
IST880I SETUP STATUS = ACTIV
IST875I ADJSSCP TOWARDS PLU = ISTAPNCP
IST875I ADJSSCP TOWARDS SLU = ISTAPNCP
IST875I ALSNAME TOWARDS PLU = RA5PCHA4  4
IST875I ALSNAME TOWARDS SLU = RASPC8A
IST933I LOGMODE=M2BSCNQ , COS=*BLANK*
IST875I APPNCOS TOWARDS PLU = #CONNECT
IST875I APPNCOS TOWARDS SLU = #CONNECT
IST314I END
```

Figure 69. Session Seen from VTAM RAS

```
              4
D NET,ID=RA5PCHA4,E
IST097I DISPLAY ACCEPTED
IST075I NAME = RA5PCHA4, TYPE = PU_T2.1 421
IST486I STATUS= ACTIV--L--, DESIRED STATE= ACTIV
IST1043I CP NAME = RAA,  5 NETID = USIBMRAG, DYNAMIC LU = YES
IST1105I RESOURCE STATUS TGN CP-CP TG CHARACTERISTICS
IST1106I RA5PCHA4 AC/R    21 YES   982D0000000000000000017100808080
IST081I LINE NAME = RA5LCHA4, LINE GROUP = RA5GCHA1, MAJNOD = RA5NCA6
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST355I LOGICAL UNITS:
IST080I RAAAT01  ACT/S----Y RAA      ACT/S----Y RAAT420  ACT/S----Y
IST314I END
```

Figure 70. Adjacent Link Station Points to Next Hop VTAM RAA

Now, let us have a look at the topology. The display in Figure 71 on page 57 **6** proves that VTAM RAS has inter-cluster links to both VTAM RAA and VTAM RA3. Both adjacent VTAMs have net IDs different from one another.

```
D NET,TOPO,ID=RAS,LIST=ALL
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = TOPOLOGY 470
IST1295I CP NAME          NODETYPE ROUTERES CONGESTION  CP-CP WEIGHT
IST1296I USIBMRAT.RAS     NN       128      NONE        *NA*  *NA*
IST1297I                  ICN/MDH  CDSERVR  RSN
IST1298I                  YES      YES      8
IST1223I                  BN       NATIVE
IST1224I                  YES      YES
IST1299I TRANSMISSION GROUPS ORIGINATING AT CP USIBMRAT.RAS
IST1357I                                                CPCP
IST1300I DESTINATION CP   TGN      STATUS   TGTYPE      VALUE WEIGHT
IST1301I USIBMRAG.RAA     21       OPER     INTERCLUST  YES   *NA*
IST1301I USIBMRA.RA3      21       OPER     INTERCLUST  YES   *NA*
IST314I END                                 6
```

Figure 71. VTAM RAS Connecting to Two Networks with Different Net IDs

In our tests of the border node function, we use various types of sessions. LU 2 sessions and LU 6.2 sessions between VTAM applications work fine across one or several extended border node connections.

We also use LU 6.2 sessions between CM/2 on one end of the session and VTAM or a second CM/2 on the other end of the session. This works, if the CM/2 is set up as an end node.

If the CM/2 is a network node, it requires a later level of software to provide interoperability with border node. If the basic CM/2 R1.1 code is used, lack of this support results in session setup failure with sense code 088C 4683. After installing the latest software update, we find that sessions can be set up as expected.

5.6 Disjoint APPN Networks

This section describes another cross-network scenario, which *cannot* be done with SNI: going from one net ID through an intermediate network to a separate network with the same net ID. This scenario will apply to a customer who replaces his trunk connections by connections to a carrier network, for example the IBM Information Network.

5.6.1 Description

As shown in Figure 72, this scenario is much the same as that described above in 5.5, "Inter-Network Channel Links" on page 52, except we are now dealing with only two net IDs. VTAM RAA and VTAM RA3 are both in net ID USIBMRA. But the connection between them is through VTAM RAS in net ID USIBMRAT.

Apart from the net ID the definitions are the same as in 5.5.2, "Definitions" on page 53, so they will not be shown here. The LU-LU sessions are between two APPC/VTAM applications: server on VTAM RAA and requester on VTAM RA3.

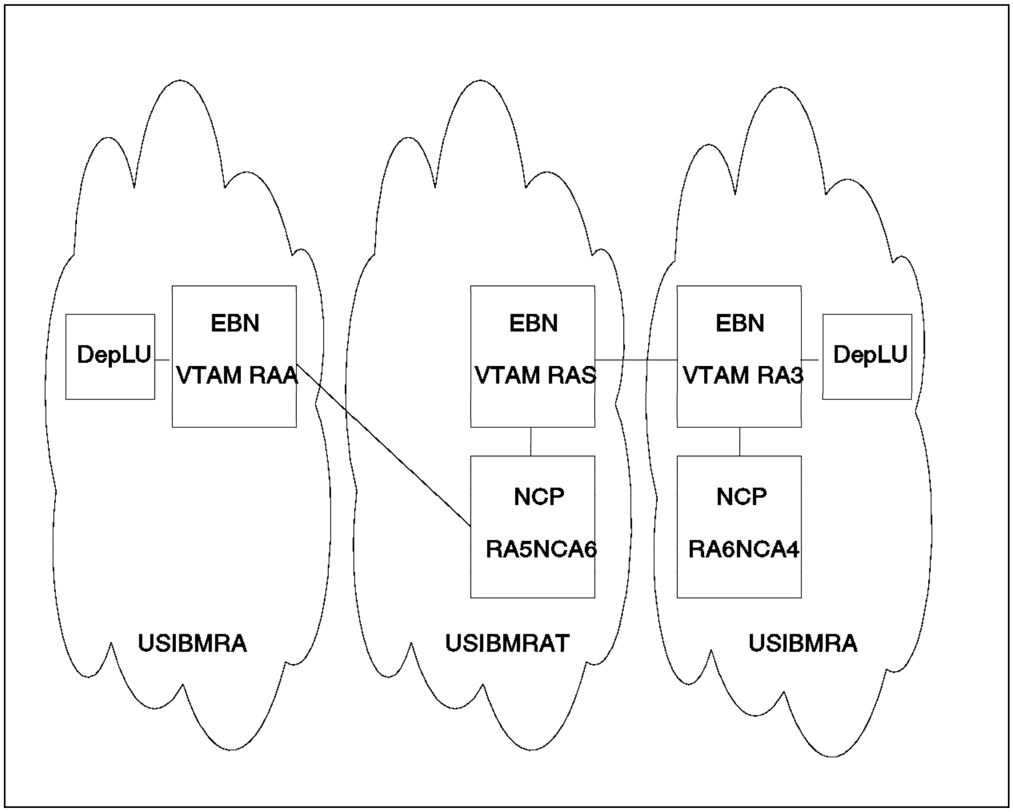

Figure 72. VTAM RAS Connecting Two Disjoint Networks with the Same Net ID

5.6.2 Test Results

The VTAM APPC requester application RA3AREQ runs on VTAM RA3. It is in session with the server application RAAASER on VTAM RAA. Both VTAMs are in the same net ID USIBMRA. But there is no connection from VTAM RAA to VTAM RA3 within net ID USIBMRA. The session path crosses two network borders: from USIBMRA into USIBMRAT, then from USIBMRAT into USIBMRA.

The VTAMs in each subnetwork (or cluster) cannot see beyond the borders. Any node on the other side of the border appears as if it were an end node attached to the adjacent border node. Thus, the directory in VTAM RA3 sees the application RAAASER as being in an end node USIBMRA.RAA served by network node server USIBMRAT.RAS. (It is valid in the architecture for an end node to have a net ID different from its network node server.) This is illustrated in Figure 73 **1**.

```
D NET,ID=RAAASER,E
IST097I DISPLAY ACCEPTED
IST075I NAME = USIBMRA.RAAASER, TYPE = CDRSC 917
IST486I STATUS= ACT/S----Y, DESIRED STATE= ACTIV
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 = RA3, VERIFY OWNER = NO
IST1184I CPNAME = USIBMRA.RAA - NETSRVR = ***NA***
IST1044I ALSLIST = ISTAPNPU
IST082I DEVTYPE = INDEPENDENT LU / CDRSC
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST171I ACTIVE SESSIONS = 0000000002, SESSION REQUESTS = 0000000000
IST206I SESSIONS:
IST1081I ADJACENT LINK STATION = RA3PCA8
IST634I NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I RA3AREQ  ACTIV-P    F6FF41646041ADDD 0000 0000  0  0 USIBMRA
IST635I RA3AREQ  ACTIV-P    F6FF41646041ADDC 0001 0001  0  0 USIBMRA
IST924I -------------------------------------------------------------
IST075I NAME = USIBMRA.RAAASER, TYPE = DIRECTORY ENTRY
IST1186I DIRECTORY ENTRY = DYNAMIC LU
IST1184I CPNAME = USIBMRA.RAA - NETSRVR = USIBMRAT.RAS 1
IST314I END
```

Figure 73. Session Partner Seen in VTAM RA3

The topology display in Figure 74 shows two inter-cluster links from VTAM RAS to VTAM RAA and VTAM RA3. Comparing to Figure 71 on page 57, we find the only difference is that VTAMs RAA and RA3 are now in the same net ID.

In fact, from the topology display it is not evident that VTAMs RAA and RA3 are not in the same subnetwork. Only Figure 73 shows that the sessions go through net ID USIBMRAT, which they would not do if VTAMs RAA and RA3 were *connected* in the same subnetwork.

```
D NET,TOPO,ID=RAS,LIST=ALL
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = TOPOLOGY 075
IST1295I CP NAME           NODETYPE ROUTERES CONGESTION  CP-CP WEIGHT
IST1296I USIBMRAT.RAS      NN       128      NONE        *NA*  *NA*
IST1297I                   ICN/MDH  CDSERVR  RSN
IST1298I                   YES      YES      8
IST1223I                   BN       NATIVE
IST1224I                   YES      YES
IST1299I TRANSMISSION GROUPS ORIGINATING AT CP USIBMRAT.RAS
IST1357I                                                 CPCP
IST1300I DESTINATION CP    TGN      STATUS   TGTYPE      VALUE WEIGHT
IST1301I USIBMRA.RAA       21       OPER     INTERCLUST  YES   *NA*
IST1301I USIBMRA.RA3       21       OPER     INTERCLUST  YES   *NA*
IST314I END
```

Figure 74. VTAM RAS Connecting Two Networks with the Same Net ID

5.7 Connecting Two Networks with the Same Net ID

A new capability with APPN, which is not possible with SNI, is to divide a single network into several subnetworks which still retain the same net ID. A reason for this might be to minimize the topology database in network nodes which do not have enough storage to hold a very large topology, for instance the 3174 with just 6 megabytes. A large APPN network comprising a VTAM with many attached 3174s, each serving a number of PS/2 end nodes, is a likely candidate for this feature. In our test we use another VTAM as the network node server because the 3174 microcode that supports connection to non-native networks is not available at the time of writing.

5.7.1 Description

Figure 75 shows the configuration we use.

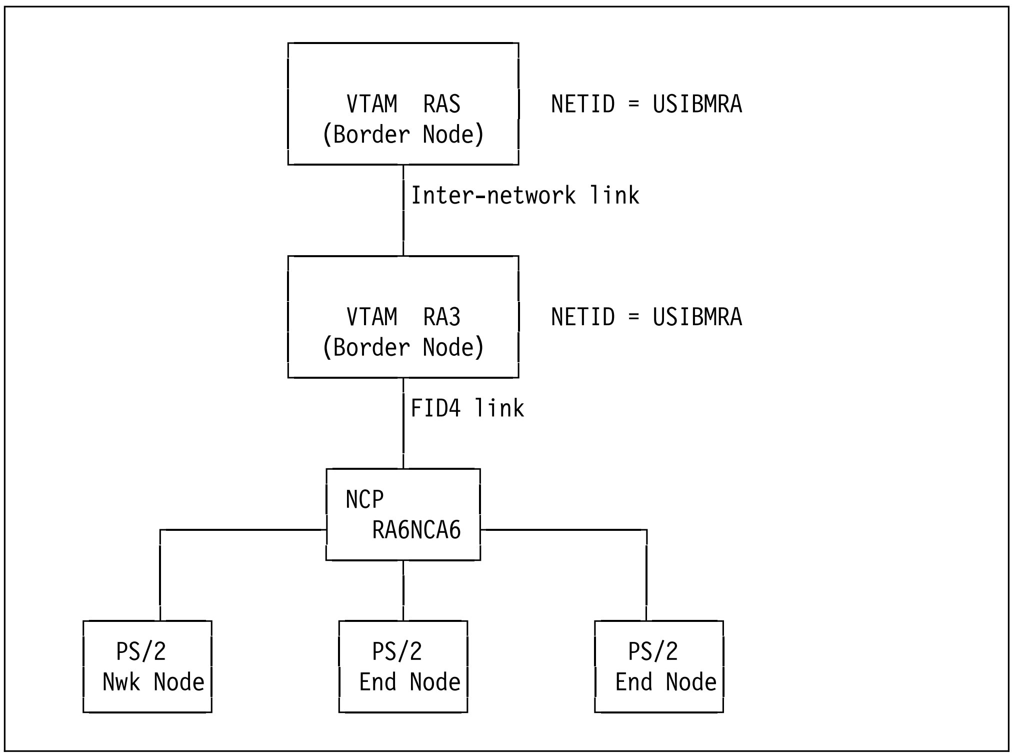

Figure 75. Border Node Configuration - Same Net ID

The two VTAMs, both defined as border nodes, are connected via an AHHC link similar to that described in Chapter 3, “APPN Host-to-Host Connection” on page 11. Every node shown has the same net ID.

In the case of the 3174, of course, the border would be a peripheral border but the principles are the same.

5.7.2 Definitions

The two VTAMs RA3 and RAS are both started with the option BN=YES, as described in 5.3.2, “Definitions” on page 38.

The TRL major nodes are exactly the same as we use in the AHHC test described in Chapter 3, “APPN Host-to-Host Connection” on page 11. The

LOCAL major nodes connecting the VTAMs, however, differ as can be seen in Figure 76 on page 62.

```
RASHER   VBUILD  TYPE=LOCAL
* APPN MPC LINK FROM RAS TO RA3
RASPC8A  PU   TRLE=RASKMPC,PUTYPE=2,XID=YES,NATIVE=NO   1


RA3HER   VBUILD  TYPE=LOCAL
* APPN MPC LINK FROM RA3 TO RAS
RA3PC8A  PU   TRLE=RA3KMPC,PUTYPE=2,XID=YES,NATIVE=NO   1
```

Figure 76. LOCAL Major Nodes - Non-Native Connection

The parameter NATIVE=NO **1** must be coded in this case to inform the VTAM border nodes that this link station represents an inter-subnetwork link. Unless this is coded, the VTAMs will recognize the same net ID at XID exchange, and will set up a “normal” intermediate link.

In the case of a peripheral border, NATIVE=NO would be coded on the border node side only (for example, on the 3174 PU definition).

5.7.3 Test Results

We begin by activating the TRL and LOCAL major nodes in each VTAM, to establish the inter-subnetwork link. Figure 77 shows that RA3 sees the link to RAS as inter-cluster **1**, even though their net IDs are the same.

```
NCCF                      N E T V I E W    RA3AO BUCZAK   02/07/94 13:42:56
* RA3AO    D NET,TOPO,ID=RA3,LIST=ALL
  RA3AO    IST097I  DISPLAY  ACCEPTED
' RA3AO
IST350I  DISPLAY TYPE = TOPOLOGY
IST1295I  CP NAME          NODETYPE ROUTERES CONGESTION  CP-CP WEIGHT
IST1296I  USIBMRA.RA3       NN       128      NONE        *NA*  *NA*
IST1297I                    ICN/MDH  CDSERVR  RSN
IST1298I                    YES      YES      14
IST1223I                    BN       NATIVE
IST1224I                    YES      YES
IST1299I  TRANSMISSION GROUPS ORIGINATING AT CP USIBMRA.RA3
IST1357I                                                  CPCP
IST1300I  DESTINATION CP    TGN      STATUS   TGTYPE      VALUE WEIGHT
IST1301I  USIBMRA.RACONNET  21       OPER     INTERM      NO    *NA*
IST1301I  USIBMRA.JERZY     21    2  OPER     ENDPT       YES   *NA*
IST1301I  USIBMRA.BERND     21    3  OPER     INTERM      YES   *NA*
IST1301I  USIBMRA.RAS       21    1  OPER     INTERCLUST  YES   *NA*
IST314I  END
```

Figure 77. Topology Display - Inter-Cluster Link with Same Net ID

We see also a PS/2 end node **2**, which will be used to establish a session to RAS; and also a network node **3** which is not aware of the presence of RAS in the adjacent network, as a display of its topology shows in Figure 78 on page 63.

```
**************************************
*        Topology Information        *
**************************************
Number of network nodes                      2

1>Network node CP name                       USIBMRA.RA3
  Route additional resistance                128
  Congested?                                 No
  Quiescing?                                 No
  ISR depleted?                              No
  Number of TGs                              3

1.1>TG partner CP name                       USIBMRA.RACONNET
    Transmission group number                21
    TG partner node type                     Virtual (connection network)
    Quiescing?                               No
    Topology                                 Network
    Effective capacity                       7.80 kilobits per second
    Cost per connect time                    0
    Cost per byte                            0
    Propagation delay                        9.22 milliseconds (telephone)
    User defined parameter 1                 128
    User defined parameter 2                 128
    User defined parameter 3                 128
    Security                                 Nonsecure

1.2>TG partner CP name                       USIBMRA.BERND
    Transmission group number                21
    TG partner node type                     Real
    Quiescing?                               No
    Topology                                 Network
    Effective capacity                       7.80 kilobits per second
    Cost per connect time                    0
    Cost per byte                            0
    Propagation delay                        9.22 milliseconds (telephone)
    User defined parameter 1                 128
    User defined parameter 2                 128
    User defined parameter 3                 128
    Security                                 Nonsecure

1.3>TG partner CP name                1      USIBMRA.RAS
    Transmission group number                21
    TG partner node type                     Real
    Quiescing?                               No
    Topology                                 Network
    Effective capacity                       7.80 kilobits per second
    Cost per connect time                    0
    Cost per byte                            0
    Propagation delay                        9.22 milliseconds (telephone)
    User defined parameter 1                 128
    User defined parameter 2                 128
    User defined parameter 3                 128
    Security                                 Nonsecure
```

Figure 78 (Part 1 of 2). APPN Topology in PS/2 Network Node

```
2>Network node CP name                      USIBMRA.BERND
  Route additional resistance               128
  Congested?                                No
  Quiescing?                                No
  ISR depleted?                             No
  Number of TGs                             2

2.1>TG partner CP name                      USIBMMK.MK34
    Transmission group number               0
    TG partner node type                    Real
    Quiescing?                              No
    Topology                                Local
    Effective capacity                      3.99 megabits per second
    Cost per connect time                   0
    Cost per byte                           0
    Propagation delay                       384.00 microseconds (local area net
    User defined parameter 1                128
    User defined parameter 2                128
    User defined parameter 3                128
    Security                                Nonsecure

2.2>TG partner CP name                      USIBMRA.RA3
    Transmission group number               21
    TG partner node type                    Real
    Quiescing?                              No
    Topology                                Network
    Effective capacity                      3.99 megabits per second
    Cost per connect time                   0
    Cost per byte                           0
    Propagation delay                       384.00 microseconds (local area net
    User defined parameter 1                128
    User defined parameter 2                128
    User defined parameter 3                128
    Security                                Nonsecure
```

Figure 78 (Part 2 of 2). APPN Topology in PS/2 Network Node

Note that the (one-way) TG from RA3 to RAS is reported **1** in the topology, but the node RAS itself is not known to BERND. The topology of RAS and anything beyond is not propagated in the subnetwork containing RA3 and BERND.

Next, we establish an LU 6.2 session from JERZY to RASASER, which is the ATELL application in RAS. The session displays show that, apart from the topology isolation, the APPN functions work in just the same way as with a native APPN connection. Figure 79 on page 65 shows the subarea session information as seen from RAS.

```
 NLDM.CON                  SESSION CONFIGURATION DATA                PAGE    1
 -------------- PRIMARY ---------------+-------------- SECONDARY --------------
 NAME JERZY     SA 0000001C  EL 0084   |   NAME RASASER   SA 0000001C  EL 006E
 --------------------------------------+---------------------------------------
 DOMAIN RASAN  C-C  PCID USIBMRA.JERZY.D54B4E4B32FEFD57          DOMAIN RASAN
               +-------------+                  +-------------+
 RAS           |   CP/SSCP   | ---          --- |   CP/SSCP   | RAS
 ISTPUS28(0000) | SUBAREA PU  |                  | SUBAREA PU  | ISTPUS28(0000)
               +------+------+                  +------+------+
                      |                                |
               +------+------+                  +------+------+
        (0079) |     CUA     |                  |     LU      | RASASER (006E)
               +------+------+                  +-------------+
                      |
               +------+------+
 RASPC8A (007A) |     PU      | LOGMODE #INTER
               +------+------+ PADJ CP RA3
                      |
               +------+------+
 JERZY   (0084) |     ILU     |
               +-------------+

SELECT PT, ST (PRI, SEC TRACE), RT (RESP TIME), P, AR, FC
CMD==>
```

Figure 79. Session Configuration Seen from RAS

The displays of non-native resources RASASER from RA3 and JERZY from RAS are shown in Figure 80 on page 66 and Figure 81 on page 66. They demonstrate that to RA3, RASASER is seen as an LU on the network node RAS **1** (which is correct); and to RAS, JERZY is seen as an end node served by RA3 **2** (again, correct).

```
NCCF                      N E T V I E W    RA3AO BUCZAK   02/07/94 14:01:42
C RA3AO    DISPLAY NET,ID=RASASER,SCOPE=ALL
  RA3AO    IST097I  DISPLAY  ACCEPTED
' RA3AO
IST075I  NAME = USIBMRA.RASASER  , TYPE = CDRSC
IST486I  STATUS= ACT/S----Y, DESIRED STATE= ACTIV
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 = RA3     , VERIFY OWNER = NO
IST1184I  CPNAME = USIBMRA.RAS       - NETSRVR = ***NA***
IST1044I  ALSLIST = ISTAPNPU
IST082I  DEVTYPE = INDEPENDENT LU / CDRSC
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST171I  ACTIVE SESSIONS = 0000000002, SESSION REQUESTS = 0000000000
IST206I  SESSIONS:
IST1081I  ADJACENT LINK STATION = RA3PC8A
IST634I  NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I  JERZY    ACTIV-P    D54B4E4B32FEFD57           0  0 USIBMRA
IST635I  JERZY    ACTIV-P    D54B4E4B31FEFD57           0  0 USIBMRA
IST924I  -------------------------------------------------------------
IST075I  NAME = USIBMRA.RASASER  , TYPE = DIRECTORY ENTRY
IST1186I  DIRECTORY ENTRY = DYNAMIC    LU                  1
IST1184I  CPNAME = USIBMRA.RAS       - NETSRVR = ***NA***
IST314I  END
```

Figure 80. Display of Non-Native Resource from RA3

```
NCCF                      N E T V I E W    RASAN BUCZAK   02/07/94 14:03:08
C RASAN    DISPLAY NET,ID=JERZY,SCOPE=ALL
  RASAN    IST097I  DISPLAY  ACCEPTED
' RASAN
IST075I  NAME = USIBMRA.JERZY    , TYPE = CDRSC
IST486I  STATUS= ACT/S----Y, DESIRED STATE= ACTIV
IST977I  MDLTAB=***NA*** ASLTAB=***NA***
IST1333I  ADJLIST = ***NA***
IST861I  MODETAB=***NA*** USSTAB=***NA*** LOGTAB=***NA***
IST934I  DLOGMOD=SNASVCMG USS LANGTAB=***NA***
IST597I  CAPABILITY-PLU ENABLED  ,SLU ENABLED  ,SESSION LIMIT NONE
IST231I  CDRSC    MAJOR NODE = ISTCDRDY
IST479I  CDRM NAME = RAS     , VERIFY OWNER = NO
IST1184I  CPNAME = USIBMRA.JERZY     - NETSRVR = ***NA***
IST1044I  ALSLIST = ISTAPNPU
IST082I  DEVTYPE = INDEPENDENT LU / CDRSC
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST171I  ACTIVE SESSIONS = 0000000002, SESSION REQUESTS = 0000000000
IST206I  SESSIONS:
IST1081I  ADJACENT LINK STATION = RASPC8A
IST634I  NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I  RASASER  ACTIV-S    D54B4E4B32FEFD57 0002 0001       USIBMRA
IST635I  RASASER  ACTIV-S    D54B4E4B31FEFD57 0001 0001       USIBMRA
IST924I  -------------------------------------------------------------
IST075I  NAME = USIBMRA.JERZY    , TYPE = DIRECTORY ENTRY
IST1186I  DIRECTORY ENTRY = DYNAMIC    EN                    2
IST1184I  CPNAME = USIBMRA.JERZY     - NETSRVR = USIBMRA.RA3
IST314I  END
```

Figure 81. Display of Non-Native Resource from RAS

5.8 Multiple Subnetwork Paths Scenario

This section describes a test scenario where multiple subnetworks can act as an intermediate subnetwork. We describe the way VTAM maintains the adjacent cluster tables and also the way VTAM searches across the inter-subnetwork links. We examine the impact of setting the subnet visit count value, and explore the option of using a predefined adjacent cluster table to better control cross subnetwork searches.

5.8.1 Description

In the test environment, we use four MVS VTAM systems and two PS/2 Communications Manager/2 systems. Figure 82 shows the schematic diagram of the test environment. The system environment contains four subnetworks and five net IDs.

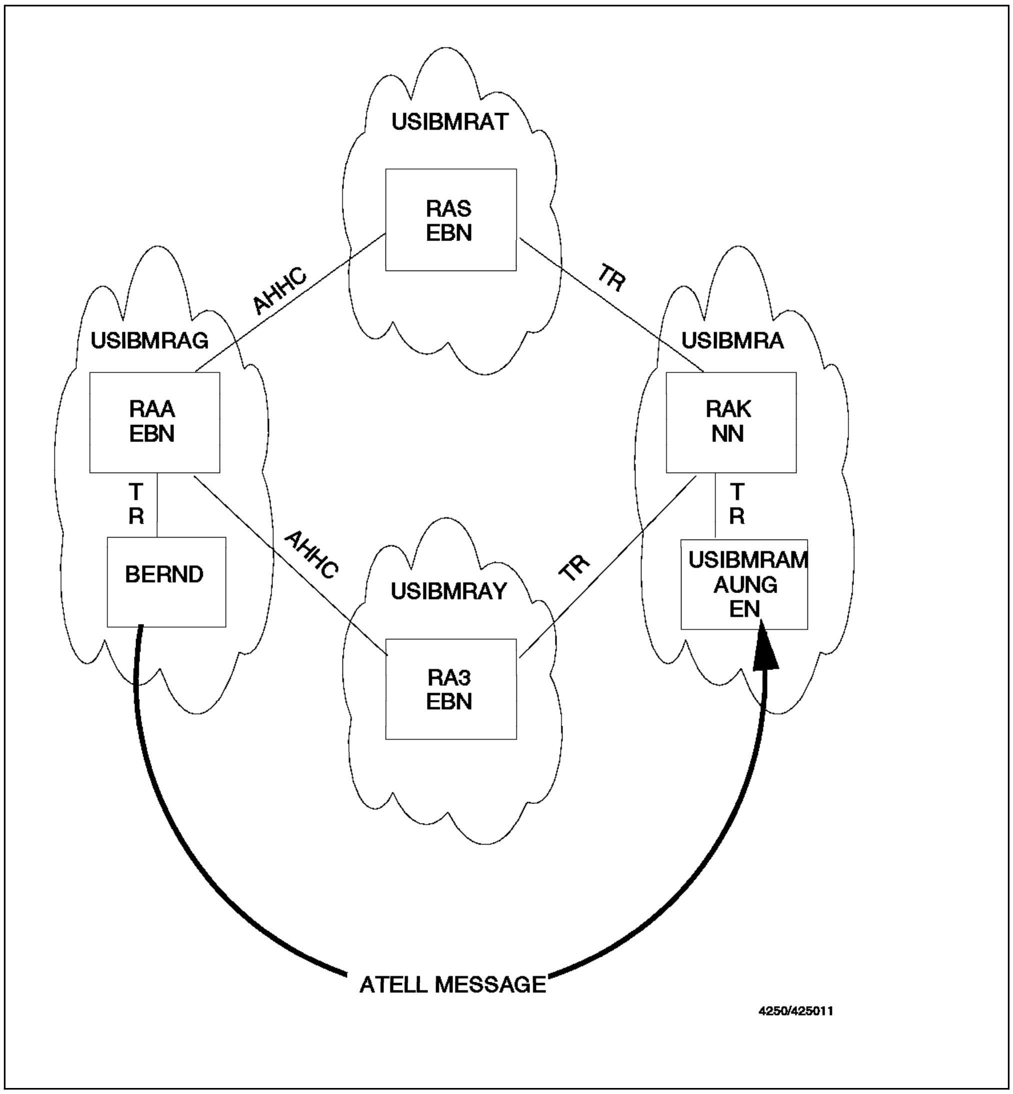

Figure 82. Schematic Diagram of the Test Environment

- RAA is a VTAM V4R2 network node with the net ID of USIBMRAG. It has MPC connections to RAS and RA3. The MPC connections use AHHC as described in Chapter 3, "APPN Host-to-Host Connection" on page 11.

  It also has a 3172, which provides the LAN connection.

  RAA is defined as a border node.

- BERND, a PS/2 running Communications Manager/2, is connected to RAA via the LAN and the 3172. It is defined as an end node and has the net ID of USIBMRAG.
- RAS is a VTAM V4R2 composite network node. It has the net ID of USIBMRAT. It connects to RAA via an AHHC connection. It connects to RAK via the token-ring LAN. Connectivity to the token-ring LAN is provided by a dedicated NCP.

  RAS is defined as a border node.

- RA3 is also a VTAM V4R2 composite network node. It has the net ID of USIBMRAY. It connects to RAA via an AHHC connection. It connects to RAK via the token-ring LAN. Connectivity to the token-ring LAN is provided by a dedicated NCP.

  RA3 is defined as a border node.

- RAK is a VTAM V4R1 composite network node. It has the net ID of USIBMRA.
- AUNG, a PS/2 running Communications Manager/2, is connected to RAK via the LAN. It is defined as an end node and has the net ID of USIBMRAM.

Figure 83 shows the inter-cluster links.

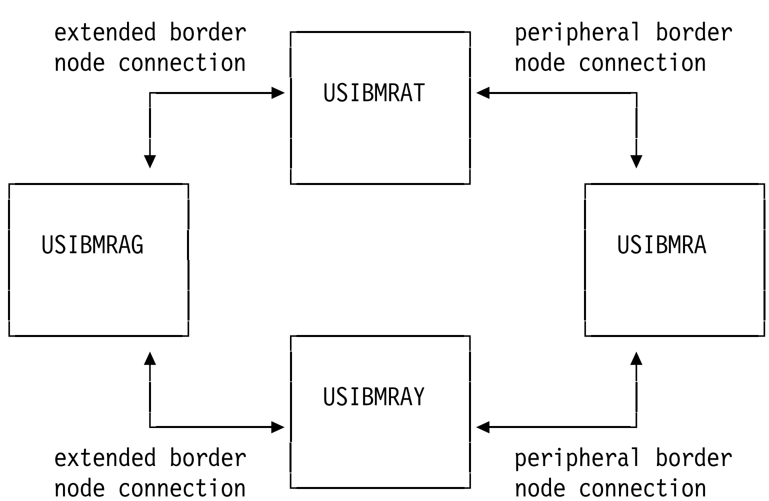

Figure 83. Inter-Cluster Links in the Test Environment

5.8.2 Definitions

This section contains the system definitions for the test environment for the multiple subnetwork paths.

5.8.2.1 RAA Definitions

In RAA, we define the XCA (External Communications Adapter) major node for the 3172 and the AHHC connections to RAS and RA3. Figure 84, Figure 85 and Figure 86 show the definitions.

```
RAA3172 VBUILD TYPE=XCA
**
RAA3172P PORT  ADAPNO=0,                                             * X
               CUADDR=EA2,                                           * X
               MEDIUM=RING,                                          * X
               SAPADDR=4,                                            * X
               TIMER=60
**
RAAXCAG1 GROUP DIAL=YES,CALL=INOUT,DYNPU=YES
*
RAAXL1   LINE
RAAXP1   PU
```

Figure 84. 3172 Definitions

```
RAAAHHC  VBUILD  TYPE=TRL
RAAMPCS  TRLE   LNCTL=MPC,READ=(C12),WRITE=(C14)
RAAMPC3  TRLE   LNCTL=MPC,READ=(C16),WRITE=(C18)
```

Figure 85. TRL Major Node in RAA

```
RAALCA   VBUILD  TYPE=LOCAL
RAAMPCSP PU   TRLE=RAAMPCS,PUTYPE=2,XID=YES
RAAMPC3P PU   TRLE=RAAMPC3,PUTYPE=2,XID=YES
```

Figure 86. LOCAL Major Node in RAA

5.8.2.2 RAS Definitions

In RAS, we define the AHHC connection to RAA and a switched major node for the dial-out connection to RAK.

```
RASAHHCA VBUILD  TYPE=TRL
RASMPCA  TRLE   LNCTL=MPC,READ=(C14),WRITE=(C12)
```

Figure 87. TRL Major Node in RAS

```
RASLCAA  VBUILD  TYPE=LOCAL
RASMPCAP PU   TRLE=RASMPCA,PUTYPE=2,XID=YES
```

Figure 88. LOCAL Major Node in RAS

```
        VBUILD TYPE=SWNET,MAXNO=20,MAXGRP=5
RASUK1   PU    ADDR=01,MAXDATA=512,MAXOUT=7,MAXPATH=5,PASSLIM=7,
               PUTYPE=2,
               CPNAME=RAK,
               NETID=USIBMRA,
               CONNTYPE=APPN,
               DYNLU=YES,
               CPCP=YES
*
RASUK1P  PATH  DIALNO=0004400001290000,            IN NCP RA03745
               GRPNM=RA5GTRL2                       IN NCP RA5NCA.
```

Figure 89. Definition for Switched Connection from RAS to RAK

5.8.2.3 RA3 Definitions

The RA3 definitions are similar to those of RAS.

5.8.2.4 RAK Definitions

A switched major node is defined for the connections from RAS, RA3 and AUNG.

5.8.2.5 AUNG and BERND Definitions

We defined each as an end node with appropriate net IDs. We also define the connection to the preferred network node server by providing the appropriate MAC address.

5.8.3 Test Results

The aim of the test is to establish LU 6.2 sessions from BERND to AUNG and to determine the cross-network path taken by the session. We examine the impact of setting the subnet visit count value, and explore the option of using a predefined adjacent cluster table to better control cross subnetwork searches.

5.8.3.1 Initialization

We start RAA by issuing:

```
S NETA0,,,(LIST=A0,BN=YES,BNDYN=FULL,NETID=USIBMRAG)
```

Note the start option **BNDYN=FULL** which means that entries are dynamically added to the adjacent cluster table in VTAM even if the net ID of the entry added does not match the net ID to which the table applies. This means that VTAM will search adjacent subnetworks even if the adjacent subnetworks do not have the same net ID as the resource that VTAM is searching for.

RAS and RA3 are started similarly.

RAK, which is a VTAM V4R1 system, is started without any border node option.

The AHHC connections are established by activating the appropriate major nodes in each VTAM.

The token-ring connections are established using the **VARY DIAL** command as shown below.

```
V NET,DIAL,ID=RASUK1,CPCP=YES
```

Once the VTAM systems are active, the links from the PS/2s, AUNG and BERND, are activated.

5.8.3.2 Topology Displays

Figure 90, Figure 91, Figure 92 on page 72, and Figure 93 on page 72 show the topology information in the 4 VTAMs involved in the tests. In the border nodes, the subnetwork links are shown as having the TGTYPE of INTERCLUST. Please see **1** and **2** for examples.

In contrast, the same subnetwork links are seen by RAK, which is not a border node, as links to end nodes. They are shown in the topology display for RAK as having the TGTYPE of ENDPT. Please see **3** and **4** for examples.

```
* RAAAN    D NET,TOPO,ID=RAA,LIST=ALL
  RAAAN    IST097I  DISPLAY  ACCEPTED
' RAAAN
IST350I  DISPLAY TYPE = TOPOLOGY
IST1295I  CP NAME          NODETYPE ROUTERES CONGESTION  CP-CP WEIGHT
IST1296I  USIBMRAG.RAA     NN       128      NONE        *NA*  *NA*
IST1297I                   ICN/MDH  CDSERVR  RSN
IST1298I                   YES      YES      8
IST1223I                   BN       NATIVE
IST1224I                   YES      YES
IST1299I  TRANSMISSION GROUPS ORIGINATING AT CP USIBMRAG.RAA
IST1357I                                                 CPCP
IST1300I  DESTINATION CP   TGN      STATUS   TGTYPE      VALUE WEIGHT
IST1301I  USIBMRAT.RAS     21       OPER     INTERCLUST  YES   *NA*
IST1301I  USIBMRAY.RA3     21       OPER     INTERCLUST  YES   *NA*
IST1301I  USIBMRAG.BERND   21       OPER     ENDPT       YES   *NA*
IST314I  END
```

Figure 90. Topology in RAA

```
* RASAN    D NET,TOPO,ID=RAS,LIST=ALL
  RASAN    IST097I  DISPLAY  ACCEPTED
' RASAN
IST350I  DISPLAY TYPE = TOPOLOGY
IST1295I  CP NAME          NODETYPE ROUTERES CONGESTION  CP-CP WEIGHT
IST1296I  USIBMRAT.RAS     NN       128      NONE        *NA*  *NA*
IST1297I                   ICN/MDH  CDSERVR  RSN
IST1298I                   YES      YES      42
IST1223I                   BN       NATIVE
IST1224I                   YES      YES
IST1299I  TRANSMISSION GROUPS ORIGINATING AT CP USIBMRAT.RAS
IST1357I                                                 CPCP
IST1300I  DESTINATION CP   TGN      STATUS   TGTYPE      VALUE WEIGHT
IST1301I  USIBMRAG.RAA     21       OPER     INTERCLUST  YES   *NA*
IST1301I  USIBMRA.RAK      21       OPER   1 INTERCLUST  YES   *NA*
IST314I  END
```

Figure 91. Topology in RAS

```
* RA3A0    D NET,TOPO,ID=RA3,LIST=ALL
  RA3A0    IST097I  DISPLAY  ACCEPTED
' RA3A0
IST350I  DISPLAY TYPE = TOPOLOGY
IST1295I  CP NAME           NODETYPE ROUTERES CONGESTION  CP-CP WEIGHT
IST1296I  USIBMRAY.RA3      NN       128      NONE        *NA*  *NA*
IST1297I                    ICN/MDH  CDSERVR  RSN
IST1298I                    YES      YES      14
IST1223I                    BN       NATIVE
IST1224I                    YES      YES
IST1299I  TRANSMISSION GROUPS ORIGINATING AT CP USIBMRAY.RA3
IST1357I                                                  CPCP
IST1300I  DESTINATION CP    TGN      STATUS   TGTYPE      VALUE WEIGHT
IST1301I  USIBMRAG.RAA      21       OPER     INTERCLUST  YES   *NA*
IST1301I  USIBMRA.RAK       21       OPER  2  INTERCLUST  YES   *NA*
IST314I  END
```

Figure 92. Topology in RA3

```
* RAKAN    D NET,TOPO,ID=RAK,LIST=ALL
  RAKAN    IST097I  DISPLAY  ACCEPTED
' RAKAN
IST350I  DISPLAY TYPE = TOPOLOGY
IST1295I  CP NAME           NODETYPE ROUTERES CONGESTION  CP-CP WEIGHT
IST1296I  USIBMRA.RAK       NN       128      NONE        *NA*  *NA*
IST1297I                    ICN/MDH  CDSERVR  RSN         BASEWEIGHT
IST1298I                    YES      NO       168         *NA*
IST1299I  TRANSMISSION GROUPS ORIGINATING AT CP USIBMRA.RAK
IST1357I                                                 CPCP
IST1300I  DESTINATION CP    TGN      STATUS   TGTYPE     VALUE WEIGHT
IST1301I  USIBMRA.RAP       21       OPER     INTERM     YES   *NA*
IST1301I  USIBMRAT.RAS      21       OPER  3  ENDPT      YES   *NA*
IST1301I  USIBMRAY.RA3      21       OPER  4  ENDPT      YES   *NA*
IST1301I  USIBMRAM.AUNG     21       OPER     ENDPT      YES   *NA*
IST314I  END
```

Figure 93. Topology in RAK

5.8.3.3 Adjacent Cluster Displays

VTAM provides a command to display the adjacent cluster tables that are used by VTAM to determine the adjacent CPs to which the search requests will be directed. The command, DISPLAY ADJCLUST, is fully documented in *VTAM Operation* and *VTAM Messages and Codes*.

The entries in the tables can be predefined, by using the adjacent cluster definitions, or dynamically discovered. By not coding adjacent cluster definitions and by using the VTAM start option of BNDYN=FULL, we choose to add the entries to the adjacent cluster table dynamically.

The command, **D NET,ADJCLUST,SCOPE=ALL**, will show you all the table entries for various net IDs in the table. As more network searches are made, the table will grow with additional net IDs.

DISPLAY ADJCLUST with the NETID parameter specified will show you the adjacent CPs to be used to direct the search for a specific net ID. Figure 94 on page 73 shows the command in RAA straight after the VTAM initialization. RAS and RA3 are listed as the adjacent CPs to be used by RAA in searching for the net ID of USIBMRAM. VTAM performs the search serially. That is, it will send the search to the first CP in the table (**1**), and the second entry in the table (**2**) will only be used if the search sent to the first entry fails to find the resource.

Note in Figure 94 that the entries are shown as being DYNAMIC under the TYPE field (**3**), meaning that they are dynamically discovered entries rather than predefined ones. Also note that the STATE (**4**) is shown as ACTIVE for each entry.

Note that at VTAM initialization time, the entries in the table are in alphabetical order by network-qualified CP name. VTAM will dynamically alter the order of the table so that the most recently successful adjacent CP is always at the top of the priority list. We will see this in the next few sections.

```
* RAAAN    D NET,ADJCLUST,NETID=USIBMRAM
  RAAAN    IST097I  DISPLAY  ACCEPTED
' RAAAN
IST350I  DISPLAY TYPE = ADJACENT CLUSTER TABLE
IST1325I  DYNAMIC TABLE FOR USIBMRAM       DYNAMICS = FULL
IST1326I  CP NAME           TYPE 3  STATE 4    STATUS       SNVC
IST1327I  USIBMRAG.RAA      DYNAMIC ACTIVE     *** N/A ***  N/A
IST1327I  USIBMRAT.RAS      DYNAMIC ACTIVE     NOT SEARCHED 003  1
IST1327I  USIBMRAY.RA3      DYNAMIC ACTIVE     NOT SEARCHED 003  2
IST314I  END
```

Figure 94. Adjacent Cluster Table for USIBMRAM in RAA

5.8.3.4 Cross Network Searches

Now we invoke ATELL in USIBMRAG.BERND to send a message to USIBMRAM.AUNG. This results in BERND sending a Locate Find search request to RAA, and RAA in turn sending the request to RAS, which is the first entry in RAA's adjacent cluster table. The session path is successfully found through RAS and the following displays confirm the fact.

Note in Figure 96 on page 74 that the directory entry for AUNG (**1**) in RAA shows AUNG as an end node connected to RAS. Also note in the adjacent cluster display in Figure 95 that the search for LU name USIBMRAM.AUNG is successful through RAS (**2**) and therefore RA3 (**3**) is not searched.

```
* RAAAN    D NET,ADJCLUST,NETID=USIBMRAM
  RAAAN    IST097I  DISPLAY  ACCEPTED
' RAAAN
IST350I  DISPLAY TYPE = ADJACENT CLUSTER TABLE
IST1325I  DYNAMIC TABLE FOR USIBMRAM       DYNAMICS = FULL
IST1326I  CP NAME           TYPE    STATE      STATUS       SNVC
IST1327I  USIBMRAT.RAS      DYNAMIC ACTIVE     FOUND        003  2
IST1327I  USIBMRAG.RAA      DYNAMIC ACTIVE     *** N/A ***  N/A
IST1327I  USIBMRAY.RA3      DYNAMIC ACTIVE     NOT SEARCHED 003  3
IST314I  END
```

Figure 95. Adjacent Cluster Table for USIBMRAM in RAA

```
* RAAAN    D NET,DIRECTRY,ID=USIBMRAM.AUNG,E
  RAAAN    IST097I  DISPLAY  ACCEPTED
' RAAAN
IST350I  DISPLAY TYPE = DIRECTORY
IST1186I  DIRECTORY ENTRY = DYNAMIC    EN
IST1184I  CPNAME = USIBMRAM.AUNG     - NETSRVR = USIBMRAT.RAS  1
IST314I  END
```

Figure 96. Directory Entry for AUNG in RAA

5.8.3.5 Dynamic Reordering of Adjacent Cluster Table

VTAM will not vary the order of the entries (priority list) in the adjacent cluster table as long as the search to the first entry in the table is successful. If alternative CPs are searched, then VTAM will dynamically adjust the priority list so that the successful entry moves to the top of the list.

This is tested in our test scenario by disabling the AHHC connection from RAA to RAS. As the result, RAS is deleted from the adjacent cluster table, and the table now looks as shown in Figure 97.

```
* RAAAN    D NET,ADJCLUST,NETID=USIBMRAM
  RAAAN    IST097I  DISPLAY  ACCEPTED
' RAAAN
IST350I  DISPLAY TYPE = ADJACENT CLUSTER TABLE
IST1325I  DYNAMIC TABLE FOR USIBMRAM       DYNAMICS = FULL
IST1326I  CP NAME          TYPE    STATE      STATUS       SNVC
IST1327I  USIBMRAG.RAA      DYNAMIC ACTIVE     *** N/A ***  N/A
IST1327I  USIBMRAY.RA3      DYNAMIC ACTIVE     NOT SEARCHED 003
IST314I  END
```

Figure 97. Adjacent Cluster Table for USIBMRAM in RAA

We then invoke ATELL again to send a message from USIBMRAG.BERND to USIBMRAM.AUNG. As shown in Figure 98, the directory entry for AUNG in RAA now shows AUNG as an end node connected to RA3. The adjacent cluster display in Figure 99 on page 75 shows that the search for LU name USIBMRAM.AUNG is successful through RA3, and that RA3 has now been moved to the top of the priority list.

```
* RAAAN    D NET,DIRECTRY,ID=USIBMRAM.AUNG,E
  RAAAN    IST097I  DISPLAY  ACCEPTED
' RAAAN
IST350I  DISPLAY TYPE = DIRECTORY
IST1186I  DIRECTORY ENTRY = DYNAMIC    EN
IST1184I  CPNAME = USIBMRAM.AUNG     - NETSRVR = USIBMRAY.RA3
IST314I  END
```

Figure 98. Directory Entry for AUNG in RAA

```
* RAAAN    D NET,ADJCLUST,NETID=USIBMRAM
  RAAAN    IST097I  DISPLAY  ACCEPTED
' RAAAN
IST350I  DISPLAY TYPE = ADJACENT CLUSTER TABLE
IST1325I  DYNAMIC TABLE FOR USIBMRAM       DYNAMICS = FULL
IST1326I  CP NAME          TYPE    STATE      STATUS       SNVC
IST1327I  USIBMRAY.RA3     DYNAMIC ACTIVE     FOUND        003
IST1327I  USIBMRAG.RAA     DYNAMIC ACTIVE     *** N/A ***  N/A
IST314I  END
```

Figure 99. Adjacent Cluster Table for USIBMRAM in RAA

Re-enabling the connection between RAA and RAS causes RAS to reappear in the adjacent cluster table — at the bottom of the priority list (**1**) as shown in Figure 100. Deleting the directory entry in RAA for AUNG and invoking ATELL, to again cause a search, shows that the search goes through RA3 as expected.

```
* RAAAN    D NET,ADJCLUST,NETID=USIBMRAM
  RAAAN    IST097I  DISPLAY  ACCEPTED
' RAAAN
IST350I  DISPLAY TYPE = ADJACENT CLUSTER TABLE
IST1325I  DYNAMIC TABLE FOR USIBMRAM       DYNAMICS = FULL
IST1326I  CP NAME          TYPE    STATE      STATUS       SNVC
IST1327I  USIBMRAY.RA3     DYNAMIC ACTIVE     FOUND        003
IST1327I  USIBMRAG.RAA     DYNAMIC ACTIVE     *** N/A ***  N/A
IST1327I  USIBMRAT.RAS     DYNAMIC ACTIVE     NOT SEARCHED 003  1
IST314I  END
```

Figure 100. Adjacent Cluster Table for USIBMRAM in RAA

5.8.3.6 Setting Subnet Visit Count (SNVC)

In this section we test the effect of changing the subnet visit count (SNVC) value in the various subnets in the test configuration. SNVC defines the maximum number of subnetworks that the border node will search when looking for a resource. SNVC=1 restricts the search to the current network. SNVC=2 restricts the search to the adjacent networks.

In order to test the impact of SNVC, we first change the SNVC value in RAA using the MODIFY VTAMOPTS command. RAA now has an SNVC value of 2. RA3 and RAS continue to have the default SNVC value of 3.

Given that it requires the visit of 3 subnets to reach USIBMRAM.AUNG from USIBMRAG.BERND, we expect the next ATELL test to fail. The next ATELL test fails as expected. Figure 101 on page 76 shows the result of DISPLAY ADJCLUST command after the ATELL test. It shows in **1** and **2** that both RA3 and RAS are used during the search but the search is not successful.

```
* RAAAN    D NET,ADJCLUST,NETID=USIBMRAM
  RAAAN    IST097I  DISPLAY  ACCEPTED
' RAAAN
IST350I  DISPLAY TYPE = ADJACENT CLUSTER TABLE
IST1325I  DYNAMIC TABLE FOR USIBMRAM       DYNAMICS = FULL
IST1326I  CP NAME           TYPE    STATE      STATUS       SNVC
IST1327I  USIBMRAY.RA3      DYNAMIC ACTIVE     NOT FOUND    002  1
IST1327I  USIBMRAT.RAS      DYNAMIC ACTIVE     NOT FOUND    002  2
IST1327I  USIBMRAG.RAA      DYNAMIC ACTIVE     *** N/A ***  N/A
IST314I  END
```

Figure 101. Adjacent Cluster Table for USIBMRAM in RAA

To determine the impact of changing the SNVC value in the intermediate CPs, we reset the SNVC value in RAA to 3 and set the SNVC value in RA3 to 1. We set the SNVC value in RAS to 2.

Given that to reach USIBMRAM from USIBMRAY where RA3 resides takes two subnet visits, we expect the search through RA3 to fail but the alternate search through RAS to succeed as RAS has the SNVC value set to 2.

The next ATELL test confirms our expectations. The test succeeds and the directory display as shown in Figure 102 confirms that AUNG is reached through RAS. Figure 103 shows the adjacent cluster display after the test. It shows in **1** that RAS as the CP involved in the most recent successful search now moves to the top of the priority list. It shows in **2** that RA3 now moves a position down the list.

```
* RAAAN    D NET,DIRECTRY,ID=USIBMRAM.AUNG,E
  RAAAN    IST097I  DISPLAY  ACCEPTED
' RAAAN
IST350I  DISPLAY TYPE = DIRECTORY
IST1186I  DIRECTORY ENTRY = DYNAMIC    EN
IST1184I  CPNAME = USIBMRAM.AUNG     - NETSRVR = USIBMRAT.RAS
IST314I  END
```

Figure 102. Directory Entry for AUNG in RAA

```
* RAAAN    D NET,ADJCLUST,NETID=USIBMRAM
  RAAAN    IST097I  DISPLAY  ACCEPTED
' RAAAN
IST350I  DISPLAY TYPE = ADJACENT CLUSTER TABLE
IST1325I  DYNAMIC TABLE FOR USIBMRAM       DYNAMICS = FULL
IST1326I  CP NAME           TYPE    STATE      STATUS       SNVC
IST1327I  USIBMRAT.RAS      DYNAMIC ACTIVE     FOUND        003   1
IST1327I  USIBMRAY.RA3      DYNAMIC ACTIVE     NOT FOUND    003   2
IST1327I  USIBMRAG.RAA      DYNAMIC ACTIVE     *** N/A ***  N/A
IST314I  END
```

Figure 103. Adjacent Cluster Table for USIBMRAM in RAA

5.8.3.7 Using Adjacent Cluster Definitions in VTAM

VTAM V4R2 provides you with an option to predefine entries in the adjacent cluster table. This is done by using adjacent cluster definitions in VTAM. Full details on the definition statements involved can be found in *VTAM Resource Definition Reference*.

Figure 104 shows the predefined adjacent cluster table that we will use. Note in **1** that by defining **BNDYN=NONE** for the default network we override the start parameter. As shown in **2**, we specify **BNDYN=FULL** for the net ID of USIBMRAM. Using the NEXTCP statement as shown in **3** and **4**, we define two entries for the net ID of USIBMRAM.

```
RAAACLST VBUILD  TYPE=ADJCLUST
DEFAULT  NETWORK BNDYN=NONE 1
RANETP   NETWORK NETID=USIBMRAM,SNVC=5,BNDYN=FULL 2
RACP1P   NEXTCP  CPNAME=USIBMRAT.RAS 3
RACP2P   NEXTCP  CPNAME=USIBMRAY.RA3 4
```

Figure 104. Predefined Adjacent Cluster Table in RAA

Figure 105 shows the adjacent table display after the predefined table has been activated. You can see in **1** and **2** that no dynamic entries are added to the adjacent cluster table except for the net ID of USIBMRAM. You can also see in **3** and **4** the two predefined entries for the net ID of USIBMRAM.

```
* RAAAN    D NET,ADJCLUST,SCOPE=ALL
  RAAAN    IST097I  DISPLAY  ACCEPTED
' RAAAN
IST350I  DISPLAY TYPE = ADJACENT CLUSTER TABLE
IST1325I  DEFINED TABLE FOR DEFAULT_NETID  DYNAMICS = NONE 1
IST924I  -------------------------------------------------------------
IST1325I  DEFINED TABLE FOR USIBMRAM       DYNAMICS = FULL 2
IST1326I  CP NAME          TYPE    STATE      STATUS       SNVC
IST1327I  USIBMRAT.RAS     DEFINED ACTIVE     NOT SEARCHED 005 3
IST1327I  USIBMRAY.RA3     DEFINED ACTIVE     NOT SEARCHED 005 4
IST1327I  USIBMRAG.RAA     DYNAMIC ACTIVE     *** N/A ***  N/A
IST924I  -------------------------------------------------------------
IST1325I  DEFAULT TABLE FOR USIBMRAT       DYNAMICS = NONE 1
IST924I  -------------------------------------------------------------
IST1325I  DEFAULT TABLE FOR USIBMRAY       DYNAMICS = NONE 1
IST314I  END
```

Figure 105. Adjacent Cluster Display in RAA

We now perform an ATELL test from USIBMRAG.BERND to USIBMRAM.AUNG. The search goes through RAS and the session is set up through RAS as we expect. Figure 106 on page 78 confirms the fact.

```
* RAAAN    D NET,ADJCLUST,NETID=USIBMRAM
  RAAAN    IST097I  DISPLAY  ACCEPTED
' RAAAN
IST350I  DISPLAY TYPE = ADJACENT CLUSTER TABLE
IST1325I  DEFINED TABLE FOR USIBMRAM       DYNAMICS = FULL
IST1326I  CP NAME           TYPE    STATE      STATUS       SNVC
IST1327I  USIBMRAT.RAS      DEFINED ACTIVE     FOUND        005
IST1327I  USIBMRAY.RA3      DEFINED ACTIVE     NOT SEARCHED 005
IST1327I  USIBMRAG.RAA      DYNAMIC ACTIVE     *** N/A ***  N/A
IST314I  END
```

Figure 106. Adjacent Cluster Display for USIBMRAM in RAA

To assess the flexibility of adjacent cluster tables, we change the configuration so that only RA3 is predefined. This is done by activating an alternate adjacent cluster table in VTAM. Figure 107 shows the alternate adjacent cluster table that we use.

```
RAAACLST VBUILD  TYPE=ADJCLUST
DEFAULT  NETWORK BNDYN=NONE
RANETP   NETWORK NETID=USIBMRAM,SNVC=3,BNDYN=FULL
RACP2P   NEXTCP  CPNAME=USIBMRAY.RA3
```

Figure 107. Alternate Adjacent Cluster Table in RAA

Figure 108 shows the adjacent table display after the alternate table has been activated. You can see in **1** that RA3 has been added as a predefined entry, and in **2** that RAS has been added as a dynamic entry.

```
* RAAAN    D NET,ADJCLUST,NETID=USIBMRAM
  RAAAN    IST097I  DISPLAY  ACCEPTED
' RAAAN
IST350I  DISPLAY TYPE = ADJACENT CLUSTER TABLE
IST1325I  DEFINED TABLE FOR USIBMRAM       DYNAMICS = FULL
IST1326I  CP NAME           TYPE    STATE      STATUS       SNVC
IST1327I  USIBMRAY.RA3      DEFINED ACTIVE     NOT SEARCHED 003  1
IST1327I  USIBMRAG.RAA      DYNAMIC ACTIVE     *** N/A ***  N/A
IST1327I  USIBMRAT.RAS      DYNAMIC ACTIVE     NOT SEARCHED 003  2
IST314I  END
```

Figure 108. Adjacent Cluster Display in RAA

To determine the impact of a CP-CP connection failure, we disable the connection between RAA and RAS. Figure 109 on page 79 shows the adjacent cluster display after the connection is disabled. Note in **3** that RAS is still kept in the display but with a STATE of NOT ACTIVE.

```
* RAAAN    D NET,ADJCLUST,NETID=USIBMRAM
  RAAAN    IST097I  DISPLAY  ACCEPTED
' RAAAN
IST350I  DISPLAY TYPE = ADJACENT CLUSTER TABLE
IST1325I  DEFINED TABLE FOR USIBMRAM       DYNAMICS = FULL
IST1326I  CP NAME           TYPE    STATE      STATUS       SNVC
IST1327I  USIBMRAY.RA3      DEFINED NOT ACTIVE FOUND        003  3
IST1327I  USIBMRAG.RAA      DYNAMIC ACTIVE     *** N/A ***  N/A
IST1327I  USIBMRAT.RAS      DYNAMIC ACTIVE     NOT SEARCHED 003
IST314I  END
```

Figure 109. Adjacent Cluster Display in RAA

We next invoke ATELL to send a message from BERND to AUNG and as expected the search and the session go through RAS as shown in **4** in Figure 110.

```
* RAAAN    D NET,ADJCLUST,NETID=USIBMRAM
  RAAAN    IST097I  DISPLAY  ACCEPTED
' RAAAN
IST350I  DISPLAY TYPE = ADJACENT CLUSTER TABLE
IST1325I  DEFINED TABLE FOR USIBMRAM       DYNAMICS = FULL
IST1326I  CP NAME           TYPE    STATE      STATUS       SNVC
IST1327I  USIBMRAY.RA3      DEFINED NOT ACTIVE FOUND        003
IST1327I  USIBMRAT.RAS      DYNAMIC ACTIVE     FOUND        003  4
IST1327I  USIBMRAG.RAA      DYNAMIC ACTIVE     *** N/A ***  N/A
IST314I  END
```

Figure 110. Adjacent Cluster Display in RAA

Re-enabling the connection from RAA to RA3 reverts the STATE of RA3 in the adjacent cluster table to ACTIVE as shown in **5** in Figure 111.

```
* RAAAN    D NET,ADJCLUST,NETID=USIBMRAM
  RAAAN    IST097I  DISPLAY  ACCEPTED
' RAAAN
IST350I  DISPLAY TYPE = ADJACENT CLUSTER TABLE
IST1325I  DEFINED TABLE FOR USIBMRAM       DYNAMICS = FULL
IST1326I  CP NAME           TYPE    STATE      STATUS       SNVC
IST1327I  USIBMRAY.RA3      DEFINED ACTIVE  5  FOUND        003
IST1327I  USIBMRAT.RAS      DYNAMIC ACTIVE     FOUND        003
IST1327I  USIBMRAG.RAA      DYNAMIC ACTIVE     *** N/A ***  N/A
IST314I  END
```

Figure 111. Adjacent Cluster Display in RAA

We now activate another adjacent cluster table that has BNDYN=NONE coded for the net ID of USIBMRAM as shown in **1** in Figure 112 on page 80. As the result of this, we expect the next adjacent cluster display to contain **no** dynamic entry for the net ID of USIBMRAM.

```
RAAACLST VBUILD  TYPE=ADJCLUST
DEFAULT  NETWORK BNDYN=NONE
RANETP   NETWORK NETID=USIBMRAM,SNVC=3,BNDYN=NONE 1
RACP1P   NEXTCP  CPNAME=USIBMRAT.RAS
```

Figure 112. Alternate Adjacent Cluster Table in RAA

Figure 113 shows the result. Note that, as expected, the display shows no dynamic entry that has been added.

```
* RAAAN    D NET,ADJCLUST,NETID=USIBMRAM
  RAAAN    IST097I  DISPLAY  ACCEPTED
' RAAAN
IST350I  DISPLAY TYPE = ADJACENT CLUSTER TABLE
IST1325I  DEFINED TABLE FOR USIBMRAM       DYNAMICS = NONE
IST1326I  CP NAME           TYPE    STATE      STATUS       SNVC
IST1327I  USIBMRAT.RAS      DEFINED ACTIVE     NOT SEARCHED 003
IST314I  END
```

Figure 113. Adjacent Cluster Display in RAA

5.9 Hints and Tips

Connecting networks with extended border node connections is easy, compared to SNI. (Few parameters have to be set.) You get an extended border node connection by starting as border nodes both VTAMs at the connection end points.

Use BNDYN=FULL to make use of the dynamic search functions across network borders. Only large networks require more control by using adjacent cluster definitions and BNDYN=LIMITED.

Increase the SNVC (subnet visit count) to the appropriate value. It must be large enough to cover all backup search paths. It can be changed by operator command.

Define a CDRM major node containing the local VTAM with parameters CDRDYN=YES and CDRSC=OPT. If the CDRM major node consists of the local VTAM only, NETID need not be coded.

Use architected COS names. If user defined COS names must be used, at least use identical COS definitions within one topology subnetwork.

Use APPNCOS=#CONNECT as a start option.

You cannot use multiple link TGs with border node connections, because border node connections are always FID2 connections.

In any configuration where sessions are to be set up across a border, you must ensure that the network nodes involved can interoperate with a border node. That is, they implement the APPN base function set 1013.

We suggest you use **extended** border node connections.

Peripheral border node connections must be used when communicating with VTAM V4R1 network nodes or network nodes other than VTAM. You can set up sessions across **one** peripheral border in any configuration.

You can use the peripheral border node function to inter-connect two APPN subnetworks that do not have border node capabilities, via an APPN network with extended border node capabilities — that is, a subnetwork that has VTAM V4R2 extended border nodes at connection points to the other subnetworks.

When there are two peripheral border connections in a search path, they can only be at both ends of the search path. One or more extended border node connections must be between them.

When setting up a peripheral border connection between VTAM and PS/2, do not define VTAM as network node in the PS/2. During connection setup, VTAM changes its role from network node to end node and, therefore, must be defined as an unknown node type in the PS/2.

Chapter 6. Virtual Route-Based Transmission Group

In this chapter we describe how VTAM V4R2 optimizes route selection by mapping APPN routing on to the subarea network. This support is called "virtual route-based transmission group", or VR-TG.

6.1 Benefits

1. In a network consisting of both APPN and subarea portions, VR-TG allows route selection to be performed across the *whole* topology subnetwork rather than within the individual (subarea and APPN) portions. This means that, for a given session, the best possible route according to the defined class of service (COS) is much more likely to be chosen.

   Without VR-TG, although the session route *within each portion* may be the optimum, the route as a whole is not optimized in any way.

2. Network management is made easier because the *entire* topology is known to each node in the network; it is not necessary to consult multiple management agents for a view of the session path.

3. Networks with multilink TGs between NCPs (not supported in the APPN architecture) can take advantage of APPN connectivity.

6.2 How VR-TG Works

A "virtual route-based transmission group" is an APPN representation of the subarea connectivity between two VTAM domains. As such it encompasses *all* possible subarea paths (VRs) between those VTAM domains. Both the topology of an APPN network, and the session establishment methods, are affected by VR-TG.

6.2.1 Topology with VR-TG

To illustrate the difference between the APPN topologies with and without VR-TG, consider the diagram in Figure 114 on page 84.

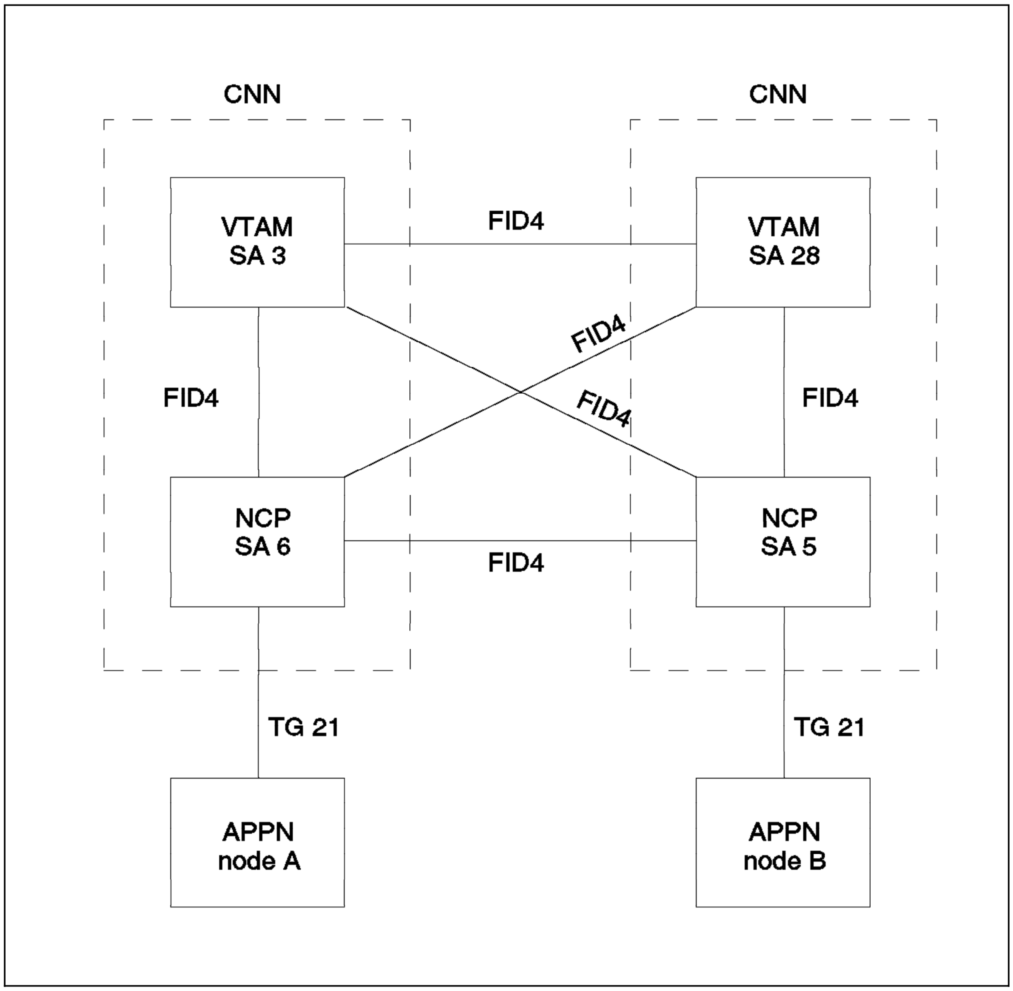

Figure 114. Two VTAM Domains with VR-TG

If VTAM 3 owns NCP 6, and VTAM 28 owns NCP 5, and both are APPN network nodes, the network comprises two composite network nodes. If all the links shown are FID4, we have full subarea connectivity in the network; however without VR-TG there is no APPN connection between the two VTAMs and there is no topology exchange between them. Without VR-TG the APPN topology is disjoint, and looks like the diagram in Figure 115 on page 85.

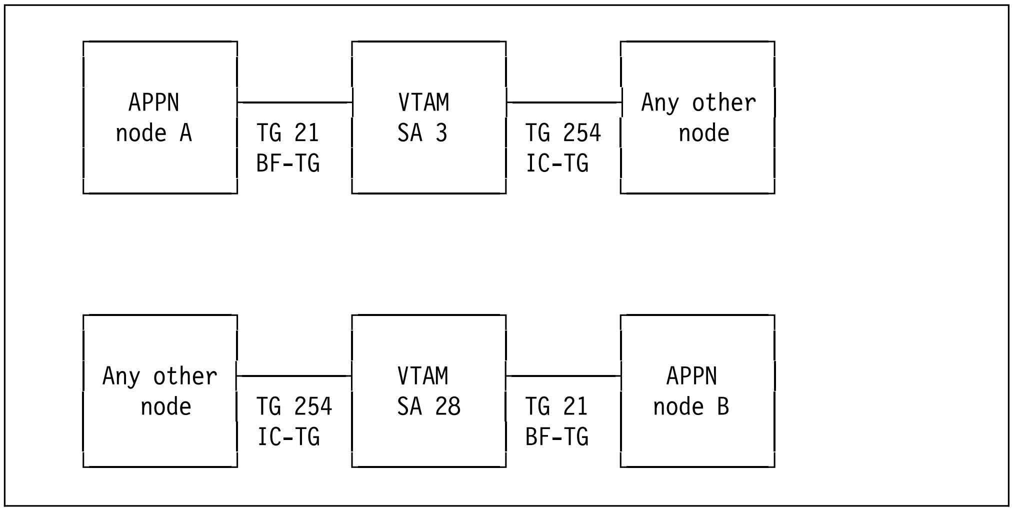

Figure 115. APPN View of Topology - IC-TG

The TGs which attach the APPN network to a subarea network are known by VTAM as "boundary function transmission groups" (BF-TGs). The links which appear to terminate the APPN networks are called "interchange transmission groups" (IC-TGs), and always take the number 254. Since they are endpoint links, they will not be found in the network topology database.

A route calculated by a network node in the APPN network (for instance, A) to a node in or across the subarea network (for instance, B or an application in subarea 28) will be of the following form:

A - (TG21) - VTAM 3 - (TG254) - B.

At the subarea boundary VTAM 3 takes over and selects a virtual route across the subarea network to VTAM 28. Finally VTAM 28 finishes off the route across the APPN network, which now looks like the following:

A - (TG254) - VTAM 28 - (TG21) - B.

Thus the route is computed in three separate stages, each independent of the other two.

With VR-TG, that connection across the subarea network is treated as an APPN transmission group (now always numbered 255), and reported accordingly in the topology database exchanges. The APPN view of the entire network looks like Figure 116.

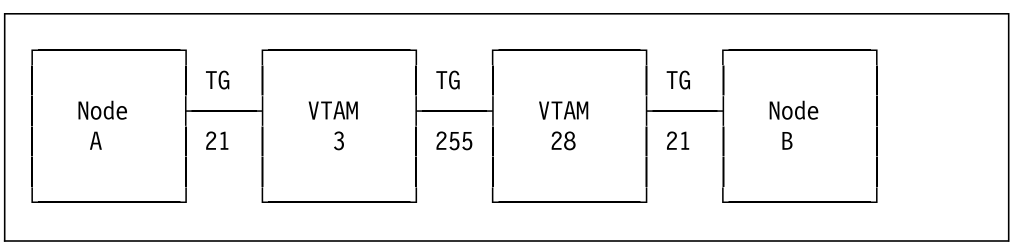

Figure 116. APPN View of Topology - VR-TG

Node A is now able to decide - based on correct information - between this route via the subarea network, and some other route that may exist between A and B. If the subarea network is chosen, the SSCP at the entry point (subarea 3 in this case) chooses the best route (to subarea 5 in this case) between all the possible paths linking subareas 3, 5, 6 and 28; TG 255 represents all those paths.

6.2.2 Sessions with VR-TG

With VR-TG - as in the case where there are separate APPN and subarea connections between VTAMs - both SSCP-SSCP and CP-CP sessions exist. The difference is that with VR-TG, the CP-CP sessions take a subarea path. A comparison between the five possible cases will illustrate the point:

1. If two VTAMs have only a subarea connection (without VR-TG), there are no CP-CP sessions and the SSCP sessions take the subarea route.
2. If two VTAMs have only an APPN connection, there are no SSCP sessions and the CP-CP sessions take an APPN route. The adjacent CP is seen by each VTAM as an independent LU on the adjacent link station that connects to the other VTAM.
3. If two VTAMs have separate APPN and subarea connections (without VR-TG), both (1) and (2) are true - all the sessions exist.
4. If two VTAMs have only a VR-TG connection, the SSCP sessions take a subarea route and the CP-CP sessions also take a subarea route (which may be different from the SSCP route). The adjacent CP is seen by each VTAM as a cross-domain resource belonging to the other VTAM′s CDRM.
5. If two VTAMs have both a VR-TG and an APPN connection, the SSCP sessions take a subarea route. The CP-CP sessions may take either route, depending on which TG VTAM chooses for them.

Remember also that there are two CP-CP sessions, which may take different routes. It is even possible for one to take a subarea path and the other to take an APPN path, in case (5) above.

VR-TGs do not affect how resources are searched for. If a target resource is found in a VR-TG "partner" VTAM′s domain, the SSCP session is used to set up the LU-LU session over the correct virtual route.

6.2.3 VR-TG in Practice

With multiple VTAMs, as with multiple network nodes in an APPN network, it is not necessary (indeed, not desirable) to have CP-CP sessions between each pair. It is possible to have a VR-TG without CP-CP sessions, in which case there will be no direct APPN searches between the two VTAMs, but session setup using TG 255 will be possible. It is even possible to set up such a session without any VR-TG at all, provided there is a VR-TG path via other VTAMs. In this case, the two (or more) TG 255s will be combined into one. This process is called RSCV pruning, and is performed by the first VTAM to see the RSCV - either the one that calculates it in the first place, or the first one to receive it after some other node has calculated it. RSCV pruning also occurs when the session path must traverse a VR-TG *and* an adjacent IC-TG, in which case a single TG 254 is used for the session path.

Figure 117 on page 87 illustrates these points.

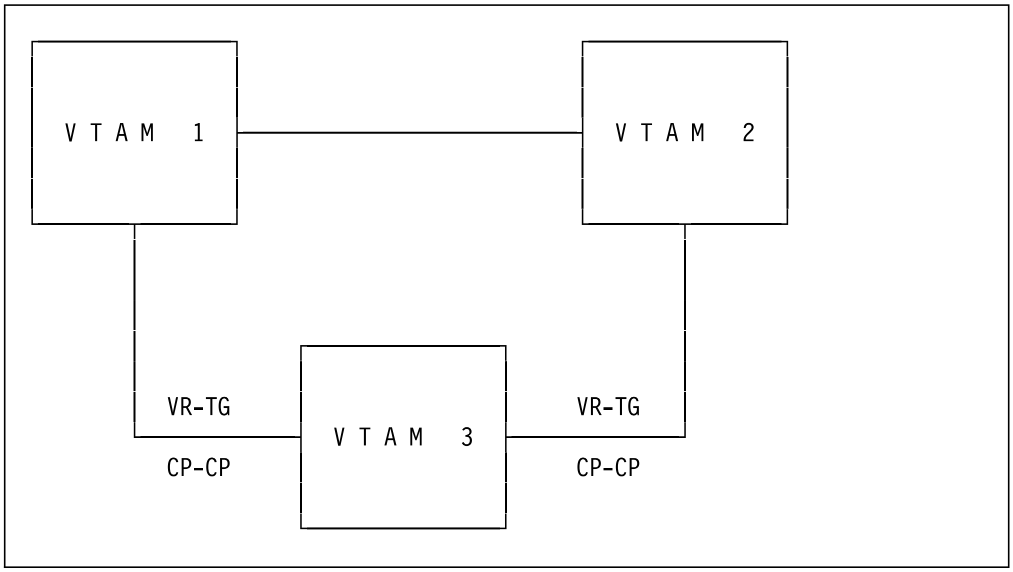

Figure 117. VTAM Connections with and without VR-TG

In this diagram, the links between VTAM 3 and its two companions are VR-TGs, with CP-CP sessions.

If the link between VTAM 1 and VTAM 2 is a VR-TG *without* CP-CP sessions, no APPN searches are sent on it; however, a session between the domains of VTAM 1 and VTAM 2 will be set up using TG 255. This is exactly the equivalent of two APPN nodes having a link, but no CP-CP session.

If the link between VTAM 1 and VTAM 2 is *not a VR-TG at all*, then it will not appear in the topology database and no APPN route will be calculated across it. Any route between VTAM 1 and VTAM 2 will be calculated via VTAM 3, using two TG 255s in the RSCV. However, VTAM 1 must ensure that the subarea portion of the route is a single VR. It calculates the best VR to VTAM 2, then prunes the RSCV so that it contains a single TG 255 directly between VTAM 1 and VTAM 2. Therefore, a direct route can be set up between VTAM 1 and VTAM 2 **even though there is no path in the topology database** between them.

To summarize the rules regarding VR-TG:

- You **must have** meshed SSCP-SSCP connectivity within the subarea network. In other words, every VTAM must have an SSCP session with every other VTAM.
- You **need not have** meshed VR-TG or meshed CP-CP connectivity within the subarea network.
- You **must have** CP-CP connectivity to all VTAMs in the subarea network. In other words, every VTAM must be able to reach every other VTAM via a path (preferably more than one path) of CP-CP sessions.
- You **must not have** meshed VR-TG or meshed CP-CP connectivity in a large subarea network, with more than a dozen or so VTAMs. This is because excessive TDUs and search requests will flow through the network.

6.3 VR-TG for Optimum Routing

In this section we describe how VR-TG aids in selecting the best route in a complex network comprising both subarea and APPN portions.

6.3.1 Description

Figure 118 shows a complex arrangement of VTAMs and NCPs, with both subarea and APPN connections linking them. In practice, the two composite network nodes would not be connected in this manner unless they were remote from each other, but our test equipment is all in the same room; we use a VM virtual channel-to-channel connection to illustrate VR-TG where a remote 3172 (or, more likely, an NCP - NCP link) would be used in a real installation.

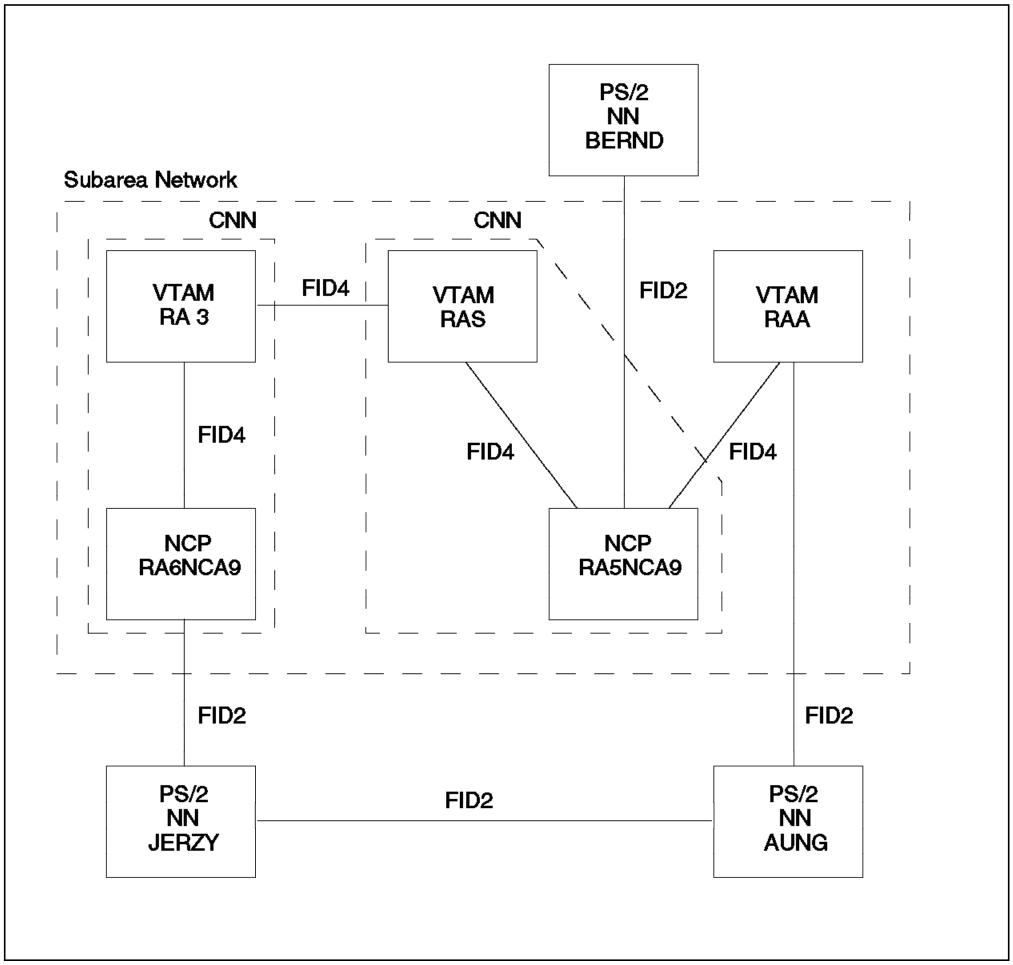

Figure 118. VR-TG Network Diagram

In the figure, two composite network nodes (RA3 and RAS) and one VTAM interchange node (RAA) are all connected via FID4 links. The three PCs are all network nodes in order to allow them to calculate routes; in practice they are more likely to be network node servers for a number of end nodes. The direct connection between RAA and AUNG is via a 3172.

The APPN view of this network is shown in Figure 119 on page 89.

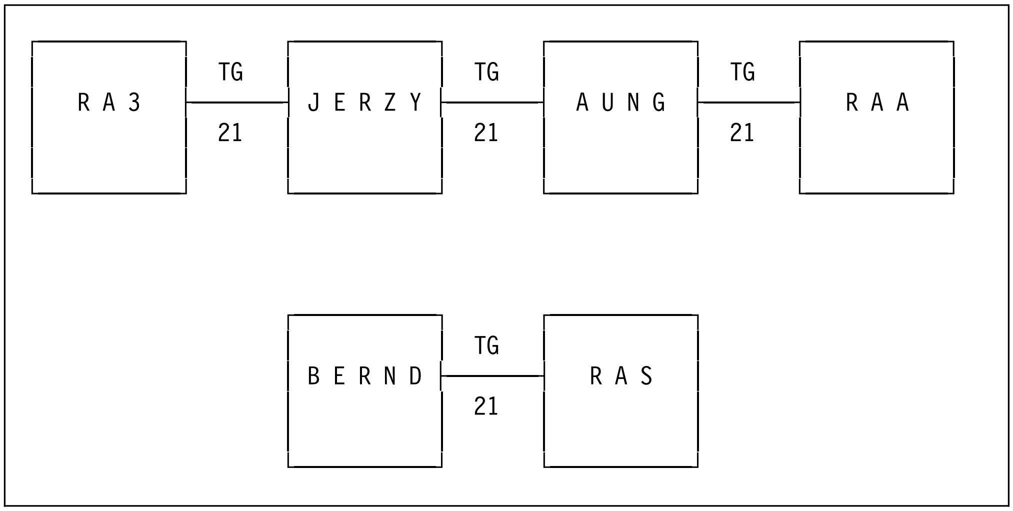

Figure 119. APPN View of Configuration without VR-TG

To demonstrate VR-TG, we:

- Set up the subarea network so that RAS searches RA3 before RAA
- Establish a session between BERND and AUNG
- Restart the network with VR-TG switched on
- Establish another session between BERND and AUNG

6.3.2 Definitions

The VTAM and NCP definitions we use before switching on VR-TG are straightforward subarea definitions, and can be seen in Appendix A, "Definitions" on page 193:

- NCPs RA6NCA9 and RA5NCA9 (Figure 256 on page 206)
- XCA major node RAA3172 (Figure 258 on page 212)
- Subarea MPC major nodes RASKRA3 and RA3KRAS (Figure 259 on page 213)
- CA major node RAACRX5 (Figure 260 on page 213)
- VTAM path tables RAAEV1, RA3EV1, RASEV1 (Figure 261 on page 214)

Of particular interest are the CDRM major nodes and the adjacent SSCP tables. The adjacent SSCP table we use in RAS ensures that RAS will search RA3 before RAA when it receives a session request, as shown in Figure 120. RAS will search RA3 whether the destination LU has a net ID specified or not.

```
        VBUILD TYPE=ADJSSCP
*       ADJSSCP TABLE TO ROUTE SEARCHES VIA RA3 BEFORE RAA
RA3    ADJCDRM
RAA    ADJCDRM
       NETWORK NETID=USIBMRA
RA3    ADJCDRM
RAA    ADJCDRM
```

Figure 120. ADJSSCP Table in RAS

In accordance with the usual practice, we use the same CDRM major node for all three VTAMs, as shown in Figure 121 on page 90.

```
       VBUILD TYPE=CDRM
*     CDRM MAJOR NODE FOR VR-TG TEST 1
       NETWORK NETID=USIBMRA
RAA      CDRM  SUBAREA=10,CDRDYN=YES,CDRSC=OPT
RAS      CDRM  SUBAREA=28,CDRDYN=YES,CDRSC=OPT
RA3      CDRM  SUBAREA=3,CDRDYN=YES,CDRSC=OPT
```

Figure 121. CDRM Major Node - No VR-TG

There are two ways to turn on VR-TG:

- Code VRTG=YES on the VTAM start parameters, in which case VTAM will attempt to use VR-TG with all its CDRM partners.
- Code VRTG=YES on the individual CDRM definitions. This method is recommended because it allows fine tuning of VR-TG design.

We adopt the second approach, although it means using individual CDRM major nodes for the three VTAMs, because we do not want a VR-TG between RAA and RA3. These two VTAMs are not physically adjacent and there is no need for a VR-TG. Figure 122 shows the CDRM definitions we use.

```
       VBUILD TYPE=CDRM
*     CDRM MAJOR NODE FOR VRTG TEST 1 PART 2 (RAA)
       NETWORK NETID=USIBMRA
RAA      CDRM  SUBAREA=10,CDRDYN=YES,CDRSC=OPT
RAS      CDRM  SUBAREA=28,CDRDYN=YES,CDRSC=OPT,VRTG=YES
RA3      CDRM  SUBAREA=3,CDRDYN=YES,CDRSC=OPT

       VBUILD TYPE=CDRM
*     CDRM MAJOR NODE FOR VRTG TEST 1 PART 2 (RAS)
       NETWORK NETID=USIBMRA
RAS      CDRM  SUBAREA=28,CDRDYN=YES,CDRSC=OPT
RAA      CDRM  SUBAREA=10,CDRDYN=YES,CDRSC=OPT,VRTG=YES
RA3      CDRM  SUBAREA=3,CDRDYN=YES,CDRSC=OPT,VRTG=YES

       VBUILD TYPE=CDRM
*     CDRM MAJOR NODE FOR VRTG TEST 1 PART 2 (RA3)
       NETWORK NETID=USIBMRA
RA3      CDRM  SUBAREA=3,CDRDYN=YES,CDRSC=OPT
RAA      CDRM  SUBAREA=10,CDRDYN=YES,CDRSC=OPT
RAS      CDRM  SUBAREA=28,CDRDYN=YES,CDRSC=OPT,VRTG=YES
```

Figure 122. CDRM Major Nodes - with VR-TG

6.3.3 Test Results

When the network is initialized and all the connections and CP-CP sessions are established, we display the topology to verify that it matches Figure 119 on page 89. Figure 123 on page 91, as an example, illustrates RAA′s view of its adjacent nodes.

```
NCCF                      N E T V I E W    RAAAN BUCZAK   02/04/94 13:37:30
* RAAAN    D NET,TOPO,ID=RAA,LIST=ALL
  RAAAN    IST097I  DISPLAY  ACCEPTED
' RAAAN
IST350I  DISPLAY TYPE = TOPOLOGY
IST1295I  CP NAME           NODETYPE ROUTERES CONGESTION  CP-CP WEIGHT
IST1296I  USIBMRA.RAA       NN       128      NONE        *NA*  *NA*
IST1297I                    ICN/MDH  CDSERVR  RSN
IST1298I                    YES      NO       30
IST1223I                    BN       NATIVE
IST1224I                    NO       YES
IST1299I  TRANSMISSION GROUPS ORIGINATING AT CP USIBMRA.RAA
IST1357I                                                  CPCP
IST1300I  DESTINATION CP    TGN      STATUS   TGTYPE      VALUE WEIGHT
IST1301I  USIBMRA.AUNG 1    21       OPER     INTERM      YES   *NA*
IST1301I  USIBMRA.RAK       21       INOP     INTERM      YES   *NA*
IST1301I  USIBMRA.RA3       21       INOP     INTERM      YES   *NA*
IST1301I  USIBMRA.RAS       21       INOP     INTERM      YES   *NA*
IST1301I  USIBMRA.RACONNET  21       INOP     INTERM      NO    *NA*
IST314I  END
```

Figure 123. Topology from RAA - No VR-TG

As usual, there are entries from previous tests in the topology; the only operative link is that to AUNG **1** which corresponds with Figure 119 on page 89. RA3 and RAS are known as CDRMs, not as CPs.

We then establish the session between BERND and AUNG, using ATELL. Figure 124 shows the session configuration.

```
NLDM.CON                  SESSION CONFIGURATION DATA                PAGE    1
-------------- PRIMARY ---------------+-------------- SECONDARY --------------
NAME BERND     SA 00000005  EL 009C   |   NAME AUNG     SA 00000006  EL 00A5
--------------------------------------+---------------------------------------
DOMAIN RASAN       PCID USIBMRA.BERND.E287CA233FBF9106          DOMAIN RASAN
               +-------------+                  +-------------+
               |             | ---          --- |   CP/SSCP   | AUNG
RA5NCA9 (0000) | SUBAREA PU  |  APPN TP 02      | LOCAL DATA  |
     2         +------+------+  SUBA TP 00      +------+------+
                      |              VR 02             |
               +------+------+       ER 05 3    +------+------+
J0005013       |    LINK     |      RER 02      |     LU      | AUNG    (00A5)
               +------+------+                  +-------------+
                      |       APPNCOS #INTER
               +------+------+
W32256  (001E) |     PU      | LOGMODE #INTER
               +------+------+ PADJ CP BERND
                      |
               +------+------+
BERND   (009C) |     ILU     |
     1         +-------------+

SELECT PT, ST (PRI, SEC TRACE), RT (RESP TIME), P, ER, VR, AR
CMD==>
```

Figure 124. Session Configuration - No VR-TG

We can see that the session path goes from BERND **1** via NCP RA5NCA9 **2**, then via ER 5 **3** to its destination AUNG. Figure 125 on page 92 shows where ER 5 goes.

```
 NLDM.ER                   SPECIFIC ER CONFIGURATION                  PAGE    1
 ------------------------------------------------------------------------------
 SUBAREA1 00000005   SUBAREA2 00000006  ER  05 | NODES (TOTAL/MIGRATION): 04/00
 ------------------------------------------------------------------------------
                              (A)
                               V
  +-----+ NAME: RA5NCA9     +-----+ NAME: N/A
  | INN |   SA: 00000005    | INN |   SA: 00000003
  +--+--+ SSCP: RAS      1  +--+--+ SSCP: N/A      3
     |                         |
 1) TG001                  3) TG001
     |                         |
  +--+--+ NAME: ISTPUS28    +--+--+ NAME: N/A
  | INN |   SA: 0000001C    | INN |   SA: 00000006
  +--+--+ SSCP: RAS      2  +--+--+ SSCP: N/A      4
     |
 2) TG001
     |
     V
    (A)

END OF DATA
ENTER SEL# (FOR TG DETAIL)
CMD==>
```

Figure 125. Explicit Route - No VR-TG

The route goes from subarea 5 **1** to subarea 28 (RAS) **2** to subarea 3 (RA3) **3** to subarea 6 (RA6NCA9) **4**. A similar session display taken in RA3 shows that the destination LU AUNG is reached via a link station on RA6NCA9.

The complete route taken is thus as follows:

BERND - RA5NCA9 - RAS - RA3 - RA6NCA9 - JERZY - AUNG.

This is because:

- BERND, which calculates the first part of the APPN route, does not know of the subarea link between RAS and RAA, nor of the APPN link between RAA and AUNG.
- RAS, which calculates the subarea route, does not know of the APPN link between AUNG and RAA.

Next, we terminate the session between BERND and AUNG, delete the (dynamically created) CDRSC definitions in RAS, and inactivate all the CDRM sessions between the VTAMs. Thus we start with a "clean" network for the next stage.

Figure 126 on page 93 shows what happens when we activate the new CDRM definitions in the last of the three VTAMs, which happens to be RAA.

```
NCCF                      N E T V I E W    RAAAN BUCZAK   02/04/94 13:55:09
* RAAAN    ACT RAACDRX
C RAAAN    VARY NET,ACT,ID=RAACDRX
  RAAAN    IST097I  VARY     ACCEPTED
  RAAAN    IST093I  RAACDRX  ACTIVE
  RAAAN    IST093I  RA3      ACTIVE
  RAAAN    IST1086I  APPN CONNECTION FOR USIBMRA.RAS       IS ACTIVE - TGN =
           255   1
  RAAAN    IST093I  RAS      ACTIVE
  RAAAN    IST1096I  CP-CP SESSIONS WITH USIBMRA.RAS       ACTIVATED 2
```

Figure 126. Activation of CDRM with VR-TG

Activation of RA3 proceeds as before; activation of RAS immediately results in activation of the APPN connection **1** using TG 255, and the CP-CP sessions **2**. A display of the topology in RAA now shows the new link, as in Figure 127.

```
NCCF                      N E T V I E W    RAAAN BUCZAK   02/04/94 13:56:30
* RAAAN    D NET,TOPO,ID=RAA,LIST=ALL
  RAAAN    IST097I  DISPLAY  ACCEPTED
' RAAAN
IST350I  DISPLAY TYPE = TOPOLOGY
IST1295I  CP NAME           NODETYPE ROUTERES CONGESTION  CP-CP WEIGHT
IST1296I  USIBMRA.RAA       NN       128      NONE        *NA*  *NA*
IST1297I                    ICN/MDH  CDSERVR  RSN
IST1298I                    YES      NO       34
IST1223I                    BN       NATIVE
IST1224I                    NO       YES
IST1299I  TRANSMISSION GROUPS ORIGINATING AT CP USIBMRA.RAA
IST1357I                                                  CPCP
IST1300I  DESTINATION CP    TGN      STATUS   TGTYPE      VALUE WEIGHT
IST1301I  USIBMRA.AUNG      21       OPER     INTERM      YES   *NA*
IST1301I  USIBMRA.RAK       21       INOP     INTERM      YES   *NA*
IST1301I  USIBMRA.RACONNET  21       INOP     INTERM      NO    *NA*
IST1301I  USIBMRA.RA3       21       INOP     INTERM      YES   *NA*
IST1301I  USIBMRA.RAS       21       INOP     INTERM      YES   *NA*
IST1301I  USIBMRA.RAS    1  255      OPER     INTERM VRTG YES   *NA*
IST314I  END
```

Figure 127. Topology from RAA - with VR-TG

The subarea link to RAS now shows up as an APPN link using TG 255 **1**, with the extra indication “VRTG” under TGTYPE to confirm its nature. There is no operative APPN link to RA3. Figure 128 on page 94 shows the full APPN topology, now complete.

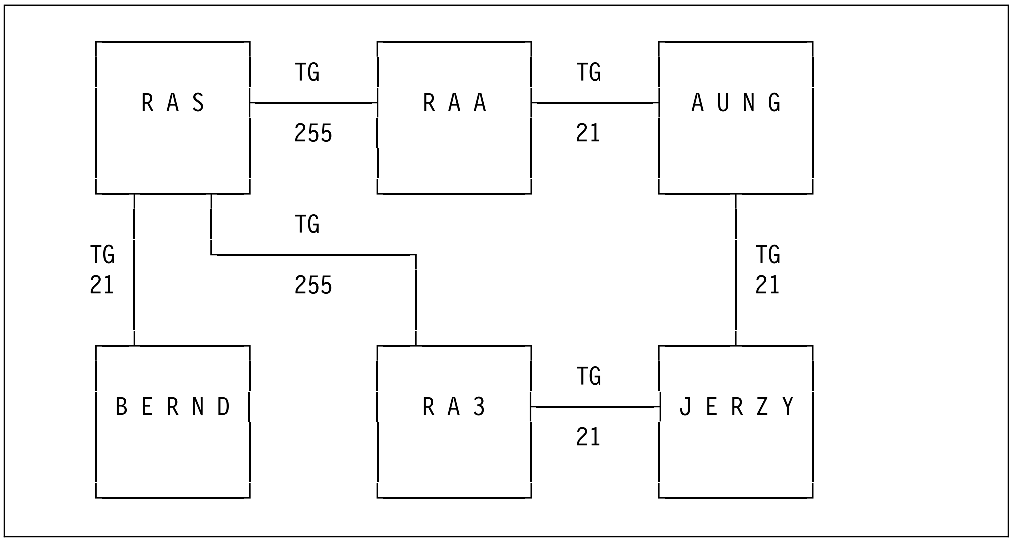

Figure 128. APPN View of Configuration with VR-TG

Now we set up the LU 6.2 session between BERND and AUNG again. This time, the display of the session route looks exactly like Figure 124 on page 91, except that the explicit route is shown as ER 0 to subarea 10, not ER 5 to subarea 6. Figure 129 shows a display of ER 0.

```
 NLDM.ER                    SPECIFIC ER CONFIGURATION                PAGE    1
 ------------------------------------------------------------------------------
 SUBAREA1 00000005   SUBAREA2 0000000A   ER  00 | NODES (TOTAL/MIGRATION): 02/00
 ------------------------------------------------------------------------------


  +-----+ NAME: RA5NCA9
  | INN |   SA: 00000005
  +--+--+ SSCP: RAS
     |
 1) TG001     1
     |
  +--+--+ NAME: N/A
  | INN |   SA: 0000000A
  +--+--+ SSCP: N/A







END OF DATA
ENTER SEL# (FOR TG DETAIL)
CMD==>
```

Figure 129. Explicit Route - with VR-TG

This time, the session has taken the shorter route from RAS to RAA (subarea 10) via subarea TG 1 **1**. The APPN session route can now be seen in Figure 130 on page 95.

```
 NLDM.AR               APPN SESSION ROUTE CONFIGURATION                PAGE   1
 -- PRIMARY ---+-- SECONDARY --+--------------- PCID ---------------+- DOMAIN -
 NAME BERND    | NAME AUNG     | USIBMRA.BERND.E287CA2341BF9106     | RASAN
 --------------+---------------+------------------------------------+----------
                                       TG021 |
  +---------+                           +----+----+
  |   CP    |                           |   CP    |
  |BERND    |                           |AUNG     |
  +----+----+                           +---------+
 TG021 |
  +----+----+
  | CP(ICN) |
  |RAS      |
  +----+----+
 TG255 |       1
  +----+----+
  | CP(ICN) |
  |RAA      |
  +----+----+
       |

END OF DATA
SELECT PAR, SAR
CMD==>
```

Figure 130. APPN Session Route - with VR-TG

Here we can see that the session goes between RAS and RAA using APPN TG 255 **1**. Thus the complete session path is:

BERND - RA5NCA9 - RAA - AUNG.

The shorter session route has been chosen because:

- BERND knows of the APPN link between RAS and RAA, and is aware that the route through those nodes has a lower weight than the alternative route through RA3. In fact, all the TGs have the same characteristics in this network so the route with the smallest number of hops is always selected.
- RAS, as the SSCP responsible for the subarea boundary with BERND, has determined the best subarea route between RA5NCA9 and RAA.

6.4 VR-TG Over a Multilink Transmission Group

Many customers' networks include multilink transmission groups between NCPs, either for granularity in bandwidth or for greater availability. The APPN architecture, however, does not support multilink TGs; therefore such TGs cannot be converted to FID2 parallel TGs without some loss of resilience.

VR-TG provides the ability to use multilink TGs for APPN *connectivity*, although not for APPN *dynamics* since subarea path tables are still required.

6.4.1 Description

Figure 131 shows two composite network nodes connected via a two-link TG between their NCPs.

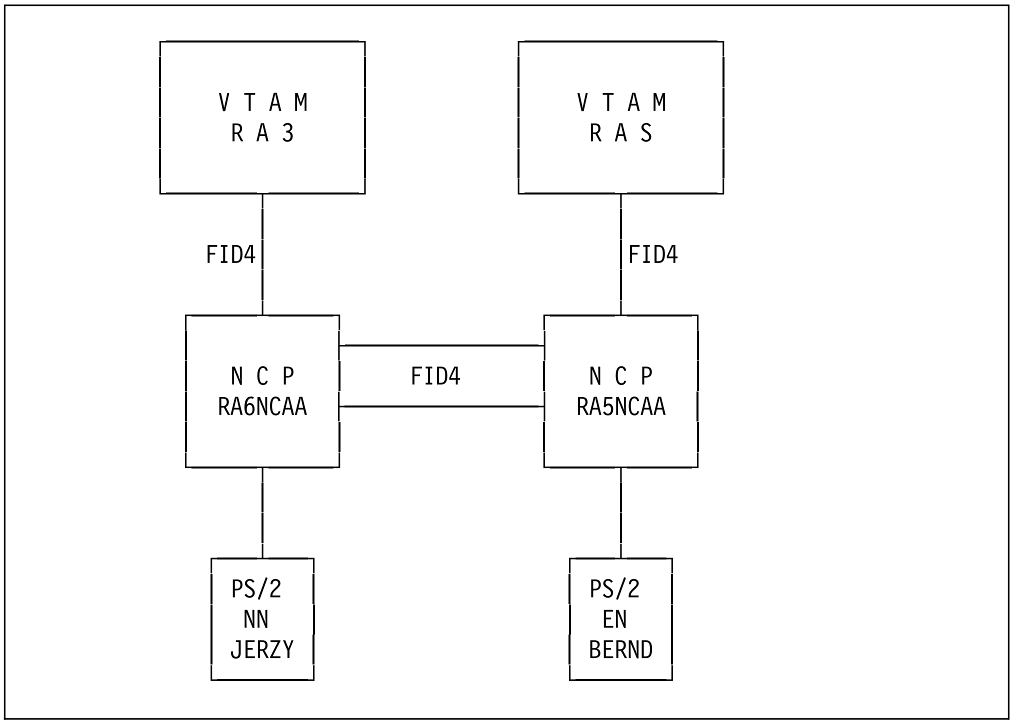

Figure 131. VR-TG with a Multilink Transmission Group

The two links between the NCPs are both token-ring subarea connections. With VR-TG, the APPN topology looks like Figure 132.

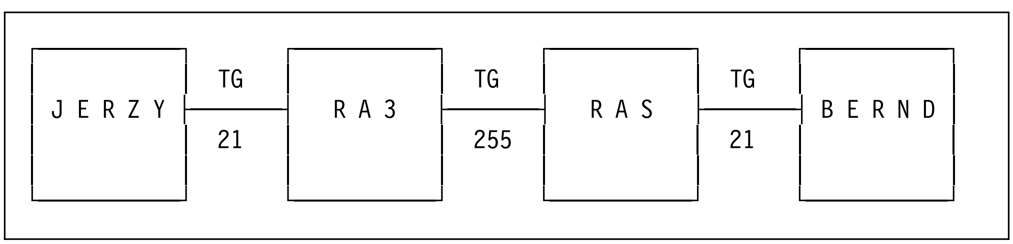

Figure 132. APPN Topology - VR-TG with Multilink Transmission Group

6.4.2 Definitions

Most of the definitions we use in this test are not APPN related, and can be seen in Appendix A, "Definitions" on page 193:

- NCPs RA6NCAA and RA5NCAA (Figure 257 on page 208)
- VTAM path tables RA3EV2 and RASEV2 (Figure 262 on page 216)

The CDRM definitions are the same as in the previous VR-TG example (RASCDRX and RA3CDRX) and can be seen in Figure 122 on page 90.

6.4.3 Test Results

We set up all the connections and verify the topology as shown in Figure 133.

```
NCCF                      N E T V I E W    RA3AO BUCZAK   02/09/94 14:46:43
* RA3AO    D NET,TOPO,ID=RA3,LIST=ALL
  RA3AO    IST097I  DISPLAY  ACCEPTED
' RA3AO
IST350I  DISPLAY TYPE = TOPOLOGY
IST1295I  CP NAME           NODETYPE ROUTERES CONGESTION  CP-CP WEIGHT
IST1296I  USIBMRA.RA3       NN       128      NONE        *NA*  *NA*
IST1297I                    ICN/MDH  CDSERVR  RSN
IST1298I                    YES      NO       4
IST1223I                    BN       NATIVE
IST1224I                    NO       YES
IST1299I  TRANSMISSION GROUPS ORIGINATING AT CP USIBMRA.RA3
IST1357I                                                  CPCP
IST1300I  DESTINATION CP    TGN      STATUS   TGTYPE      VALUE WEIGHT
IST1301I  USIBMRA.RAS       255      OPER     INTERM VRTG YES   *NA*
IST1301I  USIBMRA.JERZY     21       OPER     INTERM      YES   *NA*
IST314I  END
```

Figure 133. Topology Display from RA3

Next, we establish a session between JERZY and BERND, and verify both the APPN and subarea views of it. Figure 134 on page 98 shows the APPN route, which corresponds to Figure 132 on page 96.

```
NLDM.AR              APPN SESSION ROUTE CONFIGURATION               PAGE   1
-- PRIMARY ---+-- SECONDARY --+--------------- PCID ---------------+- DOMAIN -
NAME JERZY    | NAME BERND    | USIBMRA.JERZY.D54B4E4B334BD70D     | RASAN
--------------+---------------+------------------------------------+----------
                                       TG021 |
 +---------+                            +----+----+
 |   CP    |                            |   CP    |
 |JERZY    |                            |BERND    |
 +----+----+                            +---------+
TG021 |
 +----+----+
 | CP(ICN) |
 |RA3      |
 +----+----+
TG255 |
 +----+----+
 | CP(ICN) |
 |RAS      |
 +----+----+
      |

END OF DATA
SELECT PAR, SAR
CMD==>
```

Figure 134. APPN View of Session

Figure 135 shows the subarea route taking TG 6 **1** between the two NCPs (subareas 6 and 5).

```
NLDM.ER                  SPECIFIC ER CONFIGURATION                   PAGE    1
------------------------------------------------------------------------------
SUBAREA1 00000006   SUBAREA2 00000005   ER  00 | NODES (TOTAL/MIGRATION): 02/00
------------------------------------------------------------------------------


 +-----+ NAME: RA6NCAA
 | INN |   SA: 00000006
 +--+--+ SSCP: RA3
    |
1) TG006    1
    |
 +--+--+ NAME: N/A
 | INN |   SA: 00000005
 +--+--+ SSCP: N/A






END OF DATA
ENTER SEL# (FOR TG DETAIL)
CMD==>
```

Figure 135. Explicit Route for Session

Finally, a display of the link stations on NCP RA5NCAA (Figure 136 on page 99) shows that TG 6 does, indeed, have two links **1** .

```
* RASAN    D NET,STATIONS,ID=RA5NCAA
  RASAN    IST097I  DISPLAY  ACCEPTED
' RASAN
IST350I  DISPLAY TYPE = STATIONS
IST393I  PU T4/5 MAJOR NODE RA5NCAA , SUBAREA =         5
IST396I  LNKSTA    STATUS     CTG GTG   ADJNODE      ADJSA     NETID
IST397I  RA5P132   NEVAC        8   8                   0
IST610I                             LINE RA5L132  - STATUS NEVAC
IST397I  RA5UA881  ACTIV----E   6   6   RA6NCAA         6   USIBMRA
IST610I                       1     LINE RA5IA881 - STATUS ACTIV----E
IST397I  RA5UA891  ACTIV----E   6   6   RA6NCAA         6   USIBMRA
IST610I                             LINE RA5IA891 - STATUS ACTIV----E
IST397I  RA5PCHA3  ACTIV----E   1   1  ISTPUS28        28   USIBMRA
IST610I                             LINE RA5LCHA3 - STATUS ACTIV----E
IST314I  END
```

Figure 136. Display of TG 6

6.5 Hints and Tips

VR-TG brings significant benefits and is very easy to implement, so we recommend that it be used between all eligible VTAM V4R2 APPN nodes. “Eligible” means a VTAM with both subarea and APPN connections, so VR-TG is available only on interchange nodes and migration data hosts.

There are, however, some points worth remembering:

- VR-TG connections are not possible across SNI boundaries.
- VR-TG does not remove the need for path table definitions; APPN-only connections are required to do away with path tables, but such connections are not practical in many cases such as multilink TGs and backup SSCPs.
- Two migration data hosts linked via a VR-TG will not have CP-CP sessions as they are APPN end nodes.
- Make sure that VRTG=YES is specified on the CDRM definition at both ends of a connection; otherwise, the CDRM session will be set up without any APPN connection or CP-CP sessions.
- If two VTAMs are connected via both VR-TG and FID2 links, care needs to be taken over defining the VR-TG characteristics, in order to ensure the desired load balance between FID4 and FID2 links. The VR-TG characteristics (TGP, COSTBYTE, CAPACITY and so on) can be defined on the CDRM statement.
- If the network is large, do not let all the parameters default but consider carefully where you need VR-TG and where you need CP-CP sessions. The following advice is summarized from the recommendations given in *VTAM V4R2 for MVS/ESA Network Implementation Guide*:

  1. Define CDRM connectivity between *every* pair of VTAMs.
  2. Define CP-CP connectivity (and therefore VR-TG) such that *every* VTAM can be reached via an APPN search path.
  3. Do *not* define VR-TG and CP-CP connectivity between VTAMs whose domains are not physically adjacent. RSCV pruning will take care of sessions passing through an intermediate VTAM.
  4. If there are multiple CMCs in a network, define sufficient VR-TGs and CP-CP sessions between the CMCs to meet the network′s availability requirements. In other words, build in just enough redundancy to ensure that each CMC can be reached by an APPN search if a resource fails. Do *not* define VR-TGs in addition to this minimum requirement.
  5. Define a VR-TG (with CP-CP sessions) between each migration data host and its network node server. (If a migration data host has more than one potential server, define a VR-TG to each such node.) Do *not* define VR-TGs between MDHs since RSCV pruning makes them unnecessary.

Chapter 7. CNN Routing

This chapter describes a new function that is transparent to the user. For certain configurations of mixed APPN and subarea networks, CNN routing provides a shorter route through the subarea network. We explain when and how it is done. We describe several cases where CNN routing automatically improves the route. We also show scenarios where it does not improve the route, and a non-optimal route is chosen. In these cases you may want to overrule the automatic process by giving preference to certain links. Examples for the necessary definitions are given.

7.1 Benefits

1. Whenever possible, VTAM calculates an APPN route which takes advantage of the optimal route in the subarea network (zero-hop route).
2. Whenever possible, VTAM reroutes an incoming BIND to take advantage of the optimal route in the subarea network (zero-hop route).

7.2 How CNN Routing Works

In general, the interchange node function in VTAM provides seamless connection between the subarea part and the APPN part of an SNA network. For all the functions you have in subarea, you get an equivalent or better function in APPN. In VTAM V4R1, the exception to this rule is "piece-wise optimal" routing. This means that for certain configurations you do not get the optimal route end to end.

In VTAM V4R2, the answers to this problem are VR-TG (covered in Chapter 6, "Virtual Route-Based Transmission Group" on page 83) and CNN Routing. These functions provide solutions for most configurations. Where they do not, in some cases a non-optimal route may be acceptable. In others, defining different APPNCOS weights may help. Again in others, this may be a reason to retain subarea networking.

In VTAM V4R1, non-optimal routes can be chosen when parallel TGs exist between an APPN node and a composite network node. A typical example could be a CMC where a VTAM owns two NCPs, as shown in Figure 137 on page 103. Towards a VTAM data host (APPN end node), there are channel connections from each NCP and a channel connection from the CMC VTAM. So in total, there are three channel connections. From the APPN point of view, which is shown in Figure 138 on page 103, these are three TGs between two nodes. For session establishment one TG is selected at random. This may not be optimal.

The CNN routing function in VTAM V4R2 can become effective in two different ways during sessions setup:

"Route calculation," when VTAM acts as the network node server

"BIND rerouting," when another node calculates the route

7.2.1 Route Calculation

When VTAM in the composite network node is the network node server of the OLU, VTAM has the task of calculating the route to the DLU.

If the OLU is connected to only one subarea node of the composite network node, and also the DLU is connected to only one subarea node of the composite network node, the CNN routing function has no effect. (If OLU and DLU connect to different subarea nodes, the VR between these subareas is chosen according to the subarea COS table. If OLU and DLU connect to the same subarea, a zero-hop route is chosen in the subarea network.)

If the OLU or the DLU or both are connected with TGs to more than one subarea node of the composite network node and if these TGs have equal weight for a particular COS, then VTAM compares the subarea numbers. Whenever the subarea number of the entry TG (OLU side) and the subarea number of the exit TG (DLU side) match, these TGs will be preferred for route selection. This results in a zero-hop route in the subarea network.

7.2.2 BIND Rerouting

When VTAM in the composite network node is not the network node server of the OLU, some other network node calculates the route to the DLU without any knowledge of the composite network node′s subarea network. Figure 137 on page 103 gives an example where the route is calculated by network node BERND. The CNN routing function becomes active, when a BIND reaches the composite network node. The composite network node has to leave the entry TG unchanged. But it may modify the exit TG.

If the DLU is connected to only one subarea of the composite network node, the CNN routing function does not take effect.

If the DLU is connected with TGs to more than one subarea of the composite network node an alternate TG will be substituted if all of the following conditions are met:

- An alternate TG originates at the same subarea as the entry TG. (For example, both PS/2 BERND and VTAM RA3 have TGs connecting to the same NCP.)
- The alternate TG terminates at the same CP as the exit TG. (For example, the channel links from both NCPs connect to VTAM RA3.)
- The characteristics of the alternate TG must match those of the original exit TG. (In our example, all three links are channel links and by default have the same characteristics.)

BIND rerouting can change the RSCV contained in the BIND. So, in a trace the RSCV in the BIND request can be different from the RSCV in the BIND response. The modified RSCV is reported to NetView.

An NCP owned by two or more VTAMs (shared NCP) is seen as belonging to several CPs. If the entry TG and the exit TG are controlled by different VTAMs, the CNN routing function does not take effect.

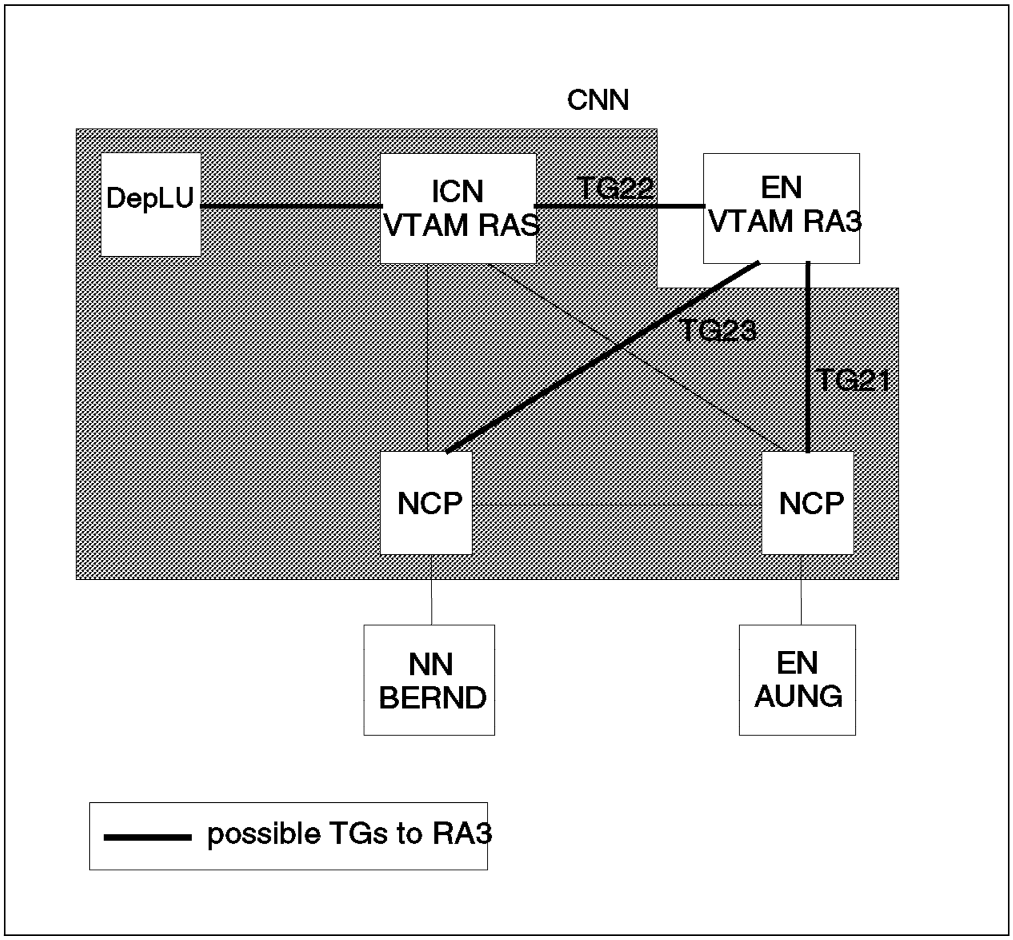

Figure 137. Network, where CNN Routing Takes Effect

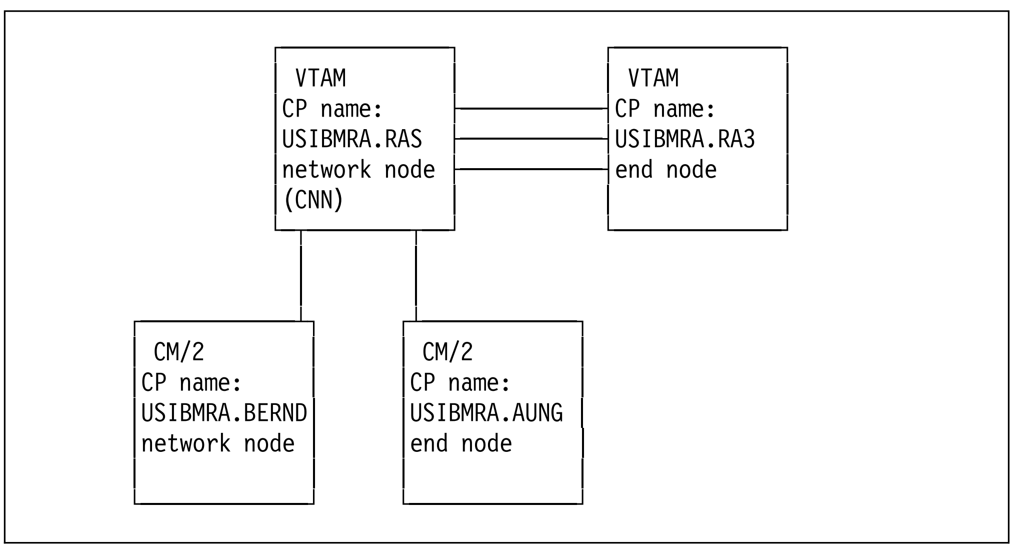

Figure 138. The APPN View of the Network

7.3 Optimal Routes

This section describes a combined subarea and APPN network in which the CNN routing function provides optimal routes end to end. Examples for both “route calculation” and “BIND rerouting” are given.

7.3.1 Description

As shown in Figure 137 on page 103, in this scenario there is a composite network node consisting of VTAM RAS (subarea 28) and two NCPs (RA5NCA3, RA6NCB1). There is an application host with VTAM RA3, an APPN end node. It is channel-connected to VTAM RAS and both NCPs. So there are three TGs between the composite network node RAS and end node RA3. The local terminal RAST421 belongs to subarea 28. An end node AUNG is connected to NCP RA6NCB1. A network node BERND is connected to NCP RA5NCA3. Everything is in the same net ID (USIBMRA).

7.3.2 Definitions

There are no parameters to control CNN routing. The function is always there. Of course, it only takes effect under the circumstances described above in 7.2, “How CNN Routing Works” on page 101.

Figure 140 shows the relevant parameters in the start list of an end node. Figure 139 contains the typical start parameters for an interchange node. For test purposes we want to make sure that routing trees are not reused. Each route calculation should pick one of the TGs between RA3 and RAS randomly. We achieve this by coding RESUSAGE=1 in the VTAM start parameters. Figure 141, Figure 142, and Figure 143 on page 105 contain the definitions of the channel connections between RA3 and RAS. Although they are the same as in previous scenarios, they are shown here to document the PU names. Later in this section, the PU names are used to identify a particular TG and to prove that CNN routing works as expected.

```
SSCPID=28,GWSSCP=YES,
HOSTSA=28,
CONFIG=SC,SUPP=NOSUP,
NETID=USIBMRA,HOSTPU=ISTPUS28,SSCPNAME=RAS,
NCPBUFSZ=2048,
PPOLOG=YES,DYNLU=YES,
NODETYPE=NN,
RESUSAGE=1,                                          CNN1
CPCP=YES,
```

Figure 139. Start Parameters for Composite Network Node

```
CONFIG=3C,SUPP=NOSUP,
NETID=USIBMRA,HOSTPU=ISTPUS03,SSCPNAME=RA3,
NODETYPE=EN,
DYNLU=YES,
CPCP=YES,
```

Figure 140. Start Parameters for End Node

```
RA3RA5  VBUILD TYPE=LOCAL           CHANNEL LINK VTAM3 TO NCP5
RA3LE28  PU    CUADDR=E28,                                         X
               PUTYPE=2,                                           X
               VPACING=0,                                          X
               XID=YES
```

Figure 141. Definition in EN RA3 for Link to NCP Subarea 5

```
RA3RA6  VBUILD TYPE=LOCAL           CHANNEL LINK VTAM3 TO NCP6
RA3LE29  PU    CUADDR=E29,
               MAXBFRU=34,
               PUTYPE=2,
               XID=YES
```

Figure 142. Definition in EN RA3 for Link to NCP Subarea 6

```
RA3PUTIN VBUILD  TYPE=LOCAL
* APPN MPC LINK FROM RA3 TO RAS
RA3PCA8  PU   TRLE=RA3KMPC,PUTYPE=2,XID=YES
```

Figure 143. Definition in EN RA3 for Link to VTAM Subarea 28

7.3.3 Test Results

Before we actually establish any sessions, we want to make sure that the three TGs between RAS and RA3 really have the same weight. The weight is always related to a COS. We will find that the LU 2 uses COS #CONNECT, the LU 6.2 in the PS/2s uses COS #INTER. Figure 144 and Figure 145 on page 106 show that all three TGs have equal weights for the same COS. The weight of one COS is different from that of the other COS (210 and 240, respectively).

```
D NET,TOPO,ID=RA3,APPNCOS=#CONNECT,LIST=ALL
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = TOPOLOGY 280
IST1295I CP NAME          NODETYPE ROUTERES CONGESTION  CP-CP WEIGHT
IST1296I USIBMRA.RA3      EN       *NA*     ***NA***    *NA*  32767
IST1297I                  ICN/MDH  CDSERVR  RSN
IST1298I                  NO       NO       2
IST1223I                  BN       NATIVE
IST1224I                  NO       *NA*
IST1299I TRANSMISSION GROUPS ORIGINATING AT CP USIBMRA.RA3
IST1357I                                                CPCP
IST1300I DESTINATION CP   TGN      STATUS   TGTYPE      VALUE WEIGHT
IST1301I USIBMRA.RAS      21       OPER     ENDPT       YES   210
IST1301I USIBMRA.RAS      22       OPER     ENDPT       YES   210
IST1301I USIBMRA.RAS      23       OPER     ENDPT       YES   210
IST314I END
```

Figure 144. Checking the Weight of Links from RA3 and RAS

```
D NET,TOPO,ID=RA3,APPNCOS=#INTER,LIST=ALL
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = TOPOLOGY 284
IST1295I CP NAME           NODETYPE ROUTERES CONGESTION  CP-CP WEIGHT
IST1296I USIBMRA.RA3       EN       *NA*     ***NA***    *NA*  32767
IST1297I                   ICN/MDH  CDSERVR  RSN
IST1298I                   NO       NO       2
IST1223I                   BN       NATIVE
IST1224I                   NO       *NA*
IST1299I TRANSMISSION GROUPS ORIGINATING AT CP USIBMRA.RA3
IST1357I                                                 CPCP
IST1300I DESTINATION CP    TGN      STATUS   TGTYPE      VALUE WEIGHT
IST1301I USIBMRA.RAS       21       OPER     ENDPT       YES   240
IST1301I USIBMRA.RAS       22       OPER     ENDPT       YES   240
IST1301I USIBMRA.RAS       23       OPER     ENDPT       YES   240
IST314I END
```

Figure 145. Checking the Weight of Links from RA3 and RAS

7.3.3.1 Session from LU 2 in VTAM′s Domain

We use a local LU on VTAM RAS to log on to TSO on VTAM RA3. In VTAM V4R1, the session may use any of the three TGs between composite network node RAS and end node RA3. The selection is done randomly. Now, let′s see if VTAM V4R2 will do any better!

Not being quite sure about the LU name which is logged on to TSO on RA3, we type in a new display command D NET,RSCLIST which allows the use of wildcard names **1** and displays all kinds of resource types **2** (Figure 146).

```
                    1
D NET,RSCLIST,ID=RAS*,A
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = RSCLIST 931      2
IST1417I NETID    NAME     STATUS     TYPE            MAJNODE
IST1418I USIBMRA  RAS      ACTIV      ADJCP           ISTADJCP
IST1418I USIBMRA  RAST421  ACT/S----Y CDRSC           ISTCDRDY
IST1418I USIBMRA  RAS      ACT/S----Y ADJACENT CP     ISTCDRDY
IST314I END
```

Figure 146. Display with Wildcard ID

After finding the terminal′s name, on RA3 we display RAST421. Figure 147 on page 107 shows it is connected to TSO in RA3, and actually uses the adjacent link station named RA3PCA8 **3**. Looking back to Figure 137 on page 103 and to the channel link definitions on page 105, we find that the session is routed on the optimal path. CNN routing works again.

In the same display, we also find the ALSLIST **4**. It gives the list of possible adjacent link stations. ISTAPNPU in this case means: any APPN capable link. We also get the session identifier to be used in the next display **5**. It proves our previous assumption that the session uses COS #CONNECT **6**.

Figure 148 on page 107 shows that TG 22 is used.

```
D NET,ID=RAST421,E
IST097I DISPLAY ACCEPTED
IST075I NAME = USIBMRA.RAST421, TYPE = CDRSC 060
IST486I STATUS= ACT/S----Y, DESIRED STATE= ACTIV
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 = RA3, VERIFY OWNER = NO
IST1184I CPNAME = USIBMRA.RAS - NETSRVR = ***NA***
IST1044I ALSLIST = ISTAPNPU  4
IST082I DEVTYPE = INDEPENDENT LU / CDRSC
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST171I ACTIVE SESSIONS = 0000000001, SESSION REQUESTS = 0000000000
IST206I SESSIONS:
IST1081I ADJACENT LINK STATION = RA3PCA8  3
IST634I NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I RA3AT01  ACTIV-P    F6FF416461762020 0002 0008  0  0 USIBMRA
IST314I END                    5
D NET,SESSION,SID=F6FF416461762020
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = SESSIONS 067
IST879I PLU/OLU REAL = USIBMRA.RA3AT01   ALIAS = ***NA***
IST879I SLU/DLU REAL = USIBMRA.RAST421   ALIAS = ***NA***
IST880I SETUP STATUS = ACTIV
IST875I ADJSSCP TOWARDS SLU = ISTAPNCP
IST875I ALSNAME TOWARDS SLU = RA3PCA8
IST933I LOGMODE=M2BSCNQ , COS=*BLANK*
IST875I APPNCOS TOWARDS SLU = #CONNECT  6
IST314I END
```

Figure 147. LU 2 Uses Optimal Route

```
NLDM.AR               APPN SESSION ROUTE CONFIGURATION               PAGE   1
-- PRIMARY ---+-- SECONDARY --+--------------- PCID ---------------+- DOMAIN -
NAME RA3AT01  | NAME RAST421  | USIBMRA.RA3.F6FF416461762020       | RASAN
--------------+---------------+------------------------------------+----------

 +---------+
 |   CP    |
 |RA3      |
 +----+----+
TG022 |
 +----+----+
 |   CP    |
 |RAS      |
 +---------+




END OF DATA
SELECT PAR, SAR
CMD==>
```

Figure 148. NLDM Shows Session across TG 22

7.3.3.2 Session with VTAM as NN Server

We use a PS/2 end node AUNG (connected to NCP RA6NCB1 of composite network node RAS) to establish an LU 6.2 session to an application in VTAM RA3. The application ATELL is used in the PS/2, an ATELL server named RA3ASER is used in VTAM RA3.

The point of this test is to show that the sessions from AUNG to RA3ASER are indeed established across the shortest path.

The display of the session partner in Figure 149 is done in RA3. It shows that adjacent link station RA3LE29 **1** is used. Looking back to Figure 137 on page 103 and to the channel link definitions on page 105, we find that the session is routed on the optimal path. CNN routing works again.

```
 D NET,ID=AUNG,E
 IST097I DISPLAY ACCEPTED
 IST075I NAME = USIBMRA.AUNG, TYPE = CDRSC 074
 IST486I STATUS= ACT/S----Y, DESIRED STATE= ACTIV
 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 = RA3, VERIFY OWNER = NO
 IST1184I CPNAME = USIBMRA.AUNG - NETSRVR = ***NA***
 IST1044I ALSLIST = ISTAPNPU
 IST082I DEVTYPE = INDEPENDENT LU / CDRSC
 IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
 IST171I ACTIVE SESSIONS = 0000000002, SESSION REQUESTS = 0000000000
 IST206I SESSIONS:
 IST1081I ADJACENT LINK STATION = RA3LE29  1
IST634I NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I RA3ASER  ACTIV-S    E5D3656D42B93AD8 0002 0001  0  0 USIBMRA
IST635I RA3ASER  ACTIV-S    E5D3656D41B93AD8 0001 0001  0  0 USIBMRA
IST314I END
```

Figure 149. LU in End Node Uses Optimal Route

There are several ways of finding the number of the transmission group which is used by the session. One method is doing a display of the adjacent link station as shown in Figure 68 on page 56. Using NLDM provides a more convenient way as the information is given for the whole session. Figure 150 on page 109 shows the NLDM display.

```
 NLDM.AR               APPN SESSION ROUTE CONFIGURATION                PAGE   1
 -- PRIMARY ---+-- SECONDARY --+--------------- PCID ---------------+- DOMAIN -
 NAME AUNG     | NAME RA3ASER  | USIBMRA.AUNG.E5D3656D41B93AD8      | RASAN
 --------------+---------------+------------------------------------+----------

  +---------+
  |   CP    |
  |AUNG     |
  +----+----+
 TG021 |
  +----+----+
  | CP(ICN) |
  |RAS      |
  +----+----+
 TG021 |
  +----+----+
  |   CP    |
  |RA3      |
  +---------+


END OF DATA
SELECT PAR, SAR
CMD==>
```

Figure 150. NLDM Shows Session across TG 21

7.3.3.3 Session with Non-VTAM Network Node Server

The previous sections (7.3.3.1, "Session from LU 2 in VTAM's Domain" and 7.3.3.2, "Session with VTAM as NN Server") give examples of the CNN routing function called "route calculation." This section shows the CNN routing function "BIND rerouting." The diagram we refer to is still Figure 137 on page 103.

The end node of the previous scenario is replaced by a network node. We use a PS/2 network node BERND (connected to NCP RA5NCA3 of composite network node RAS) to establish an LU 6.2 session to an application in VTAM RA3. In contrast to the previous scenario, the route is not calculated by the network node server RAS, but by network node BERND.

If network node BERND were a VTAM V4R2 CNN, it would have knowledge of the subarea domain of VTAM RAS, even if VTAM RAS were a VTAM V4R1. In our scenario, network node BERND is a PS/2 without the CNN routing function and consequently has no knowledge whatsoever of the topology of the subarea network. So, it will not be able to select the optimal route end to end. It will randomly pick one of the three TGs between composite network node RAS and end node RA3.

Figure 151 on page 110 shows that TG 23 is chosen. TG 23 is not chosen in the two previous scenarios. So it must be the optimal one. Figure 152 on page 110 proves it. RA3LE28 **1** is indeed the channel link to NCP RA5NCA3, as defined in Figure 141 on page 105.

Is this a coincidence? No. The mystery is resolved in the trace shown in Figure 153 and Figure 154 on page 112.

```
NLDM.AR              APPN SESSION ROUTE CONFIGURATION              PAGE   1
-- PRIMARY ---+-- SECONDARY --+--------------- PCID ---------------+- DOMAIN -
NAME BERND    | NAME RA3ASER  | USIBMRA.BERND.E287CA234389E599     | RASAN
--------------+---------------+------------------------------------+----------

 +---------+
 |   CP    |
 |BERND    |
 +----+----+
TG021 |
 +----+----+
 | CP(ICN) |
 |RAS      |
 +----+----+
TG023 |
 +----+----+
 |   CP    |
 |RA3      |
 +---------+


END OF DATA
SELECT PAR, SAR
CMD==>
```

Figure 151. NLDM Shows Session across TG 23

```
D NET,ID=BERND,E
IST097I DISPLAY ACCEPTED
IST075I NAME = USIBMRA.BERND, TYPE = CDRSC 005
IST486I STATUS= ACT/S----Y, DESIRED STATE= ACTIV
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 = RA3, VERIFY OWNER = NO
IST1184I CPNAME = USIBMRA.BERND - NETSRVR = ***NA***
IST1044I ALSLIST = ISTAPNPU
IST082I DEVTYPE = INDEPENDENT LU / CDRSC
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST171I ACTIVE SESSIONS = 0000000002, SESSION REQUESTS = 0000000000
IST206I SESSIONS:
IST1081I ADJACENT LINK STATION = RA3LE28  1
IST634I NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I RA3ASER  ACTIV-S    E287CA234289E599 0002 0001  0  0 USIBMRA
IST635I RA3ASER  ACTIV-S    E287CA234189E599 0001 0001  0  0 USIBMRA
IST314I END
```

Figure 152. LU in Network Node Uses Optimal Route

The CM/2 trace tool is used in network node BERND. It provides an excellent editing function.

The trace in Figure 153 contains the BIND request. The parts less relevant to this test are left out. The PLU and SLU are shown **2**. The COS name is #INTER **3**, as we have assumed. The FQPCID **4** proves network node BERND is the originator of the session. At the end of the BIND we find the route selection control vector (RSCV). It shows that network node BERND has calculated a route which takes TG 21 from itself to RAS, and TG 21 **5** from RAS to RA3. Not optimal.

```
Line:    172 Send MU
Time stamp: 13:47:08.78
DLC type: IBMTRNET
Adapter number: 00
Destination address: 40000105000004
ALS ID: 040027727601DB24
 ||TH: FID2, Exp, OIS, LFSID=0x10202, SNF=0x0005
 ||RH: RQ, SC, FI, OIC, RQD1
 ||BIND rq
 ||   Type = Negotiable
 ||   FM profile = 19
 ||   TS profile = 7
 ||   FM usage - primary:
  ...
 ||   FM usage - secondary:
  ...
 ||   FM Usage - common:
  ...
 ||   TS usage:
  ...
 ||   PS profile:
  ...
 ||   Cryptography options:
 ||      Private cryptography support = No
 ||      Session-level cryptography support = No
 ||   Primary LU name = USIBMRA.BERND 2
 ||   Structured user data:
 ||      Mode name = #INTER
 ||      Session instance identifier = 0x0102db240202db24
 ||      Network-qualified PLU name = USIBMRA.BERND
 ||   Secondary LU name = USIBMRA.RA3ASER 2
 ||   Fully qualified PCID control vector:
 ||      PCID = 0xe287ca234289e599 4
 ||      Network qualified CP name = USIBMRA.BERND
 ||   COS/TPF control vector:
 ||      Transmission priority = High
 ||      COS name = #INTER 3
 ||   Route selection control vector:
 ||      Maximum hop count = 2
 ||      Current hop count = 1
 ||      TG descriptor control vector
 ||         TG identifier control vector
 ||            TG number = 21
 ||            Partner name = USIBMRA.RAS
 ||            Partner node type = Real
 ||      TG descriptor control vector
 ||         TG identifier control vector
 ||            TG number = 21 5
 ||            Partner name = RA3
 ||            Partner node type = Real
```

Figure 153. BIND Request (before BIND Rerouting)

The BIND request travels across the network to reach the application in VTAM RA3. About a tenth of a second later the BIND response comes back, as you can see from the time stamp in Figure 154. The RSCV has changed **6**. CNN routing in VTAM RAS “reroutes” the BIND to take the optimal exit TG out of VTAM RAS.

The diligent reader will notice that several tests are done for this scenario. In fact, all tests show that network node BERND picks one TG at random out of the three possible TGs between VTAM RAS and VTAM RA3. But the final result remains the same: CNN routing in VTAM RAS makes sure the session takes the optimal route. We can see that it is the BIND *response* that matters. Its RSCV reflects the actual session path. This RSCV is reported to NetView, as the NLDM display in Figure 151 on page 110 proves.

```
Line:    186 Recv MU
Time stamp: 13:47:08.91
DLC type: IBMTRNET
Adapter number: 00
Destination address: 40000105000004
ALS ID: 040027727601DB24
 ||TH: FID2, Exp, OIS, LFSID=0x10202, SNF=0x0005
 ||RH: +RSP, SC, FI, RQD1
 ||BIND +rsp
 ||   Type = Negotiable
 ||   FM profile = 19
 ||   TS profile = 7
 ||   FM usage - primary:
  ...
 ||   FM usage - secondary:
  ...
 ||   FM Usage - common:
  ...
 ||   TS usage:
  ...
 ||   PS profile:
  ...
 ||   Cryptography options:
 ||      Private cryptography support = No
 ||      Session-level cryptography support = No
 ||   Primary LU name =
 ||   Structured user data:
 ||      Mode name = #INTER
 ||      Session instance identifier = 0x0287ca234289e599
 ||      Network-qualified SLU name = USIBMRA.RA3ASER
 ||   Fully qualified PCID control vector:
 ||      PCID = 0xe287ca234289e599
 ||      Network qualified CP name = USIBMRA.BERND
 ||   Route selection control vector:
 ||      Maximum hop count = 2
 ||      Current hop count = 2
 ||      TG descriptor control vector
 ||         TG identifier control vector
 ||            TG number = 21
 ||            Partner name = USIBMRA.RAS
 ||            Partner node type = Real
 ||      TG descriptor control vector
 ||         TG identifier control vector
 ||            TG number = 23  6
 ||            Partner name = RA3
 ||            Partner node type = Real
```

Figure 154. BIND Response (after BIND Rerouting)

7.4 Non-Optimal Routes

This section describes the test scenario where the CNN routing function in VTAM **cannot** ensure that the optimal subarea route within the composite network node is chosen. We also describe the solution, which involves assigning a lower line weight to the preferred TG.

7.4.1 Description

In the test environment, we use two MVS VTAM V4R2 systems and one PS/2 Communications Manager/2 system. Figure 155 shows the schematic diagram of the test environment.

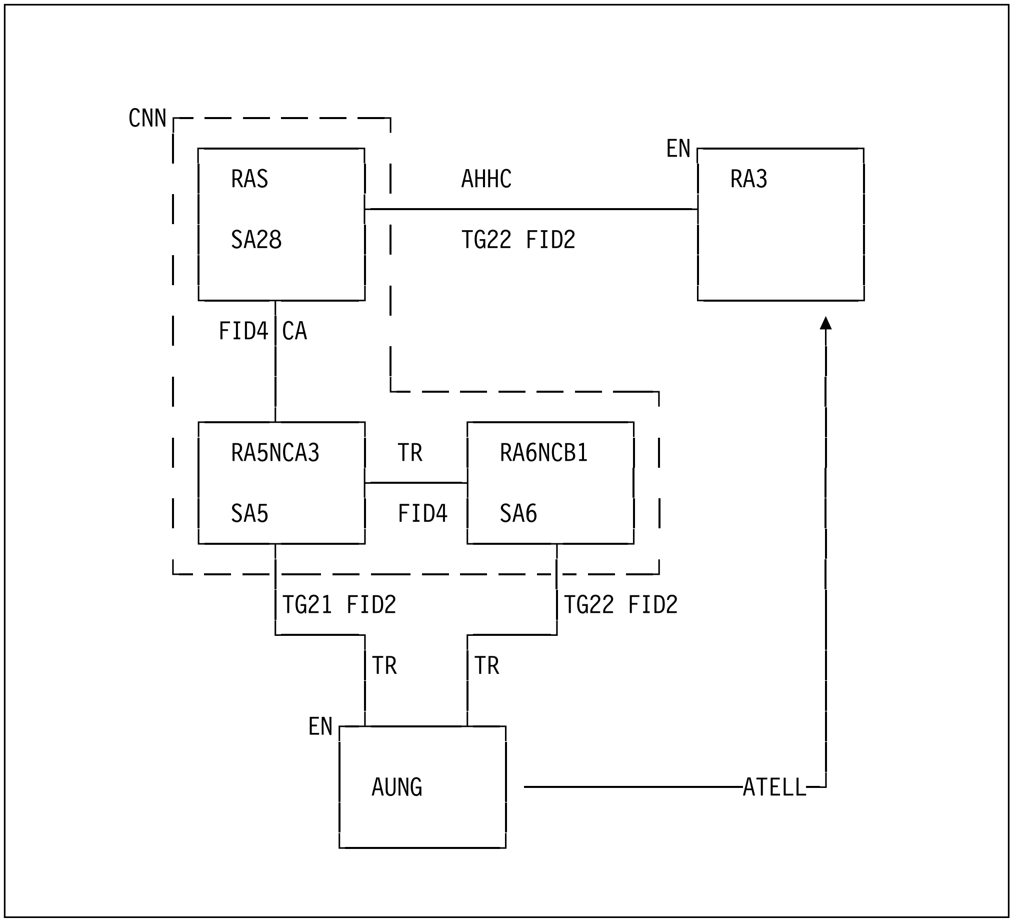

Figure 155. Schematic Diagram of the Test Environment

- RAS, a VTAM V4R2 composite network node, has an MPC connection to RA3. The MPC connections use AHHC as described in Chapter 3, "APPN Host-to-Host Connection" on page 11.

  It also has two NCPs that are interconnected as shown in Figure 155.
- RA3 is defined as a VTAM V4R2 end node. It connects to RAS via the AHHC connection.
- AUNG, a PS/2 running Communications Manager/2, is defined as an end node. It has two LAN connections to RAS via the two separate NCPs.

Figure 156 on page 114 shows the APPN view of the test environment.

We intend to show that, for sessions between AUNG and RA3:

1. When the default mode and COS are used, either TG 21 or TG 22, between AUNG and RAS, is selected randomly, even though TG 21 provides the better route within the composite network node.
2. This issue can be addressed by using a new APPNCOS to force the preferred route.

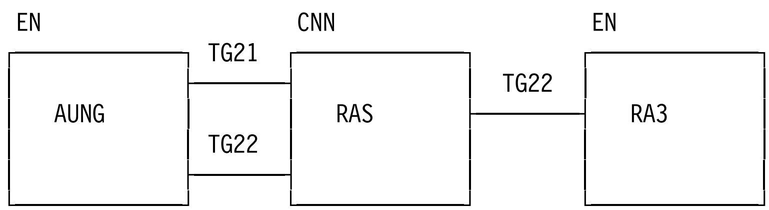

Figure 156. APPN View of the Test Environment

7.4.2 Definitions

This section contains the system definitions for the test environment. We intend to perform the test in two phases. The first phase will use default mode and default APPNCOS tables resulting in either of the two TGs from AUNG to RAS being used. We will set up additional mode and COS tables. The second phase will use them in order to always favor one TG, the TG with the optimal subarea route (TG 21 as shown in Figure 155 on page 113).

7.4.2.1 Link Definitions

In RAS, we define the major nodes for the AHHC connections to RA3. We also define the two NCPs. RA5NCA3 will be loaded and activated via the FID4 channel connection shown in Figure 155 on page 113. RA6NCB1 will also be loaded and activated via a FID4 channel but the channel will be disabled after the activation so that the only possible route from subarea 6 to subarea 28 is to go via subarea 5. In the NCPs, we define the token-ring INN connection between the two NCPs.

We also define the switched major node for the token-ring LAN connections from AUNG to RAS. The definitions are shown in Figure 157 on page 115. Note in **1** that the second PU has a pointer to a TG profile which is shown in Figure 158 on page 115.

This means that as seen from RAS, TG 21 will have a "user-defined parameter 1" (UDP1) value of 0 whereas TG 22 will have the default value which is 128. Given that the default APPNCOS table ignores the user defined parameters this will have no effect on the route selection if the default APPNCOS tables are used, allowing both TGs to be selected.

```
        VBUILD TYPE=SWNET,MAXNO=20,MAXGRP=5
 RASUP1   PU    ADDR=01,MAXDATA=521,MAXOUT=7,MAXPATH=5,PASSLIM=7,      X
                PUTYPE=2,                                              X
                CPNAME=AUNG,                                           X
                NETID=USIBMRA,                                         X
                CONNTYPE=APPN,                                         X
                DYNLU=YES,                                             X
                CPCP=YES
 *
 RASUP2   PU    ADDR=01,MAXDATA=521,MAXOUT=7,MAXPATH=5,PASSLIM=7,      X
                TGP=UPARMLOW,  1                                       X
                PUTYPE=2,                                              X
                CPNAME=AUNG,                                           X
                NETID=USIBMRA,                                         X
                CONNTYPE=APPN,                                         X
                DYNLU=YES,                                             X
                CPCP=YES
```

Figure 157. Definitions for Switched Connections from AUNG to RA3

```
* TG PROFILE TO CHANGE UPARM1 VALUE
*
UPARMLOW TGP   UPARM1=0      1
```

Figure 158. TG Profile Definition in RAS

We define AUNG as an end node. We define LINK0001, the logical link to subarea 6, and LINK0002, the logical link to subarea 5, by providing the appropriate MAC addresses. Figure 159 on page 116 shows the link definitions in the NDF file. Note that LINK0002, which connects to subarea 5, has been assigned the USER_DEFINED_1 value of 0 as shown in **3** whereas LINK0001, which connects to subarea 6, has been assigned the USER_DEFINED_1 value of 255 as shown in **4**. This means that, as seen from AUNG, TG 21 will have the UDP1 value of 0, and TG 22 will have 255. We will confirm this during the testing.

Note again that if the default APPNCOS tables are used specifying the UDP1 value will have no effect on the route selection, allowing both TGs to be selected.

```
DEFINE_LOGICAL_LINK  LINK_NAME(LINK0001)
                     ADJACENT_NODE_TYPE(NN)
                     PREFERRED_NN_SERVER(NO)
                     DLC_NAME(IBMTRNET)
                     ADAPTER_NUMBER(0)
                     DESTINATION_ADDRESS(X'40000106000004')
                     ETHERNET_FORMAT(NO)
                     CP_CP_SESSION_SUPPORT(YES)
                     SOLICIT_SSCP_SESSION(NO)
                     ACTIVATE_AT_STARTUP(YES)
                     USE_PUNAME_AS_CPNAME(NO)
                     LIMITED_RESOURCE(NO)
                     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)
                   4 USER_DEFINED_1(255)
                     USER_DEFINED_2(USE_ADAPTER_DEFINITION)
                     USER_DEFINED_3(USE_ADAPTER_DEFINITION);

DEFINE_LOGICAL_LINK  LINK_NAME(LINK0002)
                     ADJACENT_NODE_TYPE(NN)
                     PREFERRED_NN_SERVER(YES)
                     DLC_NAME(IBMTRNET)
                     ADAPTER_NUMBER(0)
                     DESTINATION_ADDRESS(X'400001050000')
                     ETHERNET_FORMAT(NO)
                     CP_CP_SESSION_SUPPORT(YES)
                     SOLICIT_SSCP_SESSION(NO)
                     ACTIVATE_AT_STARTUP(YES)
                     USE_PUNAME_AS_CPNAME(NO)
                     LIMITED_RESOURCE(NO)
                     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)
                   3 USER_DEFINED_1(0)
                     USER_DEFINED_2(USE_ADAPTER_DEFINITION)
                     USER_DEFINED_3(USE_ADAPTER_DEFINITION);
```

Figure 159. Link Definitions in AUNG

7.4.2.2 New Mode and APPNCOS Definitions

In the VTAM V4R2 systems and CM/2 in AUNG, we need to define a new class of service definition that will use the user defined parameter 1 to assign line weights to TGs. We also need to define a new mode entry that will point to the new COS definitions. By using this mode for our sessions, we can favor one of the TGs. Detailed explanations of this approach can be found in the Appendix F of *VTAM Network Implementation Guide*. We define the new mode and COS definitions in all three systems so that there is consistency across all three systems.

Figure 160 shows the new APPNCOS definition which is placed in a VTAMLST member. Note that this is a simplistic example where the decision about the line weight is made based solely on the UDP1 and all other parameters are ignored. The table simply gives a line weight of 30 to any line with the UDP1 value of under 100 and a line weight of 240 to any line with the UDP1 value of 100 or over.

We also define a new mode that points to the new COS definition. The new VTAM mode table entry — which is assembled into a copy of the default VTAM mode table, ISTINCLM — is shown in Figure 161. **3** shows how the mode points to the APPNCOS to be used.

```
    3
UPARMCOS APPNCOS  PRIORITY=HIGH             TRANSMISSION PRIORITY
         LINEROW  WEIGHT=30,                line row weight           *
               NUMBER=1,                    line row number           *
               UPARM1=(0,99),               user defined char 1       *
               UPARM2=(0,255),              user defined char 2       *
               UPARM3=(0,255),              user defined char 3       *
               CAPACITY=(MINIMUM,MAXIMUM),  line speed                *
               COSTTIME=(0,255),            cost per connect time     *
               COSTBYTE=(0,255),            cost per byte transmitted *
               PDELAY=(MINIMUM,MAXIMUM),    propagation delay         *
               SECURITY=(UNSECURE,MAXIMUM)  security level for TG
         LINEROW  WEIGHT=240,               line row weight           *
               NUMBER=2,                    line row number           *
               UPARM1=(99,255),             user defined char 1       *
               UPARM2=(0,255),              user defined char 2       *
               UPARM3=(0,255),              user defined char 3       *
               CAPACITY=(MINIMUM,MAXIMUM),  line speed                *
               COSTTIME=(0,255),            cost per connect time     *
               COSTBYTE=(0,255),            cost per byte transmitted *
               PDELAY=(MINIMUM,MAXIMUM),    propagation delay         *
               SECURITY=(UNSECURE,MAXIMUM)  security level for TG
         NODEROW  NUMBER=2,                 node row number           *
               WEIGHT=160,                  node row weight           *
               CONGEST=(LOW,HIGH),          congestion                *
               ROUTERES=(0,255)             route addition resistance
```

Figure 160. APPNCOS Definition in RAS and RA3

```
         TITLE 'CNN2MODE'
CNN2MODE MODEENT LOGMODE=CNN2MODE,FMPROF=X'13',TSPROF=X'07',          *
               ENCR=B'0000',SSNDPAC=7,                                *
               SRCVPAC=7,PSNDPAC=7,APPNCOS=UPARMCOS  3
```

Figure 161. New Mode Table Entry

Figure 162 on page 118 shows the new COS definition in the NDF file for AUNG. This entry is edited into the NDF file. It is not possible to define a COS entry using the CM/2 Setup facility. Explanations of CM/2 NDF commands and examples can be found in the OS/2 file **CMLIB\CMVERIFY.TXT**. Note the similarities between the CM/2 definition and the VTAM definition.

Also note that this definition is not necessary for our scenario to work as it is the network node (RAS) which determines the route. If AUNG is a network node then this COS definition is mandatory. We have included the COS definition to show what is necessary to add a COS definition in CM/2. We also encourage you to keep the same mode and COS tables in all the nodes in the network.

We also edit into the NDF file a new mode which is shown in Figure 163. This definition is necessary for our test scenario as mode name, COS name and TG profiles are passed in the Locate Find request sent to the network node server.

```
DEFINE_COS  COS_NAME(UPARMCOS)
            DESCRIPTION(Sample COS definition)
            TRANSMISSION_PRIORITY(NETWORK)
            NODE_ROW( WEIGHT(160)
                      CONGESTION_RANGE(YES,YES)
                      ROUTE_ADDITION_RES_RANGE(0,255))
            TG_ROW( WEIGHT(30)
                    EFFECTIVE_CAPACITY_RANGE(0,255)
                    COST_PER_CONNECT_TIME_RANGE(0,255)
                    COST_PER_BYTE_RANGE(0,255)
                    SECURITY_RANGE(NONSECURE,GUARDED_RADIATION)
                    PROPAGATION_DELAY_RANGE(MINIMUM,MAXIMUM)
                    USER_DEFINED_1_RANGE(0,99)
                    USER_DEFINED_2_RANGE(0,255)
                    USER_DEFINED_3_RANGE(0,255))
            TG_ROW( WEIGHT(240)
                    EFFECTIVE_CAPACITY_RANGE(0,255)
                    COST_PER_CONNECT_TIME_RANGE(0,255)
                    COST_PER_BYTE_RANGE(0,255)
                    SECURITY_RANGE(NONSECURE,GUARDED_RADIATION)
                    PROPAGATION_DELAY_RANGE(MINIMUM,MAXIMUM)
                    USER_DEFINED_1_RANGE(100,255)
                    USER_DEFINED_2_RANGE(0,255)
                    USER_DEFINED_3_RANGE(0,255));
```

Figure 162. COS Definition in NDF File of AUNG

```
DEFINE_MODE  MODE_NAME(CNN2MODE)
             COS_NAME(UPARMCOS)
             DEFAULT_RU_SIZE(YES)
             RECEIVE_PACING_WINDOW(4)
             MAX_NEGOTIABLE_SESSION_LIMIT(32767)
             PLU_MODE_SESSION_LIMIT(8)
             MIN_CONWINNERS_SOURCE(0)
             COMPRESSION_NEED(PROHIBITED)
             PLU_SLU_COMPRESSION(NONE)
             SLU_PLU_COMPRESSION(NONE);
```

Figure 163. Mode Definition in NDF File of AUNG

7.4.3 Test Results

The aim of the test is to establish LU 6.2 sessions between AUNG and RA3 and to show that:

1. When the default mode and COS are used, either TG 21 or TG 22, between AUNG and RAS, is selected randomly, even though TG 21 provides the better route within the composite network node.

   This test shows that the CNN routing function fails when the entry point and the exit point into the composite network node are in two different subareas.

2. When CNN2MODE is used, TG 21 is always selected.

   This test shows a possible solution to non-optimal subarea routes.

7.4.3.1 Initialization

We start the VTAMs by issuing:

```
S NET28,,,(LIST=S0,RESUSAGE=1)
S NET03,,,(LIST=30,NODETYPE=EN)
```

Note the start option **RESUSAGE** which means that the Topology and Routing Services in RAS rebuilds its routing trees for every new session request. When trees are rebuilt, choices are made randomly among equal-weighted possibilities. This ensures that, in our testing, a route selection will be made randomly whenever we establish a new ATELL session between AUNG and RA3.

The AHHC connections are established by activating the appropriate major nodes in each VTAM.

We also activate the TG profile definition shown in Figure 158 on page 115 and the new APPNCOS definition shown in Figure 160 on page 117.

The NCPs are loaded and the FID4 channel to RA6NCB1 is disabled, forcing the sessions to go from subarea 6 to subarea 28 via subarea 5 within RAS, the composite network node. This is clearly the less optimal route within the composite network node.

The links from AUNG to RAS are activated in the desired order so that TG 22 gets the first PU in the switched major node as its link station and TG 21 gets the second PU. Both TGs are already in the topology of RAS before the activation of the links from AUNG using the CM/2 subsystem management facility. The topology display in RAS after the link activation is shown in Figure 164 on page 120.

```
* RASAN    D NET,TOPO,ID=RAS,LIST=ALL
  RASAN    IST097I  DISPLAY  ACCEPTED
' RASAN
IST350I  DISPLAY TYPE = TOPOLOGY
IST1295I  CP NAME           NODETYPE ROUTERES CONGESTION  CP-CP WEIGHT
IST1296I  USIBMRA.RAS       NN       128      NONE        *NA*  *NA*
IST1297I                    ICN/MDH  CDSERVR  RSN
IST1298I                    YES      NO       2
IST1223I                    BN       NATIVE
IST1224I                    NO       YES
IST1299I  TRANSMISSION GROUPS ORIGINATING AT CP USIBMRA.RAS
IST1357I                                                  CPCP
IST1300I  DESTINATION CP    TGN      STATUS   TGTYPE      VALUE WEIGHT
IST1301I  USIBMRA.RA3       21       INOP     ENDPT       YES   *NA*
IST1301I  USIBMRA.RA3       22       OPER     ENDPT       YES   *NA*
IST1301I  USIBMRA.AUNG      21       OPER     ENDPT       YES   *NA*
IST1301I  USIBMRA.AUNG      22       OPER     ENDPT       YES   *NA*
IST350I  DISPLAY TYPE = TOPOLOGY
IST314I  END
```

Figure 164. Topology Display in RAS

We display the switched PUs, as shown in Figure 165 and Figure 166 on page 121, to find out the TG number associated with each PU. We find in Figure 166 on page 121 that TG 21 is associated with RASUP2, which is the one with the TGP parameter pointing to the profile with UDP1 value of 0. Therefore, when we display the TGs we should find that TG 21 has the UDP1 value of 0 and TG 22 has the default value, which is 128.

```
* RASAN    D NET,ID=RASUP1,E
  RASAN    IST097I  DISPLAY  ACCEPTED
' RASAN
IST075I  NAME = RASUP1           , TYPE = PU_T2.1
IST486I  STATUS= ACTIV--L--, DESIRED STATE= ACTIV
IST1043I  CP NAME = AUNG    , CP NETID = USIBMRA , DYNAMIC LU = YES
IST1105I  RESOURCE STATUS TGN CP-CP TG CHARACTERISTICS
IST1106I  RASUP1   AC/R    22 YES   982D0000000000000000017100808080
IST136I  SWITCHED SNA MAJOR NODE = RASDSWP2
IST081I  LINE NAME = J000600F, LINE GROUP = RA6GTRL2, MAJNOD = RA6NCB1
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST355I  LOGICAL UNITS:
IST080I  AUNG     ACT/S----Y
IST314I  END
```

Figure 165. Display of Switched PU in RAS

```
* RASAN    D NET,ID=RASUP2,E
  RASAN    IST097I  DISPLAY  ACCEPTED
' RASAN
IST075I  NAME = RASUP2           , TYPE = PU_T2.1
IST486I  STATUS= ACTIV     , DESIRED STATE= ACTIV
IST1043I  CP NAME = AUNG    , CP NETID = USIBMRA , DYNAMIC LU = YES
IST1105I  RESOURCE STATUS TGN CP-CP TG CHARACTERISTICS
IST1106I  RASUP2   AC/R    21 YES   982D0000000000000000017100008080
IST136I  SWITCHED SNA MAJOR NODE = RASDSWP2
IST081I  LINE NAME = J000500F, LINE GROUP = RA5GTRL2, MAJNOD = RA5NCA3
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST172I  NO LOGICAL UNITS         EXIST
IST314I  END
```

Figure 166. Display of Switched PU in RAS

We now display in RAS the TGs towards AUNG, using the DISPLAY TOPO command. As shown in **1** in Figure 167 and **2** in Figure 168, the UDP1 values in the displays confirm our expectations: 0 for TG 21 and 128 for TG 22.

```
* RASAN    D NET,TOPO,ORIG=RAS,DEST=AUNG,TGN=21
  RASAN    IST097I  DISPLAY  ACCEPTED
' RASAN
IST350I  DISPLAY TYPE = TOPOLOGY
IST1299I  TRANSMISSION GROUPS ORIGINATING AT CP USIBMRA.RAS
IST1357I                                              CPCP
IST1300I  DESTINATION CP    TGN      STATUS   TGTYPE      VALUE WEIGHT
IST1301I  USIBMRA.AUNG      21       OPER     ENDPT       YES   *NA*
IST1163I                    RSN
IST1164I                    72
IST1302I                    CAPACITY PDELAY   COSTTIME    COSTBYTE
IST1303I                    8K       TERRESTR 0           0
IST1304I                    SECURITY UPARM1   UPARM2      UPARM3
IST1305I                    UNSECURE 0 1      128         128
IST314I  END
```

Figure 167. Display TG 21 to AUNG in RAS

```
* RASAN    D NET,TOPO,ORIG=RAS,DEST=AUNG,TGN=22
  RASAN    IST097I  DISPLAY  ACCEPTED
' RASAN
IST350I  DISPLAY TYPE = TOPOLOGY
IST1299I  TRANSMISSION GROUPS ORIGINATING AT CP USIBMRA.RAS
IST1357I                                              CPCP
IST1300I  DESTINATION CP    TGN      STATUS   TGTYPE      VALUE WEIGHT
IST1301I  USIBMRA.AUNG      22       OPER     ENDPT       YES   *NA*
IST1163I                    RSN
IST1164I                    74
IST1302I                    CAPACITY PDELAY   COSTTIME    COSTBYTE
IST1303I                    8K       TERRESTR 0           0
IST1304I                    SECURITY UPARM1   UPARM2      UPARM3
IST1305I                    UNSECURE 128 2    128         128
IST314I  END
```

Figure 168. Display TG 22 to AUNG in RAS

We now check in AUNG that AUNG's view of the links matches the definitions. We use the subsystem management facility in CM/2 for this purpose. Figure 169 on page 122 shows that, from AUNG's perspective, TG 22 has the UDP1 value of 255 (**3**) as defined in the NDF file. Figure 170 on page 123 shows that, from AUNG's perspective, TG 21 has the UDP1 value of 0 (**4**) as defined in the NDF file.

```
**************************************
*      Active Links Information      *
**************************************
2>Link name                                  LINK0001
  DLC name                                   IBMTRNET
  Adapter number                             0
  Destination DLC address                    X'40000106000004'
  Link activated                             Locally
  Link state                                 Active
  Deactivating link                          No
  Active and activating sessions             2
  Max frame data (BTU) size                  1929
  Adjacent node CP name                      USIBMRA.RAS
  Adjacent node type                         Network node
  CP-CP session support                      Yes
  Connection type                            Peer
  Link station role                          Secondary
  Line type                                  Switched
  Transmission group number                  22
  Effective capacity                         4000000 bits per second
  Cost per connect time                      0
  Cost per byte                              0
  Propagation delay                          384.00 microseconds
  User defined parameter 1                   255  3
  User defined parameter 2                   128
  User defined parameter 3                   128
  Security                                   Nonsecure
  Physical unit (PU) name                    AUNG
```

Figure 169. Display Active Links in AUNG

```
3>Link name                                 LINK0002
  DLC name                                  IBMTRNET
  Adapter number                            0
  Destination DLC address                   X'40000105000004'
  Link activated                            Locally
  Link state                                Active
  Deactivating link                         No
  Active and activating sessions            0
  Max frame data (BTU) size                 1929
  Adjacent node CP name                     USIBMRA.RAS
  Adjacent node type                        Network node
  CP-CP session support                     Yes
  Connection type                           Peer
  Link station role                         Secondary
  Line type                                 Switched
  Transmission group number                 21
  Effective capacity                        4000000 bits per second
  Cost per connect time                     0
  Cost per byte                             0
  Propagation delay                         384.00 microseconds
  User defined parameter 1                  0 4
  User defined parameter 2                  128
  User defined parameter 3                  128
  Security                                  Nonsecure
  Physical unit (PU) name                   AUNG
```

Figure 170. Display Active Links in AUNG

Having verified the test configuration, we now proceed with the test. We invoke ATELL in AUNG to send a message to RA3 as shown below.

```
atell usibmra.ra3aser Hello There!!
```

Because we do not specify the mode, ATELL uses its default mode which is #INTER. Mode #INTER is associated with COS #INTER which does not use the user defined parameter values in determining what weight to assign to a TG. Therefore we expect that a succession of ATELL tests will be randomly distributed amongst TG 21 and TG 22 from AUNG to RAS.

We terminate the session after every ATELL test by using the VARY TERM command. Invoking the ATELL test 10 times confirms our expectations. The sessions use both TG 21 and TG 22 in a random manner.

We now proceed with the second phase of the test. The aim is to use the new mode and the new COS to force the sessions to go through TG 21 which provides the better subarea route within the composite network node. We invoke ATELL in AUNG as shown below, specifying the mode.

```
atell -m cnn2mode usibmra.ra3aser Hello There!!
```

The mode name of CNN2MODE, and its associated COS name, UPARMCOS, will be sent, along with the TG characteristics as seen by AUNG, in the Locate Find request issued by AUNG to RAS. RAS will consult its APPNCOS table to assign the weights to the TGs and to build the routing tree. Because UPARMCOS in RAS will consistently assign a lower weight to TG 21 based on the UDP1 value we expect that the sessions will always take TG 21.

We repeat the ATELL test 10 times, terminating the session after every test. Each test establishes a session through TG 21 only, as expected.

Using the DISPLAY SESSION command, we show the session parameters of one of the sessions as shown in Figure 171 on page 124. Note in **1** the mode being used and in **2** the APPNCOS being used for the session.

```
* RASAN    D NET,SESSION,SID=E5D3656D49DB424C
  RASAN    IST097I  DISPLAY  ACCEPTED
' RASAN
IST350I  DISPLAY TYPE = SESSIONS
IST879I  PLU/OLU REAL = USIBMRA.AUNG      ALIAS = ***NA***
IST879I  SLU/DLU REAL = USIBMRA.RA3ASER   ALIAS = ***NA***
IST880I  SETUP STATUS = ACTIV
IST875I  ADJSSCP TOWARDS PLU = ISTAPNCP
IST875I  ADJSSCP TOWARDS SLU = ISTAPNCP
IST875I  ALSNAME TOWARDS PLU = RASUP2
IST875I  ALSNAME TOWARDS SLU = RASPC8A
IST933I  LOGMODE=CNN2MODE, COS=*BLANK*  1
IST875I  APPNCOS TOWARDS PLU = UPARMCOS 2
IST875I  APPNCOS TOWARDS SLU = UPARMCOS 2
IST314I  END
```

Figure 171. Display Session between AUNG and RA3 in RAS

Using NLDM in NetView in RAS, we display the session path as shown in Figure 172. Also using NLDM in NetView in RAS, we display the APPN path for the session as shown in Figure 173 on page 125.

```
 NLDM.CON                  SESSION CONFIGURATION DATA                PAGE    1
 -------------- PRIMARY ---------------+-------------- SECONDARY --------------
 NAME AUNG      SA 00000005  EL 00D5   |   NAME RA3ASER   SA 0000001C  EL 0079
 --------------------------------------+---------------------------------------
 DOMAIN RASAN       PCID USIBMRA.AUNG.E5D3656D49DB424C        C-C DOMAIN RASAN
               +-------------+                  +-------------+
               |             | ---          --- |   CP/SSCP   | RAS
 RA5NCA3 (0000) | SUBAREA PU  |  APPN TP 03      | SUBAREA PU  | ISTPUS28(0000)
               +------+------+  SUBA TP 00      +------+------+
                      |              VR 00             |
               +------+------+       ER 01      +------+------+
 J000500F      |    LINK     |      RER 00      |     CUA     |
               +------+------+                  +------+------+
                      |       APPNCOS UPARMCOS         |
               +------+------+                  +------+------+
 RASUP2  (001C) |     PU      | LOGMODE CNN2MODE |     PU      | RASPC8A (0062)
               +------+------+ PADJ CP AUNG     +------+------+
                      |       SADJ CP RA3              |
               +------+------+                  +------+------+
 AUNG    (00D5) |     ILU     |                  |     ILU     | RA3ASER (0079)
               +-------------+                  +-------------+

SELECT PT, ST (PRI, SEC TRACE), RT (RESP TIME), P, ER, VR, AR
CMD==>
```

Figure 172. NLDM Display for the ATELL Session

```
 NLDM.AR              APPN SESSION ROUTE CONFIGURATION               PAGE   1
-- PRIMARY ---+-- SECONDARY --+--------------- PCID ---------------+- DOMAIN -
NAME AUNG     | NAME RA3ASER  | USIBMRA.AUNG.E5D3656D49DB424C      | RASAN
--------------+---------------+------------------------------------+----------

 +---------+
 |   CP    |
 |AUNG     |
 +----+----+
TG021 |
 +----+----+
 | CP(ICN) |
 |RAS      |
 +----+----+
TG022 |
 +----+----+
 |   CP    |
 |RA3      |
 +---------+


END OF DATA
SELECT PAR, SAR
CMD==>
```

Figure 173. APPN Route Display for the ATELL Session

7.5 Hints and Tips

The CNN routing may improve the route in the subarea part of the network. It never makes the route worse. There is no need and no parameter to control the function externally.

CNN routing applies to all networks, where data hosts are APPN end nodes and are connected to two or more subareas.

Configurations with remote NCPs may use a non-optimal route.

Configurations with shared NCPs may use a non-optimal route.

Define APPNCOS and link characteristics to enforce an optimal route, if CNN routing is not able to provide it.

In some cases, subarea routing may provide better routes than APPN.

The RSCV in the BIND *response* describes the actual route.

Chapter 8. Dependent LU Requester/Server

In this chapter we relate how dependent LUs are freed from the requirements that they be attached directly to a subarea node, and that all their sessions be routed via that node. The support for this in VTAM is called "dependent LU server" (DLUS), while the corresponding code at the dependent LU end is called "dependent LU requester" (DLUR).

This DLUR/S function adds to the support already provided in VTAM V4R1, which allows a dependent LU to have sessions across an APPN network provided that it is attached to a subarea in the traditional way.

8.1 Benefits

1. Sessions from dependent LUs can now be routed *via any adjacent APPN node*, rather than via the subarea node (NCP or VTAM) through which they happen to have SSCP-LU sessions. This will, in many cases, shorten session paths and improve performance.
2. Dependent LU requesters such as the 3174 can be attached *anywhere in the network*, rather than directly to a subarea node, without affecting the centralized management that VTAM provides. This can eliminate the need for passthrough or duplicate links in locations where both subarea and APPN connectivity is desired.

8.2 How DLUR/S Works

All dependent LUs, and the PUs which 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, cannot cross domain boundaries, let alone network boundaries. 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.

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 the dependent LU (or PU) and its SSCP is now encapsulated within an LU 6.2 "pipe". This "CP-SVR pipe" consists of a pair of sessions between the CPs in the DLUR and DLUS nodes, using the mode name CPSVRMGR. The "pipe" can carry a number of SSCP-PU and SSCP-LU sessions, and need not be between adjacent CPs. The "pipe" can even cross network boundaries.
- LU-LU session routing is now performed by the APPN function, not the subarea 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 server for the dependent LU,

but will show the correct CP name (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). In some cases (where the DLUR supports cross-network CPSVRMGR sessions) the DLUR may respond to a search, but it still indicates itself as the owning CP.

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 174 and Figure 175 on page 129 show the differences between SSCP operation and DLUR/S operation.

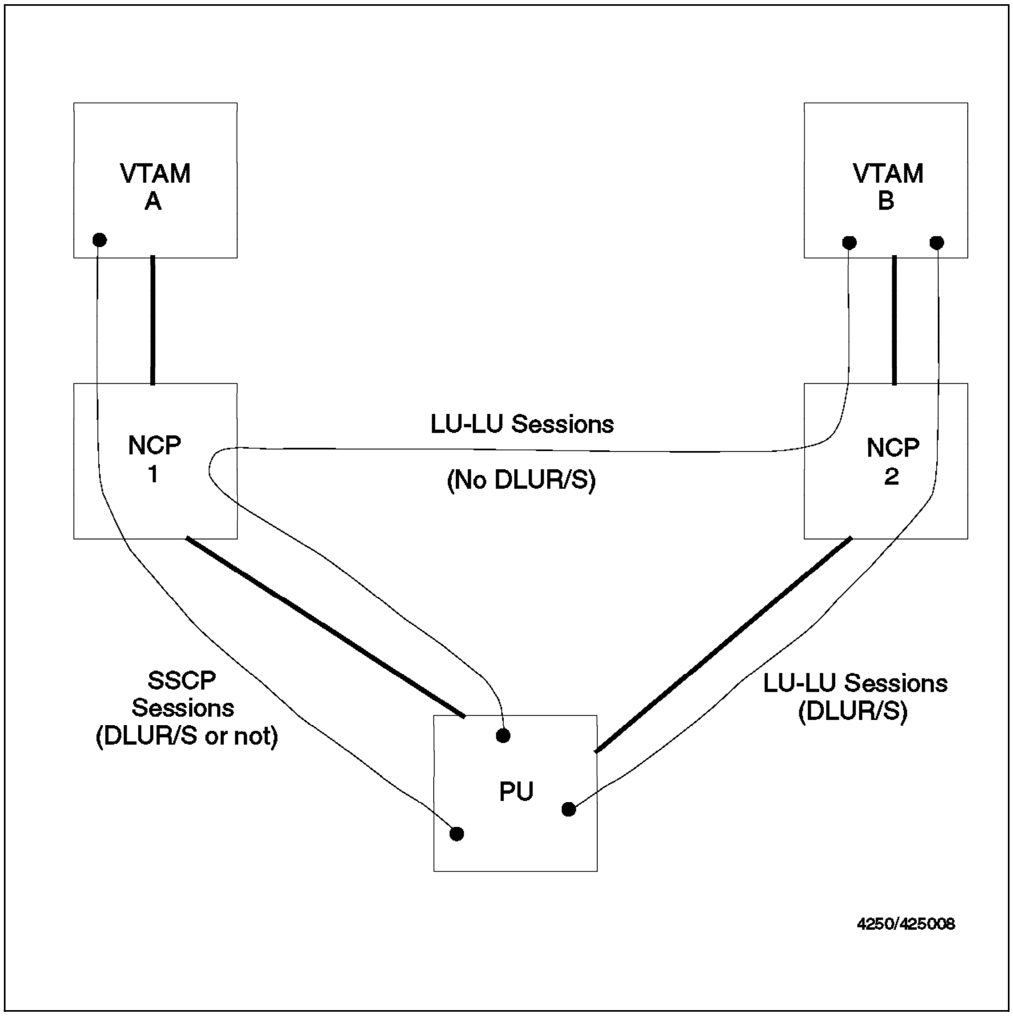

Figure 174. SSCP and DLUR Operation - Routing

Figure 174 shows a PU 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) which 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 CP-SVR 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.

Note that if the PLU is in the subarea network the session path is still via the DLUS′s domain since the DLUS acts as an interchange node. If, therefore, VTAM A and VTAM B were connected only by subarea links the LU-LU session would still take the long route.

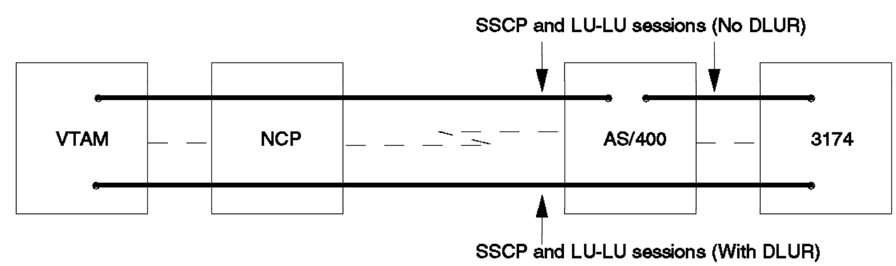

Figure 175. SSCP and DLUR Operation - Resource Utilization

Figure 175 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 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 both VTAM and the 3174 use the APPN formats for transmission headers and BIND requests.

If the session is initiated by the SLU and the PLU is in the APPN network, the PLU must support Session Services Extensions. The AS/400 does not presently have this support, so if the 3174′s LUs want to log on to the AS/400:

- The "Primary LU Support" in OS/400* V2R3 must be used.
- The AS/400 must then be in the subarea network.
- The sessions will take the roundabout route via the NCP (but not via VTAM).

8.2.1 DLUR/S Sessions

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 PUs served by a gateway), it may have multiple sessions with a number of DLU servers. Each DLUS handles one or more PUs complete with their associated LUs.

The content of the RUs carried on the CPSVRMGR sessions is as follows:

- ATTACH Service Transaction Program X′22F0F0F6′.

- GDS Variable X′1500′ containing the complete PIU that would have been sent on the SSCP session. This includes the FID2 transmission header, the RH and the RU. The transmission header is in T2.0 format, and carries the local address used to identify the LU or PU.
- Other information and control vectors appropriate to the request being sent. For instance, the encapsulated REQACTPU and ACTPU requests carry the IDBLK and IDNUM from the PU definition in the switched major node; these are used by both VTAM and the DLUR to identify the correct PU.

CPSVRMGR sessions, like CPSVCMG sessions, can never flow over a pure subarea link - they must be completely APPN (or VR-TG).

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 176 shows the relationship of PUs, LUs and sessions in a DLUR/S environment.

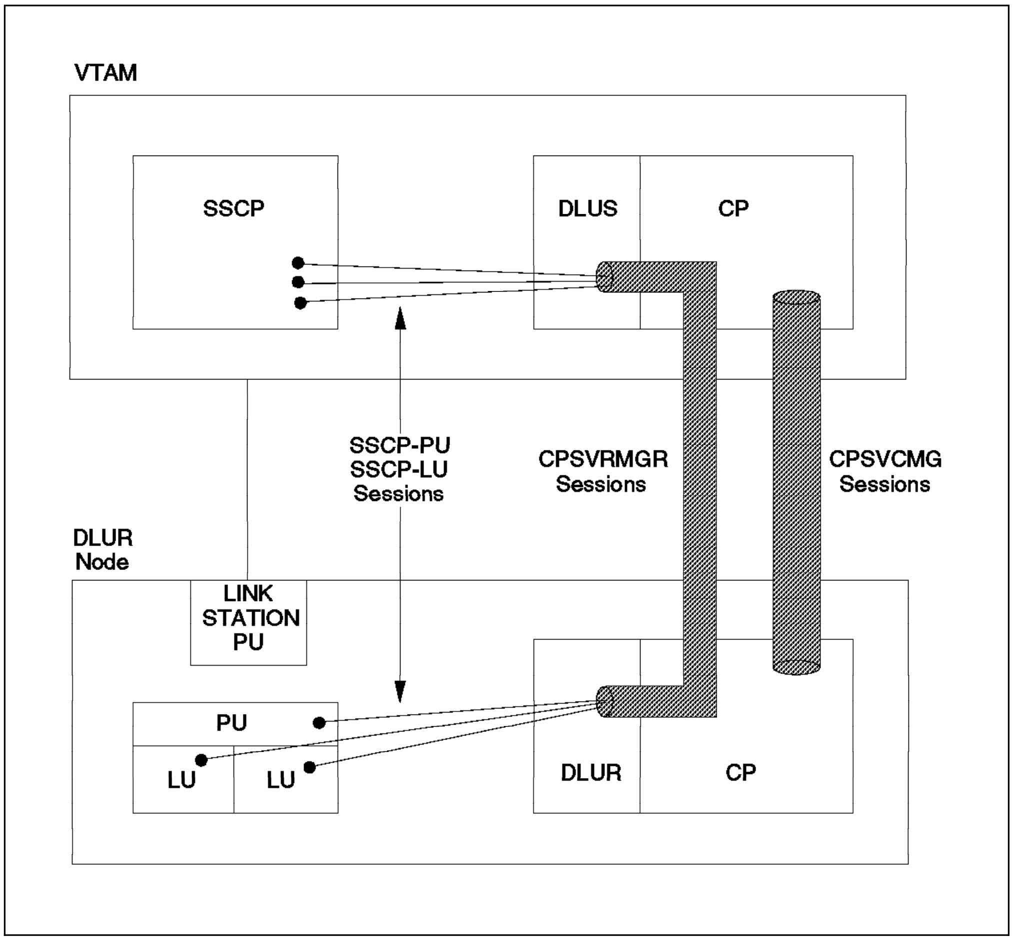

Figure 176. DLUR/S Network Resources and Sessions

The CPSVRMGR 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 CPSVCMG sessions are present only if the node is directly connected, and if the two nodes agree to establish them.

8.3 DLUR and Connection Network

This section describes a basic DLUR/S setup where the PS/2 that contains the DLUR is adjacent to the VTAM V4R2 system that contains the DLUS. We also describe the use of the connection network for sessions involving the dependent LUs.

8.3.1 Description

In the test environment, we use three MVS VTAM V4R2 systems and one PS/2 Communications Manager/2 system. Figure 177 represents the APPN view of the test environment, showing the logical connections through the LAN.

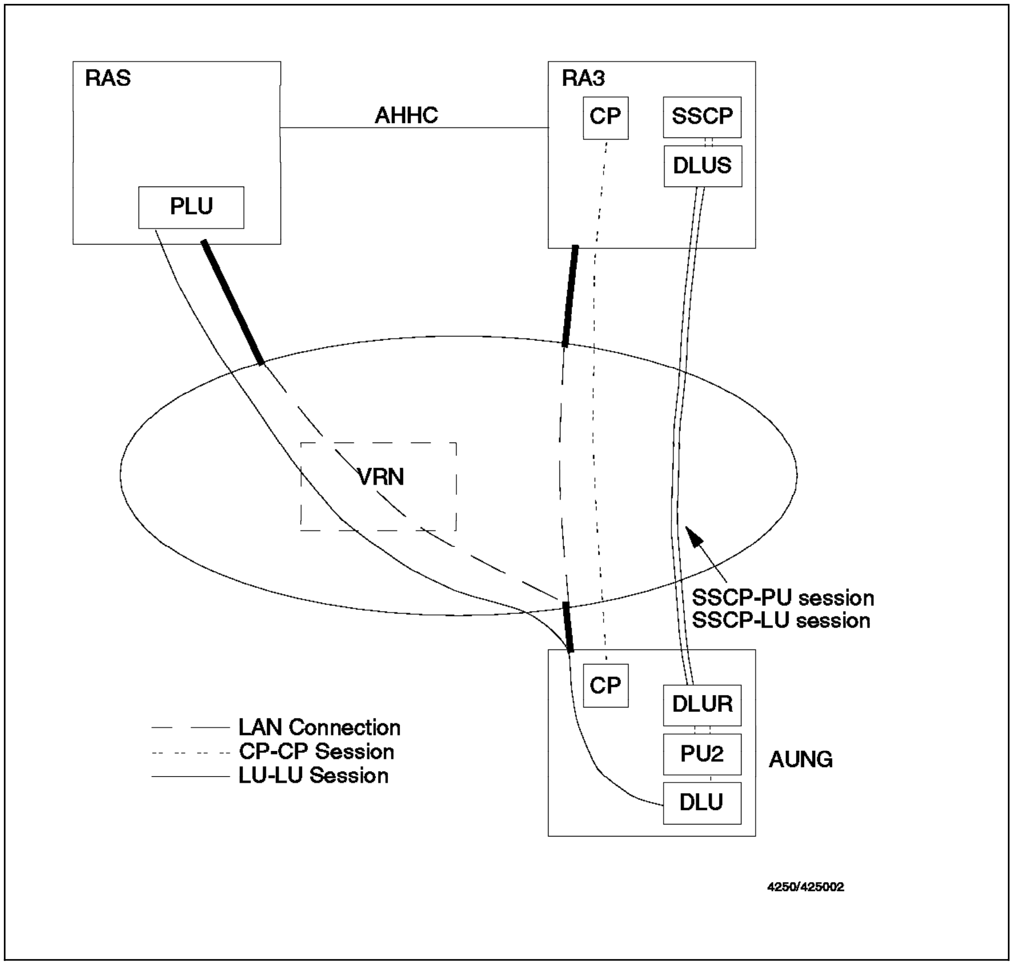

Figure 177. Schematic Diagram of the Test Environment

- RAS, a VTAM V4R2 composite network node, has an MPC connection to RA3. The MPC connections use AHHC as described in Chapter 3, “APPN Host-to-Host Connection” on page 11. It connects to the token-ring LAN via a dedicated NCP. Through the token-ring LAN, RAS is connected to the virtual routing node RACONNET.
- RA3, also a VTAM V4R2 composite network node, has an AHHC connection to RAS. It also connects to the token-ring LAN via a dedicated NCP. It also has a logical connection to the virtual routing node RACONNET.
- AUNG, a PS/2 running Communications Manager/2, is defined as an end node. Its network node server is RA3, which is connected via the token-ring

LAN. AUNG also has a logical connection to the virtual routing node RACONNET.

The testing in this section will verify that:

1. RA3 can "own" a dependent PU and a dependent LU in AUNG by establishing the SSCP-PU and SSCP-LU sessions via the CP-SVR pipe between RA3 and AUNG.

2. LU-LU session setup involving dependent LUs in the DLUR use APPN flows, and therefore the sessions can use the connection network.

8.3.2 Definitions

This section contains the system definitions for the test environment.

VTAM V4R2 requires no definition to perform the DLUS function. We choose to define a switched major node in RA3 for the dependent PU in AUNG. Details concerning the definitions for AHHC and connection networks can be found in Chapter 3, "APPN Host-to-Host Connection" on page 11 and Chapter 4, "Connection Network" on page 21.

Communications Manager/2 requires a link definition to define the "virtual link" between the DLUS and the DLUR. The session between the DLUS and the DLUR in CM/2 is seen by CM/2 as a "virtual link." This virtual link is selected when we define in CM/2 Setup the 3270 terminals attached to the DLUS. We can also see this virtual link when we use the CM/2 Subsystem Management facility to view the active links.

To create the definitions for AUNG, we take an NDF file from an existing configuration, and edit and verify it. We are not able to use the CM/2 Setup facility as no online panels are available, at the time of writing, to create the definitions regarding the link to DLUS. We then use CM/2 Setup to define the 3270 terminals because the definitions for the 3270 emulation do not appear in the NDF file. Figure 178 shows an extract from the NDF file for the configuration used in AUNG during the test. Note in **1** the node ID for AUNG. We will see in Figure 181 on page 133 that we define a different PU for dependent PU functions and that the dependent PU in AUNG has a different node ID.

```
DEFINE_LOCAL_CP  FQ_CP_NAME(USIBMRA.AUNG      )
                 CP_ALIAS(AUNG    )
                 NAU_ADDRESS(INDEPENDENT_LU)
                 NODE_TYPE(EN)
                 NODE_ID(X'05D32223')  1
                 NW_FP_SUPPORT(NONE)
                 HOST_FP_SUPPORT(YES)
                 HOST_FP_LINK_NAME(HOST$1  )
                 MAX_COMP_LEVEL(NONE)
                 MAX_COMP_TOKENS(0);
```

Figure 178. Local CP Definitions in AUNG

Figure 179 on page 133 shows the connection network definition in the NDF file. Note in **2** the virtual routing node name.

```
DEFINE_CONNECTION_NETWORK  FQ_CN_NAME(USIBMRA.RACONNET )  2
                           ADAPTER_INFO( DLC_NAME(IBMTRNET)
                                         ADAPTER_NUMBER(0));
```

Figure 179. Connection Network Definition in AUNG

Figure 180 shows the link definition in the NDF file for the link to RA3. Note in **3** that this is an APPN only link. No SSCP-PU session will be established on this link.

```
DEFINE_LOGICAL_LINK  LINK_NAME(HOST0002)
                     ADJACENT_NODE_TYPE(NN)
                     PREFERRED_NN_SERVER(NO)
                     DLC_NAME(IBMTRNET)
                     ADAPTER_NUMBER(0)
                     DESTINATION_ADDRESS(X'400001060000')
                     ETHERNET_FORMAT(NO)
                     CP_CP_SESSION_SUPPORT(YES)
                     SOLICIT_SSCP_SESSION(NO)  3
                     ACTIVATE_AT_STARTUP(NO)
                     USE_PUNAME_AS_CPNAME(NO)
                     LIMITED_RESOURCE(NO)
                     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 180. Link Definition to RA3

Figure 181 shows the definition in the NDF file for the virtual link (pipe) between the DLUS in RA3 and DLUR in AUNG. The link name as shown in **4** is the link name that we select when we define a 3270 terminal (dependent LU) to RA3 in CM/2 Setup. In **5**, we define the CP name of the VTAM where the DLUS for this link resides. The PU name parameter, as shown in **6**, is not used during communications with VTAM. It is a CM/2 internal parameter. The Node ID, as defined in **7**, is used during the XID exchange with VTAM. This matches the IDBLK **1** and IDNUM **2** parameters in the VTAM switched major node definition as shown in Figure 182 on page 134.

```
DEFINE_DEPENDENT_LU_SERVER  LINK_NAME(RA3LINK )  4
                            FQ_ADJACENT_CP_NAME(USIBMRA.RA3  5   )
                            PU_NAME(RASPU2   )  6
                            NODE_ID(X'05D33331')  7
                            ACTIVATE_AT_STARTUP(NO);
```

Figure 181. DLUS Link Definition in AUNG

```
        VBUILD TYPE=SWNET,MAXNO=20,MAXGRP=5
RA3PP1   PU     ADDR=01,MAXDATA=521,MAXOUT=7,MAXPATH=5,PASSLIM=7,       X
                PUTYPE=2,                                               X
                IDBLK=05D,  1                                           X
                IDNUM=33331  2
*
RA3PPL12 LU     LOCADDR=2,DLOGMOD=M2SDLCQ,SSCPFM=USSSCS,                *
                MODETAB=AMODETAB,USSTAB=US327X
```

Figure 182. Switched Major Node Definition in RA3

8.3.3 Test Results

The aim of the test is to establish the SSCP-PU and SSCP-LU sessions using the pipe between DLUS and DLUR. We also aim to test that the connection network is used for the LU-LU sessions.

We start the VTAMs and ensure that the NCPs are loaded and active. We also establish the connections between RAA and RAS, and between RAS and RA3. The topology display in RA3 after the link activation is shown in Figure 183. The topology display in RAS is shown in Figure 184 on page 135.

```
* RA3AO    D NET,TOPO,ID=RA3,LIST=ALL
  RA3AO    IST097I  DISPLAY  ACCEPTED
' RA3AO
IST350I  DISPLAY TYPE = TOPOLOGY
IST1295I  CP NAME          NODETYPE ROUTERES CONGESTION  CP-CP WEIGHT
IST1296I  USIBMRA.RA3      NN       128      NONE        *NA*  *NA*
IST1297I                   ICN/MDH  CDSERVR  RSN
IST1298I                   YES      NO       2
IST1223I                   BN       NATIVE
IST1224I                   NO       YES
IST1299I  TRANSMISSION GROUPS ORIGINATING AT CP USIBMRA.RA3
IST1357I                                                 CPCP
IST1300I  DESTINATION CP    TGN      STATUS   TGTYPE     VALUE WEIGHT
IST1301I  USIBMRA.RACONNET  21       OPER     INTERM     NO    *NA*
IST1301I  USIBMRA.RAS       21       OPER     INTERM     YES   *NA*
IST1301I  USIBMRA.AUNG      21       OPER     ENDPT      YES   *NA*
IST314I  END
```

Figure 183. Topology Display in RA3

```
* RASAN    D NET,TOPO,ID=RAS,LIST=ALL
  RASAN    IST097I  DISPLAY  ACCEPTED
' RASAN
IST350I  DISPLAY TYPE = TOPOLOGY
IST1295I  CP NAME           NODETYPE ROUTERES CONGESTION  CP-CP WEIGHT
IST1296I  USIBMRA.RAS       NN       128      NONE        *NA*  *NA*
IST1297I                    ICN/MDH  CDSERVR  RSN
IST1298I                    YES      NO       6
IST1223I                    BN       NATIVE
IST1224I                    NO       YES
IST1299I  TRANSMISSION GROUPS ORIGINATING AT CP USIBMRA.RAS
IST1357I                                                  CPCP
IST1300I  DESTINATION CP    TGN      STATUS   TGTYPE      VALUE WEIGHT
IST1301I  USIBMRA.RACONNET  21       OPER     INTERM      NO    *NA*
IST1301I  USIBMRA.RA3       21       OPER     INTERM      YES   *NA*
IST1301I  USIBMRA.RAA       21       OPER     INTERM      YES   *NA*
IST314I  END
```

Figure 184. Topology Display in RAS

We now activate the switched major node in RA3. We then start the link to RA3 in AUNG using the CM/2 subsystem management facilities. We observe that the CP-CP sessions are established. We also observe in CM/2 that two more sessions are established with the mode of CPSVRMGR. These two sessions represent the pipe between DLUS and DLUR. The virtual link (RA3LINK) then becomes active with two active sessions: the SSCP-PU session and the SSCP-LU session. We only have one SSCP-LU session because we choose to define only one LU.

Figure 185 shows the display of DLURs in RA3. AUNG appears as an active DLUR as shown in **3**.

```
* RA3AO    D NET,DLURS
  RA3AO    IST097I  DISPLAY  ACCEPTED
' RA3AO
IST350I  DISPLAY TYPE = DLURS
IST1352I  DLUR NAME          DLUS CONWINNER STATE  DLUS CONLOSER STATE
IST1353I  USIBMRA.AUNG  3    ACTIVE                ACTIVE
IST314I  END
```

Figure 185. Display of DLURs in RA3

Figure 186 on page 136 shows the display of the switched PU in RA3. It shows a PU type 2, which is supported by the DLUR in AUNG **4**.

```
* RA3AO    D NET,ID=RA3PP1,E
  RA3AO    IST097I  DISPLAY  ACCEPTED
' RA3AO
IST075I  NAME = RA3PP1          , TYPE = PU_T2
IST486I  STATUS= ACTIV     , DESIRED STATE= ACTIV
IST1043I  CP NAME = ***NA***, CP NETID = USIBMRA , DYNAMIC LU = YES
IST1354I  DLUR NAME = AUNG 4          MAJNODE = RA3DSWD3
IST136I  SWITCHED SNA MAJOR NODE = RA3DSWD3
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST355I  LOGICAL UNITS:
IST080I  RA3PPL12 ACTIV      RA3PPL13 ACTIV
IST314I  END
```

Figure 186. Display of Switched PU in RA3

Figure 187 shows the display of AUNG in RA3. It shows (in **5**) RA3PP1 as a PU supported by the DLUR in AUNG. **6** and **7** show the DLUS-DLUR sessions between AUNG and RA3. Figure 188 on page 137 and Figure 189 on page 137 show the display of these sessions. Note in **1** and **2** that the CPSVRMGR mode uses different APPNCOS entries in CM/2 and VTAM. The class of service associated with CPSVRMGR depends on the implementation, and may cause the two CPSVRMGR sessions to take different APPN routes.

```
* RA3AO    D NET,ID=AUNG,E
  RA3AO    IST097I  DISPLAY  ACCEPTED
' RA3AO
IST075I  NAME = USIBMRA.AUNG     , TYPE = ADJACENT CP
IST486I  STATUS= ACT/S----Y, DESIRED STATE= ACTIV
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.AUNG      - NETSRVR = ***NA***
IST1044I  ALSLIST = ISTAPNPU
IST082I  DEVTYPE = INDEPENDENT LU / CDRSC
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST171I  ACTIVE SESSIONS = 0000000004, SESSION REQUESTS = 0000000000
IST206I  SESSIONS:
IST1081I  ADJACENT LINK STATION = W32223
IST634I  NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I  RA3    6 ACTIV/DL-S E5D3656D55F5FF9A 0004 0000  0  0 USIBMRA
IST635I  RA3      ACTIV/CP-S E5D3656D54F5FF9A 0003 0001  0  0 USIBMRA
IST635I  RA3    7 ACTIV/DL-P F6FF416462B95C69 0000 0007  0  0 USIBMRA
IST635I  RA3      ACTIV/CP-P F6FF416462B95C68 0001 0002  0  0 USIBMRA
IST1355I  PHYSICAL UNITS SUPPORTED BY DLUR USIBMRA.AUNG
IST089I  RA3PP1   TYPE = PHYSICAL UNIT    , ACTIV    5
IST924I  -------------------------------------------------------------
IST075I  NAME = USIBMRA.AUNG     , TYPE = DIRECTORY ENTRY
IST1186I  DIRECTORY ENTRY = REGISTERED EN
IST1184I  CPNAME = USIBMRA.AUNG      - NETSRVR = USIBMRA.RA3
IST314I  END
```

Figure 187. Display of AUNG in RA3

```
* RA3AO    D NET,SESSION,SID=E5D3656D55F5FF9A
  RA3AO    IST097I  DISPLAY  ACCEPTED
' RA3AO
IST350I  DISPLAY TYPE = SESSIONS
IST879I  PLU/OLU REAL = USIBMRA.AUNG      ALIAS = ***NA***
IST879I  SLU/DLU REAL = USIBMRA.RA3       ALIAS = ***NA***
IST880I  SETUP STATUS = ACTIV
IST875I  ALSNAME TOWARDS PLU = W32223
IST933I  LOGMODE=CPSVRMGR, COS=*BLANK*
IST875I  APPNCOS TOWARDS PLU = SNASVCMG 1
IST314I  END
```

Figure 188. Display of DLUS-DLUR Session in RA3

```
* RA3AO    D NET,SESSION,SID=F6FF416462B95C69
  RA3AO    IST097I  DISPLAY  ACCEPTED
' RA3AO
IST350I  DISPLAY TYPE = SESSIONS
IST879I  PLU/OLU REAL = USIBMRA.RA3       ALIAS = ***NA***
IST879I  SLU/DLU REAL = USIBMRA.AUNG      ALIAS = ***NA***
IST880I  SETUP STATUS = ACTIV
IST875I  ADJSSCP TOWARDS SLU = ISTAPNCP
IST875I  ALSNAME TOWARDS SLU = W32223
IST933I  LOGMODE=CPSVRMGR, COS=*BLANK*
IST875I  APPNCOS TOWARDS SLU = #CONNECT 2
IST314I  END
```

Figure 189. Display of DLUS-DLUR Session in RA3

We now go to the 3270 emulation screen in AUNG and log on to TSO in RAS. Given that the LU-LU session will be established using the APPN session establishment flows, and given that RAS and AUNG belong to the same connection network, we expect that the TSO session will use the connection network. We see in the CM/2 logical link display that a direct connection between RAS and AUNG is set up as the result of the logon as expected.

The display in RAS of the dependent LU in AUNG (RA3PPL12) as shown in Figure 190 on page 138 reveals in **1** that it has been entered in the directory as a dynamic LU in the domain of RA3, the network node in which the DLUS for RA3PPL12 resides.

```
* RASAN    D NET,ID=RA3PPL12,E
  RASAN    IST097I  DISPLAY  ACCEPTED
' RASAN
IST075I  NAME = USIBMRA.RA3PPL12 , TYPE = CDRSC
IST486I  STATUS= ACT/S----Y, DESIRED STATE= ACTIV
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
IST1184I  CPNAME = USIBMRA.AUNG      - NETSRVR = ***NA***
IST1044I  ALSLIST = CNV00004 ISTAPNPU CNV00005
IST082I  DEVTYPE = INDEPENDENT LU / CDRSC
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST171I  ACTIVE SESSIONS = 0000000001, SESSION REQUESTS = 0000000000
IST206I  SESSIONS:
IST1081I  ADJACENT LINK STATION = CNV00005
IST634I  NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I  RASAT01  ACTIV-P    F627D16481F5426D 0001 0005  0  0 USIBMRA
IST924I  -------------------------------------------------------------
IST075I  NAME = USIBMRA.RA3PPL12 , TYPE = DIRECTORY ENTRY
IST1186I  DIRECTORY ENTRY = DYNAMIC    LU
IST1184I  CPNAME = USIBMRA.AUNG      - NETSRVR = USIBMRA.RA3 1
IST314I  END
```

Figure 190. Display of the Dependent LU

Using NLDM in NetView in RA3, we display the session path as shown in Figure 191 on page 139. **2** shows that a direct link has been established to AUNG using the connection network via the NCP directly attached to RAS. This confirms that the boundary function in RA3 is not used because the LU-LU session establishment uses APPN flows. **3** shows the dynamically generated link station at the other end of the link. This link station represents AUNG. **4** shows RA3PPL12 as an independent LU.

Figure 192 on page 139 shows the APPN view of the same session. This display clearly confirms that the connection network is used as shown in **5**.

```
NLDM.CON                  SESSION CONFIGURATION DATA                PAGE    1
--------------- PRIMARY ----------------+--------------- SECONDARY ---------------
NAME RASAT01   SA 0000001C  EL 000A    |   NAME RA3PPL12  SA 00000005  EL 00EA
-----------------------------------------+-----------------------------------------
DOMAIN RASAN       PCID USIBMRA.RAS.F627D16481F5426D              DOMAIN RASAN
               +-------------+                 +-------------+
RAS            |   CP/SSCP   | ---          --- |             |
ISTPUS28(0000) | SUBAREA PU  |  APPN TP 01      | SUBAREA PU  | RA5NCA5 (0000)
               +------+------+  SUBA TP 00      +------+------+
                      |              VR 00             |
               +------+------+       ER 00      +------+------+
RASAT01 (000A) |     LU      |      RER 00      |    LINK 2   | J0005003
               +-------------+                  +------+------+
                               APPNCOS #CONNECT        |
                                                 +------+------+
                               LOGMODE M2SDLCQ  |     PU 3    | CNV00005(000E)
                                                 +------+------+
                               SADJ CP AUNG            |
                                                 +------+------+
                                                 |     ILU 4   | RA3PPL12(00EA)
                                                 +-------------+

SELECT PT, ST (PRI, SEC TRACE), RT (RESP TIME), P, ER, VR, AR
CMD==>
```

Figure 191. NLDM Display for the TSO Logon

```
NLDM.AR              APPN SESSION ROUTE CONFIGURATION              PAGE   1
-- PRIMARY ---+-- SECONDARY --+--------------- PCID ---------------+- DOMAIN -
NAME RASAT01  | NAME RA3PPL12 | USIBMRA.RAS.F627D16481F5426D       | RASAN
--------------+---------------+------------------------------------+----------

 +---------+
 |   CP    |
 |RAS      |
 +----+----+
TG021 |
 +----+----+
 | CP(CN)  | 5
 |RACONNET |
 +----+----+
TG001 |
 +----+----+
 |   CP    |
 |AUNG     |
 +---------+


END OF DATA
SELECT PAR, SAR
CMD==>
```

Figure 192. APPN View of the TSO Logon

8.4 Non-Adjacent DLUS

This section describes a DLUR/S setup where the APPN node that contains the DLUS is not adjacent to the APPN node that contains the DLUR.

8.4.1 Description

In the test environment, we use three MVS VTAM V4R2 systems and one PS/2 Communications Manager/2 system. Figure 193 shows the schematic diagram of the test environment. All the connections as shown in the diagram are APPN (FID2) connections.

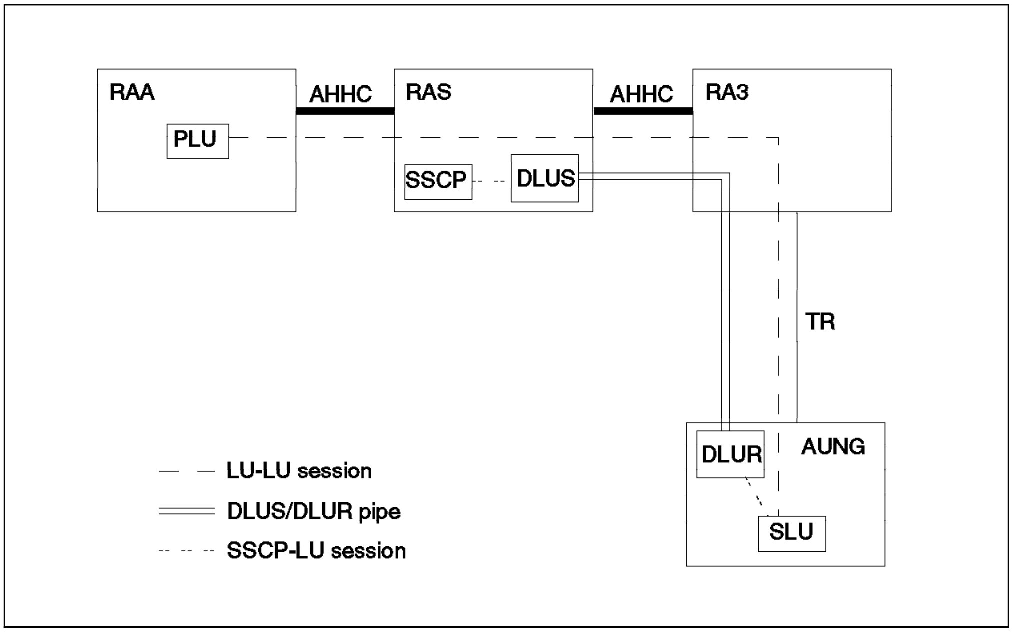

Figure 193. Schematic Diagram of the Test Environment

- RAA, a VTAM V4R2 network node, has an MPC connection to RAS. The MPC connections use AHHC as described in Chapter 3, "APPN Host-to-Host Connection" on page 11.
- RAS, also a VTAM V4R2 network node, has AHHC connections to RAA and RA3.
- RA3 is a VTAM V4R2 composite network node with a dedicated NCP. It connects to the token-ring LAN via the NCP.
- AUNG, a PS/2 running Communications Manager/2, is defined as an end node. Its network node server is RA3, which is connected via the token-ring LAN.

The testing in this section will verify that RAS can "own" a dependent PU and a dependent LU in AUNG by establishing the SSCP-PU and SSCP-LU sessions via the CP-SVR pipe set up via the APPN network.

8.4.2 Definitions

This section contains the system definitions for the test environment. VTAM V4R2 requires no definition to perform the DLUS function. Communications Manager/2 requires a link definition to define the virtual link between the DLUS and the DLUR. The CM/2 definitions here are similar to those used in 8.3, "DLUR and Connection Network" on page 131, and are created similarly.

Figure 194 shows an extract from the NDF file for the configuration used in AUNG during the test. It shows the link definition for the link to RA3. Note in **1** that this is an APPN only link. No SSCP-PU session will be established on this link.

```
DEFINE_LOGICAL_LINK  LINK_NAME(HOST0002)
                     ADJACENT_NODE_TYPE(NN)
                     PREFERRED_NN_SERVER(NO)
                     DLC_NAME(IBMTRNET)
                     ADAPTER_NUMBER(0)
                     DESTINATION_ADDRESS(X'400001060000')
                     ETHERNET_FORMAT(NO)
                     CP_CP_SESSION_SUPPORT(YES)
                     SOLICIT_SSCP_SESSION(NO)  1
                     ACTIVATE_AT_STARTUP(NO)
                     USE_PUNAME_AS_CPNAME(NO)
                     LIMITED_RESOURCE(NO)
                     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 194. Link Definition to RA3

Figure 195 shows the definition in the NDF file for the virtual link (pipe) between the DLUS in RAS and DLUR in AUNG. The link name as shown in **4** is the link name that we select when we define a 3270 terminal (dependent LU) to RA3 in CM/2 Setup. In **5**, we define the CP name of the VTAM where the DLUS for this link resides.

The Node ID, as defined in **7**, is used during the XID exchange with VTAM. This matches the IDBLK **1** and IDNUM **2** parameters in the VTAM switched major node definition in RAS as shown in Figure 196 on page 142.

```
DEFINE_DEPENDENT_LU_SERVER  LINK_NAME(RASLINK )  4
                            FQ_ADJACENT_CP_NAME(USIBMRA.RAS  5   )
                            PU_NAME(RASPU1   )
                            NODE_ID(X'05D33331')  7
                            ACTIVATE_AT_STARTUP(NO);
```

Figure 195. DLUS Link Definition in AUNG

```
        VBUILD TYPE=SWNET,MAXNO=20,MAXGRP=5
RASPP1   PU    ADDR=01,MAXDATA=521,MAXOUT=7,MAXPATH=5,PASSLIM=7,         X
               PUTYPE=2,                                                 X
               IDBLK=05D, 1                                              X
               IDNUM=33331 2
*
RASPPL12 LU    LOCADDR=2,DLOGMOD=M2SDLCQ,SSCPFM=USSSCS,                  *
               MODETAB=AMODETAB,USSTAB=US327X
```

Figure 196. Switched Major Node Definition in RAS

8.4.3 Test Results

The aim of the test is to establish the SSCP-PU and SSCP-LU sessions using the pipe between DLUS and DLUR.

We start the VTAMs and ensure that all the required resources are active. We also establish the connections between RAA and RAS, and between RAS and RA3. The topology is displayed in the three network nodes to ensure that it is as shown in Figure 193 on page 140.

We activate the switched major node in RAS. We then start the link to RA3 in AUNG using the CM/2 subsystem management facilities. We observe that the CP-CP sessions are established on the link. We also observe in CM/2 that two more sessions are established on the link. The session partner is RAS and the mode is CPSVRMGR. These two sessions represent the pipe between DLUS in RAS and DLUR in AUNG. The virtual link (RASLINK) then becomes active with two active sessions: the SSCP-PU session and the SSCP-LU session. We use the DISPLAY DLURS command in RAS to verify that AUNG is known by RAS as a DLUR.

We now display the DLUS-DLUR sessions in RAS using the NLDM functions in NetView. Figure 197 on page 143 and Figure 198 on page 143 show the display of one of these sessions. They have similar information to the DISPLAY SESSION command shown as an example in 8.3, "DLUR and Connection Network" on page 131. Note the logmode, CPSVRMGR, as shown in **1**.

```
NLDM.CON                  SESSION CONFIGURATION DATA                PAGE    1
--------------- PRIMARY ----------------+--------------- SECONDARY ---------------
NAME RAS        SA 0000001C  EL 0007    |  NAME AUNG      SA 0000001C  EL 0073
----------------------------------------+-----------------------------------------
DOMAIN RASAN        PCID USIBMRA.RAS.F627D16481F7FE56         C-C DOMAIN RASAN
                +-------------+                  +-------------+
RAS             |  CP/SSCP    | ---         ---  |  CP/SSCP    | RAS
ISTPUS28(0000)  | SUBAREA PU  |  APPN TP 01      | SUBAREA PU  | ISTPUS28(0000)
                +------+------+                  +------+------+
                       |                                |
                +------+------+                  +------+------+
RAS      (0007) |     LU      |                  |     CUA     |
                +-------------+                  +------+------+
                               APPNCOS #CONNECT         |
                                                 +------+------+
                             1 LOGMODE CPSVRMGR  |     PU      | RASMPC3P(0066)
                                                 +------+------+
                               SADJ CP RA3              |
                                                 +------+------+
                                                 |     ILU     | AUNG    (0073)
                                                 +-------------+

SELECT PT, ST (PRI, SEC TRACE), RT (RESP TIME), P, AR, FC
CMD==>
```

Figure 197. Subarea View in RAS of DLUS-DLUR Session

```
NLDM.AR              APPN SESSION ROUTE CONFIGURATION               PAGE   1
-- PRIMARY ---+-- SECONDARY --+--------------- PCID ---------------+- DOMAIN -
NAME RAS      | NAME AUNG     | USIBMRA.RAS.F627D16481F7FE56       | RASAN
--------------+---------------+------------------------------------+----------

 +---------+
 |   CP    |
 |RAS      |
 +----+----+
TG021 |
 +----+----+
 |   CP    |
 |RA3      |
 +----+----+
TG021 |
 +----+----+
 |   CP    |
 |AUNG     |
 +---------+


END OF DATA
SELECT PAR, SAR
CMD==>
```

Figure 198. APPN View in RAS of DLUS-DLUR Session

Figure 199 on page 144 shows the display of the SSCP-LU session between RAS and the dependent LU (RASPPL12) using NLDM in RAS. RAS, which owns the dependent LU, sees it as any other dependent LU - except that it keeps a note of its DLUR as shown at **3**.

```
NLDM.CON                  SESSION CONFIGURATION DATA                PAGE    1
-------------- PRIMARY ---------------+-------------- SECONDARY --------------
NAME RAS       SA 0000001C  EL 0001   |   NAME RASPPL12  SA 0000001C  EL 0074
--------------------------------------+---------------------------------------
DOMAIN RASAN       PCID USIBMRA.RAS.F627D16481F7FE58             DOMAIN RASAN
               +-------------+                  +-------------+
RAS            |    SSCP     | ---          --- |    SSCP     | RAS
ISTPUS28(0000) | SUBAREA PU  |                  | SUBAREA PU  | AUNG 3
               +-------------+                  +------+------+
                                                       |
                                                +------+------+
                                                |     PU      | RASPP1  (0072)
                                                +------+------+
                                                       |
                              SUBACOS ISTVTCOS +------+------+
                              LOGMODE N/A      |     LU      | RASPPL12(0074)
                                                +-------------+



SELECT PT, ST (PRI, SEC TRACE), RT (RESP TIME), P
CMD==>
```

Figure 199. NLDM Display of the Dependent LU in RAS

We now go to the 3270 emulation screen in AUNG and log on to TSO in RAA. We expect the TSO session to be set up from AUNG to RAA via RA3 and RAS.

The display in RAA of the dependent LU in AUNG (RASPPL12) as shown in Figure 200 on page 145 reveals interesting information. It is shown as a dynamic CDRSC entry. It is in session with TSO in RAA as shown in **4**. **5** shows correctly that the LU resides in the node with CP name of AUNG. **6** is interesting. Even though AUNG is served by RA3, the dependent LU is served by RAS. That is, given that RAS owns the dependent LU, it responds to any search request for the dependent LU. This shows that the DLUS is always the network node server for the associated dependent LUs owned via the DLUR.

```
* RAAAN    D NET,ID=RASPPL12,E
  RAAAN    IST097I  DISPLAY  ACCEPTED
' RAAAN
IST075I  NAME = USIBMRA.RASPPL12 , TYPE = CDRSC
IST486I  STATUS= ACT/S----Y, DESIRED STATE= ACTIV
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.AUNG      - NETSRVR = ***NA***
IST1044I  ALSLIST = ISTAPNPU
IST082I  DEVTYPE = INDEPENDENT LU / CDRSC
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST171I  ACTIVE SESSIONS = 0000000001, SESSION REQUESTS = 0000000000
IST206I  SESSIONS:
IST1081I  ADJACENT LINK STATION = RAAMPCSP
IST634I  NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I  RAAAT01  ACTIV-P  4 F7FF61646FC8E302 0001 0004  0  0 USIBMRA
IST075I  NAME = USIBMRA.RASPPL12 , TYPE = CDRSC
IST924I  -------------------------------------------------------------
IST075I  NAME = USIBMRA.RASPPL12 , TYPE = DIRECTORY ENTRY
IST1186I  DIRECTORY ENTRY = DYNAMIC    LU
IST1184I  CPNAME = USIBMRA.AUNG  5   - NETSRVR = USIBMRA.RAS  6
IST314I  END
```

Figure 200. Display of the Dependent LU in RAA

Using NLDM in NetView in RAS, we display the session path as shown in Figure 201. The APPN view as shown in Figure 202 on page 146 confirms the APPN path taken by the session.

```
NLDM.CON                  SESSION CONFIGURATION DATA                PAGE    1
-------------- PRIMARY ---------------+-------------- SECONDARY --------------
NAME RAAAT01   SA 0000001C  EL 0077   |   NAME RASPPL12  SA 0000001C  EL 0076
--------------------------------------+---------------------------------------
DOMAIN RASAN  C-C  PCID USIBMRA.RAA.F7FF61646FC8E302        C-C DOMAIN RASAN
               +-------------+                 +-------------+
RAS            |   CP/SSCP   | ---         --- |   CP/SSCP   | RAS
ISTPUS28(0000) | SUBAREA PU  |  APPN TP 01     | SUBAREA PU  | ISTPUS28(0000)
               +------+------+                 +------+------+
                      |                               |
               +------+------+                 +------+------+
        (0061) |     CUA     |                 |     CUA     |
               +------+------+                 +------+------+
                      |       APPNCOS #CONNECT        |
               +------+------+                 +------+------+
RASMPCAP(0062) |     PU      | LOGMODE M2SDLCQ |     PU      | RASMPC3P(0066)
               +------+------+ PADJ CP RAA     +------+------+
                      |        SADJ CP RA3            |
               +------+------+                 +------+------+
RAAAT01 (0077) |     ILU     |                 |     ILU     | RASPPL12(0076)
               +-------------+                 +-------------+

SELECT PT, ST (PRI, SEC TRACE), RT (RESP TIME), P, AR, FC
CMD==>
```

Figure 201. NLDM Display for the TSO Logon

```
 NLDM.AR               APPN SESSION ROUTE CONFIGURATION                PAGE   1
 -- PRIMARY ---+-- SECONDARY --+--------------- PCID ---------------+- DOMAIN -
 NAME RAAAT01  | NAME RASPPL12 | USIBMRA.RAA.F7FF61646FC8E302       | RASAN
 --------------+---------------+------------------------------------+----------
                                       TG021 |
  +---------+                           +----+----+
  |   CP    |                           |   CP    |
  |RAA      |                           |AUNG     |
  +----+----+                           +---------+
 TG021 |
  +----+----+
  |   CP    |
  |RAS      |
  +----+----+
 TG021 |
  +----+----+
  |   CP    |
  |RA3      |
  +----+----+
       |

END OF DATA
SELECT PAR, SAR
CMD==>
```

Figure 202. APPN View of the TSO Logon

8.5 DLUR/S and VR-TG

DLUR/S sessions may not traverse a subarea network. However, a VR-TG connection is an APPN connection in this context, so it will support a DLUR/S session. Therefore, a subarea network *can* be used to carry a DLUR/S session as long as the interchange node on the DLUR′s boundary side is VTAM V4R2. The DLUS node, of course, has to be V4R2 in any case.

8.5.1 Description

Figure 203 shows the configuration we use in this test; the VTAM and NCP setup is the same as used in 6.4, "VR-TG Over a Multilink Transmission Group" on page 96.

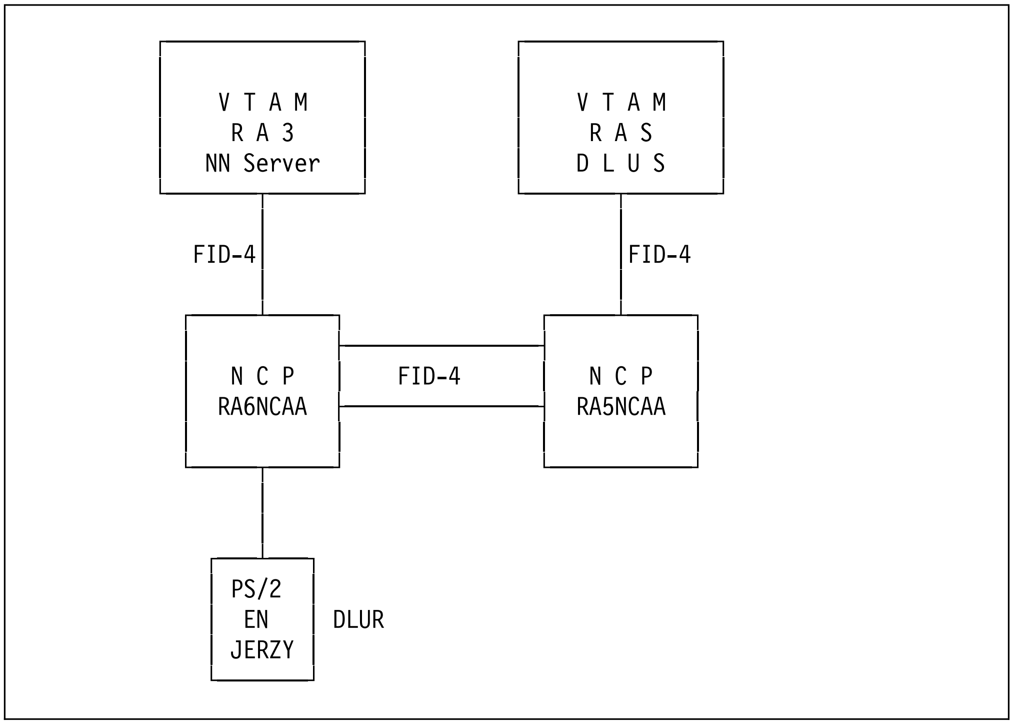

Figure 203. VR-TG with DLUR/S

RA3 and RAS are connected via a VR-TG. The end node JERZY has RA3 as its network node server, but its PU has RAS defined as its DLU server. We shall allow JERZY to initiate the DLUR/S sessions, then log on from its dependent LUs to applications in both RA3 and RAS.

8.5.2 Definitions

Since the DLUR will initiate the CPSVRMGR sessions, the definitions required in VTAM are simple. When the DLUR sends an encapsulated REQACTPU complete with IDBLK and IDNUM, VTAM uses switched logic to identify the PU; in other words, it looks it up in a switched major node. The switched major node we use is shown in Figure 204 on page 148; there are no new parameters in it relating to DLUR/S.

```
      VBUILD TYPE=SWNET,MAXDLUR=1
*      SWITCHED MAJOR NODE FOR PS/2 AS DLU REQUESTER
RASPS2   PU    ADDR=97,MAXPATH=1,IDBLK=05D,IDNUM=06503    1
RASTPS00 LU    LOCADDR=2,DLOGMOD=M2SDLCQ,SSCPFM=USSSCS,              *
               MODETAB=AMODETAB,USSTAB=US327X
RASTPS01 LU    LOCADDR=3,DLOGMOD=M2SDLCQ,SSCPFM=USSSCS,              *
               MODETAB=AMODETAB,USSTAB=US327X
RASTUS02 LU    LOCADDR=4,DLOGMOD=M2SDLCQ,SSCPFM=USSSCS,              *
               MODETAB=AMODETAB,USSTAB=US327X
RASTUS03 LU    LOCADDR=5,DLOGMOD=M2SDLCQ,SSCPFM=USSSCS,              *
               MODETAB=AMODETAB,USSTAB=US327X
```

Figure 204. Switched Major Node for DLUR/S

As with token-ring LAN attached PUs, the station address (97) is ignored; the IDBLK and IDNUM **1** are matched against those sent in by the DLUR. The LUs are defined exactly as they would be if the PU was boundary attached in the "usual" way.

The CM/2 definitions are included here (see Figure 205) as they serve to illustrate how DLUR/S works.

```
DEFINE_LOGICAL_LINK  LINK_NAME(LINK0001)              5
                     ADJACENT_NODE_TYPE(NN)
                     DLC_NAME(IBMTRNET)
                     ADAPTER_NUMBER(0)
                     DESTINATION_ADDRESS(X'400001060000')
                     ETHERNET_FORMAT(NO)
                     CP_CP_SESSION_SUPPORT(YES)
                     SOLICIT_SSCP_SESSION(NO)         6
                     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);

DEFINE_DEPENDENT_LU_SERVER  LINK_NAME(HOST0001)              1
                            FQ_ADJACENT_CP_NAME(USIBMRA.RAS) 2
                            ACTIVATE_AT_STARTUP(YES)
                            PU_NAME(RASP74)                  3
                            NODE_ID(X'05D06503');            4
```

Figure 205. CM/2 Definitions for DLUR/S

The 3270 emulation sessions defined in CM/2 point to a host link, which would normally be a real link; in this case the host link HOST0001 **1** is defined as a CP-SVR pipe, whereas the real link **5** is an APPN link, with no support **6** for SSCP sessions.

The CP-SVR pipe shows the DLUS′s CP name **2**, the PU name **3**, and the node ID **4** consisting of the IDBLK and IDNUM. This node ID can be anything you like, as long as it matches the switched major node definition (see Figure 204) and is not the same as the real node ID that is exchanged with the partner node (RA3, here) over the real link.

8.5.3 Test Results

We initialize the PS/2, wait for the USS10 messages to appear on the emulation sessions, and log on to the applications. LU RASTPS00 logs on to NetView in RA3, while the second LU RASTPS01 logs on to NetView in RAS. Thus one session is cross-domain (but should take a direct route to the application), while the other is same-domain (but must take a longer route to its application). This is the reverse of what normally happens!

First, we inspect the topologies which show that JERZY is known to RA3 as a served end node **1** (Figure 206) but is not known *as an adjacent APPN node* to its DLU server RAS (Figure 207 on page 150). RA3 and RAS are connected via a VR-TG **2**.

```
NCCF                       N E T V I E W    RA3AO BUCZAK   02/15/94 20:19:26
* RA3AO    D NET,TOPO,ID=RA3,LIST=ALL
  RA3AO    IST097I  DISPLAY  ACCEPTED
′ RA3AO
IST350I  DISPLAY TYPE = TOPOLOGY
IST1295I  CP NAME           NODETYPE ROUTERES CONGESTION  CP-CP WEIGHT
IST1296I  USIBMRA.RA3       NN       128      NONE        *NA*  *NA*
IST1297I                    ICN/MDH  CDSERVR  RSN
IST1298I                    YES      NO       4
IST1223I                    BN       NATIVE
IST1224I                    NO       YES
IST1299I  TRANSMISSION GROUPS ORIGINATING AT CP USIBMRA.RA3
IST1357I                                                  CPCP
IST1300I  DESTINATION CP    TGN      STATUS   TGTYPE      VALUE WEIGHT
IST1301I  USIBMRA.RACONNET  21       INOP     INTERM      NO    *NA*
IST1301I  USIBMRA.RAS       21       INOP     INTERM      YES   *NA*
IST1301I  USIBMRA.JERZY  1  21       OPER     ENDPT       YES   *NA*
IST1301I  USIBMRA.AUNG      21       INOP     ENDPT       YES   *NA*
IST1301I  USIBMRA.RAS    2  255      OPER     INTERM VRTG YES   *NA*
IST314I  END
```

Figure 206. Topology Display in NN Server

```
NCCF                      N E T V I E W    RASAN BUCZAK   02/15/94 20:19:53
* RASAN    D NET,TOPO,ID=RAS,LIST=ALL
  RASAN    IST097I  DISPLAY  ACCEPTED
' RASAN
IST350I  DISPLAY TYPE = TOPOLOGY
IST1295I  CP NAME           NODETYPE ROUTERES CONGESTION  CP-CP WEIGHT
IST1296I  USIBMRA.RAS       NN       128      NONE        *NA*  *NA*
IST1297I                    ICN/MDH  CDSERVR  RSN
IST1298I                    YES      NO       8
IST1223I                    BN       NATIVE
IST1224I                    NO       YES
IST1299I  TRANSMISSION GROUPS ORIGINATING AT CP USIBMRA.RAS
IST1357I                                                  CPCP
IST1300I  DESTINATION CP    TGN      STATUS   TGTYPE      VALUE WEIGHT
IST1301I  USIBMRA.RACONNET  21       INOP     INTERM      NO    *NA*
IST1301I  USIBMRA.RA3       21       INOP     INTERM      YES   *NA*
IST1301I  USIBMRA.AUNG      21       INOP     ENDPT       YES   *NA*
IST1301I  USIBMRA.RA3174    21       OPER     INTERM      YES   *NA*
IST1301I  USIBMRA.RA3    2  255      OPER     INTERM VRTG YES   *NA*
IST1301I  USIBMRAT.BERND    21       OPER     ENDPT       YES   *NA*
IST314I  END
```

Figure 207. Topology Display in DLU Server

Next, we display the resource JERZY in each of the two VTAMs. Figure 208 shows how RA3 (its network node server) sees it.

```
 NCCF                      N E T V I E W    RA3AO BUCZAK   02/15/94 20:21:33

C RA3AO    DISPLAY NET,ID=JERZY,SCOPE=ALL
  RA3AO    IST097I  DISPLAY  ACCEPTED
' RA3AO
IST075I  NAME = USIBMRA.JERZY    , TYPE = ADJACENT CP   1
IST486I  STATUS= ACT/S----Y, DESIRED STATE= ACTIV
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.JERZY     - NETSRVR = ***NA***
IST1044I  ALSLIST = ISTAPNPU
IST082I  DEVTYPE = INDEPENDENT LU / CDRSC
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST171I  ACTIVE SESSIONS = 0000000004, SESSION REQUESTS = 0000000000
IST206I  SESSIONS:
IST1081I  ADJACENT LINK STATION = W32258    3
IST634I  NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I  RAS  4   ACTIV/DL-S D54B4E4B354AE8BE           5  0 USIBMRA
IST635I  RA3      ACTIV/CP-S D54B4E4B344AE8BE 0003 0001  0  0 USIBMRA
IST635I  RAS  4   ACTIV/DL-P F627D16481A33D48           5  0 USIBMRA
IST635I  RA3      ACTIV/CP-P F6FF41646266A9B2 0001 0003  0  0 USIBMRA
IST924I  -------------------------------------------------------------
IST075I  NAME = USIBMRA.JERZY    , TYPE = DIRECTORY ENTRY
IST1186I  DIRECTORY ENTRY = DYNAMIC    EN    2
IST1184I  CPNAME = USIBMRA.JERZY     - NETSRVR = USIBMRA.RA3    5
IST314I  END
```

Figure 208. Display DLUR CP from NN Server

The display is as expected: the resource is an adjacent CP **1**, and is stored in the directory as a dynamically defined end node **2** with a server of RA3 **5**.

In addition to the expected CP-CP sessions. we see the two DLUR/S sessions to RAS **4**, denoted by the status "ACTIV/DL". They use VR 5, which (as we shall see later) uses the VR-TG and does not pass through RA3.

Figure 209 shows RAS′s view of JERZY.

```
 NCCF                     N E T V I E W    RASAN BUCZAK   02/15/94 20:22:10

C RASAN    DISPLAY NET,ID=JERZY,SCOPE=ALL
  RASAN    IST097I  DISPLAY  ACCEPTED
' RASAN
IST075I  NAME = USIBMRA.JERZY    , TYPE = CDRSC     1
IST486I  STATUS= ACT/S----Y, DESIRED STATE= ACTIV
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 = RAS     , VERIFY OWNER = NO
IST1184I  CPNAME = USIBMRA.JERZY     - NETSRVR = ***NA***
IST082I  DEVTYPE = INDEPENDENT LU / CDRSC
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST171I  ACTIVE SESSIONS = 0000000002, SESSION REQUESTS = 0000000000
IST206I  SESSIONS:
IST634I  NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I  RAS  2  ACTIV/DL-S D54B4E4B354AE8BE 002E 0000  5  0 USIBMRA
IST635I  RAS  2  ACTIV/DL-P F627D16481A33D48 0000 0033  5  0 USIBMRA
IST1355I  PHYSICAL UNITS SUPPORTED BY DLUR USIBMRA.JERZY    3
IST089I  RASPS2   TYPE = PHYSICAL UNIT    , ACTIV
IST924I  -------------------------------------------------------------
IST075I  NAME = USIBMRA.JERZY    , TYPE = DIRECTORY ENTRY
IST1186I  DIRECTORY ENTRY = DYNAMIC    EN   4
IST1184I  CPNAME = USIBMRA.JERZY     - NETSRVR = USIBMRA.RA3   5
IST314I  END
```

Figure 209. Display DLUR CP from DLU Server

Here we see that JERZY is a CDRSC **1** (since it is neither an LU owned by RAS, nor an adjacent CP), and its directory entry shows it to be a dynamically defined end node **4** (since it only became known to RAS when it issued a search for RAS). Its network node server is shown correctly at **5**.

The DLUR/S sessions are again shown at **2** (this time there are no CP-CP sessions), but in addition the message at **3** indicates that JERZY is a DLU requester on behalf of the PU RASPS2.

There are in fact two PUs associated with the node JERZY:

1. The link station used by RA3 to connect to the node, which is known only to RA3. Its name is W32258 and it can be seen displayed at **3** in Figure 208 on page 150. It has (in this case) been created dynamically, but could have been defined in a switched major node by a "PU CPNAME=JERZY" statement.
2. The internal PU used by RAS to handle the dependent LUs, which is known only to RAS. Its name is RASPS2 and it can be seen displayed in Figure 210

on page 152. It has (in this case) been defined in the switched major node shown in Figure 204 on page 148, but could have been created dynamically.

```
 NCCF                   N E T V I E W    RASAN BUCZAK   02/15/94 20:22:38

C RASAN    DISPLAY NET,ID=RASPS2,SCOPE=ALL
  RASAN    IST097I  DISPLAY  ACCEPTED
' RASAN
IST075I  NAME = RASPS2           , TYPE = PU_T2
IST486I  STATUS= ACTIV     , DESIRED STATE= ACTIV
IST1043I  CP NAME = ***NA***, CP NETID = USIBMRA , DYNAMIC LU = YES
IST1354I  DLUR NAME = JERZY      1      MAJNODE = RASCAL
IST136I  SWITCHED SNA MAJOR NODE = RASCAL
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST355I  LOGICAL UNITS:
IST080I  RASTPS00 ACT/S      RASTPS01 ACT/S      RASTUS02 ACTIV
IST080I  RASTUS03 ACTIV
IST314I  END
```

Figure 210. Display of Served PU from DLUS

Here RASPS2 is displayed exactly as a T2.0 PU, except that the DLUR name is shown at **1**. A complete list of served DLURs can be displayed using the D NET,DLUR command as seen in Figure 211.

```
 NCCF                   N E T V I E W    RASAN BUCZAK   02/15/94 20:28:59

* RASAN    D NET,DLURS
  RASAN    IST097I  DISPLAY  ACCEPTED
' RASAN
IST350I  DISPLAY TYPE = DLURS
IST1352I  DLUR NAME          DLUS CONWINNER STATE  DLUS CONLOSER STATE
IST1353I  USIBMRA.JERZY      ACTIVE                ACTIVE
IST314I  END
```

Figure 211. Display of Served DLURs

Now we display RASTPS00 from RA3. RASTPS00 is a dependent LU owned by RAS, which has logged on to NetView in RA3; we therefore expect RA3 to know about the LU. Figure 212 on page 153 shows that RA3 indeed knows about it.

```
 NCCF                      N E T V I E W    RA3A0 BUCZAK   02/15/94 20:20:34

C RA3A0    DISPLAY NET,ID=RASTPS00,SCOPE=ALL
  RA3A0    IST097I  DISPLAY  ACCEPTED
' RA3A0
IST075I  NAME = USIBMRA.RASTPS00 , TYPE = CDRSC    1
IST486I  STATUS= ACT/S----Y, DESIRED STATE= ACTIV
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 = RA3     , VERIFY OWNER = NO
IST1184I  CPNAME = USIBMRA.JERZY     - NETSRVR = ***NA***
IST1044I  ALSLIST = ISTAPNPU
IST082I  DEVTYPE = INDEPENDENT LU / CDRSC
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST171I  ACTIVE SESSIONS = 0000000001, SESSION REQUESTS = 0000000000
IST206I  SESSIONS:
IST1081I  ADJACENT LINK STATION = W32258  3
IST634I  NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I  RA3A0004 ACTIV-P 2 F6FF41646266A9B8 002E 0078  0  0 USIBMRA
IST924I  -------------------------------------------------------------
IST075I  NAME = USIBMRA.RASTPS00 , TYPE = DIRECTORY ENTRY
IST1186I  DIRECTORY ENTRY = DYNAMIC    LU
IST1184I  CPNAME = USIBMRA.JERZY 5 - NETSRVR = USIBMRA.RAS  4
IST314I  END
```

Figure 212. Dependent LU Displayed from Session Partner

We can see that RASTPS00 is, indeed, viewed as a cross-domain resource **1**, and has a session with RA3's NetView **2**. Its adjacent link station is shown correctly as W32258 **3** as is its CP name JERZY **5**. However, the directory entry **4** shows that RAS is known as the network node server. Here we have an LU served by RAS on an end node served by RA3! This information has been supplied by RAS, which replies to a search request on behalf of its served dependent LU. RA3, knowing that RASTPS00 resides on CP JERZY, has been able to calculate the direct route.

A display of RASTPS01 from RA3 is also interesting. RASTPS01 has logged on to NetView in the same domain RAS, so in theory RA3 should not be concerned. However, Figure 213 on page 154 shows otherwise.

```
 NCCF                      N E T V I E W    RA3AO BUCZAK   02/15/94 20:21:06

C RA3AO    DISPLAY NET,ID=RASTPS01,SCOPE=ALL
  RA3AO    IST097I  DISPLAY  ACCEPTED
' RA3AO
IST075I  NAME = USIBMRA.RASTPS01 , TYPE = CDRSC
IST486I  STATUS= ACT/S----Y, DESIRED STATE= ACTIV
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 = ***NA***, VERIFY OWNER = NO
IST1184I  CPNAME = USIBMRA.JERZY     - NETSRVR = ***NA***
IST1044I  ALSLIST = ISTAPNPU
IST082I  DEVTYPE = INDEPENDENT LU / CDRSC
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST171I  ACTIVE SESSIONS = 0000000001, SESSION REQUESTS = 0000000000
IST206I  SESSIONS:
IST1081I  ADJACENT LINK STATION = W32258
IST634I  NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I  RASAN009 ACTIV-P    F627D16481A33D75            5  0 USIBMRA
IST314I  END
```

Figure 213. Display of Dependent LU from Intermediate Node

RA3 is aware of this LU and its session only because the session passes through RA3's domain, and RA3 has been involved in its setup. RA3 has no directory entry for RASTPS01 because no search took place for it.

Displays of the same two LUs from their owner RAS show them simply as dependent LUs on the switched PU RASPS2. This is no different from how they would appear with "real" (not DLUR/S) SSCP sessions.

Displays of all the sessions via NetView confirm that all the session routes take the expected paths. Figure 214 on page 155 shows the configuration of one of the DLUR/S sessions.

```
 NLDM.CON                  SESSION CONFIGURATION DATA                 PAGE    1

-------------- PRIMARY ---------------+-------------- SECONDARY --------------
NAME RAS       SA 0000001C  EL 0007   |   NAME JERZY 2  SA 00000006  EL 00D4
--------------------------------------+---------------------------------------
DOMAIN RASAN       PCID USIBMRA.RAS.F627D16481A33D48             DOMAIN RASAN
               +-------------+                 +-------------+
RAS            |   CP/SSCP   | ---         --- |   CP/SSCP   | JERZY
ISTPUS28(0000) | SUBAREA PU  |  APPN TP 01     | LOCAL DATA  |
               +------+------+  SUBA TP 00     +------+------+
                      |              VR 05            |
               +------+------+       ER 06     +------+------+
RAS     (0007) |     LU      |      RER 07     |     LU      | JERZY   (00D4)
               +-------------+                 +-------------+
                              APPNCOS #CONNECT

                              LOGMODE CPSVRMGR  1




SELECT PT, ST (PRI, SEC TRACE), RT (RESP TIME), P, ER, VR, AR
CMD==>
```

Figure 214. DLUR/S Session Configuration

We can see that the mode CPSVRMGR is indeed used **1**, and that VR 5 (ER 6) is taken across the subarea network. Rather than work out what ER 6 is, we note that JERZY is seen as belonging to subarea 6 **2**, so the session flows from RAS across the VR-TG to RA3's NCP RA6NCAA. Figure 215 confirms that the VR-TG is used.

```
NLDM.AR               APPN SESSION ROUTE CONFIGURATION               PAGE   1
-- PRIMARY ---+-- SECONDARY --+--------------- PCID ---------------+- DOMAIN -
NAME RAS      | NAME JERZY    | USIBMRA.RAS.F627D16481A33D48       | RASAN
--------------+---------------+------------------------------------+----------

 +---------+
 |   CP    |
 |RAS      |
 +----+----+
TG255 |
 +----+----+
 | CP(ICN) |
 |RA3      |
 +----+----+
TG021 |
 +----+----+
 |   CP    |
 |JERZY    |
 +---------+


END OF DATA
SELECT PAR, SAR
CMD==>
```

Figure 215. APPN View of DLUR/S Session

Here we have SSCP-LU sessions passing through the domain of a “foreign” VTAM!

If we display the same-domain session between RASTPS01 and RAS′s NetView, exactly the same picture is seen. The session path is from RAS via ER 6 (TG 255 in the APPN view) to NCP RA6NCAA, then across APPN TG 21 to JERZY.

The cross-domain session between RASTPS00 and RA3′s NetView can be seen in Figure 216.

```
 NLDM.CON                    SESSION CONFIGURATION DATA                 PAGE    1
 -------------- PRIMARY ---------------+-------------- SECONDARY --------------
 NAME RA3A0004  SA 00000003  EL 0010   |   NAME RASTPS00  SA 00000006  EL 00D8
 --------------------------------------+---------------------------------------
 DOMAIN RA3A0       PCID USIBMRA.RA3.F6FF41646266A9B8              DOMAIN RA3A0
               +-------------+                  +-------------+
 RA3           |   CP/SSCP   | ---         ---  |             |
 ISTPUS03(0000) | SUBAREA PU  |                  | SUBAREA PU  | RA6NCAA (0000)
               +------+------+  SUBA TP 00      +------+------+
                      |              VR 00             |
               +------+------+       ER 01      +------+------+
 RA3A0004(0010) |     LU      |      RER 00      |    LINK     | J0006013
               +-------------+                  +------+------+
                                                       |
                                                +------+------+
                               LOGMODE M2SDLCQ  |     PU      | W32258  (0020)
                                                +------+------+
                               SADJ CP JERZY           |
                                                +------+------+
                                                |     ILU     | RASTPS00(00D8)
                                                +-------------+

SELECT PT, ST (PRI, SEC TRACE), RT (RESP TIME), P, ER, VR
CMD==>
```

Figure 216. Cross-Domain Session Configuration

This shows that the session path is from RA3 directly via the NCP RA6NCAA (the reader can verify from the path tables that ER 1 is indeed the direct route). The cross-domain session path is shorter than the same domain session path.

8.6 CP-SVR Pipe Initiated by VTAM

Activation of dependent PUs and LUs is usually driven automatically when the PU is powered on or connected to the network. Therefore, most of the examples in this book describe the CP-SVR pipe being initiated by the DLUR node. However, in some cases it is desirable that activation is performed from the DLUS side. In this test we describe how that is done.

8.6.1 Description

We use the same configuration of PS/2, VTAM and NCP as shown in Figure 203 on page 147, except that the links between the VTAMs are FID2 rather than VR-TG. The only difference this makes is that we see TG21 or TG22 rather than TG255 in the APPN session routes.

8.6.2 Definitions

Instead of the switched major nodes used in 8.5.2, "Definitions" on page 147, we need a dial-out definition as shown in Figure 217.

```
        VBUILD TYPE=SWNET,MAXDLUR=1
*       SWITCHED MAJOR NODE FOR PS/2 AS DLU REQUESTER
RASPS2   PU    ADDR=97,MAXPATH=1,DWACT=YES,CPNAME=RUBBISH  1
         PATH  DLURNAME=JERZY,PID=1,DLCADDR=(1,C,INTPU),   2         *
               DLCADDR=(2,X,05D06503)
RASTPS00 LU    LOCADDR=2,DLOGMOD=M2SDLCQ,SSCPFM=USSSCS,              *
               MODETAB=AMODETAB,USSTAB=US327X
RASTPS01 LU    LOCADDR=3,DLOGMOD=M2SDLCQ,SSCPFM=USSSCS,              *
               MODETAB=AMODETAB,USSTAB=US327X
RASTPS02 LU    LOCADDR=4,DLOGMOD=M2SDLCQ,SSCPFM=USSSCS,              *
               MODETAB=AMODETAB,USSTAB=US327X
RASTPS03 LU    LOCADDR=5,DLOGMOD=M2SDLCQ,SSCPFM=USSSCS,              *
               MODETAB=AMODETAB,USSTAB=US327X
```

Figure 217. Switched Major Node - DLUS Initiated Connection

Once again, we see **1** that the station address and the CP name *on the PU* are ignored. We code DWACT=YES on the PU statement to ensure that VTAM initiates the connection as soon as the major node is activated. To perform the "dial-out", we need a PATH statement **2** with a few differences from the traditional coding:

- The CP name of the DLUR is specified in the DLURNAME parameter.
- The dial-out path is given in the DLCADDR parameters rather than the DIALNO parameter, and it contains the IDBLK and IDNUM that will be used to identify the PU we wish to activate. This is an example of the new "expanded dial number" function introduced in VTAM V4R2.

If both dial-in and dial-out operation is required for a DLUR PU, then *both* DLCADDR (on the PATH statement) *and* IDBLK / IDNUM (on the PU statement) must be coded. VTAM will not use the DLCADDR on the PATH statement to match an incoming node ID.

8.6.3 Test Results

Since CM/2 will always attempt to set up the CP-SVR pipe when it is started (regardless of the ACTIVATE_AT_STARTUP parameter), we allow it to do so and then terminate the sessions. While CM/2 is sitting idle, we activate the switched major node RASCASSE shown in Figure 217 on page 157. There are no console messages indicating that the dial-out process has succeeded, but the emulator sessions on the PS/2 suddenly display the USS10 messages.

We see the same displays on RAS and RA3 as we saw in 8.5, "DLUR/S and VR-TG" on page 147. The only significant difference is in the NetView displays, which show that RAS, not JERZY, establishes the first CPSVRMGR session.

The contents of the DLCADDR parameters have been sent to the DLUR as subvectors 91 (INTPU) and 92 (node ID) of the TG Descriptor control vector (46), in the encapsulated ACTPU request. This information is identical to that sent by the DLUR in the encapsulated REQACTPU request when the DLUR initiates the pipe.

We have found that, **for the CM/2 implementation**, the VTAM PU name (RASPS2) is not checked against the CM/2 PU name, nor is the DLCADDR parameter INTPU checked by CM/2. This may not, of course, be true in other product implementations such as 3174.

8.7 DLUR/S with Multiple PUs and Multiple Servers

The DLUR/S architecture allows a DLUR to act on behalf of more than one PU, and these PUs need not all have the same DLU server. The circumstances in which this might be useful include:

- More than 255 LUs within a single node. For example, a 3174 can accommodate 64 coax ports plus 24 ASCII ports, each of which may have several sessions. The T2.0 architecture does not permit more than 255 LUs on a single PU.
- Downstream PUs. For example, a 3174 can act as a gateway for 250 PUs attached via token-ring LAN, Ethernet or ISDN. Note, however, that the current releases of 3174 microcode (C5) and CM/2 (the interim DLUR code supplied with VTAM) do not support DLUR for downstream PUs.
- Resilience. For example, a 3174 with two PUs connected to two hosts will still allow half its LUs access to SSCP services if one host fails.

8.7.1 Description

Figure 218 shows the configuration we use to demonstrate two DLU servers serving two PUs in the same CM/2 node.

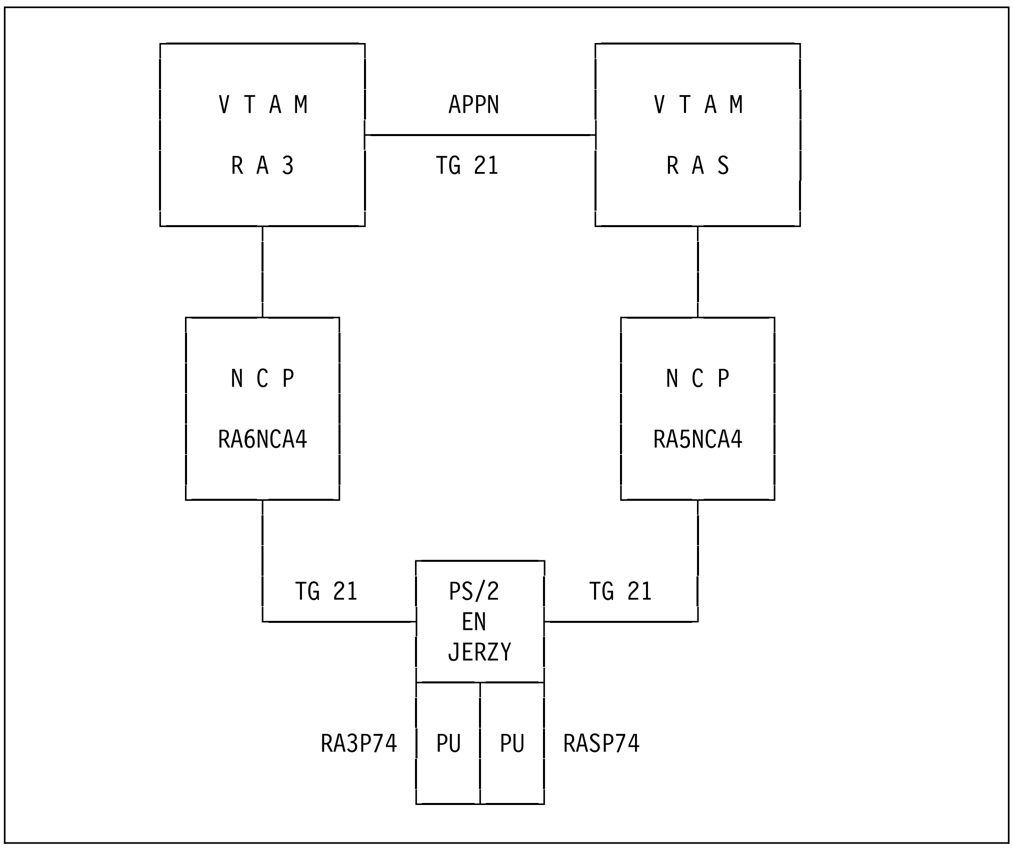

Figure 218. DLUR/S with Two Servers

The two VTAMs are both network nodes, RA3 serving the end node JERZY. The connection between VTAMs is an APPN multipath channel link, while the connections from JERZY to each VTAM are via the token-ring LAN. JERZY has two PUs defined, each with two emulator sessions. RA3 and RAS are each the DLU server for one PU.

8.7.2 Definitions

The CM/2 definitions for the host links (these define the PUs) are shown in Figure 219. The real links to the two NCPs are also shown for contrast.

```
DEFINE_LOGICAL_LINK  LINK_NAME(LINK0001)
                     ADJACENT_NODE_TYPE(NN)
                     DLC_NAME(IBMTRNET)
                     ADAPTER_NUMBER(0)
                     DESTINATION_ADDRESS(X'400001060000')
                     ETHERNET_FORMAT(NO)
                     CP_CP_SESSION_SUPPORT(YES)
                     SOLICIT_SSCP_SESSION(NO)   4
                     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 219 (Part 1 of 3). CM/2 Definitions - Two DLU Servers

```
DEFINE_LOGICAL_LINK  LINK_NAME(LINK0002)
                     ADJACENT_NODE_TYPE(NN)
                     DLC_NAME(IBMTRNET)
                     ADAPTER_NUMBER(0)
                     DESTINATION_ADDRESS(X'40000105000004')
                     ETHERNET_FORMAT(NO)
                     CP_CP_SESSION_SUPPORT(YES)
                     SOLICIT_SSCP_SESSION(NO)   4
                     NODE_ID(X'05D32258')       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 219 (Part 2 of 3). CM/2 Definitions - Two DLU Servers

```
DEFINE_DEPENDENT_LU_SERVER  LINK_NAME(HOST0001)    1
                            FQ_ADJACENT_CP_NAME(USIBMRA.RAS)
                            ACTIVATE_AT_STARTUP(YES)
                            PU_NAME(RASP74)
                            NODE_ID(X'01706503');

DEFINE_DEPENDENT_LU_SERVER  LINK_NAME(HOST0002)    2
                            FQ_ADJACENT_CP_NAME(USIBMRA.RA3)
                            ACTIVATE_AT_STARTUP(YES)
                            PU_NAME(RA3P74)
                            NODE_ID(X'01796503');
```

Figure 219 (Part 3 of 3). CM/2 Definitions - Two DLU Servers

Of the four 3270 emulation session definitions, two specify HOST0001 as their host link, and two specify HOST0002. These are defined as DLUR/S links **1** and **2**. Note in each case the DLUS CP name, and the node ID (IDBLK and IDNUM); the node IDs must be different because they identify which PU a VTAM will activate. Note also that an IDBLK different from the built-in one is permitted (017 is usually a 3174, 05D belongs to CM/2).

By contrast, the real link definitions have the real node ID of the PS/2 (one by definition **3** and one by default from the local node definition); both specify that SSCP sessions are not to be solicited **4**. No network node server is defined, so RA3 becomes the server since the link to it (LINK0001) is the first to be activated.

The VTAM definitions are shown in Figure 220; the switched major nodes RASMALAI and RA3MALAI identify their respective PUs to RAS and RA3.

```
        VBUILD TYPE=SWNET
*       SWITCHED MAJOR NODE FOR PS/2 AS DLU REQUESTER (RASMALAI)
RASP74   PU     ADDR=C1,MAXPATH=1,IDBLK=017,IDNUM=06503
RAST7400 LU     LOCADDR=2,DLOGMOD=M2SDLCQ,SSCPFM=USSSCS,                *
                MODETAB=AMODETAB,USSTAB=US327X
RAST7401 LU     LOCADDR=3,DLOGMOD=M2SDLCQ,SSCPFM=USSSCS,                *
                MODETAB=AMODETAB,USSTAB=US327X


        VBUILD TYPE=SWNET
*       SWITCHED MAJOR NODE FOR PS/2 AS DLU REQUESTER (RA3MALAI)
RA3P74   PU     ADDR=C1,MAXPATH=1,IDBLK=017,IDNUM=96503
RA3T7400 LU     LOCADDR=4,DLOGMOD=M2SDLCQ,SSCPFM=USSSCS,                *
                MODETAB=AMODETAB,USSTAB=US327X
RA3T7401 LU     LOCADDR=5,DLOGMOD=M2SDLCQ,SSCPFM=USSSCS,                *
                MODETAB=AMODETAB,USSTAB=US327X
```

Figure 220. VTAM Definitions - Two DLU Servers

These definitions are no different from what they would be if the two PUs were 3174s dialling in over a LAN.

8.7.3 Test Results

We activate the VTAM switched major nodes and start CM/2. The four LUs are activated by their DLU servers, and we establish a same-domain and a cross-domain session from the two LUs on each PU. Next, we display JERZY from each VTAM. Figure 221 shows how RAS sees it.

```
 NCCF                      N E T V I E W    RASAN BUCZAK   02/18/94 10:31:32

C RASAN    DISPLAY NET,ID=JERZY,SCOPE=ALL
  RASAN    IST097I  DISPLAY  ACCEPTED
' RASAN
IST075I  NAME = USIBMRA.JERZY    , TYPE = ADJACENT CP   1
IST486I  STATUS= ACT/S----Y, DESIRED STATE= ACTIV
IST977I  MDLTAB=***NA*** ASLTAB=***NA***
IST1333I  ADJLIST = ***NA***
IST861I  MODETAB=***NA*** USSTAB=***NA*** LOGTAB=***NA***
IST934I  DLOGMOD=CPSVRMGR USS LANGTAB=***NA***
IST597I  CAPABILITY-PLU ENABLED  ,SLU ENABLED  ,SESSION LIMIT NONE
IST231I  CDRSC    MAJOR NODE = ISTCDRDY
IST1184I  CPNAME = USIBMRA.JERZY     - NETSRVR = ***NA***
IST1044I  ALSLIST = ISTAPNPU
IST082I  DEVTYPE = INDEPENDENT LU / CDRSC
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST171I  ACTIVE SESSIONS = 0000000002, SESSION REQUESTS = 0000000000
IST206I  SESSIONS:
IST1081I  ADJACENT LINK STATION = CN000003
IST634I  NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I  RAS   3  ACTIV/DL-S D54B4E4B37B3AE1B 0011 0000  0  0 USIBMRA
IST635I  RAS   3  ACTIV/DL-P F627D16481F80421 0000 0010  0  0 USIBMRA
IST1355I  PHYSICAL UNITS SUPPORTED BY DLUR USIBMRA.JERZY
IST089I  RASP74   TYPE = PHYSICAL UNIT    , ACTIV   4
IST924I  -------------------------------------------------------------
IST075I  NAME = USIBMRA.JERZY    , TYPE = DIRECTORY ENTRY
IST1186I  DIRECTORY ENTRY = DYNAMIC    EN
IST1184I  CPNAME = USIBMRA.JERZY     - NETSRVR = USIBMRA.RA3  2
IST314I  END
```

Figure 221. View of JERZY from RAS

We see that JERZY is an adjacent CP **1** (since it has a direct link to RA5NCA4), and that its network node server is RA3 **2**. The only sessions RAS has with JERZY are the two DLUR/S sessions **3**, which serve the PU RASP74 **4**.

Figure 222 on page 163 shows how RA3 sees JERZY.

```
NCCF                    N E T V I E W    RA3AO BUCZAK   02/18/94 10:32:02
C RA3AO    DISPLAY NET,ID=JERZY,SCOPE=ALL
  RA3AO    IST097I  DISPLAY  ACCEPTED
' RA3AO
IST075I  NAME = USIBMRA.JERZY    , TYPE = ADJACENT CP
IST486I  STATUS= ACT/S----Y, DESIRED STATE= ACTIV
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.JERZY     - NETSRVR = ***NA***
IST1044I  ALSLIST = ISTAPNPU
IST082I  DEVTYPE = INDEPENDENT LU / CDRSC
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST171I  ACTIVE SESSIONS = 0000000004, SESSION REQUESTS = 0000000000
IST206I  SESSIONS:
IST1081I  ADJACENT LINK STATION = CN000007
IST634I  NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I  RA3   2  ACTIV/DL-S D54B4E4B38B3AE1B 0010 0000  0  0 USIBMRA
IST635I  RA3   1  ACTIV/CP-S D54B4E4B35B3AE1B 0006 0001  0  0 USIBMRA
IST635I  RA3   2  ACTIV/DL-P F6FF416462BC27DC 0000 0011  0  0 USIBMRA
IST635I  RA3   1  ACTIV/CP-P F6FF416462BC27D8 0001 0005  0  0 USIBMRA
IST1355I  PHYSICAL UNITS SUPPORTED BY DLUR USIBMRA.JERZY
IST089I  RA3P74   TYPE = PHYSICAL UNIT    , ACTIV    3
IST924I  -------------------------------------------------------------
IST075I  NAME = USIBMRA.JERZY    , TYPE = DIRECTORY ENTRY
IST1186I  DIRECTORY ENTRY = DYNAMIC    EN
IST1184I  CPNAME = USIBMRA.JERZY     - NETSRVR = USIBMRA.RA3
IST314I  END
```

Figure 222. View of JERZY from RA3

Here we can see both the CP-CP sessions **1** and the DLUR/S sessions **2**, which serve the PU RA3P74 **3**.

A display of both PUs in RAS (see Figure 223) shows that RA3P74 is unknown **1**, and that RASP74 is a served DLUR PU **2**. A similar display for RA3 shows the opposite results, as expected.

```
 NCCF                    N E T V I E W    RASAN BUCZAK   02/18/94 10:32:47

C RASAN    DISPLAY NET,ID=RA3P74,SCOPE=ALL    1
  RASAN    IST453I  ID       PARAMETER VALUE INVALID
C RASAN    DISPLAY NET,ID=RASP74,SCOPE=ALL
  RASAN    IST097I  DISPLAY  ACCEPTED
' RASAN
IST075I  NAME = RASP74           , TYPE = PU_T2
IST486I  STATUS= ACTIV     , DESIRED STATE= ACTIV
IST1043I  CP NAME = ***NA***, CP NETID = USIBMRA , DYNAMIC LU = YES
IST1354I  DLUR NAME = JERZY       2    MAJNODE = RASMALAI
IST136I  SWITCHED SNA MAJOR NODE = RASMALAI
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST355I  LOGICAL UNITS:
IST080I  RAST7400 ACT/S      RAST7401 ACT/S
IST314I  END
```

Figure 223. Display of PUs from RAS

The application sessions we have set up are as follows:

- LUs owned by RAS:
  - RAST7400 same domain to NetView RASAN
  - RAST7401 cross domain to NetView RA3AO
- LUs owned by RA3:
  - RA3T7400 cross domain to NetView RASAN
  - RA3T7401 same domain to NetView RA3AO

Displaying the LUs from RAS and RA3 shows very similar results, since their session configurations mirror each other. As seen from RA3:

- RA3T7400 and RA3T7401 are owned LUs.
- RAST7400 is not known, since it has no session passing through any part of RA3′s domain. Nor does it have a directory entry since RA3 is not aware of any search involving it.
- RAST7401 is known as an independent LU, since it has a session with an application owned by RA3. It also has a directory entry, since RA3 has participated in a search involving it. The directory entry shows RAST7401 as being a dynamically defined LU on CP JERZY, as can be seen in Figure 224.

```
NCCF                      N E T V I E W    RA3AO BUCZAK   02/18/94 10:36:04
C RA3AO    DISPLAY NET,ID=RAST7400,SCOPE=ALL
  RA3AO    IST453I  ID       PARAMETER VALUE INVALID
C RA3AO    DISPLAY NET,ID=RAST7401,SCOPE=ALL
  RA3AO    IST097I  DISPLAY  ACCEPTED
′ RA3AO
IST075I  NAME = USIBMRA.RAST7401 , TYPE = CDRSC
IST486I  STATUS= ACT/S----Y, DESIRED STATE= ACTIV
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 = RA3     , VERIFY OWNER = NO
IST1184I  CPNAME = USIBMRA.JERZY     - NETSRVR = ***NA***
IST1044I  ALSLIST = ISTAPNPU
IST082I  DEVTYPE = INDEPENDENT LU / CDRSC
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST171I  ACTIVE SESSIONS = 0000000001, SESSION REQUESTS = 0000000000
IST206I  SESSIONS:
IST1081I  ADJACENT LINK STATION = CN000007
IST634I  NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I  RA3AO004 ACTIV-P    F6FF416462BC27E2 0001 0002  0  0 USIBMRA
IST924I  -------------------------------------------------------------
IST075I  NAME = USIBMRA.RAST7401 , TYPE = DIRECTORY ENTRY
IST1186I  DIRECTORY ENTRY = DYNAMIC    LU
IST1184I  CPNAME = USIBMRA.JERZY     - NETSRVR = USIBMRA.RAS   1
IST314I  END
```

Figure 224. Display of RAS′s LUs from RA3

We can see at **1** that the network node server of RAST7400 is shown as RAS, although we know RA3 to be the network node server of JERZY. The directory entry for a served dependent LU always shows the DLU server rather than the real network node server. This reflects the fact that the DLU server is responsible for sessions involving the dependent LU, and replies to search requests on its behalf. The CP name is shown correctly, enabling network nodes to route sessions directly to the dependent LU.

Displaying session routes via NetView shows that, in every case, the direct session path is taken. Regardless of which VTAM actually owns an LU, the session always flows via the NCP owned by the application partner's VTAM, not the NCP owned by the LU's owner VTAM. Even the DLUR/S sessions use the optimum route, as shown in Figure 225.

```
 NLDM.CON                    SESSION CONFIGURATION DATA                 PAGE    1
 -------------- PRIMARY ---------------+-------------- SECONDARY --------------
 NAME RAS       SA 0000001C  EL 0007   |   NAME JERZY     SA 00000005  EL 00B2
 --------------------------------------+---------------------------------------
 DOMAIN RASAN       PCID USIBMRA.RAS.F627D16481F80421             DOMAIN RASAN
                +-------------+                  +-------------+
 RAS            |   CP/SSCP   | ---          --- |             |
 ISTPUS28(0000) | SUBAREA PU  |  APPN TP 01      | SUBAREA PU  | RA5NCA4 (0000)
                +------+------+  SUBA TP 00      +------+------+
                       |              VR 00             |
                +------+------+       ER 00      +------+------+
 RAS     (0007) |     LU      |      RER 00      |    LINK     | J0005011
                +-------------+                  +------+------+
                              APPNCOS #CONNECT          |
                                                 +------+------+
                              LOGMODE CPSVRMGR |     PU      | CN000003(001C)
                                                 +------+------+
                              SADJ CP JERZY             |
                                                 +------+------+
                                                 |     ILU     | JERZY   (00B2)
                                                 +-------------+

SELECT PT, ST (PRI, SEC TRACE), RT (RESP TIME), P, ER, VR, AR
CMD==>
```

Figure 225. Session Configuration for DLUR/S

This figure shows one of the DLUR/S sessions between RAS and JERZY. Even though JERZY's CP-CP sessions go to RA3 via the NCP in subarea 6, this session flows directly via subarea 5. The APPN display confirms the route as going from RAS via TG 21 to JERZY.

8.8 Hints and Tips

The DLUR/S function greatly improves the flexibility available to the network designer, by offering new options for both routing and connectivity. However, certain points need to be considered:

- More than any other function described in the book, DLUR/S demands a good knowledge of products other than VTAM. Successful implementation of DLUR requires a good understanding of how products such as 3174 and CM/2 are customized, and how DLUR affects that customization. The essence of applying DLUR to these products lies in divorcing the LU definitions from real host links, and attaching them to "virtual" DLUS links. With CM/2, our experience indicates that this process is straightforward provided you understand CM/2 setup. (We have used the interim CM/2 code provided with VTAM V4R2.) We have no experience to relate of 3174 (whose code is not available at the time of writing).
- Do not forget that the CP-SVR pipe needs an APPN search path with no subarea hops. For complete flexibility this means using VR-TG across the subarea network.
- 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.
  - If the DLUS and DLUR are in different APPN subnetworks, an extended border node is required in the DLUS′s subnetwork.
  - 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 PLU is in the subarea network, the session will take the traditional route via the DLUS as interchange node. Therefore, use VR-TG where possible between VTAMs so that APPN routing takes effect.
- For an SLU-initiated session, an APPN PLU must support Session Services Extensions. Today, only VTAM supports this.
- Those familiar only with PU T2.0 operation will need to understand that several PUs are now present in a single node, and that they have two distinct functions.
- We have used the interim CM/2 DLUR code provided with VTAM V4R2. This does not support the following functions:
  - Downstream PUs
  - SSCP takeover and giveback
  - Cross-subnetwork CP-SVR pipe
  - ANS=CONT
  - Self-defining dependent LUs

  IBM intends to provide an enhanced version of the DLUR function in a future release of CM/2.

Chapter 9. Automatic Logon Enhancements

Automatic logon (autologon) is an existing function in VTAM. This section describes the enhancements to make autologon more robust. We give examples for the parameters needed to use the functions. Our test results are reported and some hints are given.

9.1 Benefits

The automatic logon enhancements

1. Increase reliability of autologon in APPN, subarea and mixed networks.
2. Remove the constraints for CLSDST PASS to menu applications.
3. Drive the automatic logon even for USERVAR name changes.
4. Reduce overhead by doing only one search for all LUs queued to the same application.

9.2 Autologon Functions

Autologon is used to establish a controlling application for a dependent LU. This can be done in two different ways:

- Code LOGAPPL= on the LU definition. VTAM tries to establish a session from the LU to the controlling application when the LU is activated.
- Issue the operator command VARY LOGON= or VARY ACT,LOGON=. VTAM immediately tries to establish a session.

A detailed description can be found in *VTAM Network Implementation Guide*. To describe the enhancements, we have a look at the existing functions first.

9.2.1 Autologon Functions in Previous Releases of VTAM

The result of VTAM′s attempt to establish a session with the controlling application depends on the state of the application: it must be known, active and reachable. These states are explained below:

- An application is made known to VTAM by APPL or CDRSC definition. The CDRSC may be predefined or learned (dynamic).
- An application is active if its ACB is OPEN and accepts logons.
- An application is reachable if the session request of the SLU can be rerouted to the VTAM of the application. This can be done over SSCP-SSCP sessions or CP-CP sessions.

If the application is known, active and reachable, the session can be established.

If the application is known, reachable, but not active, the session request is queued by the VTAM of the application. When it becomes active, the session can be established.

If the application is known, but not reachable, a session request can be queued by the VTAM of the SLU. When a control session to an adjacent VTAM becomes active, the LU-LU session establishment is retried. The control session to the adjacent VTAM can be SSCP-SSCP or CP-CP.

9.2.2 The New Functions Provided by the Autologon Enhancements

Certain cases are not covered by the autologon function in previous VTAM releases. They include applications using CLSDST PASS, changing USERVARs, cross-net logons, and multi-hop session setup paths (even in the same network). This last case is typical for an APPN network, as the CP-CP sessions are not normally fully meshed as is the case for same-net SSCP-SSCP sessions.

The autologon enhancements remove the limitations of the previous autologon function. This is done in two ways:

- AUTORTRY provides some improvement by using a larger number of events to trigger the retry of session establishment.
- AUTOTI takes a new approach. It is no longer event driven, but uses a timer to trigger the retry of session establishment.

Both parameters can be set as VTAM start options or changed dynamically by operator command.

Network overhead is always a concern with timer driven processes. Each installation can choose the timer interval long enough to prevent the network from being flooded with autologon requests. Additionally there are provisions in VTAM to reduce the overhead: for one application, only one request is sent when the timer expires.

9.3 Experiences with Automatic Logon

We do tests to compare the effect of the old and the new functions.

9.3.1 Description

Figure 226 on page 169 shows the scenario for our test of the effect of the new autologon parameter AUTOTI. VTAM RAA with application TSO RAAAT is channel connected to VTAM RAS. VTAM RAS is connected to both VTAM RAA and VTAM RA3. Dependent LUs on both VTAM RA3 and VTAM RAS are used. They have RAAAT defined as their controlling application, as shown in Figure 227 on page 169. Only VTAM RAS has a direct CP-CP session to VTAM RAA. VTAM RA3 is not adjacent to VTAM RAA, there is no CP-CP session between the two. This will explain a difference in behavior between VTAMs RAS and RA3.

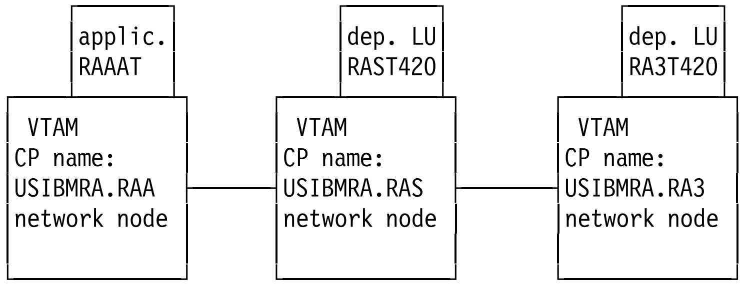

Figure 226. VTAM RAA with Adjacent VTAM and Non-Adjacent VTAM

```
RA3T420  LOCAL CUADDR=420,TERM=3277,LOGAPPL=USIBMRA.RAAAT,            X
               MODETAB=AMODETAB,USSTAB=US3270,                        X
               ISTATUS=ACTIVE,SPAN=(SPH03L),DLOGMOD=M2BSCNQ
```

Figure 227. The Usual LOGAPPL Definitions

9.3.2 Test Results

VTAM RAS has CP-CP sessions to both VTAM RAA and VTAM RA3. The local major nodes with LOGAPPL definitions are activated. As TSO in VTAM RAA is not started, no LU-LU sessions are established. We start TSO, and the LUs from both VTAMs instantly go into session with TSO.

Things are different when the link between VTAMs RAA and RAS is inactivated and re-activated. The sessions are terminated, of course. The sessions from terminals on VTAM RAS are re-established instantly. No sessions from terminals on VTAM RA3 are re-established, as it is not notified when a non-adjacent CP-CP session is established. So it does not retry the autologon. This demonstrates the limitations of autologon up to VTAM V4R1.

We continue the test, switching the autologon enhancements on by issuing the operator command to modify the VTAMOPTS to AUTOTI=1. This causes a retry every minute. (The default is 0.) After about one minute, the sessions from terminals on VTAM RA3 are also re-established.

Although several terminals with controlling application RAAAT are active in VTAM RA3, only one search flows across the network to check if the application RAAAT is active. This was verified in a separate test. When the link between VTAMs RAA and RAS is inactive and the AUTOTI option is set on by operator command, the retry message appears only for the first LU in the queue. This is shown in Figure 228 on page 170.

```
F NET03,VTAMOPTS,AUTOTI=1
IST097I MODIFY ACCEPTED
IST223I MODIFY COMMAND COMPLETED
IST899I  RETRY OF AUTOLOGON(S) TO USIBMRA.RAAAT     IN PROGRESS
IST663I  INIT OTHER REQUEST               FAILED  , SENSE=087D0001
IST664I  REAL  OLU=USIBMRAI.RA3T421    REAL  DLU=USIBMRA.RAAAT
IST889I  SID = F6FF418A62411A17
IST890I  AUTOLOGON SESSION SETUP FAILED
IST894I  ADJSSCPS TRIED  FAILURE SENSE   ADJSSCPS TRIED  FAILURE SENSE
IST895I     ISTAPNCP       08400007       RA3            08260000
IST314I  END
```

Figure 228. Setting On the Timer Controlled Autologon Retry

The test is repeated with a cross network configuration. The results are the same as for one network. The timer controlled autologon retry works for all kinds of SNA networks (subarea and APPN) and independent of the topology. Figure 229 shows that sessions are active across network borders. It also demonstrates the flexibility of the new display command D RSCLIST.

```
D NET,RSCLIST,ID=*.RA?T42*,EXCLUDE=RAA*
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = RSCLIST 068
IST1417I NETID    NAME     STATUS     TYPE             MAJNODE
IST1418I USIBMRAI RA3T421  ACT/S----Y CDRSC            ISTCDRDY
IST1418I USIBMRAI RA3T422  ACT/S----Y CDRSC            ISTCDRDY
IST1418I USIBMRAI RA3T420  ACT/S----Y CDRSC            ISTCDRDY
IST1418I USIBMRA  RAST420  ACT/S----Y CDRSC            ISTCDRDY
IST314I END
```

Figure 229. Display of All Cross-Domain Local Terminals

9.4 Hints and Tips

Set the AUTOTI value high enough to avoid unnecessary overhead in your network. As a rule, the value should be higher than 2 (minutes).

Use the default (AUTOCAP) for AUTORTRY. This gives the functions comparable to VTAM V3 R4 with subarea networking and avoids overhead caused by unproductive retries.

Do not use AUTORTRY=ALL, if there are CPs which do not support SLU-initiated sessions. This will trigger unproductive retries. Currently, only VTAM supports SLU-initiated sessions.

Additional queued sessions require additional storage for control blocks.

Chapter 10. Compression

Compression reduces line utilization at the expense of CPU load. If your network is cost optimized, has to rely on expensive lines, and some of the lines are heavily loaded, then compression is valuable to you.

Due to sophisticated flow control, SNA lines can be loaded up to almost 100% of the theoretical bandwidth. By using SNA′s transmission priorities, you can achieve consistent response times for interactive traffic even on those heavily loaded lines.

Some networking protocols require investment in bandwidth when the line load reaches about 50%. SNA defers this investment up to the time when the line reaches about 90%. Compression pushes this limit even further. You can achieve a data throughput between applications of up to 150% or more of the link capacity.

Compression was introduced in VTAM in V3R4.1. It is enhanced in VTAM V4R2. It is also available in CM/2 R1.1. Not all of the compression enhancements in VTAM V4R2 can be tested in our given environment. We use compression between VTAM applications. We document the externals for compression between a VTAM application (TSO) and an LU 2 in the PS/2, and between two APPC applications (one in VTAM, one in the PS/2).

In this section we explain how to activate compression in VTAM and the CM/2. We show ways to check the effectiveness of compression. Considerations about possible drawbacks of compression are added.

10.1 Benefits

VTAM compression has the following benefits for the customer:

1. Improves response times on slow lines
2. Reduces line cost by reducing the number of leased lines
3. Reduces line cost by reducing the traffic volume
4. Can be used selectively on a session basis
5. Is transparent to the application and the terminal LU

Additional benefits provided by the enhancements in VTAM V4R2 are:

1. Control over compression effectiveness and CPU cycle usage
2. Improved performance by using static compression
3. Improved performance by using hardware compression

10.2 How Compression Works

For a complete overview, please read the *VTAM V4R2 Network Implementation Guide*. It gives a detailed introduction to VTAM Compression.

10.2.1 Compression in VTAM V3R4.1

Five levels of compression are selectable on a session basis:

- Level 0 - No compression is done.
- Level 1 - RLE compression replaces strings of identical bytes with 1- or 2-byte code sequences. It does not require tables.
- Level 2 - LZ9 compression uses tables to replace byte strings by 9-bit codes. The tables are continuously adapted. Up to VTAM V4R1, compression always works in “adaptive” mode.
- Level 3 - LZ10 uses 10-bit codes and big tables.
- Level 4 - LZ12 uses 12-bit codes and even bigger tables.

RLE stands for run length encoding. LZ stands for Lempel-Ziv.

10.2.2 Compression Enhancements in VTAM V4R2

- The adaptive compression and decompression are done more efficiently due to fine tuning of the code.
- The tables can be periodically frozen. This eliminates the table update. VTAM has the choice of doing compression in frozen mode or, as previously, in adaptive mode. Switching between frozen and adaptive is controlled by a new parameter CMPMIPS.
- On machines with the new compression instruction LZ compression can be done by hardware. Decompression continues to be done by software.

10.3 Compression on LU 2 Sessions

Compression is available for all types of LUs, provided both session partners support it. Local non-SNA terminals, for example, do not support compression. In our test we use a 3270 emulation LU on CM/2 as SLU, and TSO as PLU.

10.3.1 Description

Compression is controlled by five different parameters, which define the maximum level of compression allowed. The actual level is found by negotiation at BIND time. Out of the five parameters, the lowest compression level defined is actually chosen. The parameters are contained in:

VTAM start options in the application host

VTAM application major node

Logmode in the SLU′s VTAM

CM/2 local node characteristics

CM/2 3270 emulation options

The VTAM start option CMPVTAM can be set to levels 0 to 4. This can be done when VTAM is started, or dynamically as shown in Figure 230 on page 173.

```
F NET03,VTAMOPTS,CMPVTAM=1
IST097I MODIFY ACCEPTED
IST223I MODIFY COMMAND COMPLETED
D NET,VTAMOPTS,OPTION=CMPVTAM
IST097I DISPLAY ACCEPTED
IST1188I ACF/VTAM V4R2 STARTED AT 11:41:50 ON 02/17/94 926
IST1349I COMPONENT ID IS 5695-11701-201
IST1348I VTAM STARTED AS INTERCHANGE NODE
IST1189I CMPVTAM  = 1
IST314I END
```

Figure 230. Setting and Displaying VTAM's Allowed Compression Level

There are parameters in the VTAM application major node to set the allowed compression level separately for input and output. Figure 231 shows part of a VTAM application major node for TSO. The GROUP statement is used to set CMPAPPLI and CMPAPPLO for all TSO address spaces. No changes in TSO are needed. Compression is transparent to the application.

```
        VBUILD TYPE=APPL
RA3ATG  GROUP  CMPAPPLI=2,CMPAPPLO=2
RA3AT   APPL AUTH=(NOACQ,PASS,NVPACE,TSO,NOPO),                    *
               EAS=1,ACBNAME=TSO
RA3AT01  APPL AUTH=(NOACQ,PASS,NVPACE,TSO,NOPO),                   *
               EAS=1,ACBNAME=TSO0001
RA3AT02  APPL AUTH=(NOACQ,PASS,NVPACE,TSO,NOPO),                   *
               EAS=1,ACBNAME=TSO0002
```

Figure 231. Setting the Compression Level for TSO

The change in the logmode table is done in the VTAM which owns the SLU. See Figure 232. COMPRES=REQD means that the SLU desires compression. The result is subject to negotiation. We have chosen to copy the source of the standard logmode table and modify it, using the standard logmodes. This is better than inventing new non-standard logmodes.

```
M2BSCNQ  MODEENT LOGMODE=M2BSCNQ,FMPROF=X'02',TSPROF=X'02',          X
               PRIPROT=X'71',SECPROT=X'40',COMPROT=X'2000',          X
               COMPRES=REQD,                                         X
               RUSIZES=X'0000',PSERVIC=X'000000000000000000000200'
M2BSCQ   MODEENT LOGMODE=M2BSCQ,FMPROF=X'02',TSPROF=X'02',           X
               PRIPROT=X'71',SECPROT=X'40',COMPROT=X'2000',          X
               RUSIZES=X'0000',PSERVIC=X'008000000000000000000200'
M2SDLCNQ MODEENT LOGMODE=M2SDLCNQ,FMPROF=X'03',TSPROF=X'03',         X
               PRIPROT=X'B1',SECPROT=X'90',COMPROT=X'3080',          X
               RUSIZES=X'8587',PSERVIC=X'020000000000185000007E00'
M2SDLCQ  MODEENT LOGMODE=M2SDLCQ,FMPROF=X'03',                       X
               TSPROF=X'03',                                         X
               PRIPROT=X'B1',SECPROT=X'90',COMPROT=X'3080',          X
               COMPRES=REQD,                                         X
               RUSIZES=X'8587',PSERVIC=X'028000000000185000007E00'
```

Figure 232. Enabling Compression for the SLU

In CM/2, changes have to be done in two different places. We use the setup menus to do both changes. The menus we went through for the first change are shown here in hierarchical order:

```
Communications Manager Setup
  Communications Manager Configuration Definition
    APPC APIs - Advanced
      Communications Manager Profile List
        Local Node Characteristics
          Local Node Options
```

We set the maximum compression level to LZ9 and allocate 100 “tokens” of storage. The help function in CM/2 explains the term “tokens” in this context. This change is reflected in the NDF file as shown in Figure 233 **1**.

```
 DEFINE_LOCAL_CP  FQ_CP_NAME(USIBMRA.BERND    )
                  CP_ALIAS(WTR32256)
                  NAU_ADDRESS(INDEPENDENT_LU)
                  NODE_TYPE(NN)
                  NODE_ID(X'05D32256')
                  NW_FP_SUPPORT(NONE)
                  HOST_FP_SUPPORT(YES)
                  MAX_COMP_LEVEL(LZ9)  1
                  MAX_COMP_TOKENS(100);  1
...
```

Figure 233. Enabling Compression in CM/2

The following menus are selected for the second change:

```
Communications Manager Setup
  Communications Manager Configuration Definition
    3270 Emulation
      Communications Manager Profile List
        3270 Emulation
          Long Session/LU Name
            3270 Logical Terminal - Options
              SNA Data Compression
```

We select to enable compression, save, and restart CM/2. This (second) change for CM/2 does not appear in the NDF file.

10.3.2 Test Results

In this test we set the CMPVTAM to 1. This value is lower than those on the TSO application major node. That is why the actual compression level is 1, as seen in Figure 234 **1**.

```
D NET,SESSION,SID=F6FF416462BC230A
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = SESSIONS 989
IST879I PLU/DLU REAL = USIBMRA.RA3AT01   ALIAS = USIBMRA.TSO0001
IST879I SLU/OLU REAL = USIBMRA.RA3UCLU1  ALIAS = ***NA***
IST880I SETUP STATUS = ACTIV
IST933I LOGMODE=M2SDLCQ , COS=*BLANK*
IST1048I COMPRESSION LEVEL - INPUT = 1, OUTPUT = 1  1
IST1049I PERCENT REDUCTION - INPUT = 0, OUTPUT = 31
IST314I END
```

Figure 234. Displaying the Compression Level of a Session

In a second test, the CMPVTAM value is raised. The display of the session in Figure 235 **2** shows that the input compression level is still 1, but for output has gone up to 2. The maximum levels we can reach with the 3270 emulation are 1 for input and 2 for output. The reason is that the "old fashioned" LUs use a non-extended BIND which does not support control vectors. It allows only a limited negotiation. Even if VTAM sends an extended BIND, for the benefit of the LU it is converted to non-extended BIND by the boundary function. Compression is transparent for the LU as it is for the application.

For this test we browse the MVS log. Figure 235 **3** shows that even with compression level 2 (LZ9) we achieve an impressive 65% reduction. The figure 65 shown in the VTAM display has been calculated since the previous display command. There is a similar display in CM/2 subsystem management under "LU 6.2 session details." It shows accumulated values since the start of the session.

```
D NET,SESSION,SID=F6FF416462BC22E9
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = SESSIONS 934
IST879I PLU/DLU REAL = USIBMRA.RA3AT01   ALIAS = USIBMRA.TSO0001
IST879I SLU/OLU REAL = USIBMRA.RA3UCLU1  ALIAS = ***NA***
IST880I SETUP STATUS = ACTIV
IST933I LOGMODE=M2SDLCQ , COS=*BLANK*
IST1048I COMPRESSION LEVEL - INPUT = 1, OUTPUT = 2  2
IST1049I PERCENT REDUCTION - INPUT = 0, OUTPUT = 65  3
IST314I END
```

Figure 235. Maximum Compression Level for Emulation

10.4 Compression on LU 6.2 Sessions

Using LU 6.2, compression can be done between VTAM applications. The highest compression level can be used in this case. The definitions are symmetrical and can be deduced from the following example.

Our test uses an LU 6.2 session between ATELL in the PS/2 and an ATELL server (an APPC/VTAM application).

10.4.1 Description

Compression is controlled by five different parameters. The lowest level defined decides which compression level is actually chosen. The parameters are contained in:

- VTAM start options in the application host
- VTAM application major node
- Logmode in the PLU′s VTAM
- CM/2 local node characteristics
- CM/2 SNA features (modes)

The VTAM start option, application major node, and logmode table parameters are equivalent to 10.3.1, "Description" on page 172. Figure 236 **1** shows the definition change for the ATELL server application.

```
RA3ASER  APPL  APPC=YES,ACBNAME=KSER,CMPAPPLI=2,CMPAPPLO=2  1
RA3AREQ  APPL  APPC=YES,ACBNAME=KREQ
RA3ALU2  APPL  APPC=YES,ACBNAME=ACBLU2
```

Figure 236. Setting Compression Level for APPC/VTAM Application

```
...
 DEFINE_MODE  MODE_NAME(#INTER  )
              COS_NAME(#INTER  )
              DEFAULT_RU_SIZE(YES)
              RECEIVE_PACING_WINDOW(7)
              MAX_NEGOTIABLE_SESSION_LIMIT(8)
              PLU_MODE_SESSION_LIMIT(8)
              MIN_CONWINNERS_SOURCE(4)
              COMPRESSION_NEED(REQUESTED)  2
              PLU_SLU_COMPRESSION(LZ9)  2
              SLU_PLU_COMPRESSION(LZ9);  2
...
```

Figure 237. Setting the Logmode for APPC Compression in CM/2

In the PS/2, changes are required in two places. The first one is the same as for LU 2 compression. It is described in 10.3, "Compression on LU 2 Sessions" on page 174. The second change is shown here in hierarchical sequence:

Communications Manager Setup
  Communications Manager Configuration Definition
    APPC - API
      Communications Manager Profile List
        SNA Features
          Modes
            Mode Definition

The change is reflected in the NDF file as seen in Figure 237 on page 176 **2**. This change becomes unnecessary if logmodes with compression are provided as standard with the products. We urge you to use standard logmodes whenever possible.

10.4.2 Test Results

Figure 238 **3** informs us about the compression level allowed for this pair of LUs. Figure 239 **4** tells us how effective compression actually is for this session.

```
D NET,ID=RA3ASER,E
IST097I DISPLAY ACCEPTED
IST075I NAME = USIBMRA.RA3ASER, TYPE = APPL 274
IST486I STATUS= ACT/S, DESIRED STATE= ACTIV
IST977I MDLTAB=***NA*** ASLTAB=***NA***
IST861I MODETAB=***NA*** USSTAB=***NA*** LOGTAB=***NA***
IST934I DLOGMOD=***NA*** USS LANGTAB=***NA***
IST597I CAPABILITY-PLU ENABLED  ,SLU ENABLED  ,SESSION LIMIT NONE
IST231I APPL MAJOR NODE = RA3BKAM
IST212I ACBNAME = KSER
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST271I JOBNAME = KAMPMANS, STEPNAME = ***NA***, DSPNAME = IST467D7
IST1050I MAXIMUM COMPRESSION LEVEL - INPUT = 2, OUTPUT = 2  3
IST171I ACTIVE SESSIONS = 0000000002, SESSION REQUESTS = 0000000000
IST206I SESSIONS:
IST634I NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I BERND    ACTIV-P    E287CA2342CFE50F 0001 0002  0  0 USIBMRA
IST635I BERND    ACTIV-P    E287CA2341CFE50F 0001 0001  0  0 USIBMRA
IST314I END
```

Figure 238. Display of Application with Compression Level

```
D NET,SESSION,SID=E287CA2342CFE50F
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = SESSIONS 278
IST879I PLU/OLU REAL = USIBMRA.BERND     ALIAS = ***NA***
IST879I SLU/DLU REAL = USIBMRA.RA3ASER   ALIAS = ***NA***
IST880I SETUP STATUS = ACTIV
IST875I ADJSSCP TOWARDS PLU = ISTAPNCP
IST875I ALSNAME TOWARDS PLU = RA3UC1
IST933I LOGMODE=#INTER  , COS=*BLANK*
IST875I APPNCOS TOWARDS PLU = #INTER
IST1048I COMPRESSION LEVEL - INPUT = 2, OUTPUT = 2
IST1049I PERCENT REDUCTION - INPUT = 47, OUTPUT = 0
IST314I END                          4
```

Figure 239. Display of Session with Compression Level

10.5 Compression in Frozen Mode

As mentioned earlier, the LZ compression uses tables. Using the normal (adaptive) mode, the tables for compression and decompression are updated continuously. The sender "learns" the most recently most frequently used character strings. After some time, the tables will be near the optimum. In adaptive mode, the sender still continues to check for new frequently used character strings.

In frozen mode, the sender avoids this check. For a certain time, it uses the tables as they are. The receiver does decompression in either adaptive or frozen mode, as directed by the sender. The new parameter CMPMIPS can be used to alternate between adaptive and frozen mode. This can be used to reduce the CPU load for compression.

10.5.1 Description

We do a series of tests with an LU 2 session. Session partners are TSO and a 3270 emulation LU in the PS/2. We repeatedly use a file transfer program to transport the same amount of data from TSO to the PS/2.

10.5.2 Test Results

The CMPMIPS parameter is changed dynamically. The change becomes effective immediately. This is different from the other compression parameters. (CMPVTAM, CMPAPPLI, CMPAPPLO, COMPRES only become effective at session establishment time.) Figure 240, Figure 241, and Figure 242 each display the current setting of CMPMIPS and the resulting reduction in the transmission of the same file from TSO to the PS/2. All three session displays refer to the same session.

```
F NET03,VTAMOPTS,CMPMIPS=0
IST097I MODIFY ACCEPTED
IST223I MODIFY COMMAND COMPLETED
D NET,SESSION,SID=F6FF416462BC23AB
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = SESSIONS 244
IST879I PLU/DLU REAL = USIBMRA.RA3AT01   ALIAS = USIBMRA.TSO0001
IST879I SLU/OLU REAL = USIBMRA.RA3UCLU1  ALIAS = ***NA***
IST880I SETUP STATUS = ACTIV
IST933I LOGMODE=M2SDLCQ , COS=*BLANK*
IST1048I COMPRESSION LEVEL - INPUT = 1, OUTPUT = 2
IST1049I PERCENT REDUCTION - INPUT = 0, OUTPUT = 0
IST314I END
```

Figure 240. Switching Compression Off

```
F NET03,VTAMOPTS,CMPMIPS=100
IST097I MODIFY ACCEPTED
IST223I MODIFY COMMAND COMPLETED
D NET,SESSION,SID=F6FF416462BC23AB
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = SESSIONS 204
IST879I PLU/DLU REAL = USIBMRA.RA3AT01   ALIAS = USIBMRA.TSO0001
IST879I SLU/OLU REAL = USIBMRA.RA3UCLU1  ALIAS = ***NA***
IST880I SETUP STATUS = ACTIV
IST933I LOGMODE=M2SDLCQ , COS=*BLANK*
IST1048I COMPRESSION LEVEL - INPUT = 1, OUTPUT = 2
IST1049I PERCENT REDUCTION - INPUT = 0, OUTPUT = 65  2
IST314I END
```

Figure 241. Switching Compression On in Adaptive Mode

We can see that CMPMIPS=0 switches compression off. Although other CMPMIPS values only concern LZ compression, CMPMIPS=0 switches off both LZ and RLE compression.

CMPMIPS=100 is the default, as it reflects the function of previous VTAMs.

CMPMIPS=50 (Figure 242 **1**) is almost as efficient as CMPMIPS=100 (Figure 241 **2**), but uses less CPU cycles. CMPMIPS=50 is the recommended initial value for any tuning activities.

The session displays, as seen in previous figures, inform you about the state of compression for a single session. A new command is available to show us the overall use of compression in a particular VTAM. An example of this command and the resulting messages is given in Figure 243 on page 180.

```
F NET03,VTAMOPTS,CMPMIPS=50
IST097I MODIFY ACCEPTED
IST223I MODIFY COMMAND COMPLETED
D NET,SESSION,SID=F6FF416462BC23AB
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = SESSIONS 173
IST879I PLU/DLU REAL = USIBMRA.RA3AT01   ALIAS = USIBMRA.TSO0001
IST879I SLU/OLU REAL = USIBMRA.RA3UCLU1  ALIAS = ***NA***
IST880I SETUP STATUS = ACTIV
IST933I LOGMODE=M2SDLCQ , COS=*BLANK*
IST1048I COMPRESSION LEVEL - INPUT = 1, OUTPUT = 2
IST1049I PERCENT REDUCTION - INPUT = 0, OUTPUT = 59  1
IST314I END
```

Figure 242. Switching Compression On in Frozen Mode

```
D NET,STATS,TYPE=COMPRESS
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = STATS,TYPE=COMPRESS 066
IST1435I LEVEL     INPUT                 OUTPUT
IST1176I                               BASIC    FROZEN
IST1177I   0         14                 15    **NA**
IST1177I   1          1                  0    **NA**
IST1177I   2          0                  0         0
IST1177I   3          0                  0         0
IST1177I   4          0                  0         0
IST314I END
```

Figure 243. Display the Effect of Compression

10.6 Hints and Tips

Use compression with care. It is beneficial for heavily loaded expensive lines. Use of compression can be controlled on a session basis.

Take into account the increased demand for CPU cycles.

Use frozen compression. CMPMIPS=50 is a good initial value.

Use hardware compression where possible.

When using hardware compression, make sure the maximum RU sizes contained in the BIND are realistic. Unnecessarily large RU size values in the BIND can cause excessive storage usage.

Chapter 11. Search Reduction

This chapter describes a new feature in ACF/VTAM* Version 4 Release 2 called "search reduction" that allows VTAM to reject subsequent requests to resources it has found previously to be unreachable in the network. As the result, these requests that do not have any chance of success do not unnecessarily consume network resources.

We describe how this function can be controlled using VTAM-provided threshold values. In our testing, we also explore the operator feedback provided by VTAM on this feature.

11.1 Benefits

1. It reduces network traffic.

   When a resource becomes unreachable in the network, there may be a great number of subsequent requests for that resource. None of these requests have any chance of success. The search reduction function keeps track of the resources that have been found to be unreachable in the network, and rejects subsequent requests to these resources, therefore saving network resources.

2. The user can have full control.

   The search reduction function provides the user the ability to control how long a resource should be considered unreachable before searching again. By having a timer threshold as well as a request counter threshold, the user is provided with maximum flexibility.

11.2 How Search Reduction Works

Search reduction is turned on in VTAM by using the VTAM option **SRCHRED**. The SRCHRED value can be set (to either ON or OFF) as a VTAM start option or using the MODIFY VTAMOPTS command.

The search reduction function uses a special search type to determine if a resource is reachable in the network. This special search, known as Resource Discovery Search (RDS), is required because VTAM needs to determine why the resource is unreachable — because it is unknown to the network or because access to it is denied by a node within the network. RDS is not associated with a specified session, an origin resource, or a class of service. Instead, it is identified as originating from the CP of the VTAM issuing the RDS. Given that it is not associated with a session or an origin resource or a COS, the only possible reason for a "not found" response to an RDS is because the resource is unreachable.

For APPN searching, an RDS is initiated if no information is known about the resource, or if a directed search has failed to find the resource. For subarea searching, an RDS is initiated when the last non-RDS request fails to locate the resource.

If a negative reply is returned for an RDS and if search reduction is turned on, a "search reduction entry" is created. Having a search reduction entry means the resource is known to VTAM as being unreachable. When a subsequent request is received by VTAM for the same resource, VTAM will know by checking its tables that it is an unreachable resource and reject the request. A failed APPN search will place a search reduction entry in the directory. A failed subarea search will create a search reduction entry as a CDRSC.

But what will happen if the resource becomes reachable after the search reduction entry has been created? VTAM provides thresholds to age the search reduction entry. The search reduction entry is deleted once the associated thresholds are exceeded. This means the resource becomes unknown to VTAM once the thresholds are exceeded. As a result, VTAM issues another search when it receives the next request for the resource. We will see this in our testing.

VTAM provides two threshold control operands for the search reduction function. They are coded as VTAM start options or specified on CDRSC definitions for specific target resources. The thresholds are maintained for each search reduction entry.

- **SRTIMER** is the time to wait after the first RDS failure before a new RDS is initiated. During this period, requests for the resource will be rejected unless the request counter threshold has been exceeded. Note that expiration of the timer results in the search reduction entry for the resource being deleted. It does not result in a search for the resource. The next request for the resource will result in VTAM issuing a new search because the resource is not known.

  The default value for SRTIMER is 30 seconds. If a zero value if coded, this means that no timer threshold is to be maintained.

- **SRCOUNT** is the number of searches to be rejected because of an existing search reduction entry before the search reduction entry is deleted. Note that the search reduction entry will be deleted when **one** of the two thresholds expires. As described before, not having a search reduction entry for the resource will result in VTAM issuing a new RDS for the next request for the resource.

  The default value for SRCOUNT is 10 searches. If a zero value is coded, this means that no request count threshold is maintained.

We suggest that you use both SRTIMER and SRCOUNT values. Using both will give you maximum control.

11.3 When Does VTAM Create Search Reduction Entries?

A search reduction entry is not created for every failed search request. It is a node role specific function. It is created only when the VTAM system is assuming one the following node roles:

- SSCP(OLU)

  This is when VTAM owns the OLU or when VTAM is an interchange node that is servicing the request entering from the APPN side of the interchange node.

- NNS(OLU)

This is when VTAM is the network node server for the OLU or when VTAM is an interchange node that is servicing the request entering from the subarea side of the interchange node.

- ODS

  This is when VTAM is the origin central directory server.

- Origin Border Node

  This is when the OLU is in the domain of the border node. In this case the search reduction entry applies only to the non-native network side.

- Entry Border Node

  This is when VTAM is a border node and the request is received across a subnetwork boundary. Note that in this case an APPN search reduction entry only limits searching within the node′s subnetwork. Therefore, if the border node has other subnetwork boundaries, cross-subnetwork searching will still be done by the border node.

Note that once the search reduction entry is created by VTAM in one of the above roles, it can be used by VTAM in other node roles *when searching the same subnetworks*. This is because the existence of the search reduction entry means the resource is known to VTAM as an unreachable resource.

11.4 Experiences with Search Reduction

This section describes the testing of the search reduction function. We describe the required definitions for the function and findings during the testing of the function.

11.4.1 Description

In the test environment, we use two MVS VTAM V4R2 systems and a PS/2 Communications Manager/2 system. Figure 244 on page 184 shows the APPN view of the test environment. All connections shown in the diagram are APPN (FID2) connections. In the testing, we will use search reduction in RA3 to eliminate the searches sent by RA3, shown as **2** in the diagram.

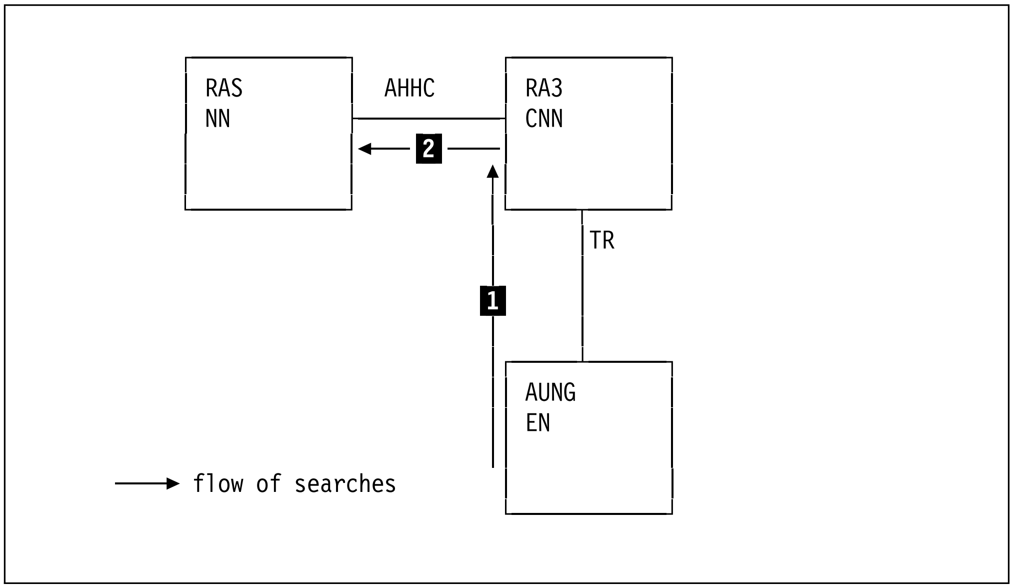

Figure 244. Schematic Diagram of the Test Environment for Search Reduction

- RAS, a VTAM V4R2 network node, has an AHHC connection to RA3. Details on AHHC connections can be found in Chapter 3, "APPN Host-to-Host Connection" on page 11.
- RA3 is a VTAM V4R2 composite network node with a dedicated NCP. It connects to the token-ring LAN via the NCP.
- AUNG, a PS/2 running Communication Manager/2, is defined as an end node in the first part of the test, and as a network node in the later part of the test. It is connected to RA3 via the token-ring LAN.

In the testing, we choose to turn on the search reduction function by using the MODIFY VTAMOPTS command. We will change SRCHRED, SRCOUNT and SRTIMER options. VTAM requires no further definition for VTAM-wide control of search reduction.

A specific CDRSC can be coded with SRCOUNT and SRTIMER parameters to control the searches for a specific target resource. We have not shown this in our testing.

In the testing, we will turn on search reduction in RA3 and examine its effect on requests issued from AUNG and RA3. We will verify that when VTAM is not the network node server of the origin LU (OLU) a search reduction entry is not created. We do this by making AUNG a network node and issuing a request from AUNG.

11.4.2 Test Results

In the testing, we use two types of session requests. One type uses pure APPN logic. This type of request will be generated by issuing an ATELL message from AUNG to an ATELL server application which initially will not be known in the network. The second type uses the interchange function in VTAM. We will issue a logon request from a local terminal attached to RA3 to an initially unknown application.

Before we commence the testing of the search reduction function, we test what happens when search reduction is not turned on. We initialize the environment as shown in Figure 244 and send an ATELL message as shown below.

```
ATELL usibmra.rasaser   Hello There!
```

This results in a LOCATE-FIND request for RASASER being sent from AUNG to RA3 which is AUNG′s network node server. Given that we haven′t activated the application major node for the ATELL server in RAS, ATELL fails with a sense code of 08400007 which means that RASASER is not found in the network. The ATELL command is repeated 9 more times with exactly the same results.

We examine the trace of RA3 so determine if the search requests for RASASER are being propagated from RA3 to RAS. The trace confirms that, for every LOCATE-FIND request received by RA3 from AUNG, one is sent to RAS from RA3.

In the testing, we come across a curious feature of ATELL — each failed execution of ATELL produces two Locate Find requests. Examination of the source code for ATELL reveals that the first search is issued by the main part of the program, but the second search is issued by the error handling routing to retrieve detailed sense code information. This information will prove useful in the next part of our testing.

We now turn search reduction on in RA3 by issuing the following commands:

```
F NET03,VTAMOPTS,SRCHRED=ON       -- to turn search reduction on
F NET03,VTAMOPTS,SRTIMER=6000     -- to set the timer threshold to a high value
F NET03,VTAMOPTS,SRCOUNT=12       -- to set the failure threshold count to 12
```

We display RASASER in both RAS and RA3 using the DISPLAY command. RASASER is not known to either VTAM.

We now execute ATELL again. This fails as expected. The sense code returned from ATELL is 087D000E which means that resource was not found because the searching was restricted due to a previous search failure for the target resource. But this is the first execution of ATELL in our test sequence so we expect to see 08400007. Recalling that ATELL error handling routine issues a second Locate Find request, we examine the CM/2 trace. It reveals that the response to the first Locate Find contains 08400007 as expected. The response to the second Locate Find, which is issued by the error handling routine, is 087D000E.

We now display RASASER in RA3 to see if a search reduction entry has been created. The result of the display command is shown in Figure 245 on page 186. It shows the search reduction entry for subarea searches as a CDRSC **1** and one for APPN searches as a directory entry **2**. This shows that both APPN and subarea entries are created even though no subarea search is performed because RA3 has no adjacent SSCPs.

The IST1401I message as shown in **3** states that a real search will be performed after 5926 seconds or for the 13th request received for this resource. That is, given that SRCOUNT is 12, VTAM will not search for the next 12 requests and will delete the entry after 12 requests. This means that the 13th request will produce a real search for RASASER by RA3. This search count number (13) has not changed since VTAM has created the subarea search reduction entry because no subarea searches have been done by VTAM.

But why is the search count number, as shown in **4**, different for the APPN search reduction entry (represented by the directory entry)? The reason is because each failed ATELL execution produces two searches. The first search from AUNG produces a search from RA3 to RAS. When it fails, RA3 creates the search reduction entries. RA3 then receives the second search, as a part of the error handling by ATELL, and returns a not found (with 087D000E). It also decrements the count by one so that it will be the 12th search from now that will cause the retry.

```
* RA3A0    D NET,ID=RASASER,E
  RA3A0    IST097I  DISPLAY  ACCEPTED
' RA3A0
IST075I  NAME = USIBMRA.RASASER  , TYPE = CDRSC 1
IST486I  STATUS= ACTIV----Y, DESIRED STATE= ACTIV
IST1402I  SRTIMER =  6000  SRCOUNT =    12
IST1401I  RESOURCE NOT FOUND-RETRY IN  5926 SEC(S) OR    13 REQUEST(S)  3
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 = ***NA***, VERIFY OWNER = NO
IST082I  DEVTYPE =             CDRSC
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST171I  ACTIVE SESSIONS = 0000000000, SESSION REQUESTS = 0000000000
IST924I  -------------------------------------------------------------
IST075I  NAME = USIBMRA.RASASER  , TYPE = DIRECTORY ENTRY 2
IST1186I  DIRECTORY ENTRY = DYNAMIC    LU
IST1184I  CPNAME = ***NA***          - NETSRVR = ***NA***
IST1402I  SRTIMER =  6000  SRCOUNT =    12
IST1401I  RESOURCE NOT FOUND-RETRY IN  5926 SEC(S) OR    12 REQUEST(S)  4
IST314I  END
```

Figure 245. Display of RASASER in RA3

We execute ATELL 3 more times. This we expect will reduce the request count threshold by 6. The display of the resource as shown in Figure 246 on page 187 confirms our expectation (**5**). Note in **6** that the request count threshold for the subarea search reduction entry still has not decreased because no subarea searches have been done.

At this point we examine the trace of RA3. The trace confirms that only one search is sent to RAS from RA3, and that the subsequent searches from AUNG to RA3 do not result in searches being sent from RA3 to RAS. Search reduction is indeed working!

```
* RA3AO    D NET,ID=RASASER,E
  RA3AO    IST097I  DISPLAY  ACCEPTED
′ RA3AO
IST075I  NAME = USIBMRA.RASASER  , TYPE = CDRSC
IST486I  STATUS= ACTIV----Y, DESIRED STATE= ACTIV
IST1402I  SRTIMER =  6000  SRCOUNT =    12
IST1401I  RESOURCE NOT FOUND-RETRY IN  5817 SEC(S) OR    13 REQUEST(S) 6
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 = ***NA***, VERIFY OWNER = NO
IST082I  DEVTYPE =            CDRSC
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST171I  ACTIVE SESSIONS = 0000000000, SESSION REQUESTS = 0000000000
IST924I  -------------------------------------------------------------
IST075I  NAME = USIBMRA.RASASER  , TYPE = DIRECTORY ENTRY
IST1186I  DIRECTORY ENTRY = DYNAMIC    LU
IST1184I  CPNAME = ***NA***          - NETSRVR = ***NA***
IST1402I  SRTIMER =  6000  SRCOUNT =    12
IST1401I  RESOURCE NOT FOUND-RETRY IN  5817 SEC(S) OR     6 REQUEST(S) 5
IST314I  END
```

Figure 246. Display of RASASER in RA3

The fact that, as shown in **5** in Figure 246, the sixth request from now will produce a retry is a bit of a dilemma for us. It means that the APPN search reduction entry will be deleted after 5 requests — that is, in the middle of the third ATELL execution from now (after the first search in that ATELL execution but before the second search). The second search — the error handling one — will produce the retry. But we won′t be able to display the resource in RA3 to show that the entry is deleted because by the end of the third ATELL we will have a new entry.

We executed ATELL 3 more times. The third ATELL returns the sense code of 08400007 which indicates that a real search has been done. A look at the CM/2 trace reveals that as expected the first search of the third ATELL execution returns the sense code of 087D000E and the second search produces the sense code 08400007.

Figure 247 on page 188 shows the display of RASASER at this point. It shows in **7** that it will take 13 more searches before a real search is made. This means 6 more ATELL executions before the search reduction entry is deleted. The seventh ATELL execution will produce a retry.

```
* RA3A0    D NET,ID=RASASER,E
  RA3A0    IST097I  DISPLAY  ACCEPTED
' RA3A0
IST075I  NAME = USIBMRA.RASASER  , TYPE = CDRSC
IST486I  STATUS= ACTIV----Y, DESIRED STATE= ACTIV
IST1402I  SRTIMER =  6000  SRCOUNT =    12
IST1401I  RESOURCE NOT FOUND-RETRY IN  5879 SEC(S) OR    13 REQUEST(S).
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 = ***NA***, VERIFY OWNER = NO
IST082I  DEVTYPE =            CDRSC
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST171I  ACTIVE SESSIONS = 0000000000, SESSION REQUESTS = 0000000000
IST924I  -------------------------------------------------------------
IST075I  NAME = USIBMRA.RASASER  , TYPE = DIRECTORY ENTRY
IST1186I  DIRECTORY ENTRY = DYNAMIC    LU
IST1184I  CPNAME = ***NA***          - NETSRVR = ***NA***
IST1402I  SRTIMER =  6000  SRCOUNT =    12
IST1401I  RESOURCE NOT FOUND-RETRY IN  5879 SEC(S) OR    13 REQUEST(S) 7
IST314I  END
```

Figure 247. Display of RASASER in RA3

At this point, we start the ATELL server in RAS by activating the required application major node and starting the program. We execute ATELL 5 more times. Even though RASASER is now known and active in RAS, because there is a search reduction entry in RA3, each execution of ATELL fails with the sense code of 087D000E. This is as expected. Figure 248 shows the directory entry at this point.

```
* RA3A0    D NET,DIRECTRY,ID=RASASER
  RA3A0    IST097I  DISPLAY  ACCEPTED
' RA3A0
IST350I  DISPLAY TYPE = DIRECTORY
IST1186I  DIRECTORY ENTRY = DYNAMIC    LU
IST1184I  CPNAME = ***NA***          - NETSRVR = ***NA***
IST1402I  SRTIMER =  6000  SRCOUNT =    12
IST1401I  RESOURCE NOT FOUND-RETRY IN  5570 SEC(S) OR     3 REQUEST(S)
IST314I  END
```

Figure 248. Directory Entry for RASASER in RA3

We execute ATELL one more time. This execution results in the search reduction entry for RASASER being deleted from the directory. Figure 249 confirms the fact.

```
* RA3A0    D NET,DIRECTRY,ID=USIBMRA.RASASER
  RA3A0    IST453I  ID       PARAMETER VALUE INVALID
```

Figure 249. Directory Entry for RASASER in RA3

The next ATELL execution produces a retry (a real search to RAS) and is successful. Figure 250 on page 189 shows the directory entry at this point.

Note that the IST1401I message, indicating a search reduction entry, does not appear. This is a directory entry for a reachable resource.

```
* RA3AO    D NET,DIRECTRY,ID=USIBMRA.RASASER
  RA3AO    IST097I  DISPLAY  ACCEPTED
' RA3AO
IST350I  DISPLAY TYPE = DIRECTORY
IST1186I  DIRECTORY ENTRY = DYNAMIC    LU
IST1184I  CPNAME = USIBMRA.RAS       - NETSRVR = ***NA***
IST1402I  SRTIMER =  6000  SRCOUNT =    12
IST314I  END
```

Figure 250. Directory Entry for RASASER in RA3

We will now move onto the next part of our testing where we will test the effect of the SRTIMER threshold value. We will now set the SRCOUNT value to a high number so that it makes no difference in our testing and we will set the SRTIMER value to a sufficiently small number that will result in the threshold being exceeded in a few minutes. We issue the following commands:

```
F NET03,VTAMOPTS,SRTIMER=300
F NET03,VTAMOPTS,SRCOUNT=500
```

We will now attempt to log on to TSO using the LU name RASAT which is not known in RAS at this stage. The first attempt produces a failure with the sense code of 08400007. Another attempt produces the sense code of 087D000E as expected. The display of the directory entry for the resource as shown in Figure 251 shows in **1** that it will be over 4 minutes before RA3 will retry the search. Note also the number of requests that remain (499). We have found that each session initiation request from a local terminal decrements the request count by two. This is due to the way the searches from a local terminal are handled internally in VTAM.

```
* RA3AO    D NET,DIRECTRY,ID=RASAT
  RA3AO    IST097I  DISPLAY  ACCEPTED
' RA3AO
IST350I  DISPLAY TYPE = DIRECTORY
IST1186I  DIRECTORY ENTRY = DYNAMIC    LU
IST1184I  CPNAME = ***NA***          - NETSRVR = ***NA***
IST1402I  SRTIMER =   300  SRCOUNT =   500
IST1401I  RESOURCE NOT FOUND-RETRY IN   250 SEC(S) OR   499 REQUEST(S) 1
IST314I  END
```

Figure 251. Directory Entry for RASAT in RA3

At this point, we start RASAT in RAS by starting the appropriate application major node and starting TSO. Another attempt at a logon from a local terminal attached to RA3 fails as expected with the sense code of 087D000E. This is because there is a search reduction entry for RASAT in RA3. We will have to wait until the timer threshold is exceeded and the search reduction entry is deleted.

Figure 252 on page 190 shows the directory entry about 4 minutes later. It shows in **2** that the search reduction entry will be deleted in seven seconds. We wait seven seconds and use the DISPLAY DIRECTRY command again to confirm that it has indeed been deleted.

```
* RA3AO    D NET,DIRECTRY,ID=RASAT
  RA3AO    IST097I  DISPLAY  ACCEPTED
' RA3AO
IST350I  DISPLAY TYPE = DIRECTORY
IST1186I  DIRECTORY ENTRY = DYNAMIC    LU
IST1184I  CPNAME = ***NA***          - NETSRVR = ***NA***
IST1402I  SRTIMER =   300  SRCOUNT =   500
IST1401I  RESOURCE NOT FOUND-RETRY IN     7 SEC(S) OR   499 REQUEST(S) 2
IST314I  END
```

Figure 252. Directory Entry for RASAT in RA3

We now attempt to log on to RASAT from a local terminal attached to RA3 and succeed.

We then proceed to the last part of our testing where AUNG is a network node. We will use ATELL again. This is to check if VTAM creates a search reduction entry if it is not the network node for the origin LU (OLU). As described in 11.3, "When Does VTAM Create Search Reduction Entries?" on page 182, we expect that VTAM will not create a search reduction entry in this scenario.

We execute ATELL to send a message to a destination LU that does not exist as shown below:

```
ATELL usibmra.fred    Hello There!
```

It returns the sense code of 08400007 which implies that a full search is performed but the resource is not found. We executed ATELL several more times. Each time the sense code is 08400007. We can infer from this that that search reduction is not being used. Examination of the trace of RA3 confirms the fact: each search request from AUNG to RA3 is propagated from RA3 to RAS.

11.5 Hints and Tips

Use both SRCOUNT and SRTIMER options so that under any condition VTAM will eventually delete the search reduction entry.

Do some testing before deciding on the SRCOUNT value to be used. As we have found in our testing, there are cases where the request threshold counter is decremented by one, there are cases where the request threshold counter is decremented by two, and there may exist some rare situations where the request threshold counter will be decremented by more than two. Remember that setting the SRTIMER value will provide a "safety net" so even if you set the SRCOUNT value too high the search reduction entry will still be deleted when the timer expires.

Chapter 12. Generic Resources

This new feature of VTAM V4R2 provides support for multiple copies of an application running on a Systems Complex (SYSPLEX), a tightly coupled group of MVS systems. Although we do not have a SYSPLEX available for testing the feature, we believe it is of sufficient importance to merit its own section in this document.

12.1 Benefits

By allowing session requests to be satisfied by any one of a group of identical applications, the Generic Resources feature:

1. Improves availability as compared with a single copy of the application. If one copy fails, only a few users lose service, and they can log on again immediately.
2. Balances the load, therefore improving overall capacity.
3. Does not require the user to be aware of multiple applications.

12.2 How the Generic Resources Feature Works

The Generic Resources feature is implemented by extensions to two existing features in VTAM V4R1: central directory server (CDS) to manage the session allocation, and USERVAR to allow mapping of generic names to real names.

A simple SYSPLEX configuration is shown in Figure 253.

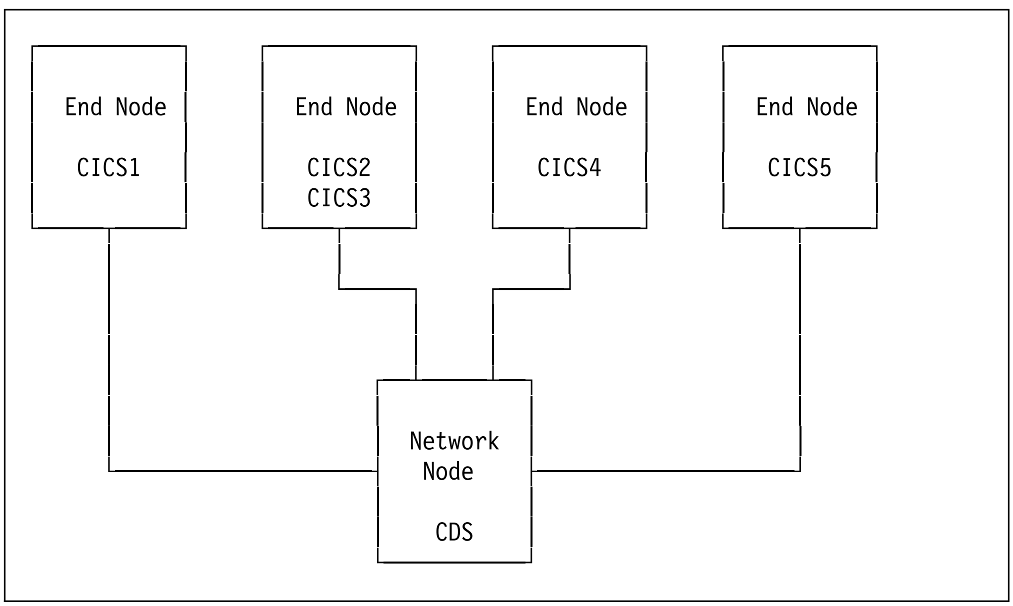

Figure 253. SYSPLEX Configuration

The diagram shows five VTAM systems residing in a SYSPLEX. Four of them are end nodes supporting five identical CICS* applications; the fifth is the network node server, the CDS, and now also the Resource Selector Node (RSN), which is responsible for directing session requests to the most appropriate available CICS system. Other options for configuring a SYSPLEX are:

- Generic resources may reside on end nodes or network nodes.
- For greater availability, more than one RSN may be configured. The RSN need not be the network node server of all (or any) of the resource owners.

Each VTAM in the SYSPLEX must be V4R2, and they must all be linked via APPN connections. The applications themselves must be levels that support generic resources. When such an application opens its ACB it uses new parameters to inform VTAM that it is a member of a generic resource group. VTAM then ensures that this information is registered at the CDS. **No** additional VTAM parameters are required in **any** host in the network.

VTAMs in the network outside the SYSPLEX need not be V4R2, and may be connected to the SYSPLEX by subarea or APPN means. The other VTAMs in the network must be at least V3R2 level with the PTF that supports "enhanced USERVAR".

When a terminal attached to a host outside the SYSPLEX issues a logon request to a SYSPLEX application, its owning VTAM searches for the application by its generic name, whether subarea or APPN searching is performed. When the search is received by the RSN (as a CDS, it will receive directed APPN searches for unknown resources), it uses its internal load balancing algorithms to determine the best real application to satisfy the session request. It then returns the *real* resource name to the requester, indicating that the resource is a "volatile USERVAR". This tells the origin VTAM that the SYSPLEX resource is a temporary alias, and is not guaranteed to retain the same real name. Therefore, the next time that resource is searched for the generic name is used again and the same load balancing mechanism comes into play.

VTAM V4R2 provides support for LU 6.2 sessions as well as LU 2, and ensures that parallel LU 6.2 sessions are set up between the same real resources.

Two points to remember about Generic Resources:

- As well as VTAM V4R2, the feature requires a supporting level of MVS/ESA, which at the time of writing had not been announced. It also requires the supporting levels of the applications, such as CICS/ESA* V4.
- Generic Resources is mutually exclusive with XRF, which provides roughly similar functions. XRF is limited to two copies of an application, but provides nondisruptive session recovery for NCP-attached LUs.

Appendix A. Definitions

Here we list the VTAM and NCP customization parameters that we used in the tests. Those directly relevant to the new features of VTAM V4R2 are shown in the appropriate places in the text; the rest are given here for reference. The full NCP definitions are shown only once, after which we simply show the incremental changes.

The CM/2 definitions relevant to the tests are shown in their proper places in the body of the book.

A.1 Base NCPs

These NCP definitions were used in all the tests. Figure 254 on page 194 shows the NCP for subarea 5, while Figure 255 on page 200 shows the one for subarea 6.


```
NCPOPT   OPTIONS NEWDEFN=(YES,ECHO,NOSUPP,REUSE),USERGEN=(FNMNDFGN)
*----------------------------------------------------------------------
*       PCCU MACRO - SA28 WILL ACTIVATE THIS NCP                     *
*----------------------------------------------------------------------
VTAM28   PCCU CUADDR=E1F,          SA28 MVS/ESA VTAM 4.2               *
               AUTODMP=YES,        ONLY ONE AUTODMP-HOST IF TWINTAIL   *
               AUTOIPL=YES,        ONLY ONE AUTOIPL-HOST IF TWINTAIL   *
               AUTOSYN=YES,        USE THE ALREADY LOADED NCP IF OK    *
               BACKUP=YES,         RESOURCE TAKEOVER PERMITTED         *
               CHANCON=COND,       CONDITIONAL CONTACT REQ. TO NCP SENT*
               DUMPDS=NCPDUMP,     DUMP DATASET                        *
               MDUMPDS=NCPDMOSS,   MOSS DUMP DATASET                   *
               CDUMPDS=NCPDCSP,    SCANNER DUMP DATASET                *
               MAXDATA=5000,                                           *
               VFYLM=YES,          VERIFY LMOD WHEN LOADING            *
               SUBAREA=28          MVS/ESA VTAM 4.2
*
*----------------------------------------------------------------------
*       BUILD MACRO - NCP/CONTROLLER INFORMATION                     *
*----------------------------------------------------------------------
         BUILD BFRS=(240),                   NCP BUFFER SIZE           *
               MEMSIZE=4M,                   4 MEGABYTES               *
               BRANCH=8000,                  BRANCH TRACE ENTRIES      *
               ADDSESS=100,                                            *
               AUXADDR=10,         ADDITIONAL PU FOR ILU  (PU T2.1)    *
               CWALL=26,           MIN. BUFFERS BEFORE SLOWDOWN        *
               CNLSQMAX=10000,     MAX. SIZE OF IP SESSIONS QUEUE      *
               CNLSQTIM=10,        TIME ALLOWED PIU IN OUT-QUEUE       *
               IPPOOL=NCP,         NCP CALCULATES IP BUFFER POOL       *
               IPRATE=(140,180),   MAX. FRAMES ACCEPTER IN .1 SECS     *
               SLODOWN=12,         SLOWDOWN AT 12% BUFFERS REMAIN      *
               DSABLTO=6.5,                                            *
               ENABLTO=6.5,        IBM 386X REQUIRE 6.5 AS MINIMUM     *
               ERLIMIT=16,         NCP SUPPORTS 16 EXPLICIT ROUTES     *
               LTRACE=4,           SIT FOR 4 LINES                     *
               LOADLIB=NCPLOAD,    NCP LOAD MODULE LIBRARY             *
               MAXSSCP=8,          8 SSCP'S CAN ACTIVATE THIS NCP      *
               MLTGORDR=MLTGPRI,   USE MLTGS IN ORDER OF PRIORITY      *
               MODEL=3745-61A,     3745-170 COMMUNICATION CONTROLLER   *
               NETID=USIBMRA,      REQUIRED                            *
               NEWNAME=RA5NCA4,    NAME OF THIS LOAD MODULE            *
               NPA=(YES,DR),       NAME OF THIS LOAD MODULE            *
               NUMHSAS=6,          6 HOSTS MAY COMMUNICATE CONCURRENTLY*
               PUNAME=RA5NCA4,     NAME OF THIS PU                     *
               SUBAREA=05,         SUBAREA ADDRESS = 05                *
               SESSACC=(YES,DEFER,100,10000,250,50,40,100,300,800,1500,*
               3000),                                                  *
               TRACE=(YES,64),     64 ADDRESS-TRACE ENTRIES            *
               TYPGEN=NCP,         CHANNEL ATTACHED  NCP               *
               TYPSYS=MVS,         MVS OPERATING SYSTEM                *
               USGTIER=5,          NCP USAGE TIER                      *
               VERSION=V7R1        NDF VERSION INDICATOR
```

Figure 254 (Part 1 of 6). Base NCP, Subarea 5

```
*
*--------------------------------------------------------------------
*       SYSCNTRL MACRO SPECIFICATIONS                               *
*--------------------------------------------------------------------
NCPSYSC  SYSCNTRL OPTIONS=(BHSASSC,ENDCALL,MODE,RCNTRL,RCOND,RECMD,RIMM*
               ,NAKLIM,SESSION,SSPAUSE,XMTLMT,STORDSP,DLRID,RDEVQ)
*--------------------------------------------------------------------
*       HOST MACRO - CHANNEL ATTACHED HOST DEFINITION               *
*--------------------------------------------------------------------
M20E28   HOST INBFRS=10,          NCP BUFFERS ALLOCATION             *
               MAXBFRU=34,        UP TO 34 VTAM BUFFERS SHIPPED      *
               UNITSZ=182,        VTAM IOBUFFER SIZE                 *
               BFRPAD=0,          BUFFER PAD                         *
               SUBAREA=28         CHANNEL ATTACHED HOST SUBAREA
*--------------------------------------------------------------------
*        DYNAMIC RECONFIGURATION POOL SPACE                         *
*--------------------------------------------------------------------
         PUDRPOOL NUMBER=8
         LUDRPOOL NUMTYP1=10,NUMTYP2=90,NUMILU=100
*--------------------------------------------------------------------
*        PATHS - SA9 NCP
*--------------------------------------------------------------------
         PATH DESTSA=3,ER0=(28,1),                                   *
               VR0=0,                                                *
               VRPWS00=(1,30)
         PATH DESTSA=28,                                             *
               ER0=(28,1),                                           *
               VR0=0,                                                *
               VRPWS00=(1,30)
         PATH DESTSA=6,                                              *
               ER0=(28,1),                                           *
               VR0=0,                                                *
               VRPWS00=(1,30)
         PATH DESTSA=20,                                             *
               ER0=(28,1),                                           *
               VR0=0,                                                *
               VRPWS00=(1,30)
*--------------------------------------------------------------------
*        SDLCST STATEMENTS FOR CONFIGURABLE LINK STATIONS           *
*--------------------------------------------------------------------
RA5SPRI1 SDLCST GROUP=RA5GPRI1,MODE=PRI,RETRIES=(7,3,5),             *
               MAXOUT=7,PASSLIM=254
RA5SSEC1 SDLCST GROUP=RA5GSEC1,MODE=SEC,RETRIES=7,MAXOUT=7,PASSLIM=254
*--------------------------------------------------------------------
*        SDLCST STATEMENTS FOR SUBAREA NTRI CONNECTION              *
*--------------------------------------------------------------------
RA5SPRI2 SDLCST GROUP=RA5GPRI2,MODE=PRI
RA5SSEC2 SDLCST GROUP=RA5GSEC2,MODE=SEC
```

Figure 254 (Part 2 of 6). Base NCP, Subarea 5

```
*----------------------------------------------------------------------
*        GROUP DEFINITION FOR SDLC SELECTION TABLE ENTRY              *
*----------------------------------------------------------------------
RA5GPRI1 GROUP LNCTL=SDLC,MODE=PRI,REPLYTO=3,ACTIVTO=180,              *
               DIAL=NO,RETRIES=(7,4,5),TYPE=NCP
RA5GSEC1 GROUP LNCTL=SDLC,MODE=SEC,REPLYTO=NONE,ACTIVTO=180,           *
               DIAL=NO,RETRIES=(7,0,0),TYPE=NCP
*----------------------------------------------------------------------
*        VIRTUAL GROUP FOR NPA                                        *
*----------------------------------------------------------------------
RA5GNPA  GROUP LNCTL=SDLC,VIRTUAL=YES,NPARSC=YES
RA5LNPA  LINE ISTATUS=ACTIVE
RA5PNPA  PU
RA5TNPA1 LU MAXCOLL=100
RA5TNPA2 LU MAXCOLL=100
*----------------------------------------------------------------------
* NTRI SUBAREA GROUP DEFINITIONS
*----------------------------------------------------------------------
RA5GPRI2 GROUP LNCTL=SDLC,ACTIVTO=240,DIAL=NO,MODE=PRI,REPLYTO=3
RA5GSEC2 GROUP LNCTL=SDLC,ACTIVTO=240,DIAL=NO,MODE=SEC,REPLYTO=3
*----------------------------------------------------------------------
*        INN LINK TO RA5NCSH                                          *
*----------------------------------------------------------------------
RA5GINN  GROUP LNCTL=SDLC,DIAL=NO,REPLYTO=1,TYPE=NCP
RA5L132  LINE  ADDRESS=(132,FULL),CLOCKNG=DIRECT,DUPLEX=FULL,NRZI=YES, *
               SDLCST=(RA5SPRI1,RA5SSEC1),SPEED=9600
RA5P132  PU    TGN=7,PUTYPE=4
*----------------------------------------------------------------------
* NTRI PHYSICAL DEFINITIONS
*----------------------------------------------------------------------
RA5GTRP1 GROUP ECLTYPE=(PHYSICAL,ANY), ALLOW SA,PERIPERAL AND IP       *
               ADAPTER=TIC2,           NTRI INTFACE TYPE USED          *
               SPEED=9600,                                             *
               MAXTSL=2044,            MAX # OF DATA XMIT IN ONE XFER  *
               MAXPU=2,              ALLOW 2 PU'S TO USE PORT          *
               NPACOLL=NO,                                             *
               PUTYPE=1
*
RA5LA88  LINE ADDRESS=(1088,FULL),                                     *
               LOCADD=400001050000,  LOCAL ADMINISTERED ADDRESS        *
               INTFACE=RA5LA88,      ASSOCIATION NAME FOR THIS PORT    *
               PORTADD=0,            ASSOCIATION # FOR LOGICAL LINES   *
               RCVBUFC=4095,         MAX # OF DATA RECV IN ONE XFER    *
               TRSPEED=4
*              STATOPT='TIC1 IP/SNA'
RA5PA881 PU ADDR=01,                 SNA NTRI ATTACHMENT ADDR          *
               INNPORT=YES,          USE THIS PORT FOR SNA-INN TRAFFIC *
               NETWORK=SNA,          THIS PU IS USED FOR SNA TRAFFIC   *
               PASSLIM=2,            SERVICE 2 FRAMES BEFORE NEXT PU   *
               ANS=CONTINUE
*              STATOPT='TIC1 SNA'
```

Figure 254 (Part 3 of 6). Base NCP, Subarea 5

```
*
RA5XA880 LU ISTATUS=INACTIVE,                                             *
               LOCADDR=0            ENABLE INDEPENDENT LU USAGE
*
RA5PA882 PU ADDR=02,                INTERNET NTRI ATTACHMENT ADDRESS  *
               NETWORK=IP,          THIS PU IS USED FOR IP TRAFFIC    *
               ARPTAB=(500,20,NOTCANON), USE NOT CANON                *
               PASSLIM=6            SERVICE 6 FRAMES BEFORE NEXT PU
*              STATOPT='TIC1 IP'
*
*----------------------------------------------------------------------
*PERIPHERAL LOGICAL GROUP  NTRI TIC 2 - CCUA                         *
*----------------------------------------------------------------------
RA5GTRL2 GROUP ECLTYPE=LOGICAL,                                       *
               AUTOGEN=10,                                            *
               CALL=INOUT,                                            *
               PHYPORT=0,NPACOLL=(YES,EXTENDED),                      *
               TYPE=NCP,                                              *
               DIAL=YES,                                              *
               DYNPU=YES,            CREATES SWITCH PU DYNAMICALLY    *
               LINEAUT=YES,                                           *
               PUTYPE=2,                                              *
               XMITDLY=NONE,                                          *
               RETRIES=(6,0,0,6)
*              STATOPT='NTRI TIC 2'
*
*----------------------------------------------------------------------
*        ETHERNET ADAPTER DEFINITION                                  *
*----------------------------------------------------------------------
RA5GETP1 GROUP ETHERNET=PHYSICAL,                                     X
               LANTYPE=DYNAMIC,                                       X
               LNCTL=SDLC
*
RA5LA62  LINE  ADDRESS=(1062,FULL), LINE ADDRESS                      X
               FRAMECNT=50,       THRESHOLD TRAFFIC ERRORS FOR ALERT  X
               INTFACE=RA5LA62    ASSOCIATION NAME OF THIS PORT
*              LOCADD=NOT-DEFINED  USE HARDWARE ADDRESS OF ESS
*
RA5PA62  PU    ANS=CONTINUE,           MUST BE CONTINUE FOR IP        X
               ARPTAB=(500,20),        ALLOW 1K ENTRIES CANONICAL     X
               PUTYPE=1                THIS AN ETHERNET-TYPE LAN PU
*
RA5LA63  LINE  ADDRESS=(1063,FULL), LINE ADDRESS                      X
               FRAMECNT=50,       THRESHOLD TRAFFIC ERRORS FOR ALERT  X
               INTFACE=RA5LA63    ASSOCIATION NAME OF THIS PORT
*              LOCADD=NOT-DEFINED  USE HARDWARE ADDRESS OF ESS
*
RA5PA63  PU    ANS=CONTINUE,           MUST BE CONTINUE FOR IP        X
               ARPTAB=(500,20),        ALLOW 1K ENTRIES CANONICAL     X
               PUTYPE=1                THIS AN ETHERNET-TYPE LAN PU
```

Figure 254 (Part 4 of 6). Base NCP, Subarea 5

```
*
*-----------------------------------------------------------------------
*       NCST LU DEFINITIONS                                            *
*-----------------------------------------------------------------------
RA5GNCS1 GROUP LNCTL=SDLC,             =SDLC REQUIRED FOR NCST=IP     *
               VIRTUAL=YES,                                           *
               NCST=IP                 USE NCST SESSION FOR IP TRAFFIC
RA5LNCS1 LINE
RA5PNCS1 PU
RA5TNCS1 LU    INTFACE=RA5TNCS1,       ASSOCIATION NAME OF THIS LU    *
               LOCADDR=1,              FIRST NCST LU                  *
               REMLU=(RAIATC1)         REMOTE LU NAME IN SA18 (RAIAN)
RA5TNCS2 LU    INTFACE=RA5TNCS2,       ASSOCIATION NAME OF THIS LU    *
               LOCADDR=2,              FIRST NCST LU                  *
               REMLU=(RAIATC2)         REMOTE LU NAME IN SA18 (RAIAN)
*
*-----------------------------------------------------------------------
*       CHANNEL ADAPTER LINK (FID4)                                    *
* ----------------------------------------------------------------------
RA5GCHA1 GROUP LNCTL=CA,                                              *
               ISTATUS=ACTIVE
*
RA5LCHA3 LINE ADDRESS=P02,        CA PHYSICAL POSITION 2              *
               CA=TYPE6,                                              *
               CASDL=120,         INTERNAL BEFORE CHANNEL SLOWDOWN    *
               DELAY=0.0,         CHAN ATTNDELAY                      *
               NCPCA=ACTIVE,      NATIVE SUBCHANNEL (NSC) ACTIVE      *
               TIMEOUT=120        INTERVAL BEFORE CHANNEL DISCONTACT
RA5PCHA3 PU PUTYPE=5,             SUBAREA CONNECTION                  *
               TGN=1
*
RA5LCHA4 LINE ADDRESS=P01,        CA PHYSICAL POSITION 1              *
               CA=TYPE6,                                              *
               CASDL=120,         INTERNAL BEFORE CHANNEL SLOWDOWN    *
               DELAY=0.0,         CHAN ATTNDELAY                      *
               NCPCA=ACTIVE,      NATIVE SUBCHANNEL (NSC) ACTIVE      *
               TIMEOUT=120        INTERVAL BEFORE CHANNEL DISCONTACT
RA5PCHA4 PU PUTYPE=2,             APPN    CONNECTION                  *
               XID=YES
```

Figure 254 (Part 5 of 6). Base NCP, Subarea 5

```
*
*-----------------------------------------------------------------------
*      IP_ADDRESS    INTERFACE    DESCRIPTION                          *
*                                                                      *
*      9.69.32.134   RA5TNCS1     SNALINK LU RA5TNCS1 TO MVS SA18      *
*      9.69.38.5     RA5LA88      TOKEN RING 1088                      *
*      9.69.38.171   RA5LA62      ESS ENET ADAPTER 1062                *
*      9.69.38.173   RA5LA63      ESS ENET ADAPTER 1063                *
*-----------------------------------------------------------------------
         IPOWNER HOSTADDR=9.69.32.135,   IP-@ TCP/IP HOST (NCPROUTE)   *
               INTFACE=RA5TNCS1          USE NCST LU NAMED RA5TNCS1
*              MAXHELLO=6                SENT ALERT IF 6 HELLOS NO ACK
*              NUMROUTE=(25,25,25)       USE DEFAULT 25,25,25
*              UDPPORT=580               USE UDP PORT 580 FOR NCPROUTE
*
         IPLOCAL INTFACE=RA5TNCS1,       SELECT INTFACE NAMED RA5TNCS1 *
               LADDR=9.69.32.134,        IP ADDRESS OF THIS INTERFACE  *
               METRIC=1,                 INTERFACE METRIC = 1          *
               P2PDEST=9.69.32.135,      IP ADDRESS OF TCP/IP HOST     *
               PROTOCOL=RIP,             MANAGED BY NCPROUTE           *
               SNETMASK=255.255.255.192  SUBNETMASK USED
*
         IPLOCAL INTFACE=RA5TNCS2,       SELECT INTFACE NAMED RA5TNCS1 *
               LADDR=9.69.38.197,        IP ADDRESS OF THIS INTERFACE  *
               METRIC=1,                 INTERFACE METRIC = 1          *
               P2PDEST=9.69.38.198,      IP ADDRESS OF TCP/IP HOST     *
               PROTOCOL=RIP,             MANAGED BY NCPROUTE           *
               SNETMASK=255.255.255.192  SUBNETMASK USED
*
         IPLOCAL INTFACE=RA5LA88,        SELECT INTFACE NAMED RA5LA88  *
               LADDR=9.69.38.5,          IP ADDRESS OF THIS INTERFACE  *
               METRIC=1,                 INTERFACE METRIC = 1          *
               PROTOCOL=RIP,             MANAGED BY NCPROUTE           *
               SNETMASK=255.255.255.192  SUBNETMASK USED
*
         IPLOCAL INTFACE=RA5LA62,        SELECT INTFACE NAMED RA5LA89  *
               LADDR=9.69.38.171,        IP ADDRESS OF THIS INTERFACE  *
               METRIC=1,                 INTERFACE METRIC = 1          *
               PROTOCOL=RIP,             MANAGED BY NCPROUTE           *
               SNETMASK=255.255.255.192 SUBNETMASK USED
*
         IPLOCAL INTFACE=RA5LA63,        SELECT INTFACE NAMED RA5LA63  *
               LADDR=9.69.38.173,        IP ADDRESS OF THIS INTERFACE  *
               METRIC=1,                 INTERFACE METRIC = 1          *
               PROTOCOL=RIP,             MANAGED BY NCPROUTE           *
               SNETMASK=255.255.255.192 SUBNETMASK USED
*
         GENEND INIT=ECLINIT,                                          *
               TMRTICK=ECLTICK,                                        *
               UGLOBAL=ECLUGBL
```

Figure 254 (Part 6 of 6). Base NCP, Subarea 5

```
NCPOPT   OPTIONS NEWDEFN=(YES,ECHO,NOSUPP,REUSE),USERGEN=(FNMNDFGN)
*----------------------------------------------------------------------
*       PCCU MACRO - SA03 WILL ACTIVATE THIS NCP                      *
*----------------------------------------------------------------------
VTAM03   PCCU CUADDR=E21,         SA03 MVS/ESA VTAM 4.2                *
               AUTODMP=YES,       ONLY ONE AUTODMP-HOST IF TWINTAIL    *
               AUTOIPL=YES,       ONLY ONE AUTOIPL-HOST IF TWINTAIL    *
               AUTOSYN=YES,       USE THE ALREADY LOADED NCP IF OK     *
               BACKUP=YES,        RESOURCE TAKEOVER PERMITTED          *
               CHANCON=COND,      CONDITIONAL CONTACT REQ. TO NCP SENT*
               DUMPDS=NCPDUMP,    DUMP DATASET                         *
               MDUMPDS=NCPDMOSS,  MOSS DUMP DATASET                    *
               CDUMPDS=NCPDCSP,   SCANNER DUMP DATASET                 *
               MAXDATA=5000,                                           *
               VFYLM=YES,         VERIFY LMOD WHEN LOADING             *
               SUBAREA=03         MVS/ESA VTAM 4.2
*
*----------------------------------------------------------------------
*       BUILD MACRO - NCP/CONTROLLER INFORMATION                      *
*----------------------------------------------------------------------
         BUILD BFRS=(240),                  NCP BUFFER SIZE            *
               MEMSIZE=4M,                  4 MEGABYTES                *
               BRANCH=8000,                 BRANCH TRACE ENTRIES       *
               ADDSESS=100,                                            *
               AUXADDR=10,        ADDITIONAL PU FOR ILU  (PU T2.1)     *
               CWALL=26,          MIN. BUFFERS BEFORE SLOWDOWN         *
               CNLSQMAX=10000,    MAX. SIZE OF IP SESSIONS QUEUE       *
               CNLSQTIM=10,       TIME ALLOWED PIU IN OUT-QUEUE        *
               IPPOOL=NCP,        NCP CALCULATES IP BUFFER POOL        *
               IPRATE=(140,180),  MAX. FRAMES ACCEPTER IN .1 SECS      *
               SLODOWN=12,        SLOWDOWN AT 12% BUFFERS REMAIN       *
               DSABLTO=6.5,                                            *
               ENABLTO=6.5,       IBM 386X REQUIRE 6.5 AS MINIMUM      *
               ERLIMIT=16,        NCP SUPPORTS 16 EXPLICIT ROUTES      *
               LTRACE=4,          SIT FOR 4 LINES                      *
               LOADLIB=NCPLOAD,   NCP LOAD MODULE LIBRARY              *
               MAXSSCP=8,         8 SSCP'S CAN ACTIVATE THIS NCP       *
               MLTGORDR=MLTGPRI,  USE MLTGS IN ORDER OF PRIORITY       *
               MODEL=3745-61A,    3745-170 COMMUNICATION CONTROLLER    *
               NETID=USIBMRA,     REQUIRED                             *
               NEWNAME=RA6NCA4,   NAME OF THIS LOAD MODULE             *
               NPA=(YES,DR),      NAME OF THIS LOAD MODULE             *
               NUMHSAS=6,         6 HOSTS MAY COMMUNICATE CONCURRENTLY*
               PUNAME=RA6NCA4,    NAME OF THIS PU                      *
               SUBAREA=06,        SUBAREA ADDRESS = 06                 *
               SESSACC=(YES,DEFER,100,10000,250,50,40,100,300,800,1500,*
               3000),                                                  *
               TRACE=(YES,64),    64 ADDRESS-TRACE ENTRIES             *
               TYPGEN=NCP,        CHANNEL ATTACHED  NCP                *
               TYPSYS=MVS,        MVS OPERATING SYSTEM                 *
               USGTIER=5,         NCP USAGE TIER                       *
               VERSION=V7R1       NDF VERSION INDICATOR
```

Figure 255 (Part 1 of 6). Base NCP, Subarea 6

```
*
*----------------------------------------------------------------------
*       SYSCNTRL MACRO SPECIFICATIONS                                  *
*----------------------------------------------------------------------
NCPSYSC  SYSCNTRL OPTIONS=(BHSASSC,ENDCALL,MODE,RCNTRL,RCOND,RECMD,RIMM*
               ,NAKLIM,SESSION,SSPAUSE,XMTLMT,STORDSP,DLRID,RDEVQ)
*----------------------------------------------------------------------
*       HOST MACRO - CHANNEL ATTACHED HOST DEFINITION                  *
*----------------------------------------------------------------------
M20E03   HOST INBFRS=10,          NCP BUFFERS ALLOCATION                *
               MAXBFRU=34,        UP TO 34 VTAM BUFFERS SHIPPED         *
               UNITSZ=182,        VTAM IOBUFFER SIZE                    *
               BFRPAD=0,          BUFFER PAD                            *
               SUBAREA=03         CHANNEL ATTACHED HOST SUBAREA
*----------------------------------------------------------------------
*        DYNAMIC RECONFIGURATION POOL SPACE                            *
*----------------------------------------------------------------------
         PUDRPOOL NUMBER=8
         LUDRPOOL NUMTYP1=10,NUMTYP2=90,NUMILU=100
*----------------------------------------------------------------------
*        PATHS - SA9 NCP
*----------------------------------------------------------------------
         PATH DESTSA=3,ER0=(03,1),                                      *
               VR0=0,                                                   *
               VRPWS00=(1,30)
         PATH DESTSA=28,                                                *
               ER0=(03,1),                                              *
               VR0=0,                                                   *
               VRPWS00=(1,30)
         PATH DESTSA=5,                                                 *
               ER0=(03,1),                                              *
               VR0=0,                                                   *
               VRPWS00=(1,30)
         PATH DESTSA=20,                                                *
               ER0=(03,1),                                              *
               VR0=0,                                                   *
               VRPWS00=(1,30)
*----------------------------------------------------------------------
*        SDLCST STATEMENTS FOR CONFIGURABLE LINK STATIONS              *
*----------------------------------------------------------------------
RA6SPRI1 SDLCST GROUP=RA6GPRI1,MODE=PRI,RETRIES=(7,3,5),                *
               MAXOUT=7,PASSLIM=254
RA6SSEC1 SDLCST GROUP=RA6GSEC1,MODE=SEC,RETRIES=7,MAXOUT=7,PASSLIM=254
*----------------------------------------------------------------------
*        SDLCST STATEMENTS FOR SUBAREA NTRI CONNECTION                 *
*----------------------------------------------------------------------
RA6SPRI2 SDLCST GROUP=RA6GPRI2,MODE=PRI
RA6SSEC2 SDLCST GROUP=RA6GSEC2,MODE=SEC
```

Figure 255 (Part 2 of 6). Base NCP, Subarea 6

```
*----------------------------------------------------------------------
*        GROUP DEFINITION FOR SDLC SELECTION TABLE ENTRY              *
*----------------------------------------------------------------------
RA6GPRI1 GROUP LNCTL=SDLC,MODE=PRI,REPLYTO=3,ACTIVTO=180,              *
               DIAL=NO,RETRIES=(7,4,5),TYPE=NCP
RA6GSEC1 GROUP LNCTL=SDLC,MODE=SEC,REPLYTO=NONE,ACTIVTO=180,           *
               DIAL=NO,RETRIES=(7,0,0),TYPE=NCP
*----------------------------------------------------------------------
*        VIRTUAL GROUP FOR NPA                                        *
*----------------------------------------------------------------------
RA6GNPA  GROUP LNCTL=SDLC,VIRTUAL=YES,NPARSC=YES
RA6LNPA  LINE ISTATUS=ACTIVE
RA6PNPA  PU
RA6TNPA1 LU MAXCOLL=100
RA6TNPA2 LU MAXCOLL=100
*----------------------------------------------------------------------
* NTRI SUBAREA GROUP DEFINITIONS
*----------------------------------------------------------------------
RA6GPRI2 GROUP LNCTL=SDLC,ACTIVTO=240,DIAL=NO,MODE=PRI,REPLYTO=3
RA6GSEC2 GROUP LNCTL=SDLC,ACTIVTO=240,DIAL=NO,MODE=SEC,REPLYTO=3
*----------------------------------------------------------------------
*        INN LINK TO RA6NCSH                                          *
*----------------------------------------------------------------------
RA6GINN  GROUP LNCTL=SDLC,DIAL=NO,REPLYTO=1,TYPE=NCP
RA6L132  LINE  ADDRESS=(132,FULL),CLOCKNG=DIRECT,DUPLEX=FULL,NRZI=YES, *
               SDLCST=(RA6SPRI1,RA6SSEC1),SPEED=9600
RA6P132  PU    TGN=7,PUTYPE=4
*----------------------------------------------------------------------
* NTRI PHYSICAL DEFINITIONS
*----------------------------------------------------------------------
RA6GTRP1 GROUP ECLTYPE=(PHYSICAL,ANY), ALLOW SA,PERIPERAL AND IP       *
               ADAPTER=TIC2,           NTRI INTFACE TYPE USED          *
               SPEED=9600,                                             *
               MAXTSL=2044,            MAX # OF DATA XMIT IN ONE XFER  *
               MAXPU=2,              ALLOW 2 PU'S TO USE PORT          *
               NPACOLL=NO,                                             *
               PUTYPE=1
```

Figure 255 (Part 3 of 6). Base NCP, Subarea 6

```
*
RA6LA88  LINE ADDRESS=(1088,FULL),                                    *
               LOCADD=400001060000,  LOCAL ADMINISTERED ADDRESS       *
               INTFACE=RA6LA88,      ASSOCIATION NAME FOR THIS PORT   *
               PORTADD=0,            ASSOCIATION # FOR LOGICAL LINES  *
               RCVBUFC=4095,         MAX # OF DATA RECV IN ONE XFER   *
               TRSPEED=4
*              STATOPT='TIC1 IP/SNA'
RA6PA881 PU ADDR=01,                 SNA NTRI ATTACHMENT ADDR         *
               INNPORT=YES,          USE THIS PORT FOR SNA-INN TRAFFIC *
               NETWORK=SNA,          THIS PU IS USED FOR SNA TRAFFIC  *
               PASSLIM=2,            SERVICE 2 FRAMES BEFORE NEXT PU  *
               ANS=CONTINUE
*              STATOPT='TIC1 SNA'
*
RA6XA880 LU ISTATUS=INACTIVE,                                         *
               LOCADDR=0             ENABLE INDEPENDENT LU USAGE
*
RA6PA882 PU ADDR=02,                 INTERNET NTRI ATTACHMENT ADDRESS *
               NETWORK=IP,           THIS PU IS USED FOR IP TRAFFIC   *
               ARPTAB=(500,20,NOTCANON), USE NOT CANON                *
               PASSLIM=6             SERVICE 6 FRAMES BEFORE NEXT PU
*              STATOPT='TIC1 IP'
*
*---------------------------------------------------------------------
*PERIPHERAL LOGICAL GROUP  NTRI TIC 2 - CCUA                         *
*---------------------------------------------------------------------
RA6GTRL2 GROUP ECLTYPE=LOGICAL,                                       *
               AUTOGEN=10,                                            *
               CALL=INOUT,                                            *
               PHYPORT=0,NPACOLL=(YES,EXTENDED),                      *
               TYPE=NCP,                                              *
               DIAL=YES,                                              *
               DYNPU=YES,             CREATE SWITCH PU DYNAMICALLY    *
               LINEAUT=YES,                                           *
               PUTYPE=2,                                              *
               XMITDLY=NONE,                                          *
               RETRIES=(6,0,0,6)
*              STATOPT='NTRI TIC 2'
*
```

Figure 255 (Part 4 of 6). Base NCP, Subarea 6

```
*----------------------------------------------------------------------
*       ETHERNET ADAPTER DEFINITION                                   *
*----------------------------------------------------------------------
RA6GETP1 GROUP ETHERNET=PHYSICAL,                                      X
               LANTYPE=DYNAMIC,                                        X
               LNCTL=SDLC
*
RA6LA62  LINE  ADDRESS=(1062,FULL), LINE ADDRESS                       X
               FRAMECNT=50,       THRESHOLD TRAFFIC ERRORS FOR ALERT   X
               INTFACE=RA6LA62    ASSOCIATION NAME OF THIS PORT
*              LOCADD=NOT-DEFINED  USE HARDWARE ADDRESS OF ESS
*
RA6PA62  PU    ANS=CONTINUE,           MUST BE CONTINUE FOR IP         X
               ARPTAB=(500,20),        ALLOW 1K ENTRIES CANONICAL      X
               PUTYPE=1                THIS AN ETHERNET-TYPE LAN PU
*
RA6LA63  LINE  ADDRESS=(1063,FULL), LINE ADDRESS                       X
               FRAMECNT=50,       THRESHOLD TRAFFIC ERRORS FOR ALERT   X
               INTFACE=RA6LA63    ASSOCIATION NAME OF THIS PORT
*              LOCADD=NOT-DEFINED  USE HARDWARE ADDRESS OF ESS
*
RA6PA63  PU    ANS=CONTINUE,           MUST BE CONTINUE FOR IP         X
               ARPTAB=(500,20),        ALLOW 1K ENTRIES CANONICAL      X
               PUTYPE=1                THIS AN ETHERNET-TYPE LAN PU
*
*----------------------------------------------------------------------
*      NCST LU DEFINITIONS                                            *
*----------------------------------------------------------------------
RA6GNCS1 GROUP LNCTL=SDLC,             =SDLC REQUIRED FOR NCST=IP      *
               VIRTUAL=YES,                                            *
               NCST=IP                 USE NCST SESSION FOR IP TRAFFIC
RA6LNCS1 LINE
RA6PNCS1 PU
RA6TNCS1 LU    INTFACE=RA6TNCS1,       ASSOCIATION NAME OF THIS LU     *
               LOCADDR=1,              FIRST NCST LU                   *
               REMLU=(RAIATC1)         REMOTE LU NAME IN SA18 (RAIAN)
RA6TNCS2 LU    INTFACE=RA6TNCS2,       ASSOCIATION NAME OF THIS LU     *
               LOCADDR=2,              FIRST NCST LU                   *
               REMLU=(RAIATC2)         REMOTE LU NAME IN SA18 (RAIAN)
*
*----------------------------------------------------------------------
*      CHANNEL ADAPTER LINK (FID4)                                    *
* ---------------------------------------------------------------------
RA6GCHA1 GROUP LNCTL=CA,                                               *
               ISTATUS=ACTIVE
*
RA6LCHA3 LINE ADDRESS=P03,        CA PHYSICAL POSITION 3               *
               CA=TYPE6,                                               *
               CASDL=120,         INTERNAL BEFORE CHANNEL SLOWDOWN     *
               DELAY=0.0,         CHAN ATTNDELAY                       *
               NCPCA=ACTIVE,      NATIVE SUBCHANNEL (NSC) ACTIVE       *
               TIMEOUT=120        INTERVAL BEFORE CHANNEL DISCONTACT
RA6PCHA3 PU PUTYPE=2,             APPN    CONNECTION                   *
               XID=YES
*
```

Figure 255 (Part 5 of 6). Base NCP, Subarea 6

```
RA6LCHA4 LINE ADDRESS=P04,        CA PHYSICAL POSITION 4            *
              CA=TYPE6,                                             *
              CASDL=120,          INTERNAL BEFORE CHANNEL SLOWDOWN  *
              DELAY=0.0,          CHAN ATTNDELAY                    *
              NCPCA=ACTIVE,       NATIVE SUBCHANNEL (NSC) ACTIVE    *
              TIMEOUT=120         INTERVAL BEFORE CHANNEL DISCONTACT
RA6PCHA4 PU PUTYPE=5,             SUBAREA CONNECTION                *
              TGN=1
*
*--------------------------------------------------------------------
*       IP_ADDRESS    INTERFACE    DESCRIPTION                     *
*                                                                  *
*       9.69.32.134   RA6TNCS1     SNALINK LU RA6TNCS1 TO MVS SA18 *
*       9.69.38.5     RA6LA88      TOKEN RING 1088                 *
*       9.69.38.171   RA6LA62      ESS ENET ADAPTER 1062           *
*       9.69.38.173   RA6LA63      ESS ENET ADAPTER 1063           *
*--------------------------------------------------------------------
         IPOWNER HOSTADDR=9.69.32.135,   IP-@ TCP/IP HOST (NCPROUTE)   *
              INTFACE=RA6TNCS1          USE NCST LU NAMED RA6TNCS1
*
         IPLOCAL INTFACE=RA6TNCS1,       SELECT INTFACE NAMED RA6TNCS1 *
              LADDR=9.69.32.134,        IP ADDRESS OF THIS INTERFACE  *
              METRIC=1,                 INTERFACE METRIC = 1          *
              P2PDEST=9.69.32.135,      IP ADDRESS OF TCP/IP HOST     *
              PROTOCOL=RIP,             MANAGED BY NCPROUTE           *
              SNETMASK=255.255.255.192  SUBNETMASK USED
*
         IPLOCAL INTFACE=RA6TNCS2,       SELECT INTFACE NAMED RA6TNCS1 *
              LADDR=9.69.38.197,        IP ADDRESS OF THIS INTERFACE  *
              METRIC=1,                 INTERFACE METRIC = 1          *
              P2PDEST=9.69.38.198,      IP ADDRESS OF TCP/IP HOST     *
              PROTOCOL=RIP,             MANAGED BY NCPROUTE           *
              SNETMASK=255.255.255.192  SUBNETMASK USED
*
         IPLOCAL INTFACE=RA6LA88,        SELECT INTFACE NAMED RA6LA88  *
              LADDR=9.69.38.5,          IP ADDRESS OF THIS INTERFACE  *
              METRIC=1,                 INTERFACE METRIC = 1          *
              PROTOCOL=RIP,             MANAGED BY NCPROUTE           *
              SNETMASK=255.255.255.192  SUBNETMASK USED
*
         IPLOCAL INTFACE=RA6LA62,        SELECT INTFACE NAMED RA6LA89  *
              LADDR=9.69.38.171,        IP ADDRESS OF THIS INTERFACE  *
              METRIC=1,                 INTERFACE METRIC = 1          *
              PROTOCOL=RIP,             MANAGED BY NCPROUTE           *
              SNETMASK=255.255.255.192 SUBNETMASK USED
*
         IPLOCAL INTFACE=RA6LA63,        SELECT INTFACE NAMED RA6LA63  *
              LADDR=9.69.38.173,        IP ADDRESS OF THIS INTERFACE  *
              METRIC=1,                 INTERFACE METRIC = 1          *
              PROTOCOL=RIP,             MANAGED BY NCPROUTE           *
              SNETMASK=255.255.255.192 SUBNETMASK USED
*
         GENEND INIT=ECLINIT,                                          *
              TMRTICK=ECLTICK,                                         *
              UGLOBAL=ECLUGBL
```

Figure 255 (Part 6 of 6). Base NCP, Subarea 6

A.2 NCP Changes for VR-TG Tests

Figure 256 shows the changes made to the NCPs for 6.3, "VR-TG for Optimum Routing" on page 88. The only changes in subarea 6 are the new path tables, whereas subarea 5 has a FID4 channel link to RAA as well as the path tables. The NCPs are now named RA5NCA9 and RA6NCA9.

```
*---------------------------------------------------------------------
*        PATHS - SA6 NCP
*---------------------------------------------------------------------
         PATH  DESTSA=3,                                              *
               ER0=(3,1),ER1=(20,1),                                  *
               VR0=0,                                                 *
               VRPWS00=(1,30),VRPWS01=(1,30),VRPWS02=(1,30),          *
               VR1=1,                                                 *
               VRPWS10=(2,60),VRPWS11=(2,60),VRPWS12=(2,60)
         PATH  DESTSA=5,                                              *
               ER0=(5,6,5000,5000,5000,20000),                        *
               ER1=(20,1),ER2=(3,1),                                  *
               VR0=0,                                                 *
               VRPWS00=(1,30),VRPWS01=(1,30),VRPWS02=(1,30),          *
               VR1=1,                                                 *
               VRPWS10=(2,60),VRPWS11=(2,60),VRPWS12=(2,60),          *
               VR2=2,                                                 *
               VRPWS20=(3,90),VRPWS21=(3,90),VRPWS22=(3,90)
         PATH  DESTSA=28,                                             *
               ER0=(20,1),ER1=(28,1),ER2=(9,9),                       *
               ER3=(28,1),ER4=(28,1),ER5=(9,6),                       *
               ER6=(9,7),                                             *
               VR0=1,                                                 *
               VRPWS00=(1,30),VRPWS01=(1,30),VRPWS02=(1,30),          *
               VR1=0,                                                 *
               VRPWS10=(2,60),VRPWS11=(2,60),VRPWS12=(2,60),          *
               VR2=5,                                                 *
               VRPWS20=(2,60),VRPWS21=(2,60),VRPWS22=(2,60),          *
               VR3=6,                                                 *
               VRPWS30=(2,60),VRPWS31=(2,60),VRPWS32=(2,60),          *
               VR4=2,                                                 *
               VRPWS40=(2,60),VRPWS41=(2,60),VRPWS42=(2,60)
```

Figure 256 (Part 1 of 3). NCP Changes - RA6NCA9 and RA5NCA9

```
*---------------------------------------------------------------------
*       PATHS - SA5 NCP
*---------------------------------------------------------------------
         PATH  DESTSA=3,                                               *
               ER0=(3,1),ER1=(11,1),ER2=(28,1),                        *
               VR0=0,                                                  *
               VRPWS00=(1,30),VRPWS01=(1,30),VRPWS02=(1,30),           *
               VR1=1,                                                  *
               VRPWS10=(2,60),VRPWS11=(2,60),VRPWS12=(2,60),           *
               VR2=2,                                                  *
               VRPWS20=(2,60),VRPWS21=(2,60),VRPWS22=(2,60)
         PATH  DESTSA=6,                                               *
               ER1=(6,6,5000,5000,5000,20000),                         *
               ER0=(20,1),ER5=(28,1),                                  *
               VR0=1,                                                  *
               VRPWS00=(1,30),VRPWS01=(1,30),VRPWS02=(1,30),           *
               VR1=0,                                                  *
               VRPWS10=(2,60),VRPWS11=(2,60),VRPWS12=(2,60),           *
               VR2=5,                                                  *
               VRPWS20=(3,90),VRPWS21=(3,90),VRPWS22=(3,90)
         PATH  DESTSA=10,                                              *
               ER0=(10,1),                                             *
               VR0=0,                                                  *
               VRPWS00=(1,30),VRPWS01=(1,30),VRPWS02=(1,30)
         PATH  DESTSA=28,                                              *
               ER0=(20,1),ER1=(28,1),                                  *
               VR0=1,                                                  *
               VRPWS00=(1,30),VRPWS01=(1,30),VRPWS02=(1,30),           *
               VR1=0,                                                  *
               VRPWS10=(2,60),VRPWS11=(2,60),VRPWS12=(2,60)
```

Figure 256 (Part 2 of 3). NCP Changes - RA6NCA9 and RA5NCA9

```
*---------------------------------------------------------------------
*      CHANNEL ADAPTER LINK (FID4) FOR NCP SA5                        *
* --------------------------------------------------------------------
*
RA5LCHA4 LINE ADDRESS=P01,         CA PHYSICAL POSITION 1              *
               CA=TYPE6,                                               *
               CASDL=120,          INTERVAL BEFORE CHANNEL SLOWDOWN    *
               DELAY=0.0,          CHAN ATTNDELAY                      *
               NCPCA=ACTIVE,       NATIVE SUBCHANNEL (NSC) ACTIVE      *
               TIMEOUT=120         INTERVAL BEFORE CHANNEL DISCONTACT
RA5PCHA4 PU PUTYPE=5,              FID4    CONNECTION                  *
               TGN=1
```

Figure 256 (Part 3 of 3). NCP Changes - RA6NCA9 and RA5NCA9

Figure 257 on page 208 shows the changes made to the NCPs for 6.4, “VR-TG Over a Multilink Transmission Group” on page 96. These comprise both the path tables and the second token-ring LAN link in TG6 between the NCPs. The NCPs are now called RA5NCAA and RA6NCAA.

```
*----------------------------------------------------------------------
*        PATHS - SA6 NCP
*----------------------------------------------------------------------
         PATH  DESTSA=3,                                               *
               ER0=(3,1),ER1=(20,1),ER3=(3,1),                         *
               VR0=0,                                                  *
               VRPWS00=(1,30),VRPWS01=(1,30),VRPWS02=(1,30),           *
               VR1=1,                                                  *
               VRPWS10=(2,60),VRPWS11=(2,60),VRPWS12=(2,60)
         PATH  DESTSA=5,                                               *
               ER0=(5,6,5000,5000,5000,20000),                         *
               ER1=(20,1),ER2=(3,1),ER3=(5,6),                         *
               VR0=0,                                                  *
               VRPWS00=(1,30),VRPWS01=(1,30),VRPWS02=(1,30),           *
               VR1=1,                                                  *
               VRPWS10=(2,60),VRPWS11=(2,60),VRPWS12=(2,60),           *
               VR2=2,                                                  *
               VRPWS20=(3,90),VRPWS21=(3,90),VRPWS22=(3,90)
         PATH  DESTSA=28,                                              *
               ER0=(20,1),ER1=(28,1),ER2=(9,9),                        *
               ER3=(28,1),ER4=(28,1),ER5=(9,6),                        *
               ER6=(9,7),ER7=(5,6),                                    *
               VR0=1,                                                  *
               VRPWS00=(1,30),VRPWS01=(1,30),VRPWS02=(1,30),           *
               VR1=0,                                                  *
               VRPWS10=(2,60),VRPWS11=(2,60),VRPWS12=(2,60),           *
               VR2=5,                                                  *
               VRPWS20=(2,60),VRPWS21=(2,60),VRPWS22=(2,60),           *
               VR3=6,                                                  *
               VRPWS30=(2,60),VRPWS31=(2,60),VRPWS32=(2,60),           *
               VR4=2,                                                  *
               VRPWS40=(2,60),VRPWS41=(2,60),VRPWS42=(2,60),           *
               VR5=7,                                                  *
               VRPWS50=(2,60),VRPWS51=(2,60),VRPWS52=(2,60)
```

Figure 257 (Part 1 of 6). NCP Changes - RA6NCAA and RA5NCAA

```
*----------------------------------------------------------------------
* NTRI PHYSICAL DEFINITIONS - SECOND TIC IN SA6
*----------------------------------------------------------------------
RA6GTRP2 GROUP ECLTYPE=(PHYSICAL,ANY), ALLOW SA,PERIPERAL AND IP      *
               ADAPTER=TIC2,          NTRI INTFACE TYPE USED          *
               SPEED=9600,                                            *
               MAXTSL=2044,           MAX # OF DATA XMIT IN ONE XFER  *
               NPACOLL=NO,                                            *
               PUTYPE=1
*
RA6LA89  LINE ADDRESS=(1089,FULL),                                    *
               LOCADD=400001060001,  LOCAL ADMINISTERED ADDRESS       *
               PORTADD=1,            ASSOCIATION # FOR LOGICAL LINES  *
               RCVBUFC=4095,         MAX # OF DATA RECV IN ONE XFER   *
               TRSPEED=4
*              STATOPT='TIC2 IP/SNA'
RA6PA891 PU ADDR=01,                 SNA NTRI ATTACHMENT ADDR         *
               INNPORT=YES,          USE THIS PORT FOR SNA-INN TRAFFIC *
               NETWORK=SNA,          THIS PU IS USED FOR SNA TRAFFIC  *
               PASSLIM=2,            SERVICE 2 FRAMES BEFORE NEXT PU  *
               ANS=CONTINUE
*              STATOPT='TIC2 SNA'
*
RA6XA890 LU ISTATUS=INACTIVE,                                         *
               LOCADDR=0             ENABLE INDEPENDENT LU USAGE
*
*----------------------------------------------------------------------
* SUBAREA LOGICAL GROUP   NTRI INN LINK        TIC #1 -CCUA          *
*----------------------------------------------------------------------
RA6GTRL3 GROUP ECLTYPE=(LOGICAL,SUBAREA),                             *
               PHYPORT=1,              MAPS PHYSICAL GROUP DEF. ABOVE *
               SDLCST=(RA6SPRI2,RA6SSEC2),                            *
               TYPE=NCP,                                              *
               DIAL=NO,                                               *
               NPACOLL=NO,                                            *
               PUTYPE=4,                                              *
               RETRIES=(6,0,0,6)
*
RA6IA891 LINE TGN=6,                                                  *
               ISTATUS=ACTIVE,                                        *
               UACB=X$L1A
*
RA6UA891 PU ADDR=04400001050001,      MAPS THE OTHER NCP (SA5) TR ADDR *
               MAXDATA=2044,                                          *
               PUTYPE=4
*
```

Figure 257 (Part 2 of 6). NCP Changes - RA6NCAA and RA5NCAA

```
*----------------------------------------------------------------------
*PERIPHERAL LOGICAL GROUP  NTRI TIC 2 - SA6                          *
*----------------------------------------------------------------------
RA6GTRL4 GROUP ECLTYPE=LOGICAL,                                        *
               AUTOGEN=10,                                             *
               CALL=INOUT,                                             *
               PHYPORT=1,NPACOLL=(YES,EXTENDED),                       *
               TYPE=NCP,                                               *
               DIAL=YES,                                               *
               DYNPU=YES,            CREATE SWITCH PU DYNAMICALLY      *
               LINEAUT=YES,                                            *
               PUTYPE=2,                                               *
               XMITDLY=NONE,                                           *
               RETRIES=(6,0,0,6)
```

Figure 257 (Part 3 of 6). NCP Changes - RA6NCAA and RA5NCAA

```
*----------------------------------------------------------------------
*        PATHS - SA5 NCP
*----------------------------------------------------------------------
         PATH  DESTSA=3,                                               *
               ER3=(6,6,5000,5000,5000,20000),                         *
               ER0=(3,1),ER1=(11,1),ER2=(28,1),                        *
               VR0=0,                                                  *
               VRPWS00=(1,30),VRPWS01=(1,30),VRPWS02=(1,30),           *
               VR1=1,                                                  *
               VRPWS10=(2,60),VRPWS11=(2,60),VRPWS12=(2,60),           *
               VR2=2,                                                  *
               VRPWS20=(2,60),VRPWS21=(2,60),VRPWS22=(2,60),           *
               VR3=3,                                                  *
               VRPWS30=(2,60),VRPWS31=(2,60),VRPWS32=(2,60)
         PATH  DESTSA=6,                                               *
               ER0=(20,1),ER1=(6,6),ER5=(28,1),                        *
               ER6=(6,6),                                              *
               VR0=1,                                                  *
               VRPWS00=(1,30),VRPWS01=(1,30),VRPWS02=(1,30),           *
               VR1=0,                                                  *
               VRPWS10=(2,60),VRPWS11=(2,60),VRPWS12=(2,60),           *
               VR2=5,                                                  *
               VRPWS20=(3,90),VRPWS21=(3,90),VRPWS22=(3,90)
         PATH  DESTSA=28,                                              *
               ER0=(20,1),ER1=(28,1),ER7=(28,1),                       *
               VR0=1,                                                  *
               VRPWS00=(1,30),VRPWS01=(1,30),VRPWS02=(1,30),           *
               VR1=0,                                                  *
               VRPWS10=(2,60),VRPWS11=(2,60),VRPWS12=(2,60)
```

Figure 257 (Part 4 of 6). NCP Changes - RA6NCAA and RA5NCAA

```
*----------------------------------------------------------------------
* NTRI PHYSICAL DEFINITIONS - SECOND TIC IN SA5
*----------------------------------------------------------------------
RA5GTRP2 GROUP ECLTYPE=(PHYSICAL,ANY), ALLOW SA,PERIPERAL AND IP      *
               ADAPTER=TIC2,           NTRI INTFACE TYPE USED         *
               SPEED=9600,                                            *
               MAXTSL=2044,            MAX # OF DATA XMIT IN ONE XFER  *
               NPACOLL=NO,                                            *
               PUTYPE=1
*
RA5LA89  LINE ADDRESS=(1089,FULL),                                    *
               LOCADD=400001050001,  LOCAL ADMINISTERED ADDRESS       *
               PORTADD=1,            ASSOCIATION # FOR LOGICAL LINES  *
               RCVBUFC=4095,         MAX # OF DATA RECV IN ONE XFER   *
               TRSPEED=4
*              STATOPT='TIC2 IP/SNA'
RA5PA891 PU ADDR=01,                 SNA NTRI ATTACHMENT ADDR         *
               INNPORT=YES,          USE THIS PORT FOR SNA-INN TRAFFIC *
               NETWORK=SNA,          THIS PU IS USED FOR SNA TRAFFIC  *
               PASSLIM=2,            SERVICE 2 FRAMES BEFORE NEXT PU  *
               ANS=CONTINUE
*              STATOPT='TIC1 SNA'
*
RA5XA890 LU ISTATUS=INACTIVE,                                         *
               LOCADDR=0             ENABLE INDEPENDENT LU USAGE
*              STATOPT='TIC2 IP'
*
*----------------------------------------------------------------------
* SUBAREA LOGICAL GROUP   NTRI INN LINK         TIC #2 -CCUA         *
*----------------------------------------------------------------------
RA5GTRL3 GROUP ECLTYPE=(LOGICAL,SUBAREA),                             *
               PHYPORT=1,              MAPS PHYSICAL GROUP DEF. ABOVE *
               SDLCST=(RA5SPRI2,RA5SSEC2),                            *
               TYPE=NCP,                                              *
               DIAL=NO,                                               *
               NPACOLL=NO,                                            *
               PUTYPE=4,                                              *
               RETRIES=(6,0,0,6)
*
RA5IA891 LINE TGN=6,                                                  *
               ISTATUS=ACTIVE,                                        *
               UACB=X$L1A
*
RA5UA891 PU ADDR=04400001060001,      MAPS THE OTHER NCP (SA6) TR ADDR *
               MAXDATA=2044,                                          *
               PUTYPE=4
*
```

Figure 257 (Part 5 of 6). NCP Changes - RA6NCAA and RA5NCAA

```
*---------------------------------------------------------------------
*PERIPHERAL LOGICAL GROUP  NTRI TIC 2 - SA5                          *
*---------------------------------------------------------------------
RA5GTRL4 GROUP ECLTYPE=LOGICAL,                                        *
               AUTOGEN=10,                                             *
               CALL=INOUT,                                             *
               PHYPORT=1,NPACOLL=(YES,EXTENDED),                       *
               TYPE=NCP,                                               *
               DIAL=YES,                                               *
               DYNPU=YES,             CREATES SWITCH PU DYNAMICALLY    *
               LINEAUT=YES,                                            *
               PUTYPE=2,                                               *
               XMITDLY=NONE,                                           *
               RETRIES=(6,0,0,6)
*              STATOPT='NTRI TIC 2'
```

Figure 257 (Part 6 of 6). NCP Changes - RA6NCAA and RA5NCAA

A.3 3172 XCA Major Node

This is used in all the tests requiring direct access from RAA to the token-ring.

```
**********************************************************************
RAA3172 VBUILD TYPE=XCA
RAA3172P PORT  ADAPNO=0,                                             * X
               CUADDR=EA2,                                           * X
               MEDIUM=RING,                                          * X
               SAPADDR=4,                                            * X
               TIMER=60
**********************************************************************
RAAXCAG1 GROUP DIAL=YES,CALL=INOUT,DYNPU=YES
*
RAAXL1   LINE
RAAXP1   PU
RAAXL2   LINE
RAAXP2   PU
RAAXL3   LINE
RAAXP3   PU
RAAXL4   LINE
RAAXP4   PU
RAAXL5   LINE
RAAXP5   PU
RAAXL6   LINE
RAAXP6   PU
RAAXL7   LINE
RAAXP7   PU
RAAXL8   LINE
RAAXP8   PU
RAAXL9   LINE
RAAXP9   PU
```

Figure 258. RAA3172 - XCA Major Node

A.4 Subarea MPC Links

These major nodes are used in 6.3, "VR-TG for Optimum Routing" on page 88 to connect RAS and RA3 via FID4 links. Contrast them with Figure 3 on page 14 and Figure 4 on page 14 which show a FID2 connection over the same link.

```
RASKRA3 VBUILD TYPE=CA
*       MPC MAJOR NODE FOR RAS - RA3 FID4 CTC CONNECTION
RASG88   GROUP LNCTL=MPC,ISTATUS=ACTIVE,MAXBFRU=3
RASLMP3  LINE  READ=(C08),WRITE=(C10)
RASPMP3  PU    PUTYPE=4,TGN=1


RA3KRAS VBUILD TYPE=CA
*       MPC MAJOR NODE FOR RA3 - RAS FID4 CTC CONNECTION
RA3G88   GROUP LNCTL=MPC,ISTATUS=ACTIVE,MAXBFRU=3
RA3LMPS  LINE  READ=(C10),WRITE=(C08)
RA3PMPS  PU    PUTYPE=4,TGN=1
```

Figure 259. Subarea MPC Connections - RASKRA3 and RA3KRAS

A.5 FID4 Channel Link between VTAM and NCP

This is used to connect RAA as a "data host" to the NCP in subarea 5.

```
RAACRX5 VBUILD TYPE=CA
*      CA MAJOR NODE FOR RAA TO NCP SA5 FID4 LINK
RAAGE28  GROUP LNCTL=NCP,ISTATUS=ACTIVE
RAALE28  LINE ADDRESS=E28,MAXBFRU=34
RAAPE28  PU   PUTYPE=4,MAXDATA=5000,TGN=1
```

Figure 260. CA Major Node - VTAM to NCP (RAACRX5)

A.6 Path Tables

The VTAM path tables used in the VR-TG tests are shown here. Figure 261 on page 214 shows those used in 6.3, "VR-TG for Optimum Routing" on page 88, while Figure 262 on page 216 shows those used in 6.4, "VR-TG Over a Multilink Transmission Group" on page 96. These path tables actually cater for a larger network than we use, so there are paths shown via subareas not connected to our test environment.

```
RAS      VPATH      NETID=USIBMRA
*        PATHS ORIGINATING AT RAS (SUBAREA 28)
         PATH  DESTSA=3,                                                *
               ER0=(3,1),ER1=(20,1),ER2=(3,1),                          *
               VR0=0,                                                   *
               VRPWS00=(1,30),VRPWS01=(1,30),VRPWS02=(1,30),            *
               VR1=1,                                                   *
               VRPWS10=(2,60),VRPWS11=(2,60),VRPWS12=(2,60)
         PATH  DESTSA=5,                                                *
               ER0=(5,1),ER1=(20,1),ER2=(5,1),                          *
               VR0=0,                                                   *
               VRPWS00=(1,30),VRPWS01=(1,30),VRPWS02=(1,30),            *
               VR1=1,                                                   *
               VRPWS10=(2,60),VRPWS11=(2,60),VRPWS12=(2,60)
         PATH  DESTSA=6,                                                *
               ER0=(20,1),ER1=(6,1),ER2=(9,1),                          *
               ER3=(9,1),ER4=(9,1),ER5=(3,1),                           *
               VR0=1,                                                   *
               VRPWS00=(1,30),VRPWS01=(1,30),VRPWS02=(1,30),            *
               VR1=0,                                                   *
               VRPWS10=(2,60),VRPWS11=(2,60),VRPWS12=(2,60),            *
               VR2=3,                                                   *
               VRPWS20=(2,60),VRPWS21=(2,60),VRPWS22=(2,60),            *
               VR3=4,                                                   *
               VRPWS30=(2,60),VRPWS31=(2,60),VRPWS32=(2,60),            *
               VR4=2,                                                   *
               VRPWS40=(2,60),VRPWS41=(2,60),VRPWS42=(2,60)
         PATH  DESTSA=10,                                               *
               ER0=(5,1),                                               *
               VR0=0,                                                   *
               VRPWS00=(2,60),VRPWS01=(2,60),VRPWS02=(2,60)
```

Figure 261 (Part 1 of 3). VTAM Path Tables

```
RA3      VPATH      NETID=USIBMRA
*        PATHS ORIGINATING AT RA3 (SUBAREA 3)
         PATH  DESTSA=5,                                               *
               ER0=(11,1),ER1=(5,1),ER2=(28,1),                        *
               VR0=1,                                                  *
               VRPWS00=(1,30),VRPWS01=(1,30),VRPWS02=(1,30),           *
               VR1=0,                                                  *
               VRPWS10=(2,60),VRPWS11=(2,60),VRPWS12=(2,60),           *
               VR2=2,                                                  *
               VRPWS20=(2,60),VRPWS21=(2,60),VRPWS22=(2,60)
         PATH  DESTSA=6,                                               *
               ER0=(20,1),ER1=(6,1),ER5=(6,1),                         *
               VR0=1,                                                  *
               VRPWS00=(1,30),VRPWS01=(1,30),VRPWS02=(1,30),           *
               VR1=0,                                                  *
               VRPWS10=(2,60),VRPWS11=(2,60),VRPWS12=(2,60)
         PATH  DESTSA=10,                                              *
               ER0=(28,1),                                             *
               VR0=0,                                                  *
               VRPWS00=(3,90),VRPWS01=(3,90),VRPWS02=(3,90)
         PATH  DESTSA=28,                                              *
               ER0=(28,1),ER1=(20,1),                                  *
               VR0=0,                                                  *
               VRPWS00=(1,30),VRPWS01=(1,30),VRPWS02=(1,30),           *
               VR1=1,                                                  *
               VRPWS10=(2,60),VRPWS11=(2,60),VRPWS12=(2,60)
```

Figure 261 (Part 2 of 3). VTAM Path Tables

```
RAA      VPATH      NETID=USIBMRA
*        PATHS ORIGINATING AT RAA (SUBAREA 10)
         PATH  DESTSA=3,                                               *
               ER2=(5,1),                                              *
               VR0=2,                                                  *
               VRPWS00=(3,90),VRPWS01=(3,90),VRPWS02=(3,90)
         PATH  DESTSA=5,                                               *
               ER0=(5,1),                                              *
               VR0=0,                                                  *
               VRPWS00=(1,30),VRPWS01=(1,30),VRPWS02=(1,30)
         PATH  DESTSA=28,                                              *
               ER1=(5,1),                                              *
               VR0=1,                                                  *
               VRPWS00=(2,60),VRPWS01=(2,60),VRPWS02=(2,60)
```

Figure 261 (Part 3 of 3). VTAM Path Tables

```
RAS      VPATH      NETID=USIBMRA
*        PATHS ORIGINATING AT RAS (SUBAREA 28)
         PATH  DESTSA=3,                                                  *
               ER0=(3,1),ER1=(20,1),ER2=(3,1),                            *
               ER3=(5,1),                                                 *
               VR0=0,                                                     *
               VRPWS00=(1,30),VRPWS01=(1,30),VRPWS02=(1,30),              *
               VR1=1,                                                     *
               VRPWS10=(2,60),VRPWS11=(2,60),VRPWS12=(2,60),              *
               VR2=3,                                                     *
               VRPWS20=(3,90),VRPWS21=(3,90),VRPWS22=(3,90)
         PATH  DESTSA=5,                                                  *
               ER0=(5,1),ER1=(20,1),ER2=(5,1),                            *
               VR0=0,                                                     *
               VRPWS00=(1,30),VRPWS01=(1,30),VRPWS02=(1,30),              *
               VR1=1,                                                     *
               VRPWS10=(2,60),VRPWS11=(2,60),VRPWS12=(2,60)
         PATH  DESTSA=6,                                                  *
               ER0=(20,1),ER1=(6,1),ER2=(9,1),                            *
               ER3=(9,1),ER4=(9,1),ER5=(3,1),                             *
               ER6=(5,1),                                                 *
               VR0=1,                                                     *
               VRPWS00=(1,30),VRPWS01=(1,30),VRPWS02=(1,30),              *
               VR1=0,                                                     *
               VRPWS10=(2,60),VRPWS11=(2,60),VRPWS12=(2,60),              *
               VR2=3,                                                     *
               VRPWS20=(2,60),VRPWS21=(2,60),VRPWS22=(2,60),              *
               VR3=4,                                                     *
               VRPWS30=(2,60),VRPWS31=(2,60),VRPWS32=(2,60),              *
               VR4=2,                                                     *
               VRPWS40=(2,60),VRPWS41=(2,60),VRPWS42=(2,60),              *
               VR5=6,                                                     *
               VRPWS50=(2,60),VRPWS51=(2,60),VRPWS52=(2,60)
```

Figure 262 (Part 1 of 2). VTAM Path Tables

```
RA3      VPATH      NETID=USIBMRA
*        PATHS ORIGINATING AT RA3 (SUBAREA 3)
         PATH  DESTSA=5,                                           *
               ER0=(11,1),ER1=(5,1),ER2=(28,1),                    *
               ER3=(6,1),                                          *
               VR0=1,                                              *
               VRPWS00=(1,30),VRPWS01=(1,30),VRPWS02=(1,30),       *
               VR1=0,                                              *
               VRPWS10=(2,60),VRPWS11=(2,60),VRPWS12=(2,60),       *
               VR2=2,                                              *
               VRPWS20=(2,60),VRPWS21=(2,60),VRPWS22=(2,60),       *
               VR3=3,                                              *
               VRPWS30=(2,60),VRPWS31=(2,60),VRPWS32=(2,60)
         PATH  DESTSA=6,                                           *
               ER0=(20,1),ER1=(6,1),ER5=(6,1),                     *
               VR0=1,                                              *
               VRPWS00=(1,30),VRPWS01=(1,30),VRPWS02=(1,30),       *
               VR1=0,                                              *
               VRPWS10=(2,60),VRPWS11=(2,60),VRPWS12=(2,60)
         PATH  DESTSA=28,                                          *
               ER0=(28,1),ER1=(20,1),ER7=(6,1),                    *
               VR0=0,                                              *
               VRPWS00=(1,30),VRPWS01=(1,30),VRPWS02=(1,30),       *
               VR1=1,                                              *
               VRPWS10=(2,60),VRPWS11=(2,60),VRPWS12=(2,60),       *
               VR2=7,                                              *
               VRPWS20=(3,90),VRPWS21=(3,90),VRPWS22=(3,90)
```

Figure 262 (Part 2 of 2). VTAM Path Tables

List of Abbreviations

ACB	Access Method Control Block
ACF	Advanced Communication Function
AHHC	APPN Host-to-Host Connection
API	Application Programming Interface
APPC	Advanced Program-to-Program Communication
APPN	Advanced Peer-to-Peer Networking*
BF-TG	Boundary Function - Transmission Group
BSC	Binary Synchronous Communication
CA	Channel Adapter
CDRM	Cross-Domain Resource Manager
CDRSC	Cross-Domain Resource
CDS	Central Directory Server
CICS	Customer Information Control System
CM	Communications Manager
CMC	Communications Management Configuration
CNN	Composite Network Node
COS	Class Of Service
CP	Control Point
CSMA/CD	Carrier Sensed Multiple Access / Collision Detect
CTC	Channel-To-Channel
DLC	Data Link Control
DLU	Destination Logical Unit
DLUR	Dependent LU Requester
DLUS	Dependent LU Server
DLUR/S	Dependent LU Requester / Server
DR	Dynamic Reconfiguration
ER	Explicit Route
ESA	Enterprise Systems Architecture
ESCON	Enterprise Systems Connection
FDDI	Fibre Distributed Data Interface
FID	Format Identifier
FQPCID	Fully Qualified Procedure Correlation Identifier
FRR	Functional Recovery Routine
GDS	Generalized Data Stream
IBM	International Business Machines Corporation
IC-TG	Interchange Transmission Group
ICP	Interconnect Controller Program
IPX	Internet Packet Exchange
ITSO	International Technical Support Organization
LAN	Local Area Network
LEN	Low-Entry Networking
LIC	Licensed Internal Code
LSAP	Link-level Service Access Point
LU	Logical Unit
MAC	Media Access Control
MPC	Multipath Channel
MVS	Multiple Virtual Storage
NAU	Network accessible unit
NCP	Network Control Program
NETBIOS	Network Basic Input Output System
NLDM	Network Logical Data Manager
NMVT	Network Management Vector Transport
NNS	Network Node Server
NPSI	NCP Packet Switching Interface
OLU	Origin Logical Unit
PC	Personal Computer
PIU	Path Information Unit
PLU	Primary Logical Unit
PU	Physical Unit
RH	Request Header, or Response Header

RSCV	Route Selection Control Vector
RSN	Resource Selector Node
RSN	Resource Sequence Number
RU	Request Unit
SLU	Secondary Logical Unit
SNA	Systems Network Architecture
SNI	SNA Network Interconnection
SNVC	Subnet Visit Count
SSCP	System Services Control Point
TCP/IP	Transmission Control Protocol / Internet Protocol
TG	Transmission Group
TIC	Token-ring Interface Coupler
TP	Transmission Priority
TRLE	Transport Resource List Entry
TSO	Time Sharing Option
USS	Unformatted System Services
VM	Virtual Machine
VR	Virtual Route
VR-TG	Virtual Route-Based Transmission Group
VRN	Virtual Routing Node
VTAM	Virtual Telecommunications Access Method
WAN	Wide Area Network
XCA	External Communications Adapter
XID	Exchange Identification
XRF	Extended Recovery Facility

Index

Numerics

A

B

C

D

E

F

G

I

L

ITSO Technical Bulletin Evaluation RED000

VTAM V4R2 Early User Experiences

Publication No. GG24-4250-00

Your feedback is very important to help us maintain the quality of ITSO Bulletins. **Please fill out this questionnaire and return it using one of the following methods:**

- Mail it to the address on the back (postage paid in U.S. only)
- Give it to an IBM marketing representative for mailing
- Fax it to: Your International Access Code + 1 914 432 8246
- Send a note to REDBOOK@VNET.IBM.COM

Please rate on a scale of 1 to 5 the subjects below.
(1 = very good, 2 = good, 3 = average, 4 = poor, 5 = very poor)

Overall Satisfaction ____

Organization of the book	____	Grammar/punctuation/spelling	____
Accuracy of the information	____	Ease of reading and understanding	____
Relevance of the information	____	Ease of finding information	____
Completeness of the information	____	Level of technical detail	____
Value of illustrations	____	Print quality	____

Please answer the following questions:

a) If you are an employee of IBM or its subsidiaries:

Do you provide billable services for 20% or more of your time? Yes____ No____

Are you in a Services Organization? Yes____ No____

b) Are you working in the USA? Yes____ No____

c) Was the Bulletin published in time for your needs? Yes____ No____

d) Did this Bulletin meet your needs? Yes____ No____

If no, please explain:

What other topics would you like to see in this Bulletin?

What other Technical Bulletins would you like to see published?

Comments/Suggestions: (THANK YOU FOR YOUR FEEDBACK!)

Name

Address

Company or Organization

Phone No.

ITSO Technical Bulletin Evaluation **RED000**

GG24-4250-00

Cut or Fold
Along Line

Fold and Tape **Please do not staple** Fold and Tape

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

IBM International Technical Support Organization
Department 985, Building 657
P.O. BOX 12195
RESEARCH TRIANGLE PARK NC
USA 27709-2195

Fold and Tape **Please do not staple** Fold and Tape

Cut or Fold
Along Line

Printed in U.S.A.

GG24-4250-00