IBM Network Utility
Description and Configuration Scenarios

Tim Kearby, Peter Gayek, Gallus Schlegel, Imre Szabo, Zhi-Yong Zhang

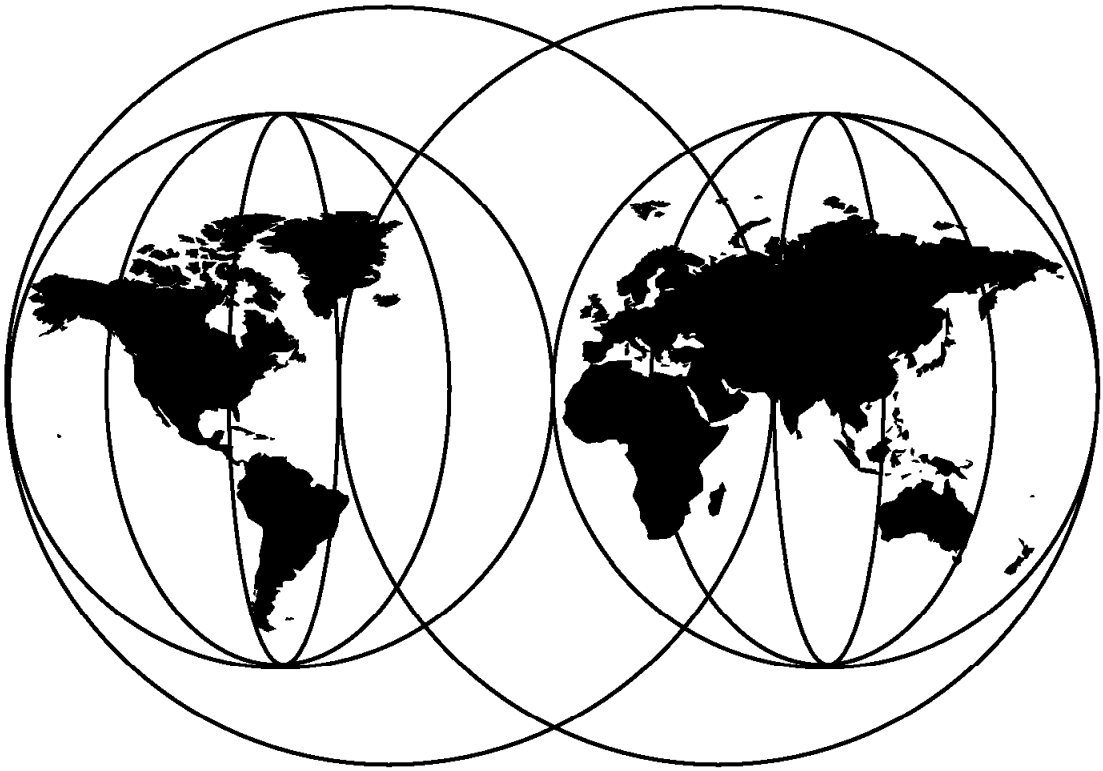

International Technical Support Organization

http://www.redbooks.ibm.com

This book was printed at 240 dpi (dots per inch). The final production redbook with the RED cover will be printed at 1200 dpi and will provide superior graphics resolution. Please see "How to Get ITSO Redbooks" at the back of this book for ordering instructions.

SG24-5289-00

International Technical Support Organization

SG24-5289-00

IBM Network Utility
Description and Configuration Scenarios

January 1999

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix C, “Special Notices” on page 193.

First Edition (January 1999)

This edition applies to the IBM Network Utility Models TN1 and TX1 at Code Level B.

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

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

Contents

Preface

The accelerating growth in networked users in today's e-business environments is driving the need for higher processing capacity in routers as they perform memory-intensive tasks such as IP-SNA integration, TN3270E Server functions, Data Link Switching, Enterprise Extender services, and/or APPN high-performance routing. IBM's new Network Utility addresses this need for higher processing capacity and increased numbers of logical connections.

This redbook provides networking professionals the information they need to quickly understand the functions of this new product so that they can begin exploiting its advanced features. It provides detailed explanations and example scenarios covering the hardware platform and the functions of the IBM Network Utility Code Levels A and B.

The example scenarios, written jointly by the International Technical Support Organization (ITSO) and the IBM Networking Hardware Division (NHD), show how to deploy the Network Utility in more than a dozen applications where the power of the product can be leveraged. These scenarios include using the Network Utility as a TN3270E server using a variety of host connections including both SNA subarea and APPN, host channel gateway configurations (both ESCON and parallel channel) for both TCP/IP and SNA, and how the Network Utility can be used as a peer router for Data Link Switching (DLSw) circuit termination.

The Highly Available Fault-Tolerant TN3270E scenario shows you a solution for scaling a TN3270 configuration in very large environments using the Network Dispatcher function to dynamically load-balance telnet sessions between multiple Network Utilities.

These scenarios provide step-by-step configuration procedures using both the command-line interface and the Configuration Program Graphical User Interface. An appendix provides detailed examples of the corresponding host definitions that are required for the Network Utility applications discussed in the scenarios.

The Team That Wrote This Redbook

This redbook was produced by a team of specialists from around the world working at the International Technical Support Organization, Raleigh Center.

Tim Kearby is an Advisory ITSO Specialist for Networking at the Systems Management and Networking ITSO Center, Raleigh. He writes redbooks and teaches workshops on local and wide area networking. Tim has held various positions in his IBM career including assignments in product development, systems engineering, and consulting. He holds a Bachelors of Science degree in Electrical Engineering from Purdue University.

Peter Gayek is a Senior Software Engineer in the IBM Networking Hardware Division in Research Triangle Park, NC. He has 17 years of experience in the development of Networking products. His areas of expertise include DLSw, APPN, SNA/IP interworking, and data link layer protocols. He holds a B.S. degree in Computer Science from Cornell University.

Gallus Schlegel is a Senior Networking Specialist in Switzerland. He has over 15 years of experience in networking. His areas of expertise include design and

implementation of LAN and WAN networks. He has written several Redbooks in this area. He holds a degree in Electronic Engineering from the ATZ in Zuerich, Switzerland.

Imre Szabo is a Staff Information Developer in the IBM Networking Hardware Division in Research Triangle Park, NC. He has 15 years of experience in Information Development and technical writing on networking products. His areas of expertise include routers, controllers, and hubs.

Zhi-Yong Zhang is a Advisory I/T Specialist in the People′s Republic of China. He has 5 years of experience in the networking field. His areas of expertise include Wide Area and ATM Campus networks. He holds both a Bachelor of Science and a Master′s degrees in Electrical Engineering & Electronics from Fudan University in Shanghai.

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

Jerzy Buczak, Bob Haimowitz, Karl Wozabal,
Gail Christensen, Shawn Walsh,
Kathryn Casamento, Linda Robinson, Mike Haley
International Technical Support Organization, Raleigh Center.

Jim Goethals
David Heath
Kevin McClain
Pat McClellan
John Averi
Dave Johnson
Renee Kovales
Howard Pearse
Jerry Sents
Dave Swingle
Jan Maher
Walt Wheeler
Gee Chia
Joe Czap
Jeff Brodd
Wayne Taylor
Jason Cornpropst
Jimmie Thomas
Jon Houghton
Jay Miller
Dave Connors
John Yarbrough
Tommy Johnson
IBM NHD
Research Triangle Park, NC

Rebecca Williams
IBM Design Center
Research Triangle Park, NC

Jimmy Weatherspoon
IBM NTSO
Gaithersburg, MD

Uttam Deedwania
IBM NSD
Research Triangle Park, NC

Comments Welcome

Your comments are important to us!

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

- Fax the evaluation form found in “ITSO Redbook Evaluation” on page 201 to the fax number shown on the form.
- Use the online evaluation form found at http://www.redbooks.ibm.com/
- Send your comments in an Internet note to redbook@us.ibm.com

Chapter 1. Introduction

To remain competitive, network users are extending their traditional internal SNA and IP networks outward to business partners, dealers, suppliers, and e-business customers. In this expanding environment, users are also searching for ways to save money and provide connectivity between their mix of SNA and TCP/IP server applications and their TCP/IP and SNA desktop client population.

In this complex, mixed environment, this significant growth in user connections requires increased bandwidth on physical links and, more important, tremendous numbers of logical connections. Traditional router and controller solutions attempt to address this requirement for high logical connectivity with multiples of products built for high-density physical connectivity and just aren't cost-effective in today's exploding logical connectivity environment.

Figure 1. The IBM Network Utility

The IBM Network Utility, shown in Figure 1, provides a simple, scalable solution for medium to large networks in this environment. It provides excellent logical connection scalability. With up to 512 MB of memory, the Network Utility is engineered to scale memory-intensive applications such as TN3270, DLSw, and APPN DLUR.

There are two models of the Network Utility:

- **Network Utility Transport (Model TX1)**

  This model is geared towards the efficient movement of network traffic with high processing requirements through concentration points. It supports both Enterprise Extender (also called HPR over IP) and Data Link Switching (DLSw) to provide traffic prioritization and non-disruptive sessions for excellent availability of SNA traffic over IP networks. It Supports High Performance Routing (HPR) in APPN networks for improved routing performance and reliability. It also provides a high-performance, low-cost channel-attached gateway. The Network Dispatcher function in the Network Utility can be used to distribute IP-based connections to multiple servers.

The Network Utility allows you to expand your IP network without a costly forklift upgrade.

- **Network Utility TN3270E Server (Model TN1)**

  This model is focused on TCP/IP connectivity to SNA mainframe hosts. It is designed for customers with SNA applications that must accommodate increasing IP user connections. It includes all of the functions of the Transport model plus a TN3270E server that can support up to 16,000 user sessions through a single machine.

  The Network Utility offers scalable TN3270 capacity by allowing you to cost-effectively add additional Network Utilities. As your network grows and exceeds the capacity of a single server, you can add another Network Utility and use the Network Dispatcher function in a Network Utility Model TN1 or TX1 (or alternatively, in the 2216 Multiaccess Connector, the 2210 Nways Multiprotocol Router, or the 3746 Multiaccess Enclosure) to balance the load among the TN3270E Servers. It is a simple, scalable and cost-effective way for you to increase the number of user connections without disrupting your data center operations.

  The Network Utility includes a Network Dispatcher Advisor for TN3270 to maximize the efficiency of multiple Network Utility TN3270E Servers.

The Network Utility hardware and operating code are combined into a single package with separately priced and orderable adapter features.

The Network Utility offers a wide range of network connectivity choices. You can choose from among the following network connections:

- LANs: token-ring, 10/100 Ethernet, or FDDI
- Channels: ESCON or Parallel
- ATM: 155 Mbps multimode or single-mode fiber
- WANs: V.24, V.35, V.36, X.21, or HSSI

Because the Network Utility models are custom-built for specific networking functions, they are easy to install, with many parameters pre-configured.

The Network Utility is an outstanding product for customers needing:

- A high-capacity, high-performance, easy-to-use TN3270E server with industry-leading cost per session
- A high-capacity, high-performance, easy-to-use SNA/IP transport device in either a pure SNA or a mixed SNA/IP environment employing Data Link Switching, APPN, or Enterprise Extender
- A high-performance, single-channel S/390 or S/370 host gateway
- A high-performance, high-speed media attach device
- High-performance load-balancing for TCP/IP servers

1.1 Hardware Overview

The hardware platform for both models of the Network Utility is the same. It is based on the IBM 2216 Model 400. It includes:

- A PowerPC processor with 256 MB (optionally 512 MB) of memory
- Two adapter slots that can be populated with most of the adapters supported by the 2216-400, including:

- LAN support for token-ring, 10/100 Ethernet, and FDDI
- ATM support for 155 Mbps OC3 multimode fiber and single-mode fiber
- Channel support for ESCON and Parallel Channels
- WAN support for V.24, V.35, V.36, X.21, and HSSI

The Network Utility is compact in size which allows for scalability through easy rack mount stacking.

For more information on the Network Utility hardware, please see Chapter 2, “The Network Utility Hardware Description” on page 7.

1.2 Software Highlights

The Network Utility provides a comprehensive set of IP and SNA routing protocols and other functions. This broad base of multiprotocol routing and transport code is based on the Multiprotocol Access Services Version 3.2 code for the 2216-400. This support includes:

- Data Link Control support for PPP, frame relay, X.25, and SDLC
- Support for LSA, LCS, and HPDT MPC (MPC+) Channel Access Methods
- Bridging, IP, DLSw, APPN, and Enterprise Extender networking protocols

In addition, the code has been tailored to provide specific networking functions unique to each model. Many of the parameters for these functions have been pre-configured. The next two sections provide the highlights for each model.

Important Note

The descriptions of code functions in this redbook pertain to the Network Utility Transport and TN3270 Code Level B support. If you purchased a Network Utility at Code Level A, you can obtain a no-charge upgrade for your model from the Network Utility Web site at URL:

`http://www.networking.ibm.com/support/downloads/networkutility`

1.2.1 The Network Utility Transport (Model TX1)

The Network Utility Transport offers:

- Code tailored to IP and SNA Transport
- Support for up to 15,000 DLSw circuits
- Support for over 16,000 LU sessions with APPN DLUR
- A single-channel gateway to S/390 servers (ESCON or Parallel Channel)
- Low-cost and high-capacity (10,000 connections/second) load balancing
- IP Passthru over the channel for simplified 3172 replacement
- Enhanced Network Dispatcher with advisors for better load balancing of news, mail, and telnet servers and with SNMP for management

The Network Utility Transport model provides tremendous value if you require:

- A cost-effective channel gateway. The Network Utility makes an ideal channel gateway in environments needing few physical connections as in, for example, a single-channel to single-LAN gateway. For high-availability

environments, multiple single-channel Network Utilities can be coupled with either VIPA support in the host for IP environments or APPN HPR and VTAM MultiNode Persistent Sessions (MNPS) for SNA and TN3270 environments.

- Load-balancing traffic among multiple TCP/IP LAN or channel-attached servers.
- High-capacity DLSw or APPN DLUR. If you have a site that requires 1,000 or more circuits to be supported by DLSw or 1,000 or more LU-LU sessions to be supported by APPN DLUR, you should consider using the Network Utility Transport instead of using multiple lower-capacity boxes.
- Network replacement of a 3720 or 3725. If you are looking to replace your 3720 or 3725 Communication Controller, you should consider an IBM Network Utility for those environments where you no longer require NCP-specific functions like NPSI, SNI, or EP. The Network Utility fits well here if you are using parallel channels with a small number of SDLC lines. (Use a 2216 to replace 3720s and 3725s with more connections.) The Network Utility can use DLSw, APPN, or Enterprise Extender to transport your SNA traffic and it also offers very competitive IP support.

Note: The following NCP-specific functions are not supported by the Network Utility:

- NCP Packet Switching Interface (NPSI) for non-QLLC protocols

  (QLLC is supported by the Network Utility.)
- SNI Back-to-Back

  (SNI Single Gateway is supported. Remote NCPs can be connected to the Network Utility via SDLC link, token-ring, QLLC, or frame relay.)
- Emulation program (EP)
- Network terminal option (NTO)
- OSI protocols
- Airlines Line Control (ALC) protocol
- Network Routing Facility (NRF)

1.2.2 The Network Utility TN3270E Server (Model TN1)

The Network Utility TN3270E Server offers:

- Code tailored for TN3270
- Support for 16,000 TN3270E sessions with a low cost per session
- Exploitation of the Network Dispatcher Advisor for TN3270 to optimize the efficiency of balancing traffic among multiple TN3270 servers
- An IP passthru mode over ESCON and Parallel Channel that allows for replacement of 3172s as IP host gateways without the added configuration complexity of IP routing
- LU pools for easier TN3270 configuration and higher availability
- TN3270 LU to IP address mapping for improved administrative control over application access
- SDDLU for dynamic creation of TN3270 LU definitions in VTAM
- Combined support for Network Dispatcher and TN3270 Server

The Network Utility TN3270E Server model provides tremendous value if you require:

- Incremental TN3270 server capacity to complement 2216 Multiaccess Connector Model 400s or 3746 Multiaccess Enclosures (MAEs). If your TN3270 session requirements would require multiple 2216s or MAEs, you should consider using the Network Utility for additional capacity. It is easy to install, and has competitive TN3270 Server function, an industry-leading price per session, and low cost for incremental capacity.
- Incremental TN3270 server capacity to complement other TN3270 servers. No matter what your current TN3270 server is, when you need additional TN3270 capacity, you should consider using the Network Utility.
- Incremental TN3270 server capacity to complement IBM 3745 NCP subarea users moving to IP desktops. Use existing 3745 or 3746 channels to carry traffic from Network Utilities, saving the additional cost of new channel connections.
- Incremental TN3270 server capacity to complement other channel gateways. No matter what your current channel gateway is, when you need additional TN3270 capacity, you should consider using the Network Utility TN3270E Server.
- Consolidation of multiple lower-capacity TN3270 servers into a higher-capacity, high-performance Network Utility.

For more information on the Network Utility software, please see Chapter 3, “Software Overview” on page 19.

Chapter 2. The Network Utility Hardware Description

This chapter provides an overview of the IBM Network Utility hardware and consolidates some of the information from the product publications that ship with the unit. For a complete reference on the Network Utility hardware, please see the following manuals:

- *2216 Nways Multiaccess Connector and Network Utility Introduction and Planning Guide*, GA27-4105
- *2216 Nways Multiaccess Connector and Network Utility Service and Maintenance Manual*, SY27-0350
- *IBM 2216 Nways Multiaccess Connector Description and Configuration Scenarios - Volume I*, SG24-4957

2.1 Base Hardware

The Network Utility is currently available in two models:

- Network Utility Transport (Model TX1)
- Network Utility TN3270E Server (Model TN1)

The base hardware is identical for both models and consists of the following:

- A 17.3″ x 18″ x 5.2″ (3U) chassis which may be placed either on a table top or installed in a rack.
- A system backplane
- A system processor card containing:
  - 233 MHz PowerPC 740 processor
  - Two Dynamic Random Access Memory (DRAM) Dual In-Line Memory Module (DIMM) slots
  - 512 KB L2 cache
  - 8 KB Non-Volatile RAM (NVRAM) for Vital Product Data (VPD)
- Hard disk
- Two LAN/WAN/Channel adapter slots
- Two PCMCIA slots for type II or Type I/II adapters
- Universal automatic voltage sensing power supply with an input range from 100VAC to 240VAC
- Two cooling fans

Note: The specifications above for the processor card describe the Network Utility Release 2 hardware. The Network Utility Release 1 hardware used the same 200 MHz processor card as the 2216-400. (However, the card shipped with the Network Utility has a slight modification that allows it to be populated with two 256 MB DIMMs if required.)

The processor card has two DIMM slots and ships with the first slot populated with one 256 MB DIMM. The second DIMM slot, if populated, must use a DIMM of the same capacity as the one in the first slot. This gives the unit a total RAM capacity of 512 MB.

Note: The feature number for the additional 256 MB DIMM is 2522.

Important Note

The Network Utility release 1 hardware shipped with two 128 MB DIMMs. If you have a release 1 Network Utility and require an upgrade to 512 MB memory you can upgrade the two 128 MB DIMMs to two 256 MB DIMMs by ordering feature number 2525. The price to upgrade to 512MB is the same for both Release 1 and Release 2 Network Utility.

The Network Utility has two PCMCIA adapter slots. These slots are type II (or I/II) PCMCIA slots. One slot is populated with an IBM Etherjet 10 Mbps Ethernet adapter.

Note: The PCMCIA Ethernet card is for service and operations purposes, such as providing a user console and transferring files. It cannot be used as a normal network routing interface.

In most countries, the other PCMCIA slot ships with a 33.6 Kbps V.34 data modem. For North American customers, the modem has an integrated DAA. For most other countries, the modem uses a programmable DAA. (See 2.4.3, “PCMCIA Modem” on page 18 for more information on using the PCMCIA modem.)

There is one fixed power subsystem on the Network Utility with a single power cord. The unit has no power ON/OFF switch.

2.2 Network Adapters

The Network Utility uses the same LAN/WAN/Channel adapters that are used on the 2216-400. However, not all the 2216 adapters are supported at this time. This section lists the adapters that are supported for the Network Utility.

Unlike the 2216-400, there are no plugging rules for the adapters in an Network Utility. (Any supported adapter may be placed in either slot, regardless of the adapter type in the other slot.)

Take Note

The adapters on the Network Utility are hot pluggable. However, you must disable all the ports before unplugging the adapter. If an adapter is removed prior to disabling the ports, a machine check can occur and the Network Utility may become idle. You can disable all the ports at the same time from the talk 5 + prompt by issuing the `disable slot #` command where # is the slot number for the adapter.

After replacing the adapter, you can issue the `enable slot #` command to re-enable the adapter. This command performs a test of each interface on the adapter and then brings them on line (if the tests pass).

The system card is *not* hot pluggable.

2.2.1 LAN Adapters

The following section lists the LAN adapters that are available for the Network Utility. The feature code for each adapter is shown in parentheses.

2.2.1.1 2-Port Token-Ring (FC 2280)

Provides for two attachments to token-ring LANs. This adapter can continually process frames of data to and from system memory and the token-ring at a speed of either 4 Mbps or 16 Mbps. It supports the use of either shielded twisted-pair or unshielded twisted-pair cable through a single connector. Cable 2713 is available for this adapter.

2.2.1.2 2-Port Ethernet (FC 2281)

Provides for two attachments to Ethernet LANs. It supports the use of either 10BASE-T cable or 10BASE2 cable. Cable 2713 is available for this adapter.

2.2.1.3 1-Port 10/100-Mbps Ethernet (FC 2288)

This adapter provides one port for 10/100 Mbps Ethernet connections using an RJ-45 connector. This attachment provides:

- 4-KB entry hardware Transparent Bridging (to support bridging at the media speed)
- IEEE 802.3 10-Mbps Ethernet
- IEEE 802.3u 100-Mbps Ethernet

Cable 2713 is available for this adapter.

2.2.1.4 1-Port FDDI (FC 2286)

This adapter provides one attachment to an FDDI dual-attaching system (DAS) connection. This attachment provides:

- One FDDI DAS port on the card
- 4-KB entry hardware Transparent Bridging (to support bridging at media speed)
- Networking support for the following Fiber Distributed Data Interface (FDDI) standards:
  - *FDDI Part 1: Token Ring Physical Layer Protocol*, ISO 9314-1, 1989
  - *FDDI Part 2: Token Ring Media Access Control (MAC)*, ISO 9314-2, 1989
  - *FDDI Part 3: Token Ring Physical Layer Medium Dependent (PMD)*, ISO/IEC 9314-3, 1990
- External cabling:
  - Fiber cabling standard: ANSI X3.T9.5
- FDDI DAS port media connectors: Multimode fiber SC
- FDDI DAS Optical Bypass Connector

Note: A cable is not provided with the Network Utility for this adapter. The IBM fiber optic cables that can be used with this adapter include the 4 m cable PN19G4864 and the 6 m cable PN19G4865.

2.2.2 WAN Adapters

The following section lists the Wide Area Network (WAN) adapters that are available for the Network Utility. The feature code for each adapter is shown in parentheses.

2.2.2.1 8-Port EIA-232E/V.24 Adapter (FC 2282)

Provides eight attachments to EIA-232E/V.24 WANs. Each attachment provides:

- Support for receiving clocking (modem-attached) at a line speed from 2.4 Kbps to 64 Kbps
- Support for providing clock (directly attached) from 9.6 Kbps to 64 Kbps
- Software selectable to receive clock (modem-attached) or provide clock (directly attached) with the appropriate cable
- A 100-pin D-shell female connector
- Support for cables FC 2701, FC 2705, and FC 2706

2.2.2.2 6-Port V.35/V.36 Adapter (FC 2290)

Provides six attachments to ITU-T V.35 or V.36 WANs. Each attachment provides:

- Support for receiving clocking (modem-attached) at a line speed from 2.4 Kbps to 2.048 Mbps
- Support for providing clock (directly attached) from 9.6 Kbps to 460.8 Kbps as well as 1.544 Mbps and 2.048 Mbps
- Software selectable to receive clock (modem-attached) or provide clock (directly attached) with the appropriate cable
- A 100-pin D-shell female connector
- Support for cables FC 2702, FC 2703, FC 2707, FC 2708, FC 2709 FC 2710, and FC 2799

With the V.35/V.36 you can use the following combination of cables: FC 2702 with FC 2707, FC 2708, and FC 2799 (for V.35) FC 2703 with FC 2709 and FC 2710 (for V.36). (See 2.3, “Cables” on page 13.)

Note: Cabling to mix V.35 and V.36 interfaces on adapter FC 2290 is not provided.

2.2.2.3 8-Port X.21 Adapter (FC 2291)

Provides eight attachments to ITU-T X.21 WANs. Each attachment provides:

- Support for receiving clocking (modem-attached) at a line speed from 2.4 Kbps to 2.048 Mbps
- Support for providing clock (directly attached) from 9.6 Kbps to 460.8 Kbps as well as 1.544 Mbps and 2.048 Mbps
- Software selectable to receive clock (modem-attached) or provide clock (directly attached) with the appropriate cable
- A 100-pin D-shell female connector
- Support for cables FC 2704, FC 2711, and FC 2712

2.2.2.4 1-Port High-Speed Serial Interface HSSI Adapter (FC 2289)

Provides one port for HSSI connection through a 5-m (16 ft, 5 in.) STP cable that is included with the adapter. The cable uses 2-row 50-pin SCSI connectors. Two types of HSSI cables can be used with the HSSI adapter: One cable for DTE-to-DCE attachment where the HSSI adapter is the DTE, and one cable for back-to-back direct connection of one HSSI adapter to another HSSI adapter. (The second cable is called a null-modem or cross-over cable). The HSSI attachment provides:

- Support for network standards:
  - *High-speed Serial Interface for Data Terminal Equipment and Data Circuit-Terminating Equipment*, ANSI/EIA/TIA 613-1993
  - *Electrical Characteristics for an Interface at Data Signaling Rates up to 52 Mbps*, ANSI/EIA/TIA 612-1993
- DTE interface for applications
- DCE interface for test
- Support for STP cable type
- Receive clocking at line speeds of 1.544 Mbps (T1 speed) to 52 Mbps
- Clocking at the following speeds:
  - 44.736 Mbps
  - 22.368 Mbps

2.2.3 ATM Adapters

The following section lists the ATM adapters that are available for the Network Utility. The feature code for each adapter is shown in parentheses.

Note

The Network Utility does not support the Fast Token-Ring (FasTR) function for the ATM adapters like the 2216-400.

2.2.3.1 1-Port ATM 155-Mbps Multimode Fiber Adapter (FC 2294)

This adapter provides one attachment to an ATM switch device over a multimode fiber-optic cable. This attachment provides:

- 8 MB of packet memory and 2 MB of control memory for high-performance support
- A specialized ATM support chip to perform the segmentation and reassembly (SAR) function for ATM Adaptation Layer 5 (AAL-5)
- SONET OC3c framing
- Nominal operating wavelength of 1300 nm using LED-based technology
- Support for a 62.5/125 micron multimode fiber
- Transceiver support for a maximum cable length of 2 km (1.24 miles)
- A multimode duplex SC connector

Notes:

1. This is the newer high-performance ATM adapter. The original Multi-Mode ATM adapter (FC 2284) is not supported in the Network Utility.
2. A cable is not provided with the Network Utility for this adapter. The IBM fiber optic cables that can be used with this adapter include the 4-m cable PN 19G4864 and the 6-m cable PN 19G4865.

2.2.3.2 1-Port ATM 155-Mbps Single-Mode Fiber Adapter (FC 2295)

This adapter provides one attachment to an ATM switch over a multimode fiber-optic cable. This attachment provides:

- 8 MB of packet memory and 2 MB of control memory for high-performance support
- A specialized ATM support chip to perform the segmentation and reassembly (SAR) function for ATM Adaptation Layer 5 (AAL-5)
- SONET OC3c framing
- Nominal operating wavelength of 1310 nm using laser-based technology
- Support for a 9/125 micron single-mode fiber
- Transceiver support for a maximum cable length of 20 km (12.4 miles)
- A single-mode polarized duplex SC connector

Notes:

1. This is the newer high-performance ATM adapter. The original Single-Mode ATM adapter (FC 2293) is not supported in the Network Utility.
2. A cable is not provided with the Network Utility for this adapter. The IBM fiber optic cables that can be used with this adapter include the 4-m cable PN 19G4864 and the 6-m cable PN 19G4865.

2.2.4 Channel Adapters

The following section lists the host channel adapters that are available for the Network Utility. The feature code for each adapter is shown in parentheses.

2.2.4.1 1-Port ESCON Channel Adapter (FC 2287)

Provides the Network Utility with access to SNA and TCP/IP host applications from LANs, WANs, and ATM over a duplex-to-duplex multimode fiber-optic cable.

The Network Utility supports up to 32 ESCON logical addresses (subchannels) per adapter for access to up to 32 hosts for Link Services Architecture (LSA) or 16 hosts for LAN Channel Station (LCS) or 16 hosts for MPC+ (assuming these types are not mixed on the adapter) when used with an IBM 9032 or 9033 ESCON Director or access to up to 15 logical host images in EMIF-capable processors operating in a logically partitioned (LPAR) mode.

Please see 3.2.12, “Channel Access Methods” on page 28 for more information on the channel access methods supported by this adapter.

2.2.4.2 1-Port Parallel Channel Adapter (FC 2299)

Provides an attachment to a mainframe through the Original Equipment Manufacturer's Information (OEMI) parallel channel interface. Gives the Network Utility access to SNA and TCP/IP host applications from LANs, WANs, and ATM with the Original Equipment Manufacturer's Information (OEMI) parallel channel interface.

The parallel channel attachment:

- Operates in direct-coupled interlock (DCI) or data-streaming mode.
- Supports data streaming rates of 3.0 and 4.5 MB per second.
- Supports extensions via connections to the 3044 Fiber-Optic Channel Extender Link Models C02 and D02.

The Parallel Channel Adapter supports up to 32 subchannels and 16 virtual net handlers per Parallel Channel Adapter.

Please see 3.2.12, "Channel Access Methods" on page 28 for more information on the channel access methods supported by this adapter.

2.3 Cables

The cables used in the Network Utility are the same as those used for the 2216. The cables for the supported adapters are listed here for your convenience.

2.3.1 WAN Cables

EIA-232E/V.24 Fanout Cable (FC 2701): This fanout cable provides eight connections (25-pin D-shell male) each of which is 1.8 m (5 ft 11 in.) in length. Each connection is suitable for connection to a EIA-232/V.24 modem. Cables 2705 and 2706 are available to complement this cable. FC 2705 attaches to the FC 2701 and extends the cable length an additional 3 m (9 ft 10 in.) for attachment to a modem. FC 2706 attaches to FC 2701 and provides clocking to allow devices to be directly attached to the Network Utility without having to use a modem. It is 3 m (9 ft 10 in.) in length and provides a female 25-pin D-shell connector.

V.35 Fanout Cable (FC 2702): This fanout cable is a 1.2-m (3 ft 10 in.) cable to a distribution box containing six 25-pin D-shell male connectors. Cables 2707 and 2708 are available to complement this cable. FC 2707 provides a 3-m (9 ft 10 in.) extension cable with a 34-pin male block connector for attachment to a modem. FC 2708 provides a 2-m (6 ft 7 in.) cable with a 34-pin female block connector for direct-device attachment.

V.36 Fanout Cable (FC 2703): This fanout cable provides six connections (37-pin D-shell male) each of which is 3 m (9 ft 10 in.) in length. Each connection is suitable for connection to a V.36 modem. Cables 2709 and 2710 are available to complement this cable. FC 2709 provides a 3-m (9 ft 10 in.) extension cable with a 37-pin male D-shell connector for attachment to a modem. FC 2710 provides a 3-m (9 ft 10 in.) cable with a 37-pin female D-shell connector for direct-device attachment.

X.21 Fanout Cable (FC 2704): This fanout cable provides eight connections (15-pin D-shell male) each of which is 1.8 m (5.9 ft) in length. Each connection is suitable for connection to an X.21 modem. Cables 2711 and 2712 are available to

complement this cable. FC 2711 provides a 3-m (9 ft 10 in.) extension cable with a 15-pin D-shell male connector for attachment to a modem. FC 2712 provides a 3-m (9 ft 10 in.) cable with a 15-pin D-shell female connector for direct-device attachment.

EIA-232E/V.24 Serial Interface Cable (FC 2705): This cable provides a 3-m (9 ft 10 in.) extension cable with a 25-pin male D-shell connector for attachment to a modem.

EIA-232E/V.24 Direct Attach Cable (FC 2706): This cable is a 3-m (9 ft 10 in.) cable with a 25-pin D-shell female connector for direct-device attachment.

V.35 Serial Interface Cable (FC 2707): This cable is a 3-m (9 ft 10 in.) extension cable with a 34-pin male block connector for attachment to a modem.

V.35 Direct Attach Cable (FC 2708): This cable is a 2-m (6.6-ft) cable with a 34-pin female block connector for direct-device attachment.

V.36 Serial Interface Cable (FC 2709): This cable is a 3 m (9 ft 10 in.) extension cable with a 37-pin male D-shell connector for attachment to a modem.

V.36 Direct Attach Cable (FC 2710): This cable is a 3-m (9 ft 10 in.) cable with a 37-pin female D-shell connector for direct-device attachment.

X.21 Serial Interface Cable (FC 2711): This cable is a 3-m (9 ft 10 in.) extension cable with a 15-pin D-shell male connector for attachment to a modem.

X.21 Direct Attach Cable (FC 2712): This cable is a 3-m (9 ft 10 in.) cable with a 15-pin D-shell female connector for direct-device attachment.

Attachment Cable for V.35 DCE (FC 2799) - for France: This cable is a 30-cm (1-ft) cable that adapts the standard V.35 34-pin male block connector to the connector required for attachment to V.35 modems in France.

2.3.2 Host Channel Cables

The sections below list the cables that are available for both the ESCON adapter and the Parallel Channel Adapter as well as some restrictions for their use. for more detailed information concerning the use of these cables, please consult the *IBM 2216 Nways Multiaccess Connector and Network Utility Introduction and Planning Guide*, GA27-4105.

2.3.2.1 Cables for the ESCON Channel Adapter

The cable for each Network Utility ESCON Channel Adapter must be ordered separately. In the U.S.A., Canada, and Latin America, order cable group #3797 or PN 14F3797 in the following standard lengths:

- 4 m (12 ft)
- 7 m (20 ft)
- 13 m (40 ft)
- 22 m (70 ft)
- 31 m (100 ft)
- 46 m (150 ft)
- 61 m (200 ft)
- 77 m (250 ft)
- 92 m (300 ft)
- 107 m (350 ft)

- 122 m (400 ft)

For additional information on planning for ESCON Channel Adapter cable installations, refer to *Fiber Optic Link Planning*, GA23-0367.

In EMEA or AP, see ESCON Cabling Information in publication number GC22-7064 to order the cables.

2.3.2.2 Cables for the Parallel Channel Adapter

The following sections describe the necessary cabling for the Parallel Channel Adapter.

V-Cable (Ships with the Parallel Channel Adapter): The Parallel Channel Adapter always requires a V-Cable that ships with the adapter.

Channel Interface-In Cable (FC 2720): You may or may not need a Channel Interface-In Cable, depending on your configuration. Use this cable if:

- The control unit immediately upstream (that is, toward the host) is not a 3172, 2216 Model 400, or Network Utility.

  or

- The control unit immediately upstream is a 3172, 2216 Model 400, or Network Utility but is too far away to cable directly to the Network Utility via V-cable.

Note: Order channel interface cables only if there is a need to connect directly to the host channel cabling. A host connection can often be achieved by cabling the adapter to some device already attached to the host channel. Such a connection is made using the V-cable that ships with the adapter.

Channel Interface-Out Cable (FC 2721): You may or may not need a Channel Interface-Out Cable, depending on your configuration. Use a channel interface-out cable if:

- The control unit immediately downstream (that is, away from the host) is not a 3172, 2216 Model 400, or Network Utility.

  Note: This is a valid configuration only at 3.0 MB per second.

  or

- The control unit immediately downstream is a 3172, 2216 Model 400, or Network Utility but is too far away to cable directly to the Network Utility via V-cable.[1]

Terminator: The S/370 parallel channel architecture requires that last device on the bus terminate the signals. The terminator that ships with the Network Utility Parallel Channel Adapter is used for this purpose if the adapter is the last device on the bus. The terminator attaches to the V-Cable on the Channel Interface-Out connector and terminates both the bus and the tag signals on the Parallel Channel.

Host Channel Cables: Host channel cables are the main cables running from the mainframe along the length of a channel. You will need to order host channel cables if those you are already using are not long enough to support adding a Parallel Channel Adapter.

---

[1] Or, in the case of a channel interface-out cable, to a standard bus-and-tag terminator.

3172 Cables: Parallel channel function in the Network Utility offers a migration path for 3172 users. However, the 3172 Interconnect Controller V-Cable is not functionally equivalent to the Network Utility V-Cable. Do not re-use the 3172 cable with the Network Utility.

2.3.2.3 Cabling Restrictions

All cables must conform to the restrictions below:

- Up to six Parallel Channel Adapters are configurable per host channel.
- Maximum instantaneous data rate of 4.5 MB per second is supported for the Parallel Channel Adapter when data streaming on a 400 ft (122 m) cable.
- Single or multiple Network Utilities (or 2216 Model 400 or 3172 units) attaching via parallel channel must be physically located at the end of a channel with no other machine types downstream.

  Note: This is true if the channel is operating in 4.5 MBps data streaming mode. If the channel adapter is operating in 3.0 MBps data streaming mode, then the Network Utility in which that adapter is seated can be located anywhere on the channel. (Remember that all devices on a channel must be configured for the same channel speed).
- The parallel channel-attached Network Utility (or 2216-400) can be extended using the 3044 Fiber-Optic Channel Extender Link Models C02 and D02 or 9034 Enterprise Systems Connection Converter extender products. However, only one parallel-channel-attached device of any type (2216 Model 400, Network Utility, 3172, or other) that attaches via twisted pair connection cables can be extended per channel using a 3044 or 9034 device.

2.3.3 Other Cables

Multipurpose RJ-45 Adapter Cable (FC 2713): This cable is a 7.6-m (25-ft) Category 5 cable with an RJ-45 connector for attachment to token-ring hubs or switches or Ethernet 10BASE-T hubs or switches.

Terminal Attachment Cable (PN 10H5569): This cable provides a 3-m (9.8 ft), 9-pin serial to 25-pin connection. Two of these cables ship as part of the base package of each Network Utility and is intended for attaching the machine to a modem or a null modem block that attaches to an ASCII terminal. When attaching to laptop computers or other machines that do not carry 25-pin ports, the second cable can be matched to the modem or null modem block to provide a 9-pin connection to the ASCII terminal. The null modem block also is included as part of the ship group as PN 10H5570.

RJ-45 Cable (No Feature Code): The Ethernet card ships with a 3 m, 10BASE-T cable with a male RJ-45 connector. This adapter-cable set is part of the base package of the Network Utility.

2.3.4 Cables That IBM Does Not Offer

The following cables are not provided as options for the Network Utility. You must obtain them, if they are required:

- Token-ring STP network adapter cable
- Ethernet 10BASE2 cable
- ATM multimode fiber adapter cable
- ATM single-mode fiber adapter cable

- FDDI cable

2.4 Controls and Indicators

This section describes the controls and indicators on the front of the Network Utility

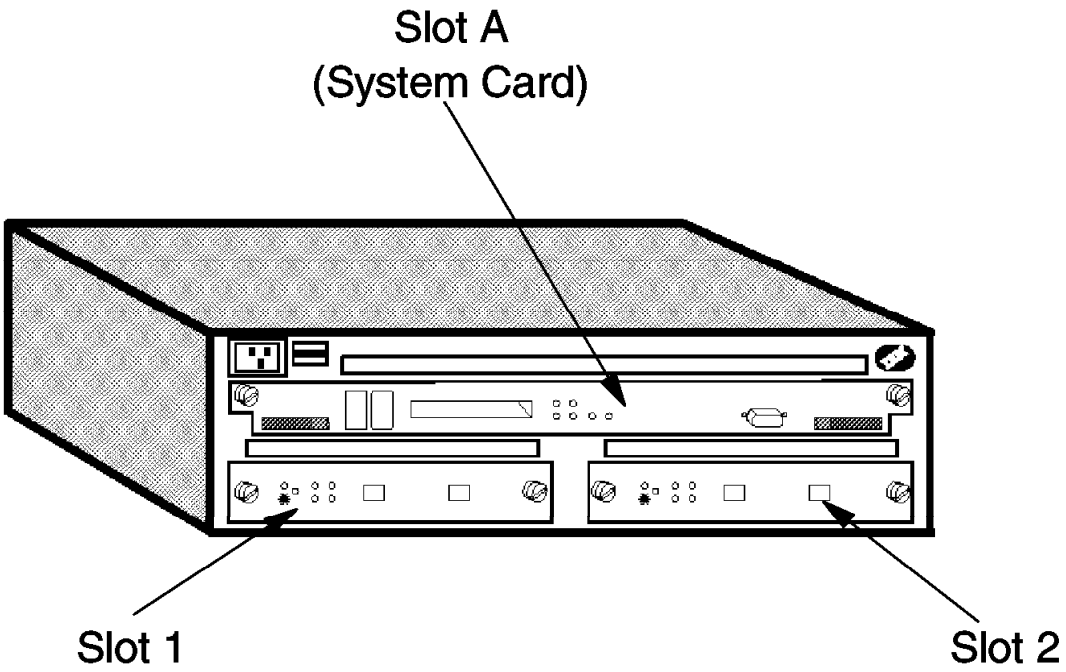

Figure 2. IBM Network Utility - Front View

Figure 2 shows the front view of the Network Utility and the slot numbering scheme. The electronic components of the Network Utility are accessed from the front of the unit. This allows for quick and easy installation and upgrades and also aids in problem determination.

2.4.1 Power

The Network Utility does not have a on/off switch. However, it does have a reset button on the CPU card.

After power-on-reset (POR), either by plugging the unit into an A/C power source or by pressing the reset button, the green and yellow LEDs will remain on until successful completion of the power-on tests (which will be under two minutes). The green LED will begin blinking, which indicates that the code is being loaded into memory. The green LED will be switched on to indicate that the system code is operational.

2.4.2 Service Port Default Settings

The default settings for the Network Utility service port are:

- Speed: 19.2 Kbps
- Parity: None
- Data bits: 8
- Stop bits: 1
- Terminal type: VT220, Monochrome

To change the settings for the serial port, follow these steps:

1. Reboot the Network Utility to the firmware main menu. To do this:
   - Perform a hardware reset or a software reload of the unit.
   - When you see the message `Starting Boot Sequence` followed by `Strike F1 key now to prematurely terminate Boot`, press `Ctrl+c` or the `F1` key immediately. Continue to hold Ctrl+c or the F1 key until you see the firmware main menu or the prompt for a supervisory password.

     Note: To make sure you do not miss this message, you can start holding down Ctrl+c at any time after the start of system board diagnostics.
2. Select Option 1, **Manage Configuration**.
3. Move the cursor to the row for the COM1 serial port and press Enter.
4. Move the cursor to the characteristic you want to change (for example, baud rate), and press Enter.
5. Select the new value and press Enter.
6. Press Esc to return to the firmware main menu.
7. If you want to continue the current boot sequence and have the operational code start using the new settings, press F9 (Start OS). If you want to reboot into the firmware and have the firmware start using the new settings, press F3 (Reboot).
8. Change the settings of your terminal or terminal emulation software to match the new settings of the Network Utility serial port.

2.4.3 PCMCIA Modem

The PCMCIA modem is a standard item that is shipped with the Network Utility in most countries. It is a 33.6 Kbps V.34 data modem, and it negotiates the data rate to be used between it and the partner modem on the other side of the telephone network. Using data compression, this modem is capable of data throughput greater than 33.6 Kbps.

The data rate between the Network Utility system and its PCMCIA modem defaults to 19.2 Kbps, but you can raise it to accommodate the higher throughput that the two modems may be able to achieve between themselves. For example, you may want to set this rate to 57.6 Kbps so that it is higher than the effective data rate of two 33.6 Kbps modems running data compression. If your modems are both faster than 19.2 Kbps, raising this rate will lower Xmodem file transfer time.

To change the data rate and any of the other settings for the PCMCIA modem, follow the same procedure given above for serial port settings, but select COM2, the PCMCIA modem, instead of the serial port.

Chapter 3. Software Overview

The Network Utility runs a modified version of Nways Multiprotocol Access Services (MAS), the software for the IBM 2216 Nways Multiaccess Connector. For the Network Utility, the code has been customized into two distinct software loads that are each bundled with the base hardware platform to create the two models, the TN1 and TX1. The software loads are referred to as the Transport and the TN3270 and support the models TX1 and the TN1, respectively.

Each software load has been tailored for the functions for which the model has been intended. For example, in the TN3270 model, the focus is on delivering cost-effective, high-capacity, and easy-to-use TN3270E server support.

However, since the Network Utility is based upon the same common code as MAS, it adopts most of the functions available in Version 3.2 of MAS. While this redbook focuses on the functions pertaining to the Transport and TN3270 functions, we provide here a brief overview of the entire Network Utility software function. For more detailed information on these functions, please reference the following MAS manuals:

- *Nways Multiprotocol Access Services Software User's Guide Version 3.2*, SC30-3886
- *Nways Multiprotocol Access Services Using and Configuring Features Version 3.2*, SC30-3993
- *Nways Multiprotocol Access Services Protocol Configuration and Monitoring Reference Vol. 1 Version 3.2*, SC30-3884
- *Nways Multiprotocol Access Services Protocol Configuration and Monitoring Reference Vol. 2 Version 3.2*, SC30-3885

3.1 Highlights

The Network Utility provides a comprehensive set of IP and SNA routing protocols and other functions. Highlights of the functions provided include:

- Routing Protocols:
  - TCP/IP with RIP, RIPv2, OSPFv2, MOSPF, DVMRPv3, BGP-4
- SNA Data Transport:
  - APPN NN, ISR, HPR, DLUR, Branch Extender, Enterprise Extender, and Extended Border Node
  - DLSw (RFC 2166) including NetBIOS support
  - SDLC, both primary and secondary capabilities
  - LAN Network Manager (LNM) support
- TN3270E server support (Model TN1 only)
- Network Dispatcher support for IP Server load balancing
  - Network Dispatcher Advisor for TN3270 to optimally balance Network Utility TN3270E Servers
- Bridging (SR, TB, SRT, SR-TB, IP Bridging Tunnel)
- Asynchronous Transfer Mode (ATM) Adapter support:

  - ATM Forum Compliant LAN Emulation
  - Classic IP
  - Native HPR over ATM
- WAN Data Link Controls (DLCs):
  - Frame relay (RFC 1490), including Boundary Access Node (BAN) support
  - Point to Point Protocol (PPP)
  - X.25, including DTE Transport (XTP) to ″carry″ X.25 traffic over FR or PPP connections and including QLLC for SNA
  - SDLC, both primary and secondary capabilities
- Security Features:
  - PAP/CHAP/RADIUS/TACACS/TACACS+ for PPP
  - TACACS+/RADIUS for telnet access to the Network Utility
  - Extensive filtering support for Bridging, DLSw, IP and IPX
- Bandwidth Reservation System (BRS) over FR and PPP DLCs Configuration Program GUI application for use on Windows, AIX, or OS/2 workstations

3.2 Functional Description

The following sections provide an overview of the Network Utility software and are organized as follows:

- Data Link Controls (DLCs) including:
  - Token-ring
  - Ethernet (10Mbps and 100Mbps)
  - FDDI
  - Point to Point Protocol (PPP)
  - Frame relay including frame relay boundary access node (BAN)
  - X.25
  - SDLC including SDLC Relay
  - ATM
- Channel Access Methods
  - High Performance Data Transfer (HPDT) MultiPath Channel (MPC) (MPC+)
  - LAN Channel Station (LCS)
  - Link Services Architecture (LSA)
- Network Protocols including:
  - Bridging including IP Bridging Tunnel
  - IPv4
  - IPv6
  - ARP
  - DLSw

  - APPN (including HPR, DLUR, Branch Extender, Enterprise Extender and Extended Border Node)
  - LAN emulation client
- Bandwidth Reservation System (BRS)
- Interactive Network Dispatcher
- TN3270E server
- Network Management including LAN Network Manager (LNM)
- Configuration Program

3.2.1 Token-Ring

The base token-ring support provides:

- IEEE 802.5/ISO 8802.5 support
- Support for 4Mbps or 16Mbps operation
- Token-ring MIB - RFC 1748

3.2.2 Ethernet

The base 10Mbps Ethernet support provides:

- Ethernet V2 or IEEE 802.3/ISO 8802.3 support
- Ethernet MIB - RFC 1650
- Ethernet locally administered MAC address can be configured to override the default burned-in address.

The base 100Mbps Ethernet support provides:

- IEEE 802.3 or IEEE 802.3u support
- Ethernet MIB - RFC 1650
- Network speed of 10Mbps or 100Mbps
- Ethernet locally administered MAC address can be configured to override the default burned-in address.

3.2.3 FDDI

The base FDDI support provides:

- ISO 9314-1, 9314-2, and 9314-3 support
- FDDI MIB - RFC 1512

3.2.4 PPP

Point to Point Protocol (PPP) is an RFC standards-based method to encapsulate various Network Layer protocols (such as IPv4 and APPN HPR) and multiplex them over a point-to-point WAN link. It is supported over all the available serial interfaces.

The PPP support is based on the following:

- RFC 1144 VJ Compressed TCP/IP Headers
- RFC 1220 PPP Extensions for Bridging
- RFC 1332 PPP IP Control Protocol

- RFC 1334 PPP Authentication Protocol
- RFC 1471 MIB objects for PPP LCP—PPP Link Group MIB
- RFC 1548 Point-to-Point Protocol
- RFC 1549 PPP in HDLC Framing
- RFC 1570 PPP LCP Extensions
- RFC 1638 Bridging Control Protocol (BCP) Interface Types
- RFC 1661 The Point-to-Point Protocol (PPP) (Obsoletes RFC 1548)
- RFC 1662 PPP in HDLC-like Framing (Obsoletes RFC 1549) (LQM option not supported)
- RFC 1962 The PPP Compression Control Protocol (CCP)
- RFC 1974 PPP Stac LZS Compression Protocol
- RFC 1994 CHAP - PPP Challenge Handshake Authentication Protocol
- RFC 2043 PPP SNA Control Protocol (SNACP)
- RFC 2063 PPP IPv6 Control Protocol
- RFC 2097 The PPP NetBIOS Frames Control Protocol (NBFCP)
- RFC 2118 Microsoft Point-To-Point Compression (MPPC) Protocol
- Bandwidth Reservation (BRS)
- Data Compression over PPP

  The IETF Compression Control Protocol (CCP) is used to negotiate compression algorithms and parameters. The compression algorithms supported are Stac LZS and Microsoft Point-to-Point Compression (MPPC). PPP data compression is negotiated by PPP at link open time; the algorithm used is established on a per-interface basis to allow for control of the (substantial) memory usage of compression dictionaries.

  PPP data compression can be used over any supported PPP interface and can be used at the same time as Bandwidth Reservation (BRS). BRS will operate on data before compression is applied. When compression is in use, all data that passes over the interface is compressed. The impact of attempting to compress already compressed traffic varies according to the algorithm in use.

  The compression achievable varies greatly according to traffic. Using the Calgary Corpus standard of binaries, text files and image files, the Stac LZS algorithm achieves a ratio of about 2:1. In general, data compression is much more effective when used in conjunction with lower speed WAN interfaces, for example less than 64 Kbps. As you increase the WAN data rate, the effective improvement in the overall bandwidth utilization diminishes.

  To analyze the specific compression being achieved, console commands are provided to display recent compression statistics.

- PAP/CHAP security for PPP links

  Provides PAP (Password Authentication Protocol) and CHAP (Challenge-Handshake Authentication Protocol) password security for authentication in bringing up PPP links. These links can be leased lines. A password for the link must be configured on each side of the link. This password is checked by the other side before allowing the link to come up.

This provides an added level of security for WAN links. For example, on a circuit without PAP/CHAP, someone could potentially masquerade as a trusted router, receiving data of a sensitive nature. With PAP/CHAP, the PPP connection will not come up unless the caller also knows what password has been configured.

- Authentication servers, TACACS, TACACS+, and RADIUS, can be used so that names and passwords need not be configured at each router. In addition, the remote authentication protocols TACACS+ and RADIUS support authorization and accounting. The code allows unique authentication, authorization and accounting servers to be specified that apply to all the PPP connections for the box. Each server may also have a unique backup server.
- PPP user names up to 64 characters in length are supported.

3.2.5 Frame Relay

The base frame relay support provides the capability to function as a frame relay DTE and therefore connect to a frame relay switch, or frame handler. The function supports:

- Frame relay architecture as defined in:
  - Frame relay architecture framework as defined in ANSI T1.606-1988 and ITU-T I.233.1
  - Frame relay core aspects as defined in ANSI T1.618-1991 and ITU-T Q.922 Annex A
  - Frame relay data link support as defined in ITU-T Q.922.
- DTE support

  DTE-to-DTE (that is, back to back) connectivity is supported to allow DTEs to be directly connected without an intermediate frame relay switch.
- 2-byte DLCIs
- Local management interface (LMI) as defined in:
  - ANSI T1.617 Annex D
  - CCITT Q.933 Annex A
  - LMI Revision 1
- Multiprotocol encapsulation as defined in:
  - RFC 1490
  - ANSI T1.617 Annex F
  - Frame Relay Forum Multiprotocol Encapsulation Implementation (MEI) agreement
- Inverse Address Resolution Protocol (InARP) (RFC 1293) is used to dynamically resolve DLCI to protocol address mappings for IP.
- Congestion control

  When the box receives a frame containing BECN (Backward Explicit Congestion Notification), it is the box′s responsibility to throttle down the PVC′s VIR (variable information rate) if either CIR monitoring or congestion monitoring is enabled. The box does this gradually as it receives consecutive frames with BECN until either the minimum IR is reached or a

frame without BECN arrives. As the box receives consecutive frames without BECN, the VIR gradually rises to the maximum IR.

Congestion can also be reduced by throttling transmission upon receipt of FECN and notifying the FECN source of congestion if the frame relay switch sends FECNs but not BECNs.

Other congestion notification mechanisms include CLLM (Consolidated Link Layer Management) messages and SNMP traps sent on receipt of CLLM, FECN, or BECN frames. CLLM messages come from the frame relay switch and the SNMP traps can be used to notify the Network Management station that congestion is occurring. CLLM support is based on:

- ANSI T1.618
- ITU-T Q.922
- ITU-T X.36

Discard eligibility (DE) can be configured on a DLCI basis for frames matching a BRS protocol or filter criteria to aid in congestion control.

- Bandwidth reservation (BRS)
- Frame relay DTE MIB (RFC 1315)
- Frame relay SVC MIB based on an internet draft
- Frame relay data compression is configurable per PVC or SVC to run over a frame relay interface and is based on Frame Relay Forum Data Compression Implementation Agreement (FRF.9). The STAC LZS algorithm is used.

Frame relay fragmentation/reassembly is not supported.

3.2.6 Frame Relay SVC

SVCs conserve network bandwidth and reduce network cost and complexity. SVCs are supported over leased access facilities for IP, Bridging, and DLSw. APPN support is provided by Enterprise Extender which uses IP transport. Multiprotocol support is provided by multiplexing over a single SVC using RFC 1490. Multiple SVCs between two boxes is supported. Bandwidth Reservation (BRS) is supported using either a configuration for a specific SVC or using a default traffic class definition used by other circuits on the interface. Any combination of PVCs and SVCs can be used on the same interface.

The SVC support conforms to the following standards:

- Frame Relay Forum FR SVC User to Network Implementation Agreement (FRF.4)
- FR SVC signalling layer as defined in ITU-T Q.933
- ISDN signalling layer in ITU-T Q.931
- FR data link layer in ITU-T Q.932

3.2.7 Frame Relay BAN

This function enables SNA PU 2.0, Type 2.1, and PU 4 endstations connected to the box to make a direct connection via frame relay to a front-end processor (FEP), such as the IBM 3745 Communication Controller attached to an IBM mainframe, an AS/400, or an APPN network node. Type 4 BAN traffic can be received only by an IBM 3746/900 Communication Controller. When using BAN, endstations function as if they are directly connected to an FEP or AS/400 via a

token-ring, Ethernet, or SDLC line as appropriate. Though their data actually passes through the box and over a frame relay network, this is transparent to the endstations.

The BAN feature works by modifying each frame received from an end station to use the RFC 1490 bridged token-ring frame format. It then transmits the modified frames on a BAN frame relay circuit. The BAN feature also filters out traffic, allowing only frames using the BAN circuit MAC address to pass over a BAN circuit to the mainframe.

There are two ways to use the BAN feature:

- Straight bridging using the Nways Multiprotocol Access Services bridging capability, where you configure BAN to bridge LLC2 frames from endstations straight into the NCP. The code does not terminate the LLC2 traffic received from attached endstations. Instead, the code converts the frames it receives to bridged token-ring format (RFC 1490) frames and bridges directly to the NCP. Thus, the box acts as a bridge between the NCP and the endstations. Endstation frames can be token-ring or Ethernet format, provided the bridge is configured to support that type of frame. SDLC endstations are supported via DLSw.
- DLSw terminated, in which BAN terminates the LLC2 connection at the box running DLSw. The DLSw box does not function as a bridge. The box terminates the LLC2 traffic received from attached endstations. At the same time, the box establishes a new LLC2 connection to the NCP over the frame relay network. Thus, though two LLC2 connections exist within the transaction, the break between them is transparent both to the NCP and the endstations. The result is a virtual LLC2 connection between NCP and the endstations.

Note: BAN support in NCP was introduced in NCP Version 7 Release 3. Users running with NCP Version 7 Release 1 or 2 will require the appropriate NCP PTF to obtain BAN support. Refer to NCP Sales Manual for more information. BAN support for INN traffic was introduced in NCP Version 7 Release 5, and is applicable only to the 3746 Model 900.

3.2.8 X.25

This code provides the function for connection to an X.25 network or a remote X.25 DTE. The X.25 support is provided on all serial adapters. The X.25 code contains the following major functions:

- QLLC - for SNA and DLSw natively
- LAP/B - X.25 link layer
- PLC - X.25 packet layer
- RFC 1356 - Multi Protocol Interconnect for X.25 (MPI)
- RFC 1256 - X.121 DTE address conversion

3.2.8.1 X.25 Base Functions

- Mod 8 or 128 both LAP/B and PLC
- PVC and SVC (IC,TW,OG) support
- 1980, 1984, 1988, and 1992 support
- Default parameters for LAP/B and PLC

- Default profile for X.25
- Department of Defense X.25 support
- X.121 to IP address conversion
- ISO 7776 and 8208 compatible
- SVC flow control negotiation
- PLC M-bit support
- DTE/DTE or DTE/DCE operation
- Configuration for most of the above functions
- On-Demand SVC call setup

3.2.8.2 X.25 VC Scalability

- Up to 4095 SVCs and/or PVCs per interface.
- SVC scalability based on available memory and the capacity of the X.25 link. Supports over 1000 SVCs depending on configuration.
- PVC scalability up to 2500 PVCs per interface.

3.2.8.3 QLLC Support

QLLC provides the Physical Services Header function for SNA entities which attach to an X.25 network. These functions are normally provided by SDLC or LLC layers. These services consist of link layer setup and reset, FRMR, XID, and TEST. QLLC also provides timeout backup for these services. QLLC acts as a primary or secondary station, or as a PEER station.

3.2.8.4 X.25 Closed User Group

Closed User Group facilities allow customers to form groups of X.25 DTEs that will send or receive traffic only to other specific groups of DTEs. The following capabilities are supported:

- Basic format (indices 0-99) and extended format (indices 100-9999) CUG facilities
- Receipt and transmission of the Closed User Group Selection facility
- Receipt and transmission of the Closed User Group with Outgoing Access Selection facility
- Receipt and transmission of the Bilateral Closed User Group Selection facility
- Receipt and transmission of the Bilateral Closed User Group with Outgoing Access Selection facility

3.2.8.5 X.25 Transport over TCP (XTP)

The X.25 Transport over TCP support is a method of transporting X.25 protocol traffic over an IP backbone, allowing customers with X.25 devices to migrate from an X.25 network (private or public) to the network of their choice while retaining their investment in the X.25 devices. X.25 devices can directly attach via HDLC lines to the box and this code transports the X.25 packet data over TCP/IP. The code uses a spoofing technique similar to the method that Data Link Switching (DLSw) uses for transporting SNA and NetBIOS over TCP/IP. The spoofing is done at the X.25 packet layer and is transparent to any logical link control (LLC) above the packet layer. TCP/IP has been chosen as the transport protocol because it guarantees the delivery of packets and maintains their order.

The connections between boxes can be composed of any media that can route IP traffic.

It is also possible to provide X.25 Transport over TCP within a single device. This allows X.25 devices attached to the same or different interfaces on a single Network Utility to communicate.

X.25 DTEs (data terminal equipment) can be directly connected to the box over one of its serial interfaces. In this configuration the box maintains the role of the DCE (data circuit-terminating equipment). A locally attached DTE establishes a virtual circuit to the box. The box locates a remote box which is attached to a remote DTE with which the locally attached DTE wishes to communicate. The remote box establishes a virtual circuit to the remote DTE so that both DTEs can communicate. Both Permanent Virtual Circuits (PVCs) and Switched Virtual Circuits (SVCs) are supported. The virtual circuit type on each side of the X.25 network must be the same.

3.2.9 SDLC

Synchronous Data Link Control (SDLC) support allows communication between SNA/SDLC endstations and the DLSw or APPN function in the box. This support includes the following:

- Primary link station support:
  - Operates in Two-Way Alternating (TWA) mode
  - Modulo 8 or 128
  - Support for the following physical and logical line types:
    - Physical point-to-point lines on non-switched, full-duplex or half-duplex facilities
    - Physical multipoint lines on non-switched, full duplex facilities
    - Logical multipoint lines. There is a logical multipoint line if the secondary station is operating as a gateway (that is, it is receiving/transmitting frames to/from the primary station for its station address as well as the station address for multiple downstream secondary stations).
- Secondary link station support

  With SDLC Secondary Support, the box can assume a secondary link station role and support the attachment of primary SDLC stations. Both the box and the respective SDLC primary station, for example a 3745 or AS/400, must have the links configured for two-way alternating mode. The physical connection can be non-switched on full or half-duplex facilities.

  SDLC secondary support for APPN and DLSw allows the box to represent multiple secondary link stations to a directly attached SDLC primary device. DLSw uses this ″logical multipoint″ capability to represent multiple endstations in the DLSw network. Each device supported on the DLSw network (attached to this or another DLSw router via any DLC type) appears to the primary SDLC device as a single secondary SDLC device on a multipoint line. To reduce non-productive idle polling in this configuration, SDLC secondary support includes the IBM 3174 Group Poll function. Support for multiple SNA PU types

  Depending on whether DLSw is using SDLC, attached devices may be SNA PU types 2.0, 2.1, 4 (for subarea INN traffic), or 4/5 (a host possibly coupled

with an FEP, running peripheral BNN traffic). SDLC itself is not sensitive to the PU type of the attached device, and allows multiple PU types attached to the same physical or logical multipoint link. For PU types 2.1 and 4 on point-to-point lines, SDLC allows full role negotiation and automatic address discovery.

3.2.10 SDLC Relay

SDLC Relay is a mechanism to connect two disjoint HDLC links, typically SDLC, without regard to the media types in between. It uses a form of IP encapsulation to pass the SNA traffic through the network so that the two SDLC link stations think they are directly connected to each other. It does not have the functionality associated with DLSw such as link acknowledgements and firewalling. Consequently, the RR polls and acknowledgements will flow from end to end through the network. If there is congestion in the network, sessions may time out. Also note this requires another Nways Multiprotocol Access Services, Nways Multiprotocol Routing Services, or IBM Nways Multiprotocol Routing Network Services box as one of the encapsulation endpoints, although it may pass through IP routers in the middle of the network. The link stations at both ends of the SNA session must be connected via SDLC and point-to-point, non-switched, full-duplex facilities to use the SDLC Relay function.

Configuration support for this capability is only provided via the command line interface.

3.2.11 ATM

The base ATM support provides:

- ATM Forum UNI 3.0 and 3.1
- ATM Forum Interim Local Management Interface (ILMI)
- PVC and SVC
- RFC 1483 encapsulation for IP and Bridging
- Standard ATM MIB - RFC 1695
- APPN/HPR (natively)
- All the supported routed protocols and native ATM bridging may be multiplexed onto a single ATM permanent virtual circuit.

3.2.12 Channel Access Methods

Multiprotocol Access Services provides High Performance Data Transfer (HPDT) MultiPath Channel (MPC) (also referred to as MPC+) for VTAM/SNA, HPDT UDP and TCP/IP host applications. Link Services Architecture (LSA) support is provided for VTAM/SNA host applications. LAN Channel Station (LCS) support is also provided for TCP/IP host applications. Up to 32 logical addresses (subchannels) are supported per ESCON or Parallel Channel Adapter.

3.2.12.1 HPDT MPC (MPC+)

Up to 16 HPDT MPC groups are supported per channel adapter. Each HPDT MPC Group consists of at least one ″read″ subchannel and one ″write″ subchannel. The channel adapter can be shared by more than one host protocol, for example LCS/LSA/HPDT MPC. (Note: Running HPDT MPC on the same channel with LCS or LSA may exhibit some performance impacts to HPDT MPC.)

APPN HPR over HPDT MPC (MPC+) is supported for both ESCON and parallel channel.

HPDT UDP extends the efficiencies of HPDT services to applications using OS/390 UNIX System Services UDP interface. HPDT reduces CPU cycle consumption and achieves a more efficient transfer of data. HPDT UDP is initially targeted for communications between DB2 on OS/390 V2R4 and Systems, Applications, Products in Data Processing (SAP) R/3 application servers. Other UNIX System Services socket applications using UDP, such as NFS and DCE, can also transparently take advantage of HPDT UDP services over the Network Utility ESCON channel. HPDT UDP is not supported on the Parallel Channel Adapter.

HPDT TCP/IP extends the efficiencies of HPDT services to IP applications using OS/390 V2R5. HPDT reduces CPU cycle consumption and achieves a more efficient transfer of data. It is supported over the ESCON and Parallel Channel Adapters. HPDT MPC supports multicast IP addresses.

3.2.12.2 LSA

The LSA access method supports up to 16 LAN appearances, any mix of token-ring, Ethernet, or FDDI per channel adapter. LSA supports up to 10,000 link stations per SAP. LSA also provides for LAN Gateway support, providing connectivity between ACF/VTAM 3.4 (or later) VM, MVS, and VSE SNA host applications and:

- Token-ring, Ethernet, and FDDI LANs
- ATM token-ring LAN emulation or Ethernet LAN emulation
- Other channels (via LSA)

3.2.12.3 LCS

The LCS access method supports up to 16 LAN appearances, any mix of token-ring, Ethernet or FDDI per ESCON or Parallel Channel Adapter. LCS supports multicast IP addresses and an IP passthru mode that bypasses the IP routing code.

For more information on the Network Utility host channel support, please see Chapter 7, “Channel Gateway” on page 111.

3.2.13 Bridging

The bridging support provided uses Adaptive Source Route Transparent (ASRT) bridging. Depending on the configuration parameters of the ASRT bridge and its interfaces, the ASRT bridge will operate in one of five modes:

- SRB - a pure source route bridge
- TB - a pure transparent bridge
- SRT - a source route transparent bridge
- SR-TB - a bridge that translates from SRB to TB
- ASRT - a bridge that supports all of the above and adapts the bridging technique depending on the source and destination devices.

In addition to the above bridging modes, the code also supports:

- Dual spanning trees
- Bridging on PPP per RFC 1220

- Bridging over ATM SVCs and PVCs per RFC 1483
- The ability to bridge IP on some interfaces and route it on others. IP packets may be routed between the bridged network and the routed network.
- A new multiaccess bridge port for source route bridging over frame relay incorporates many DLCIs in a bridge port for vastly improved scalability of data center Network Utilities. When used with branch office 2218 FRADs, the 2218 can be configured to dynamically switch between primary and backup frame relay virtual circuits without losing its LLC connections with the data center. In this configuration, separate Network Utilities must be used for the primary and backup connections. Fully meshed configurations support any-to-any connectivity and the spanning tree protocol may be used to prevent bridging loops. Non-fully meshed configurations support only branch to data center connectivity since bridging between virtual circuits on the same segment is not supported. In non-fully meshed configurations the spanning tree protocol cannot be used.
- Filter support
  - MAC Address
    - Source MAC Address (with Mask)
    - Destination MAC Address (with Mask)
  - NetBIOS Name
  - Sliding Window Filter based on:
    - Base Offset (MAC Header or I-Field)
    - Frame Type
    - Filter Offset (0 to packet size)
    - Filter Mask (32 bytes)
    - Filter Data (32 bytes)
  - Protocol Filtering
    - Service Access Point (SAP)
    - Ethertype
    - Sub-Network Access Protocol (SNAP)
- SR-TB bridge maintains the same route so that SNA connections can be maintained from TB endstations. Bridged BAN access from TB endstation is enabled to multiple DLCIs that are using the same BAN MAC address.
- MIB support for:
  - RFC 1474 ″MIB Definition for the Bridge NCP of PPP″
  - RFC 1493 ″Definitions of Managed Objects over Bridges″
  - RFC 1525 ″Definitions of Managed Objects for Source Route″

3.2.14 IP Bridging Tunnel

The bridging tunnel is a mechanism which uses IP encapsulation to get any protocol across any backbone media that is supported for IP routing by the box. This means a user can take any type of data coming from any type of bridge (that is, source route bridge (SRB), transparent bridge (TB), SRT, or SRTB) and encapsulate it into an IP ″tunnel.″ A Nways Multiprotocol Access Services, Nways Multiprotocol Routing Services IBM Nways Multiprotocol Routing Network

Services box must be at both ends of the tunnel. In between the tunnel endpoints there may be an IP router.

3.2.15 TCP/IP (Version 4)

The Internet Protocol (IPv4) packet forwarder provides packet forwarding in accordance with the forwarding rules of the TCP/IP family of network protocols. This allows the box to support communication between systems using the TCP/IP network protocols on different physical networks. The physical networks can be of the same or different type, speed, or architecture.

The IPv4 forwarding implementation supports:

- IP Version 4
- IP options:
  - End of Open List
  - No Operation
  - Loose source route
  - Strict source route
    - A configuration option is provided to disable IP source routing if it is not desirable for security or other reasons to bypass the routing tables.
  - Record route
  - Stream identifier (this option is ignored)
  - Timestamp
- IP variable-length subnetting
- Multiple IP addresses per interface (up to 32/interface)
- Same subnet address on multiple network interfaces
- Path MTU discovery (router function only)
- Broadcast
  - Configurable broadcast address (zeroes-fill or ones-fill)
  - Directed broadcast
- BOOTP/DHCP relay agent
- UDP Broadcast Forwarding
- Virtual Router Redundancy Protocol (VRRP)
- The ability route IP traffic on an interface that is also bridging IP traffic
- Resource Reservation Protocol (RSVP)
- IPv6 tunnel
- IP MTU by interface - This allows the configuration of a separate MTU for the IPv4 protocol on each interface to avoid problems with IP MTU mismatches.

3.2.15.1 ICMP

The ICMP implementation includes the following support:

- Echo (PING)
- Destination Unreachable
  - Net unreachable
  - Host unreachable
  - Protocol unreachable
  - Port unreachable
  - Fragmentation needed and DF set (including next-hop MTU for RFC-1191)
  - Source route failed
- Redirect
  - Redirect datagrams for the Host
- Disable ICMP redirects globally or for specific interfaces
- Time Exceeded
  - Time to live exceeded in transit
- Parameter Problem
- Address Mask (RFC-950)

Source Quench, Redirect datagrams for anything other than Host, Fragment Reassembly time exceeded, Timestamp, and Information are not implemented.

3.2.15.2 TCP

Support based on RFC 793 is provided.

3.2.15.3 UDP

Support based on RFC 768 is provided.

3.2.15.4 RIP

The RIP implementation conforms with RFC 1058. This protocol allows routers to dynamically find all attached networks, and the best route to each network. If one router goes down, its routes will time out, and a new route will be used.

The RIP implementation includes the following support:

- Policy Support:
  - Route Acceptance
  - Route Advertisement
  - Route Policy relating to other protocols
  - Default Origination
- RIP Version 2 (RFC 1723) adds the following features: route tags to propagate EGP information, subnet masks to support variable subnet masks, next hop addresses to support optimization of routes, authentication for a password passing, and multicasting so that multicast can be used instead of broadcast.
- Split Horizon/Poison Reverse
- Accept Host Routes (for compatibility with IBM 6611)

- RIP outage-only advertisements

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

3.2.15.5 OSPF

OSPF is an implementation of Version 2 of the Open Shortest Path First routing protocol, which is a dynamic Interior Gateway Protocol (IGP) based on link-state technology. The OSPF V2 specification is described in RFC 2178.

The OSPF V2 support includes the following:

- Area Border Router (ABR) Support
- Stub Areas
- Autonomous System Border Router (ASBR) support
- Interface Support for:
  - Numbered Point-to-Point
  - Unnumbered Point-to-Point
  - Broadcast
  - Non-Broadcast Multi-Access (NBMA)
  - Point-to-Multipoint (PMP)
  - Virtual Link
- Equal Cost multi-path route calculation
- Simple Authentication with type configurable on an interface basis
- Type of Service (TOS) zero support only
- Policy Support
  - Import of external routes configurable by following groups:
    - BGP Routes
    - RIP Routes
    - Static Routes
    - Direct Attached Interface Routes
    - Subnet Routes
    - Import as ASE type 1 or 2 routes
  - Generation and import of the Default Route
- Supernet (CIDR) route support
- Support for Demand Circuits as described in RFC 1793.
- Support for the MinLSArrival constant
- MTU size included in Database Description packets to detect MTU mismatches
- Support for area range summarization by ABRs (Area Border Routers) with support for overlapping address ranges as described in RFC 2178. There is

also an option to calculate the OSPF area ranges based on the cost of the closest (lowest cost) component network.

- Support for per neighbor cost which allows a different TOS 0 cost to be associated with each neighbor on a point-to-multipoint interface.

3.2.15.6 BGP-4

The Border Gateway Protocol (BGP) is an inter-Autonomous System (AS) routing protocol designed to be used in conjunction with IP. It was introduced to overcome the shortcomings of EGP. It has evolved significantly to address scaling problems in the Internet. This BGP implementation adheres to RFC 1654 which defines the BGP-4 Protocol Specification. It is not downward compatible with previous versions of BGP. BGP-4 is an extension of BGP-3 that provides support for routing information aggregation and reduction based on the Classless Inter-Domain Routing architecture (CIDR).

This implementation provides the following support:

- CIDR/Route aggregation support.
- Export, Import and Originate routing (Inclusive and Exclusive) policy support.
- Transit, Multihomed, and Stub AS support.
- BGP4-OSPF Interaction with TAG support.
- External and Internal BGP peer support.

BGP-4 enhancements for policies per neighbor and attributes for path selection allow customers to control inbound and outbound traffic path selection, backup path, and load balancing (based on net number advertisement).

3.2.15.7 Static Routes

Static routes are useful when destinations cannot be discovered dynamically, when the routes don’t change, or when the exchange of routing information is not desirable (dial-on-demand circuits). Static routes may be used to define a default gateway to which packets should be forwarded. For subnetted networks, a separate default gateway can be defined for each subnetted network. The customer can define up to four static routes for backup and alternative routing using an independent cost for each route. For frame relay networks, next hop awareness can be used to determine if a route is viable.

3.2.15.8 Multicast

The IP forwarder supports the routing of IP multicast datagrams, which are identified as packets whose destinations are class D IP addresses (that is, whose first byte lies in the range 224 through 239). A single multicast datagram may be delivered to multiple destinations, called a multicast group. Multicast routing is achieved through an implementation of the Multicast Extensions to OSPF (MOSPF) as documented in RFC 1584 or the Distance Vector Multicast Routing Protocol Version 3 (DVMRPv3) documented in an internet-draft. The box supports a complete implementation of the DVMRP multicast routing protocol that forwards multicast datagrams within a single Autonomous System using a broadcast and prune approach. The DVMRP protocol supports the use of IP-IP tunnels to forward multicast datagrams across non-multicast networks. An implementation of the Internet Group Management Protocol Version 2 (IGMPv2) is also included and will support any IP multicast application that uses IGMP to declare its group membership.

The box itself includes several native multicast applications. The DLSw Multicast Group support uses multicast OSPF. The ICMP ping command in the IP console can use an IP multicast address. The box can also be configured to send SNMP traps to one or more IP multicast addresses.

The box also includes commands to join and leave multicast groups, so that it can be the target (as well as the source) of multicast pings and the target of SNMP GETs.

The box will cache multicast forwarding entries to enhance performance.

3.2.15.9 Classical IP (CIP)

Classical IP and ARP refers to support for IP and ARP with ATM as a replacement for a local LAN and with routers operating as they do in ″classical″ LANs. RFC 1577 is a specification for how IP and ARP operate over ATM AAL5. RFC 1577 describes the creation of an administrative grouping called a logical IP subnet (LIS). Each LIS operates and communicates independently of other LISs on the same ATM network. ATM hosts communicate directly with each other within a LIS and via an IP router between LISs.

The Nways Multiprotocol Access Services support includes:

- RFC 1577 support
- RFC 1755 support
- RFC 1483 support
- RFC 1626 support (default MTU of 9180)
- ATM Forum UNI 3.0 and 3.1 support
- ATM PVCs and SVCs
- LLC/SNAP encapsulation for transmitting all IP and ARP packets
- Up to 32 separate LISs supported on a single ATM interface
- IP client support
- ARP server support
- Simultaneous ARP server/client attached to a single ATM endpoint
- Configurable client/server ATM addresses
- Aging of ARP entries on a per client basis
- Refresh of entries can be performed through either ARP to the ARP Server or through InATMARP on an existing channel
- Each client may specify different aging periods
- Bandwidth specification validation on incoming VCs
- Bandwidth specification on outgoing VCs
- Automatic walkdown of data rates on rejected calls
- Configurable InATMARP interval on permanent PVCs and SVCs
- Virtual ATM interfaces to support a greater number of Logical IP subnets
- Next Hop Resolution Protocol (NHRP) Server and Client to minimize the number of router hops.

3.2.15.10 IPv4 TOS/Precedence Usage

Enterprise Extender, DLSw and TN3270 all encapsulate SNA traffic in IP. There are two options for prioritizing this SNA traffic when it is encapsulated in IP and will be encrypted elsewhere within the IP network:

- Configure the use of the precedence bits in the TOS field when configuring Enterprise Extender, DLSw, or TN3270E Server. The following Precedence bit values will be used:
  - 110 - HPR Network Transmission Priority
  - 100 - HPR High Priority
  - 011 - ISR, FID2 (anything mapped to DLSw or TN3270, HPR UID exchanges
  - 010 - HPR Medium priority
  - 001 - HPR Low Priority
- Define IP filters to set the TOS field to a user-defined value for these traffic types.

Note: TN3270 packets have the priority set only in communications from the server to the client.

BRS prioritization supports both options.

3.2.15.11 Channel Support Via LCS

Supports IP routing between TCP/IP Host applications (TCP/IP for MVS, VM, and VSE; OS/390; AIX/370; and AIX/ESA) and any other interface of the box that supports IP. The use of multicast IP addresses over the channel is supported. OSPF multicast capability is new in OS/390 TCP/IP Release 3.6.

IP passthru support provides a new configuration option so that IP support over the channel can now bypass the IP routing code, offering a simplified configuration option similar in function to what the 3172 supports for easier migration when replacing 3172s with Network Utilities.

3.2.15.12 Channel Support Via MPC+

Supports IP routing between HPDT TCP/IP or UDP MVS Host applications and any other interface of the box that supports IP. The use of multicast IP addresses over the channel is supported. OSPF multicast capability is new in OS/390 TCP/IP Release 3.6.

3.2.15.13 Filter Support (Access Control)

IP Access Controls are used to permit or deny access to a private, trusted network. IP packets are inspected and are either routed, dropped, or delivered to IPSec or NAT for further processing based on configured criteria. IP access controls can be defined with the following properties:

- Filter type - Permit or deny routing of the packet
- Source Address - Specified as an IP address and mask
- Destination Address - Specified as an IP address and mask
- Protocol Range - Specified as two protocol numbers (0 through 255)
- Source port range - Specified as two port numbers. Applies only to TCP and UDP packets

- Destination port range - Specified as two port numbers. Applies only to TCP and UDP packets
- Scope - Filter applies to all packets or packets associated with a particular interface
- Direction - Inbound or outbound. Applies only to interface filters.
- TCP connection establishment - To prevent establishment of certain TCP connections
- ICMP message type and code - Includes ICMP redirects, echo requests, and echo replies (used by PING)
- Suppress packets with Record Route and Timestamp options
- Source address verification - Ensures the packet source address is consistent with the interface it was received over by checking the routing tables
- Fragment Attack protection - Drops IP fragments that attempt to overlay the TCP header as specified in RFC 1858 for Overlapping Fragment attacks.
- Precedence/TOS value - Optionally, packets passing inclusive filters can have their Precedence/TOS values set to a user-specified value.
- Policy-based routing - For inclusive inbound or global filters, the next hop can be specified, giving network administrators a powerful capability to forward packets based on pre-defined policies.

It is possible to generate an ELS message when a particular filter rule is invoked. This ELS message may be packaged in an SNMP Trap for forwarding to a network management station.

The box SYSLOG facility can be used for IPSec rule processing.

Local ELS messages and remote logging of filter events is supported.

3.2.15.14 IPv4 MIB Support

This implementation supports the following MIBs:

- RFC 1213 - MIB II
- RFC 1253 - OSPF
- RFC 1657 - BGP-4
- IETF draft - Classical IP (draft-ietf-ion-mib-02.txt)

3.2.15.15 Management Services

In order to manage, configure, load code, and monitor the basic functioning of the box, the code supports the following services:

- SNMP for management support
- Telnet server/client function - for access to the command line interface
- TFTP protocol for uploading and downloading of new code and configurations

  Trivial File Transfer Protocol (TFTP) permits the transfer of files over the Internet UDP protocol. This TFTP implementation can act as either a client (initiate the transfer) or as the server (services the transfer request).
- Disable the TFTP Server as a defense against corruption
- IP diagnostic tools of ping and traceroute

When invoked, the ping will send ICMP echo requests to a specified destination, reporting the round trip time of the responses received. The source address, packet rate, packet size, and Time-to-Live (TTL) can be specified. A summary indicates the percentage of responses received together with the range of response times. The code also supports the ability to remotely issue a Ping request from NetView for AIX to send an ICMP echo from the box to a specified IP address. The traceroute command lists the path (hop by hop) to the destination specified. For each hop the IP address and round trip time are indicated.

3.2.16 Virtual Router Redundancy Protocol (VRRP)

VRRP specifies an IETF standards-based election protocol that dynamically allows a set of routers running VRRP to backup each other on a LAN. The VRRP router controlling one or more IP addresses is called the Master router, and forwards packets sent to these IP addresses. The election process provides dynamic fail-over of the forwarding responsibility should the Master become unavailable. This allows any of the VRRP router′s IP addresses on the LAN to be used as the default first hop router by end-hosts. The advantage gained from using VRRP is a higher availability default path without requiring configuration of dynamic routing or router discovery protocols on every end-host.

VRRP is supported on Ethernet, token-ring, Fast Ethernet (100Mbps), and FDDI. It is also supported on bridged interfaces where IP has been configured to forward traffic over the bridged interface.

3.2.17 Resource Reservation Protocol (RSVP)

RSVP is a signaling mechanism that applications use on IP networks to reserve network resources in order to achieve a desired quality of service for packet delivery. Network Utility RSVP supports the Controlled Load level of service, which provides, under all conditions, service that would be received when the network was lightly loaded. RSVP is supported on the following interface types:

- ATM with point-to-point SVCs
- PPP interfaces which are leased in nature
- Frame relay interfaces with PVCs or SVCs
- Ethernet (10Mbps and 10/100Mbps)
- Token-ring
- X.25 interfaces with PVCs or SVCs

RSVP reserves overall interface bandwidth and cannot set up reservations on a per virtual circuit basis for frame relay or X.25. RSVP and BRS should not both be used on the same interface.

3.2.18 IPv6

IPv6 is a new version of the Internet Protocol (IP) and is the successor to IPv4. IPv6 includes a much larger address space, a hierarchical address scheme for efficient route aggregation, and Neighbor Discovery Protocol (NDP) for host auto-configuration. The supported routing protocols are static routes, RIPng, and PIM-DM (dense mode). Both the TCP-6 and UDP-6 transport protocols are provided along with telnet. The link types supported by IPv6 are multicast support for Ethernet (including 100 Mbps), token-ring, FDDI, and Channel (ESCON and Parallel with LCS), and point-to-point support for PPP. Both configured or

automatic tunnelling of IPv6 over IPv4 is provided. Path MTU Discovery is supported for IPv6. Multicast Listener Discovery (MLD) is also supported. The ICMPv6 protocol is supported for informational and error messages; PING-6 and Traceroute-6 are supported to test reachability. SNMP MIB support is provided for the IPv6 forwarder, ICMP6, TCP6, and UDP6.

3.2.19 Address Resolution Protocol (ARP)

The ARP implementation supports:

- User cache manipulation (add/delete permanent entries, clear cache)
- Configurable aging timer
- Auto Refresh option to reduce broadcasts
- InARP support over frame relay and ATM
- Proxy ARP (ARP subnet routing - RFC 1027)
- Caching of source route RIF information
- Primary and backup ARP servers for failure recovery
- Distributed ARP Server eliminates single point of failure using Server Cache Synchronization Protocol (SCSP)

3.2.20 Data Link Switching (DLSw)

DLSw is a methodology developed by IBM to transport SNA and NetBIOS protocol traffic across IP networks. Supported end stations attach to the box via LAN interfaces, remote-bridging WAN interfaces, and SDLC interfaces. DLSw communicates with neighbor DLSw routers using TCP/IP through any IP-enabled interface.

Key customer benefits provided by DLSw are in the areas of availability and efficient WAN utilization. Availability is improved because DLSw enables the box to internally terminate the 802.2 layer and locally acknowledge frames, while encapsulating the data in TCP/IP. Local acknowledgement minimizes time outs and the resulting session losses. Availability improvements are achieved through the automatic route switching provided by IP which re-routes traffic around failed nodes and through the reliable WAN transport provided by TCP.

Efficient WAN utilization is achieved because DLSw permits all supported protocols (including SNA and NetBIOS traffic) to be consolidated on a link. In addition, DLSw handles the link layer flow control for SNA endstations. Therefore, WAN and LAN traffic overhead is minimized because DLSw can handle the polling for SDLC and the ″keep alive″ messages for the 802.2 logical link control (LLC) on either the token-ring or Ethernet LAN.

This DLSw implementation is compliant with AIW Version 1 Standard DLSw (also published as RFC 1795), and also with the AIW Version 2 Standard DLSw (RFC 2166). This implementation of the DLSw MIB supports the IETF standard DLSw MIB, RFC 2024. DLSw interoperates with the DLSw implementation in MRNS V1 R1 and R2 for SNA traffic and with MRNS V1 R3 and later for SNA and NetBIOS.

The DLSw RFC 1795 support includes:

- Base functions
  - SNA explorer flows

  - Circuit-level pacing
  - Required capability exchange CVs
  - LF bit propagation
- Options implemented
  - SNA backward compatibility to RFC 1434 implementations
  - NetBIOS (see below)
  - Adaptive pacing
  - Version string control vector in capabilities exchange
  - Run-time capabilities exchange
  - Circuit priority (based on protocol type)
  - Single full-duplex TCP connection
  - Remote SAP list filter
  - Capabilities exchange of NetBIOS names
  - DLSw Switch-to-Switch Protocol for MAC address list support

The NetBIOS Support includes:

- AIW_V1 explorer flows
- UI-frame type filters
- Duplicate UI-frame filters
- Dynamic name caching (learning)
- Integrated cache/filter and bridge code
- Session Alive spoofing

Specific areas of note in the DLSw base RFC 1795 support are as follows:

- Capabilities Exchange

  When RFC 1795 compliant DLSw routers establish a transport connection with each other, they exchange configuration messages, which convey the supported feature set of the DLSw router, such as optional DLSw features, number of TCP connections per DLSw neighbor, and supported SAP/NetBIOS name lists.

- Improved Flow Control (circuit-level pacing)

  RFC 1795 formally defines per circuit flow control using an adaptive or fixed pacing mechanism. This implementation supports adaptive pacing, but will also work with the less powerful fixed pacing mechanism. When interoperating with RFC 1434 routers, this implementation uses the older xBUSY messages to halt and restart the flow of data on congested DLSw sessions.

- Largest Frame Size Handling

  RFC 1795 provides for the propagation of bits from the source-routing information field that indicate the largest frame that can flow on a particular source-routed path. This enables endstations on one side of the IP network to see frame size restrictions imposed by bridges on the opposite side of the IP network. This DLSw implementation includes these flows and also uses

LF bit information to help choose which of several neighbor routers to use to establish a new DLSw session.

- Circuit Prioritization

  RFC 1795 defines a general way to negotiate a data transmission priority to different circuits (DLSw sessions) within the same transport connection. This DLSw implementation uses the circuit priority mechanism to assign SNA and NetBIOS sessions different transmission priorities on the same TCP connection. In addition, the customer can now configure a range of connections if the remote RFC 1795 DLSw requests it, or when interoperation with a remote RFC 1434 DLSw product is necessary.

- SDLC PIU Segmenting

  SNA FID-2 frames are segmented when LLC frames from a host are too large to be handled by a remote SDLC controller.

- DLSw Neighbor Priority

  All DLSw Groups and TCP connections may be assigned one of three priorities, which allows transport (TCP) connections to be used in preference to other connections. This is useful for backup connections between DLSw routers. When a endstation search message is broadcast to multiple destinations and multiple responses are received for a given destination, the neighbor with the highest priority will be used. Without transport priority, or with transports of equal priority, the first neighbor to respond will be used. Transport priority works with any connected RFC 1434 or RFC 1795 DLSw router.

- Multicast Group Support

  This is an implementation feature allowing DLSw neighbors to dynamically discover each other, rather than requiring that each neighbor be statically configured. Multicast OSPF provides the multicast IP routing capability that this feature uses.

Additional DLSw functions include:

- Channel support via LSA
  - Supports local conversions (single DLSw router, no TCP/IP) between VTAM 3.4 (or later) SNA Host applications on the channel and SDLC attached SNA devices and locally attached LANs Supports remote conversions (two partner DLSw routers with TCP/IP) between VTAM 3.4 (or later) SNA Host applications on the channel and remote:
    - SDLC attached SNA devices
    - LLC
      - Bridged LAN interfaces (including token-ring, Ethernet and FDDI)
      - Bridged WAN interfaces (including frame relay)
      - ATM LAN emulation client for token-ring and Ethernet
    - APPN
    - Other channels (via LSA)
  - Supports SNA applications for ACF/VTAM 3.4 (or later) on VM, MVS and VSE
- Multicast exploration and frame forwarding (AIW Version 2 Standard DLSw, RFC 2166)

This is an implementation of the AIW standard DLSw extensions to improve the scalability of DLSw to large networks. Rather than bring up static TCP connections to configured neighbor routers, the code uses both multicast and directed UDP packets to propagate address and name resolution messages, and NetBIOS UI-frame traffic. Transport connections between multicast-capable DLSw nodes are established only on-demand, and are taken down after they are no longer needed.

These scalability enhancements are compatible with back-level DLSw nodes that have no multicast IP support, and with MRNS releases that provided for MOSPF-based discovery of DLSw partners.

- NetBIOS name list support

  This allows a user with a structured naming convention to have more control over the broadcasting of NetBIOS datagrams. DLSw builds and sends a list of local NetBIOS names (with wildcard characters), and uses the information in received lists to minimize the number of destinations to which it sends NetBIOS SSP messages.

- Switch-to-Switch Protocol for MAC address list support

  DLSw partners can exchange lists of MAC addresses they support which can be used to limit traffic on WANs. These lists, which are configured and remain static, contain a set of MAC addresses that are accessible by remote stations. The list is sent to each DLSw partner with which the route has an established TCP session. The partner uses this list to determine which SSP messages should be sent to this particular DLSw partner.

- APPN Remote Device Attachment

  This allows APPN to connect through DLSw to remote endstations.

- QLLC Device Attachment

  Support for QLLC attached devices on X.25 networks includes devices on both PVCs and SVCs, and allows dynamic call-in from non-configured QLLC devices.

- No cache aging option

  Allows the user to select ″no cache aging″. This dramatically reduces DLSw-generated search traffic in situations where IP is running over dial-up links and the location of endstation resources is fixed.

- SDLC secondary

  The ability for DLSw to act as a secondary SDLC station to an attached T2.1 or T4/5 node. The code is able to represent multiple secondary stations on a single port.

- SDLC no TEST support

  Allow the user to have PU 2.0 SDLC devices polled with SNRM rather than TEST frames. The user can also configure whether or not polling should be delayed until the remote device connects to the 2216.

- SDLC PU4 INN support

  Allow PU4s on point-to-point SDLC lines to connect through DLSw to other PU4s on LANs, SDLC, and frame relay. This does not include MLTG or remote load/dump support.

- IETF Standard DLSw MIB Support

This DLSw implementation supports the recently approved IETF standard MIB for DLSw, RFC 2024. To maintain compatibility with products that also manage MRNS V1R3 2210s, this implementation also continues to support the final IETF pre-standard draft of this MIB (draft-ietf-dlswmib-mib-09.txt).

- LNM support

  Allows LNM to manage bridges on the other side of a DLSw IP ″cloud″.

- NetBIOS hex name support

  Hex editing/entering/display of the full 16 bytes of NetBIOS names

- Local SDLC to SDLC support

  DLSw sessions between two SDLC devices attached to the same DLSw box are supported.

- Enhanced duplicate MAC address handling

  These enhancements:

  – Allow the user to define a set of MAC addresses that possess different MAC cache and explorer characteristics than the globally defined characteristics.

  – Allow the user to disable the neighbor priority feature on a global basis. When this is disabled, the DLSw partner information is not cached for the MAC address. SNA and NetBIOS explorers are always sent to all applicable DLSw partners and the first DLSw partner to respond is used to establish the DLSw session.

For more information on the Network Utility DLSw functions, please see Chapter 9, “Data Link Switching” on page 145.

3.2.21 Advanced Peer-to-Peer Networking (APPN)

The code provides the capability of being an APPN network node (NN) with intermediate routing functions and provides network services to both APPN end nodes and LEN nodes which attach to the network node. The APPN network node can also attach to T2.0 nodes using the DLUR function. The APPN network node is able to establish CP-CP sessions with its adjacent APPN network nodes to exchange network topology and resource information. Only one CP-CP session is required if the network node is just used for routing since CP-CP sessions to adjacent APPN end nodes and network nodes are optional. CP-CP sessions are required from APPN end nodes to the APPN network node that provides network services, such as locating resources in the APPN network.

3.2.21.1 APPN Base R2 (Compliance with APPN R2 Architecture)

- APPN Version 2 Network Node Base Functions

  See *SNA APPN Architecture Reference*, SC30-3422, Appendix A, ″Base and Option Sets″, for a detailed list of functions included in the Version 2 Base. APPN Option Set Towers supported:

  – Adjacent Link Station Name (#1002)

  – Adjacent Node Name Change (#1004)

  – Parallel TGs (#1007)

  – LU Name = CP Name (#1012)

  – Dependent LU Requestor (DLUR) (#1067)

- Generalized ODAI Usage (#1071)
- Preloaded Directory Cache (#1101)
- Central Resource Registration (#1107)
- Network Node Server Support for DLUS-Served LU Registration (#1116)

  This feature enables the box, when configured as a NN, to understand when a DLUR EN that supports Option Set 1116 registers its LUs.

  The Dependent LU Requestor (DLUR) architecture has been enhanced to allow DLUR end nodes to register their LUs with a network node Server so that searches do not need to take place each time. This can cut down significantly on unneeded network searches. With the Network Node Server support the box can receive the LU registration and recognize the LU as a DLUR LU.
- Tree Caching (#1200)
- AIW Standard HPR including:
  - High Performance Routing (HPR) (#1400)
  - Rapid Transport Protocol (RTP) (#1401)
  - Control Flows over RTP (#1402)
  - Native ATM DLC for HPR (#2001)
- Branch Extender (#1121)
- Border Node including:
  - Peripheral Border Node (#1014)
  - Extended Border Node (#1016)
  - HPR Border Node(#1405)
  - Prerequisites for NNS Support (#1061)
  - SS Extensions NNS Support (#1063)
  - Border Node MIB (awaiting IETF approval)

- Additional APPN functions supported:
  - Frame relay BAN

    APPN will support the capability to receive and process RFC 1490 frames with either of these encapsulations:
    - Token-ring (802.5) bridge encapsulation format
    - Q.933 encoding for APPN traffic as described in ANSI T1.617 Annex F

    Support for the RFC 1490 802.5 bridge frame format on APPN frame relay TGs is useful when the box has a frame relay connection to a bridge/DLSw BAN/gateway FRAD that is providing frame relay connectivity for token-ring LAN attached SNA devices. The frame relay BAN can be defined to a connection network which reduces configuration and network flows. The use of the bridge frame format over the frame relay PVC allows multiple LLC2 token-ring LAN connections to be multiplexed over a single frame relay DLCI/SAP. Another advantage of this format is that it allows Data-Link Switching and APPN to share the same DLCI/SAP.
  - X.25 QLLC

This support allows the APPN network node to connect to other APPN nodes and PU2.0 nodes across an X.25 network using both PVCs and SVCs. All APPN/DLUR configurations are supported on X.25. There is no X.25 support for HPR. APPN X.25 connection network and APPN X.25 Short Hold Mode are not supported.

Connections using PVCs are handled like leased lines. Connections using SVCs are of the following two types:

1. If activate at startup = yes, the connection will be treated like a PVC, and
2. If activate at startup = no, the connection will be activated on demand and will be handled as a limited resource with the session being brought down when there are no active sessions. CP-CP sessions cannot be handled as a limited resource so they should not be routed on sessions of the second type.

– LU 6.2 Security

  Enhanced LU 6.2 session-level authentication protocol, or session-security, is the LU 6.2 method for LU authentication. Two partner LUs attempting to establish a single or parallel LU 6.2 sessions must share the same password (or key) to allow sessions to be established.

– Channel support - HPR over HPDT MPC and LSA
  - Supports APPN HPR routing between VTAM 4.4 SNA Host applications on the channel and:
    - Token-ring, Ethernet and FDDI LANs (via HPR, ISR, or DLUR)
    - WANs
      – Frame relay PVCs (via HPR, ISR, or DLUR)
      – PPP (via HPR or ISR)
      – SDLC (via ISR or DLUR)
      – ISDN (PPP or frame relay over ISDN)
      – V.25bis (PPP or frame relay over V.25bis)
    - ATM token-ring LANE or Ethernet LANE (via HPR, ISR or DLUR)
    - Other channels (via HPR over HPDT MPC or ISR over LSA)
  - Supports ACF/VTAM 4.4 on MVS

– Channel support - ISR over LSA
  - Supports APPN ISR routing between VTAM 4.2 (or later) SNA Host applications on the channel and:
    - Token-ring, Ethernet and FDDI LANs (via HPR, ISR, or DLUR)
    - WANs
      – Frame relay PVCs (via HPR, ISR, or DLUR)
      – PPP (via HPR or ISR)
      – SDLC (via ISR or DLUR)
      – ISDN (PPP or frame relay over ISDN)
      – V.25bis (PPP or frame relay over V.25bis)
    - ATM token-ring LANE or Ethernet LANE (via HPR, ISR or DLUR)

    - Other channels (via HPR over HPDT MPC, HPR over LSA, or ISR over LSA)
  - Supports SNA applications for ACF/VTAM 4.2 (or later) on VM, MVS and VSE

3.2.21.2 Intermediate Session Routing (ISR)

At session endpoints it is the role of the LU, in conjunction with control point services, to establish sessions with a session partner and route session data back and forth to the partner LU. If the session partners reside on non-adjacent nodes, session data will pass through intermediate network nodes. As these intermediate nodes do not control any of the LU endpoints, LU services cannot be invoked on these nodes. It is the responsibility of the intermediate session router to forward session data to the next node along the session path.

3.2.21.3 Dependent LU Requester (DLUR)

DLUR in conjunction with Dependent LU Server (DLUS) removes the dependency for direct attachment of the dependent LU to a VTAM or NCP and allows implementation of 3270 sessions on nodes that are remote from any VTAM or NCP boundary function. DLUS/DLUR allows the migration of dependent LU devices to APPN networks.

Dependent LUs (LU 0, 1, 2, 3, and 6.2) are dependent on the services of a System Services Control Point (SSCP) in order to establish LU-LU sessions. Ownership of dependent sessions for LUs residing in T1.0 or T2.0/2.1 nodes is dependent on services provided by the VTAM or NCP boundary function. The owned node and LU must appear adjacent to their owner. LU-LU sessions must flow through the boundary function before their route can deviate from the ownership session′s route.

Dependent LU Server (DLUS) at a VTAM network node enables SSCP-PU and SSCP-LU sessions with PUs and dependent LUs on or attached to APPN end nodes and APPN network nodes that support the dependent LU Requester (DLUR) function. A DLUR box provides downstream dependent PU (DSPU) and LU support. The DLUS function is available on VTAM V4R2 or later. The DLUR node provides the boundary function. The only requirement imposed on the DLUS/DLUR connection is that it flow across an APPN network.

Once the session with DLUS has been established, the DLUR can obtain SSCP services for its dependent LUs. The dependent LUs appear to be within the domain of the SSCP. Session initiation flows are emulated from the SSCP, but session BIND and data paths are calculated directly between the DLUR node servicing the dependent LU and its session partner.

DLUR offers the following benefits:

- Dependent LUs can be placed anywhere in an APPN network without regard to the location of the owning VTAM or NCP boundary function.
- APPN dynamics are used for locating LUs and choosing optimal, class-of-service based session routes (dependency on adjacency is removed).
- Ownership of DLUR resources may span network boundaries, thus enhancing backup and recovery.
- Takeover and giveback are supported for DLUR resources.

- LU-LU sessions survive failure of SSCP control sessions.
- Dependent LUs may benefit from HPR's faster routing and nondisruptive path switch.
- Network management is improved due to full visibility of LU-LU and SSCP control sessions.

3.2.22 Branch Extender (BX)

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

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

Configuration restrictions are:

- A network node (NN) cannot be defined on a downlink of a Branch Extender.
- A node cannot be defined on both the uplink and downlink. In addition, a node serviced by a Branch Extender cannot have CP-CP sessions to the NN Server of the Branch Extender node.
- A Branch Extender may have CP-CP session with only one NN at a time.

3.2.23 High Performance Routing (HPR)

HPR consists of enhancements to APPN's speed, services, and routing techniques. The main goals of HPR are to improve APPN data routing performance, improve APPN reliability, exploit new data link technology, provide seamless migration, maintain minimal implementation costs, and preserve common network management.

A network may contain a mixture of APPN and HPR nodes since all HPR nodes fully interoperate with existing APPN nodes. All HPR nodes implement Automatic Network Routing (ANR) functions. HPR nodes which only implement the ANR functionality are called HPR base nodes and can only function as intermediate nodes of HPR connections.

Some HPR nodes additionally implement the Rapid Transport Protocol (RTP). HPR nodes with RTP support are called HPR tower nodes and can serve as an endpoint (that is, edge node) of HPR connections. This code contains the RTP tower code and can function as either an HPR intermediate or edge node.

Every HPR tower node contains a boundary function that performs the necessary mapping between APPN and HPR formats. Thus, session traffic may originate at an APPN or HPR node, traverse APPN and HPR nodes, and be destined for an APPN or HPR node. However, the maximum amount of benefit is achieved when a series of contiguous HPR nodes comprise the route between the two nodes.

HPR provides high-speed data transportation through low-level intermediate routing, as well as minimized error recovery and flow control flows. HPR minimizes the number of flows required for error recovery and flow control by performing these protocols at the HPR connection endpoints instead of at each link along the route. In addition to achieving high link utilization, HPR′s congestion control mechanism prevents congestion rather than simply recovering from it.

Beyond providing high-speed routing, HPR improves reliability by providing non-disruptive path switch. HPR′s path switch function automatically reroutes HPR connections around node and link failures without session disruptions as long as an alternate HPR route is available.

HPR′s route test function provides a means of measuring data routing performance within HPR networks. The routing performance of each hop is reported to SNMP. Network managers can use the information provided by the route test function to identify over-utilized links, determine locations requiring faster links, identify under-utilized links, verify the efficiency of the network configuration, and monitor the overall traffic load of the network.

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

3.2.23.1 Automatic Network Routing (ANR)

ANR is the low-level routing mechanism that minimizes cycles and storage requirements for routing packets through the intermediate nodes of an HPR connection. ANR is a type of source routing. Source routing refers to the general technique in which the entire path is predetermined, encoded as a sequence of routing labels, and prefixed to a data packet. SNA prioritization is maintained at all HPR nodes.

3.2.23.2 Forward Explicit Congestion Notification (FECN)

The code will detect when the frame relay FECN bit is set (on a particular frame) and indicate this to RTP at the end of the route. This will cause RTP′s adaptive rate-based flow and congestion control (ARB) algorithm to throttle back its send rate.

3.2.23.3 Rapid Transport Protocol (RTP)

Rapid Transport Protocol (RTP) is the HPR component which is responsible for error recovery, flow and congestion control, non-disruptive path switch, and segmentation/reassembly, as well as sending and receiving data.

3.2.23.4 Error Recovery

Improvements in link error rates makes it possible, and desirable, to provide end-to-end error recovery instead of performing error recovery on each link. HPR operates DLCs in non-error recovery mode in order to bypass link-level error recovery. RTP supports ″in-order″ delivery. Efficient error recovery is performed by RTP through selective retransmission. In other words, RTP only retransmits packets which fail to reach the RTP partner successfully. RTP provides data reliability by byte sequence numbering data, looking for missing sequence numbers, and selectively retransmitting missing data.

3.2.23.5 Adaptive Rate-Based Flow and Congestion Control (ARB)

ARB regulates the flow of data over an HPR connection by adaptively changing the sending RTP′s rate based on feedback from the receiving RTP. High link level utilization is achieved by considering both the traffic load and the network conditions. Unlike APPN′s adaptive pacing, ARB is designed to prevent traffic bursts and congestion from happening. It also prevents packet loss and maintains a high level of steady state throughput.

ARB is based on a rate control and feedback mechanism conducted between the two RTP edge nodes. Each RTP partner samples the rate at which it is sending, receiving, and forwarding received data. Rate messages are exchanged periodically, included on data messages, between the RTP partners. Based on the rate at which data is traversing the network, RTP either increases or decreases the current rate.

Beginning with Version 3, ARB has been enhanced to allow rate adjustments to respond more quickly to dynamic changes in the network.

3.2.23.6 Non-disruptive Path Switch (NDPS)

RTP monitors HPR connections via status exchanges and timers. After several unsuccessful retransmissions, RTP will assume a node or link failure exists within the HPR connection. RTP requests the APPN control point to compute a new HPR path to the RTP partner node. During path switch, the APPN route selection algorithm looks for a route between the HPR edge nodes comprised only of HPR nodes. If an acceptable HPR route is available, the session is rerouted. All path switch attempts are reported to SNMP.

3.2.23.7 HPR Connection Establishment and Termination

In order to establish a session in a mixed APPN/HPR network, an LU initiates a directory search to locate the destination LU. After the destination LU is located, the network node server uses APPN′s route selection algorithm to obtain the optimal route based on the application′s desired class of service. The route selection algorithm does not attempt to deliberately select HPR paths during session initiation.

3.2.23.8 Extended Border Node

If two APPN network nodes with differing network IDs connect, they are not permitted to establish CP-CP sessions and cannot, therefore, exchange network topology or perform directory services functions. In addition, base APPN provides no means to subdivide a given APPN Network ID network into smaller ″topology subnetworks″.

An APPN Extended Border Node is able to subdivide an APPN network (single Network ID) into smaller subnetworks as well as connect multiple enterprises with different network IDs. In addition to this base Extended Border Node support, the following is also supported:

- Extended Border Node support for HPR sessions that span the border node, allowing multi-subnet sessions to benefit from HPR′s better performance and non-disruptive path switching
- Network Node Server (NNS) support for Session Services Extensions (SSE), enabling the APPN network node to act as an NNS for a VTAM end node
- Session Services Extensions (SSE) to allow VTAM end nodes to establish CP-CP sessions with VTAM across the Border Node boundary

- Extended Border Node MIB as defined by the APPN Implementers' Workshop

3.2.23.9 Supported DLCs

APPN is supported on the following DLCs:

- Token-ring - ISR/HPR/DLUR
- Ethernet - ISR/HPR/DLUR
- FDDI - ISR/HPR/DLUR
- Frame relay PVC - ISR/HPR/DLUR (either bridged or routed format)
- Frame relay SVC - Enterprise Extender only
- ATM native (RFC 1493) - HPR
- Point to Point Protocol (PPP) - ISR/HPR
- Connection through DLSw to remote APPN devices - ISR/DLUR
- SDLC leased (primary, secondary, negotiable) - ISR/DLUR
- APPN over LAN emulation - ISR/HPR/DLUR
- APPN over X.25 with QLLC
- HPR over channel using HPDT MPC or LSA
- IP (Enterprise Extender) - HPR

Note: All DLC types that support APPN also support Extended Border Node

3.2.23.10 Dynamic Configuration

APPN will allow most configuration parameters to be changed dynamically without restarting the APPN node. The dynamic changes in the box configuration are allowed via the command line interface.

It is also possible to use the Configuration Program to create the revised APPN configuration and then use the command line interface to invoke the changes dynamically. In order to use this method the Configuration Program changes should be sent without specifying a restart of the box. Then the command line interface can be used to activate the new APPN configuration dynamically.

The following is the list of parameters that will cause the node to restart (that is, this is the list of non-dynamic parameters):

- Changing any of the connection network parameters
- Changing the mode name
- Changing the COS information
- Changing the NODE parameters
- Changing session and ISR accounting
- Changing the DLUR parameters at the node level (for example, enabling DLUR, changing the primary and backup DLUS)
- Changing the HPR parameters at the node level (for example, max number of HPR sessions, RTP inactivity timers, path switch timers, max RTP retries)
- Changing tuning parameters
- Changing the adjacent node type on a link station definition

3.2.24 APPN Network Management

3.2.24.1 Route Test

HPR′s route test is invoked through SNMP. Two variations of route test exist. The first tests the wrap-around time of an established HPR connection. The second tests the wrap-around time of an APPN selected route to a specified destination node.

3.2.24.2 Operator Path Switch

Network operators can force an HPR connection originating at a node to be rerouted via SNMP′s operator path switch. The operator path switch causes APPN′s route selection algorithm to obtain the best HPR route to the other HPR edge node. If one exists, the HPR path is rerouted using RTP′s path switch protocol. If the current route is returned or no alternative path is available, the HPR connection is not altered. In any case, the path switch attempt and result are reported in the HPR MIB′s path switch status table.

3.2.24.3 MIBs

- Support of Set, Get, Get_Next and Trap
- IETF Proposed Standard APPC MIB - RFC 2051
- IETF Proposed Standard APPN MIB - RFC 2155
- IETF Proposed Standard APPN HPR MIB (draft-ietf-snanau-hprmib-01.txt)
- IETF Proposed Standard APPN DLUR MIB (draft-ietf-snanau-dlurmib-01.txt)
- APPN Implementers Workshop Extended Border Node MIB
- Portions of previously available IBM Enterprise Specific APPN MIB
  - Memory and Accounting
- Portions of previously available IBM Enterprise Specific APPN HPR MIB
  - HPR NCL and Route Test
- IBM Enterprise Specific APPN Branch Extender

3.2.24.4 SNA Management Services (MS)

The SNA MS support includes the following:

- Connection with explicit Focal Point (Defined/initiated by NetView)
- Connection with implicit Focal Point (Defined/initiated by router)
- Switch to backup FP
- Route over MDS/LU6.2 (also migration support of CP/MSU)
- Send FP information to served end nodes
- Generate generic alerts (APPN and APPC)
- Forward end node MDS MU data to FP

3.2.25 Enterprise Extender (HPR over IP)

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

Enterprise Extender provides the flexibility for SNA parallel sysplex features, which are currently available in HPR networks now, to be available to users in networks that have IP backbones, or even IP clients when coupled with TN3270E server support. Enterprise Extender also makes it possible for SNA networks to use IP attachments as alternate and backup routes for the SNA network.

Enterprise Extender technology can also reduce the demands on the data center routing platforms, such as the 2216 or 3746 MAE, and, thus, provide a more cost-effective solution than other integration technologies. Enterprise Extender seamlessly routes packets through the network protocol ″edges″ eliminating the need to perform costly protocol translation and the store-and-forward associated with transport-layer functions such as DLSw.

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

3.2.26 LAN Emulation

LAN emulation protocols allow ATM networks to provide the appearance of LANs such as Ethernet or token-ring. Although LAN emulation does not exploit all of ATM′s benefits, it is useful in migrating to ATM technology and lowering network management costs. High-speed ATM links can be utilized, while still protecting software and hardware investments. Software investments are protected because application interfaces are unchanged (LAN emulation is implemented within the Data Link Control layer), and hardware investments are protected by bridging LAN and ATM networks (enabling utilization of existing adapters/wiring). LAN emulation support for both bridged and ATM-attached stations allows incremental installation of ATM adapters in stations with high-bandwidth requirements (for example, servers and engineering/multimedia workstations).

The network management benefits of ELANs come from increased flexibility in handling moves, adds and changes. Membership in an ELAN is not based on physical location; instead, logically related stations may be grouped to form an ELAN (and stations may be members of multiple ELANs). Then, as long as ELAN memberships are retained, no reconfiguration is required when stations move to new physical locations. Similarly, no wiring modifications are needed to move stations from one ELAN to another.

The Nways Multiprotocol Access Services LAN emulation client implementation complies with the ATM Forum LAN Emulation specification, Version 1.0. This specification identifies the following components to implement an ELAN: LE clients (LECs), the LE Configuration Server (LECS), the LE Server (LES) and the Broadcast and Unknown Server (BUS). The LES, BUS and LECS are collectively referred to as the LE Service components. Each ELAN has a dedicated LES and

BUS. LECs reside in end systems (ATM-attached stations or bridges) and provide a MAC-level service interface to higher layer software. Either Ethernet/802.3 or IEEE 802.5 token-ring LANs may be emulated (but all stations on an ELAN must be the same type). In bridges, the proxy function associated with the LEC interfaces between the physical LANs and associated LECs. In order to emulate a LAN, LECs request services from the LE Service components.

The LEC provides an internal interface to higher layer protocols and emulates the MAC interface of either an IEEE 802.3/Ethernet or 802.5 LAN. This is what allows existing LAN applications to use ATM services with LAN emulation. Each workstation or other device (bridge, router, etc.) connecting to an ELAN must have an LEC. The LEC performs control functions (VCC establishment, address resolution, etc.) and data forwarding.

Generic LAN emulation quality of service, or Configurable Quality of Service (QoS), allows LAN emulation networks to take advantage of ATM′s QoS capabilities. The parameters best suited for the communicating stations are picked with a negotiated algorithm for both ELAN-wide and LEC-based services.

For added failure recovery, a backup gateway for endstations on LAN emulation can be used. Default gateway IP addresses are manually configured. If the primary gateway goes down, the backup gateway automatically starts passing packets from the endstation to other subnets.

3.2.27 ELAN Summary

Nways Multiprotocol Access Services ELAN support includes:

- Support provided for multiple ATM devices
- ATM Forum UNI 3.0 and 3.1
- ATM Forum Interim Local Management Interface (ILMI)
- Support for both Ethernet/IEEE 802.3 and 802.5 ELANs
- A single instance of the LEC code supports multiple clients (on different ELANs)
- LEC with proxy support provided for both transparent and source route bridging
- A single instance of the LEC with proxy supports bridging to multiple LECs/ELANs
- Support provided for LEC configuration via the LECS

  LECS will support ELAN assignment based on ATM address, MAC address, ELAN name, or defaults. The customer will be able to define precedence relationships among the configuration options; for example, use MAC address if present, else use ELAN name if provided, otherwise use default.
- The LEC will forward unknown data frame to the BUS regardless of LAN type
- SVC support only

3.2.28 Bandwidth Reservation System (BRS)

Bandwidth reservation (BRS) is a means to decide what packets to transmit on a congested PPP or frame relay interface. The code allows reservation of percentages of the total bandwidth on a interface for specific ″classes″ of traffic. These bandwidth percentages are a guaranteed minimum in a fully loaded network. A class can far exceed its guaranteed minimum in a network under light traffic load, up to the full bandwidth. Bandwidth reservation applies to output only.

Additionally bandwidth reservation provides the ability to prioritize traffic within each class, where each bandwidth class has four priority levels to which traffic can be set—Urgent, High, Normal, and Low. The order is enforced within a class; for example, all urgent tagged messages will be sent first within the class, then followed by High, Normal and Low. However, it should be noted that the priority settings within a bandwidth class have no effect on any traffic in other bandwidth classes; for example, urgent traffic in class 1 has no effect on low traffic in class 2.

The BRS support for frame relay uses a nesting mechanism for reserving bandwidth at both the per-circuit level and at the interface level. When a frame is transmitted, BRS assigns the frame to a traffic class and a priority within that class based on the frame′s protocol type and the results of BRS filtering. Each circuit has its own set of traffic classes and each traffic class is assigned a percentage of that DLCI′s bandwidth. Circuits are grouped into circuit classes that are assigned a percentage of the interface′s bandwidth.

Bandwidth reservation guarantees bandwidth for specific types of encapsulated traffic (classes) identified by either the protocol type or a filter. BRS supports the following protocol types:

- IP
- ARP
- Bridging
- SNA/APPN-ISR
- APPN-HPR

BRS supports the following filters:

- UDP or TCP port number or range of port numbers
- MAC Address tags
- NetBIOS
- SNA/APPN-ISR
- IP Tunneling (IP)
- Multicast (IP)
- SNMP (IP)
- Rlogin (IP)
- Telnet (IP)
- DLSw (IP)
- SDLC Relay (IP)

- APPN-HPR
- HPR-Network
- HPR-High
- HPR-Medium
- HPR-Low
- X.25 Transport over TCP (IP)
- IP Precedence bits or entire IP TOS field

  BRS can be configured to filter based on IP precedence bit values that are set for different types of SNA traffic on IP or to filter on user-defined TOS values.

3.2.29 Interactive Network Dispatcher

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

The Network Dispatcher sees only the incoming packets from the client to the server. It does not need to see the outgoing packets, which significantly reduces the overhead imposed by load balancing. The client′s packet is forwarded to the chosen server exactly as it was created. Network Dispatcher is useful for many applications such as e-mail servers, Web servers, distributed parallel database queries and other TCP/IP applications. Network Dispatcher also supports stateless UDP applications.

Network Dispatcher lets you load balance news servers, mail servers, Web servers, file servers, and other servers with its protocol advisors for Network News Transfer Protocol (NNTP), Post Office Protocol (POP3), Simple Mail Transfer Protocol (SMTP), Telnet, HTTP, and FTP. Advisors query the servers and analyze the results to help determine the best distribution of incoming requests.

The box also includes an advisor for TN3270E to maximize the efficiency of the load-balancing for multiple Network Utility, 2216, or 2210 TN3270E servers. When balancing TN3270 servers, one of the servers may be in the same Network Utility as the Network Dispatcher.

When the Interactive Network Dispatcher in the Network Utility interfaces with a S/390 Parallel Sysplex, workload balancing is performed by exploiting Parallel Sysplex technology inherent in the S/390 system that allows for policy-driven workload management with the traditional benefits of high availability, scalability and security offered by S/390.

3.2.30 TN3270E Server

The IBM Network Utility Model TN1 includes the TN3270E server component that provides a TN3270 gateway function for TN3270 clients downstream from an SNA/VTAM S/390. The clients connect to the server using a TCP connection which the server then maps to a corresponding SNA LU-LU session that the TN3270 server maintains with the S/390. The Multiprotocol Access Services TN3270E server supports the capabilities defined in RFCs 1576, 1646 and 1647. The connection from the TN3270E server to the S/390 may be either a subarea SNA connection or an APPN connection (subarea connections still make use of the APPN code). Subarea connections between the TN3270 server and the S/390 are supported on Ethernet (10 Mbps or 10/100 Mbps), token-ring, FDDI, ATM LANE, LSA (channel), HPDT MPC (MPC+) (channel), and frame relay. The APPN connection is between DLUR in the Multiprotocol Access Services and DLUS in VTAM. This connection uses either APPN ISR or HPR transport and is supported, locally and remotely, over all the interfaces that support these transport protocols.

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

The TN3270E server connection function provides a multiple PU appearance to an SNA host, allowing support for greater than 253 LU sessions.

The TN3270 server support includes a partial implementation of the TN3270E MIB (including the TCP Connection Table), the TN3270 Response Time MIB, and a SNA Resource Map Table. This table allows for the mapping of the VTAM SLU name to the name known at the TN3270E server for that same NAU.

LU pooling allows SNA LUs to be grouped into named pools for easier configuration for some TN3270 networks. Pools provide:

- Balancing of traffic across the PUs, by assigning multiple local PUs to a pool.
- Improved availability by allowing LUs to be updated without having to update the clients.
- Increased reliability for when a client loses a session due to a link outage. The client can reconnect using the same pool to get an LU that uses a different link.
- Improved scalability by making it easy to allocate a pool of LUs for casual TN3270 clients to share for access, leaving more resources available for other users.

IP address to LU Name mapping allows administrators to control client access to particular LUs or LU pools based on the client's IP address.

VTAM V3R4 and later support Self-Defining Dependent LUs (SDDLU)/Dynamically Defined Dependent LUs (DDDLU). This allows the TN3270E server to send VTAM a list of dependent LU addresses for each TN3270 PU so that VTAM can dynamically create its own LU definitions.

Multiple TCP port support. This allows for the definition of TCP ports that supports base TN3270 function, versus TN3270E function, to support older clients. Also, by using LU pools to associate different SNA resources with different TCP

ports, clients can access different SNA resources by connecting to different TCP ports on the same TN3270 Server.

For more information on the Network Utility TN3270E functions, please see Chapter 5, “TN3270E Server” on page 69.

3.2.31 Network Management

Nways Multiprotocol Access Services network management support is based on the following:

- The SNMP V1 protocol
- Transportation of SNMP messages between devices with IBM multiprotocol services and Network Management Stations (NMSs)
- Configurable MIB information access
- A comprehensive set of standard and enterprise-specific SNMP MIBs for monitoring and managing resources:
  - WAN interfaces including ISDN
  - LAN interfaces
  - ATM interfaces
  - Bridging and Routing
  - Data Link Switching (DLSw)
  - APPN
  - System specific
  - Enterprise specific
  - LAN Emulation Client (LEC)
- The functions of the network management application such as:
  - Nways Manager for AIX Version 1.2
  - Nways Workgroup Manager for Windows NT Version 1.1
  - Nways Manager for HP-UX Version 1.2

The user defines a relationship between the SNMP agent in the Network Utility and the network management station (NMS) called a community. The community definition consists of:

- A community name
- One or more IP address/mask pairs that indicate the source addresses of NMSs that are allowed to communicate with the agent. These addresses also are used to specify TRAP destinations.
- An access level (Read only or Read/Write) which determines if SNMP SET requests are allowed for the community. Most of the MIB variables are implemented as read-only; however, there is a limited number that supports SETs.
- An optional MIB view definition which can act as a filter and limit the community to accessing a subset of all MIBs

A number of generic and enterprise-specific SNMP traps are provided. These traps are notifications to the NMS and are triggered by a particular event

occurring on the device (for example, an interface is disabled). The community definition also allows the user to configure the destinations that traps will be sent to and allows them to enable and disable which traps are to be sent.

3.2.31.1 Management Information

Most of the network management data is provided via SNMP. In addition, APPN provides some support for SNA Management Services as described in the APPN section. The MIB contents for SNMP are listed below. For a more complete listing of the MIB information please refer to the MIB Appendix in the *Nways Multiprotocol Access Services Protocol Configuration Monitoring Reference*, SC30-3884 and SC30-3885.

Function	MIB Information
General	RFC 1213 (MIB-II) RFC 1573 (Interface MIB)
WAN Port Interfaces	RFC 1315 (Frame Relay DTE) RFC 1317 (RS-232-like serial interface) RFC 1471 (PPP LCP) RFC 1747 (SDLC)
LAN Interfaces	RFC 1650 (Ethernet MIB) RFC 1748 (802.5 Token Ring MIB) LLC MIB (IETF Draft 01) RFC 1512 (FDDI MIB)
ATM Interfaces	RFC 1695 (AToM MIB)
LAN Emulation Client	LE Client Mgmt V1.0
Bridging Functions	RFC 1493 (Bridges) RFC 1525 (Source Route Bridges)
Routing Protocols	RFC 1253 (OSPF) RFC 1657 (BGP-4)
DLSw	RFC 2024 (Standard DLSw MIB)
APPN	IETF APPN MIB (RFC 2155) APPN HPR MIB APPN DLUR MIB APPC MIB (RFC 2051) IBM Enterprise Specific APPN Branch Extender
TN3270E	TN3270E Base MIB (partial) TN3270E Response Time MIB (partial)
Others	Routing Enterprise (for example, ELS) Box Specific (for example, CPU Utilization, memory usage, thermal sensing) Bandwidth Reservation Extensions (for example, PPP, DLSw, IP remote PING)

3.2.31.2 LAN Network Manager (LNM)

The LNM support is a Source Route (SR) bridging option that enables LAN Network Manager agents on the SR bridge. The LNM function supports the following LNM agents:

- Configuration Report Server (CRS)

  The CRS agent collects and reports MAC ring topology changes to the IBM LNM application. It will send out CRS MAC requests to query the status of other ring stations when requested by the LAN Network Manager.

- Ring Error Monitor (REM)

  The REM agent collects MAC error reports from ring stations. When thresholds are exceeded, REM forwards error information to the LAN Network Manager.

- Ring Parameter Server (RPS)

  The RPS agent services MAC requests from ring stations for ring parameter information and informs the LAN Network Manager of ring insertions.

- LAN Bridge Server (LBS)

  Limited LBS support is provided. Bridge statistics, in the form of number of packets out, may be displayed via the LAN Network Manager. However, the box configuration may not be changed by the LAN Network Manager.

The LAN Network Manager establishes an LLC-2 connection to the box LNM agent. They then communicate using the prescribed LNM management requests and responses

In addition, DLSw can be used to transport LAN Network Manager (LNM) packets. LNM is identified by Logical Link Control (LLC) Service Access Point (SAP) x′F4′.

3.2.31.3 Console Interface

The code supports a console interface that is accessible locally via an ASCII terminal or remotely via telnet. Access is controlled either through a locally administered password or through the use of a TACACS+ or RADIUS server. The TACACS+ or RADIUS server can be set up to provide authentication, authorization and accounting for telnet access. The console provides the following capabilities:

- GWCON - Displays the status and statistics of the box′s hardware and code, such as protocols, network interfaces, and event logging.
- CONFIG[2] - Provides online control of various configuration parameters, such as network addresses and event logging.
- QUICK CONFIG[3] - Provides a simple, less-detailed way of configuring bridging and routing protocols.
- MONITR[4] - Receives Event Logging System (ELS) messages and messages from the operating system, and displays then on the monitor, according to user-selected filtering criteria. Optionally, the messages can be sent via

---

2 CONFIG is also known as the talk 6 process.

3 QUICK CONFIG is accessed from the talk 6 process.

4 MONITR is also known as the talk 2 process.

UDP packets to a remote IP address. These packets will be formatted to appear as syslog packets. Users of remote logging can use the syslog facilities to direct messages to the remote host console, write them to one or more disk files, send them to one or more users, etc.

ELS manages the messages logged as a result of box activity. Using ELS commands, users can set up a configuration that filters out those messages that are unimportant to them. It is also possible to capture the ELS information only for specified interfaces. Messages can be displayed on the console terminal screen or sent to a network management station using SNMP traps. ELS also provides the ability to capture, format, and offload large volumes of ELS messages. ELS configuration is only supported via the command line interface.

3.2.31.4 Dynamic Reconfiguration

Dynamic reconfiguration includes the following capabilities:

- Activate Interface - provides the capability to configure and activate a new interface, including the protocols and features configured to run on the interface, without restarting the Network Utility. This is supported for all interfaces except for ATM and SDLC. Interfaces on which LNM, XTP or IPv6 are configured can be activated, but these protocols will not activate themselves on those interfaces. Other protocols configured for the interface will also be activated. Protocols on interfaces to be activated must already be enabled on some other interface.
- Reset Interface - provides the capability to disable an existing interface, including the protocols and features running on the interface, and then automatically reactivates the interface, including the protocols and features, using new interface-based configuration parameters. This is supported for all interfaces except for ATM, X.25, and SDLC. Interfaces on which LNM, XTP, or NHRP are configured can be reset, but these protocols will not reset themselves. Dynamically changing the hardware device type, DLC type, or enabling a new protocol on an interface is not supported dynamically.
- Reset Protocol - provides the capability to disable a protocol on an existing interface (or globally) and then automatically reactivates the protocol on the interface (or globally) using new configuration parameters. IPv4, RIP (for IP), OSPF, BGP, APPN and MAC Filters provide commands to reset part or all of their configuration parameters without having to restart the Network Utility or its interfaces. IPv4 and RIP support resetting many interface-based parameters, including IPv4 addresses. IPv4 can add, change, or delete global IPv4 filters (access controls) and IPv4 packet filters. OSPF neighbors, interfaces, areas, and AS boundary routing policy can be added, deleted, or changed using reset. BGP supports adding, deleting or changing BGP neighbors. Most APPN parameters can be reset. MAC filters can be globally reset.

In all cases, if a restart is required, it will be indicated at the console. Interfaces and protocols that can be activated or reset will be activated or reset. Protocols that cannot be activated or reset will not be. If a restart is desired, then it can be set up to occur at a more convenient time.

Please refer to the product publications for additional details on the dynamic reconfiguration support and restrictions.

3.2.32 Configuration Program

The Configuration Program runs on a workstation and allows users to use a graphical user interface to create a complete configuration of the box. The Configuration Program includes the following features:

- Enhanced configuration validation

  When the Configuration Program flags a configuration error, it puts a yellow question mark in the Navigation Window. By selecting the ″validation″ option, the Configuration Program will display a message describing the error and what to do to correct it. When the error is corrected the Configuration Program puts a ″green″ check mark in the Navigation Window indicating the configuration file is valid for sending to the box.

- Enhanced query information

  The Configuration Program is able to save a unique configuration name (default name of ′config′ is provided), model, and code release of the Configuration Program in the configuration file. However, we recommend the operator supply unique names. The information provided in the above step, and sent to the box, can now be retrieved from the box via the enhanced ″query″ operation. This procedure displays the configuration file name, the machine model, and the code release of the Configuration Program which created the configuration file.

- Configuration tutorial

  A tutorial describing how to use the Configuration Program is available from the Help pull-down menu on the Navigation Window.

- Configuration File SEND via SNMP SET

  Sends formatted configuration data for the working (currently opened) configuration. This new feature uses SNMP SETs and is the preferred method of transporting configuration data to the box′s configuration memory. Executing `tftp get` from the box to copy a configuration file into configuration memory remains an alternative.

- Configuration Files SENDs to multiple boxes via SNMP SETs

  With SEND MULTIPLE, the user′s accumulated data base(s) of configurations is the source to send the respective configurations to a multiple number of boxes. For example, a group of boxes can be identified to receive their respective configuration data in one SEND MULTIPLE action.

- Configuration File Remote TIMED ACTIVATION

  This feature allows the user to restart (activate) the box at a specified date and time. The user-specified information is used to calculate the seconds from the time of command execution to the requested restart time; this count is subsequently forwarded to the designated box. Typically this feature is used with the SEND configuration and SEND MULTIPLE configurations functions.

- Retrieve Configuration

  Retrieving a configuration via a communications parameter is supported for all three platforms. In a single operation, this command retrieves an active box configuration via SNMP GETs and places it into the program work space for immediate editing and browsing.

  With Retrieve Multiple, the user can retrieve the active box configuration from a set of boxes into configuration databases.

- Read Router Configuration

  This option ″reads″ a configuration file into the Configuration Program for browsing, editing, etc., from DASD. This function assumes a prior action has occurred where the operator has obtained a copy of the desired configuration from the box via an TFTP function.

- ASCII file

  This option is for creating an ASCII version of the configuration file. You may print the file to obtain a hard copy report. You can also use the configuration program to read configuration file in ASCII format. After updating an ASCII configuration, you can read it into the configuration program and save the configuration in binary. You can load a configuration into a multiprotocol device only in the binary format.

- Command Line Facility

  This facility provides the capability to automate configuration operations which are available in the configuration program. You can place commands after the cfg command or in an argument file named cfgargs, which the configuration uses to direct its operation.

- Drag and Drop

  Certain kinds of lists can be reordered using drag and drop. For example, see the Filter Items folder on the MAC Filter-Lists option. Filter items can be re-ordered by selecting an item in the list, then selecting the middle button (for three-button mice) or right button (two-button mice) and dragging and dropping the item to the desired position in the list.

The Configuration Program currently *does not* support configuration of the following:

- Event Logging System (ELS)
- SDLC Relay

All of these items are configurable through the command line interface. Once configured via the command line interface, the Configuration Program will preserve the configuration information in subsequent retrieves and sends of the configuration.

3.2.33 Functions Not Supported

In both software loads, certain MAS functions that are not germane to the dedicated functions for which each Network Utility model is intended have been removed. This allows each model to be streamlined and optimized for the functions that it will perform. This section lists the equivalent MAS functions that are not available with the Network Utility.

- Protocols
  - AppleTalk 2
  - Banyan Vines
  - IPX
  - DECnet IV
  - DECnet V / OSI
- Adapters/DLCs
  - ISDN
  - Fast token-ring (FasTR)

  - V.25bis
  - Dial circuit support
  - Payload encryption
- Software features
  - L2TP
  - DIALs (Remote LAN Access)
  - Network Address Translator (NAT)
  - IP Security (IPsec)
  - WAN Restoral and WAN Reroute (WRS)

Chapter 4. Overview

This chapter is an introduction to the remainder of the redbook and describes how the other chapters document some of the possible applications for the IBM Network Utility.

4.1 Major Network Utility Functions

As discussed in Chapter 3, “Software Overview” on page 19, the Network Utility uses Multiprotocol Access Services (MAS) software technology to support a variety of networking functions. The Network Utility is specifically designed for CPU and memory-intensive functions at network positions requiring a small number of physical interfaces.

Key applications of Network Utility by model include:

Model TN1 - Network Utility TN3270E Server

The TN3270E Server function provides SNA host application access to IP desktop users.

One or more Network Utilities can be positioned at a regional office or host data center to provide access for medium-to-large numbers of TN3270 clients distributed throughout an IP network.

Network Utility Model TN1 also supports all the functions of Model TX1.

Model TX1 - Network Utility Transport

Data Link Switching (DLSw)

DLSw provides native SNA endstation (workstation, controller, FEP, or host) connectivity across IP backbone networks. It also performs DLC type conversion like that done in FRAD and X.25 PAD products.

One or more Network Utilities can be positioned at a regional office or host data center to terminate TCP connections from smaller DLSw routers in many branch offices.

Advanced Peer to Peer Networking (APPN)

APPN provides native SNA endstation (workstation, controller, FEP, or host) connectivity across SNA backbone networks. The Enterprise Extender feature allows this same connectivity across IP backbone networks.

Network Utilities can be positioned wherever a high-capacity APPN network node is required. You can place one at the edge of an IP network to receive traffic from other Enterprise Extender products. A Network Utility could also provide extended border node function when connecting two different APPN networks.

Channel Gateway

Network Utility supports both ESCON (fiber-optic cable) and Parallel Channel (bus and tag cable) adapters. Using one of these adapters, a Network Utility can serve as a gateway routing both SNA and IP traffic from a S/390 host to local LANs, an ATM network, or to a high-speed serial line.

Network Dispatcher

This function allows a number of IP-based application servers (for example, TN3270 servers, HTTP web servers, or FTP servers) to present a single IP address appearance to client workstations on an intranet or on the Internet. The network dispatcher function fields TCP connection requests from these clients and routes them to an available server. It provides both load balancing among the servers, and high ″logical server″ availability by bypassing failed physical servers.

The Network Utility can be placed at a host data center in front of hosts providing these server functions, or in front of multiple Network Utility Model-TN1s that are providing TN3270E Server function.

High-speed media conversion

Network Utility can serve as a high-speed bridge between interfaces on its supported adapters.

In this book, we have selected a key subset of the above functions for expanded discussion and example configurations. The chapters that follow cover:

- TN3270E Server, optionally with Network Dispatcher in front of multiple servers
- Channel Gateway, for both SNA and IP traffic
- Data Link Switching, with both TCP termination and local DLC conversion

For help in understanding and configuring Network Utility functions other than these, consult the core software publications:

- *Nways Multiprotocol Access Services Protocol Configuration and Monitoring Reference Vol. 1 Version 3.2*, SC30-3884
- *Nways Multiprotocol Access Services Protocol Configuration and Monitoring Reference Vol. 2 Version 3.2*, SC30-3885
- *Nways Multiprotocol Access Services Software User's Guide Version 3.2*, SC30-3886
- *Nways Multiprotocol Access Services Using and Configuring Features Version 3.2*, SC30-3993

You may also find configuration help in the following IBM Redbooks. Although they are specific to the IBM 2216 Model 400, some of the configuration scenarios may apply.

- *IBM 2216 Nways Multiaccess Connector Description and Configuration Scenarios Volume 1*, SG24-4957
- *IBM 2210 Nways Multiprotocol Router and IBM 2216 Nways Multiaccess Connector Description and Configuration Scenarios Volume 2*, SG24-4956

4.2 Chapter Layout and Conventions

The remaining chapters of this redbook are organized as follows.

4.2.1 Chapter Layout

Each of the three key functions (TN3270E Server, Channel Gateway, and Data Link Switching) is covered by two chapters:

- An introductory chapter that:
  - Summarizes the supported function
  - Discusses example network configurations
  - Introduces how to manage the function
- Three chapters of example configuration details, containing:
  - Labelled diagrams of key example configurations
  - Matching tables with configuration parameters for both Configuration Program users and command-line users

  The configurations shown and described in the three chapters are actual working configurations. Binary configuration files matching these configurations are downloadable from the World Wide Web. To access these files, follow the Support and Downloads links from:

  `http://www.networking.ibm.com/networkutility`

In addition, Appendix A, “Sample Host Definitions” on page 167 provides detailed examples for configuring IBM host software products to match Network Utility configurations.

4.2.2 Example Configuration Table Conventions

The configuration parameter tables used in the three example configuration chapters all follow the same format. Table columns and conventions are as follows:

Configuration Program Navigation
The sequence of folder and panel names to follow until you reach the panel where you enter parameter values.

Configuration Program Values
Parameter names and their values.

If the configuration program panel shows parameters not listed in the table, we used their default values. ***Your configuration must be for a Network Utility and not a 2216-400 to have the correct default values.***

Command-Line Commands
The commands you type to configure the same parameters using the command-line interface, as follows:

- Command sequences start from the talk 6 prompt `Config>`. Where needed, the initial command shows how to get to the right place in the menu system and the resulting command prompt.
- Commands without parameters specified will either ask for the input values or have no parameters. Parameter prompts from the system are shown in `this font`.
- Where the value prompts and values you type are self-explanatory, the details are not shown.
- ″Accept other defaults″ means that there are other parameter prompts for which you should accept the default values (by pressing Enter).

Notes

Numbers that reference comments at the bottom of each table.

Chapter 5. TN3270E Server

This chapter introduces TN3270 and summarizes the TN3270E server function implemented in Network Utility.

5.1 What Is TN3270?

Many companies today are considering the consolidation of their WAN traffic onto a single IP-only backbone. At the same time, other companies are simplifying their workstation configurations and attempting to run only the TCP/IP protocol stack at the desktop. However, most of these companies still require access to SNA application hosts.

TN3270 meets these requirements by allowing you to run IP from the desktop over the network and attach to your SNA host through a TN3270 server. The clients connect to the server using a TCP connection. The server provides a gateway function for the downstream TN3270 clients by mapping the client sessions to SNA-dependent LU-LU sessions that the server maintains with the SNA host. The TN3270 server handles the conversion between the TN3270 data stream and an SNA 3270 data stream.

To deploy a TN3270 solution, you install TN3270 client software on desktop workstations[5] and TN3270 server software in one of several places discussed below. Client software is available from IBM and many other vendors, and runs on top of the TCP/IP stack in the workstation. A given client product provides one of two possible levels of standards support:

- Base TN3270 client

  These clients conform to RFC 1576 (TN3270 Current Practices) and/or RFC 1646 (TN3270 Extensions for LU Name and Printer Selection).

- TN3270E client

  These clients conform to RFC 1647 (TN3270 Enhancements), and RFC 2355 (TN3270 Enhancements).

A server implementation that can support TN3270E clients is called a TN3270E server.

5.2 Placement of the TN3270 Server Function

The TN3270 server function can be placed in a variety of products and positions within a network, including:

- In the SNA host itself

  IBM and several other vendors provide host TN3270 server software that sits on top of the host TCP/IP stack and connects within the host to VTAM.

- In a router or Network Utility in the network

  IBM and other vendors provide TN3270 server function in networking hardware products. You can place these products directly adjacent to the

[5] You can also find small, dedicated TN3270 client products that represent printers.

SNA host, or at any position in the network where you have SNA connectivity to the host. If you are using IBM routers (2210 or 2216) or Network Utilities, and your host is running APPN, you can use Enterprise Extender technology to place the server at any position where you have IP connectivity to the host.

- In a software product in the network

  IBM and other vendors provide TN3270 server software products that you install on mid-range servers that use operating systems such as AIX, OS/2, or Windows/NT. You can place these products at any position in the network where you have SNA connectivity to the application host.

The choice of TN3270 server product and network position is a complex one, involving such factors as:

- Host capacity and cycle impact
- Price for performance and capacity
- Availability
- Impact of server failure
- Scalability

Network Utility provides a high-performing TN3270E server implementation that scales to large networks. By combining it with the Network Dispatcher feature, you can implement server redundancy and load sharing in large TN3270 installations. You can also place a Network Utility out into an SNA or IP network away from the data center and get the same advantages of scalability, incremental addition, and reduced impact of server failure.

5.3 Network Utility TN3270E Server Function

5.3.1.1 Standards Compliance

The Network Utility implementation of TN3270E server supports these RFCs:

RFC 1576 - TN3270 Current Practices
RFC 1646 - TN3270 Extensions for LU names and Printers
RFC 1647 - TN3270 Enhancements
RFC 2355 - TN3270 Enhancements (obsoletes RFC 1647)

It can handle both base TN3270 and TN3270E clients at the same time.

5.3.1.2 Host Connectivity

As mentioned above, the path from a TN3270 client to the SNA host consists of two pieces:

- A TCP connection over IP from the client to the server
- An SNA LU-LU session from the server to the host

The form of the SNA connection from the server to the host depends on how the server represents PUs and dependent LUs. When you are using Network Utility as your TN3270 server, you can configure either of two different ways to establish links and represent PUs and LUs to VTAM:

- Using SNA subarea links

  You set up Network Utility this way when you are not running APPN at the host. You configure a separate DLC-layer link to the host for every PU (maximum of 253 LUs). Multiple PUs require multiple parallel host links.

SNA frames arriving at Network Utility on one of these links flow directly to the corresponding internal PU.

Subarea host links must be a single DLC-layer hop to the product providing the SNA subarea boundary function. Typically, this product is either NCP running in a FEP, or is VTAM itself in the host. The subarea link from the Network Utility can traverse bridges or other DLC-layer forwarding mechanisms (such as protocol converters or external DLSw routers). Network Utility supports the following link types for subarea host attachment:

- Token-ring: physical, ATM LAN emulation, or channel LSA
- Ethernet: physical, ATM LAN emulation, or channel LSA
- FDDI: physical only
- Frame relay PVCs: bridged or routed RFC 1490/2427 formats

- Using an APPN dependent LU requester (DLUR) link

  You set up Network Utility this way when you are running APPN with its Dependent LU Server (DLUS) function at the host. You configure one DLC-layer link to the host to carry the DLUR-DLUS ″pipe,″ even if you are defining multiple local PUs. SNA frames arriving at Network Utility on this link flow to the DLUR function, which redirects them to the correct internal PU.

  When you are using DLUR, you can route through an APPN network using either ISR or HPR routing to reach the host. Network Utility supports the following link types as the ″first hop″ APPN link to the host:

  - Token-ring: physical, ATM LAN emulation, or channel LSA
  - Ethernet: physical, ATM LAN emulation, or channel LSA
  - FDDI: physical only
  - Frame relay PVCs: bridged or routed RFC 1490/2427 formats
  - ATM (native, not LAN emulation): HPR only
  - Channel MPC+: HPR only
  - PPP
  - SDLC: ISR only
  - X.25: ISR only
  - DLSw: ISR only
  - IP (Enterprise Extender): HPR only

  Note especially that when using DLUR and HPR routing, you can place a Network Utility TN3270E server across an IP network from the SNA application host. Enterprise Extender maintains session-level class of service and transmission priority across the IP network.

5.4 General TN3270E Server Configuration

This section covers general information about configuring Network Utility TN3270 server support. For specific example configurations, see page 73.

5.4.1 Configuring TN3270 Subarea under the APPN Protocol

In the Network Utility implementation of TN3270 server, all SNA functions are bundled within the APPN protocol. This means that *even when you are configuring SNA subarea attachment and your SNA host is not running APPN*, you must use the configuration and console services of the APPN protocol. In particular:

- You must go through the APPN protocol at the command line and the Configuration Program to configure ports, links, and TN3270 server functions.
- You must go through the APPN protocol at the command line to use TN3270 monitoring commands.
- You must still configure APPN at the node level.

When you configure SNA subarea support, Network Utility does in fact still function as an APPN network node, but only on links to other APPN nodes. If the *only* ports and links you configure are those for SNA subarea host attachment, then the APPN function serves no purpose.

5.4.2 Configuring in the APPN Environment

APPN and TN3270 server are fully configurable both from the Configuration Program and from the command line. From the Configuration Program, the TN3270 configuration parameters are always available. If you create a TN3270 configuration and download it to a Network Utility Model TX1, which does not support TN3270 server function, the Network Utility ignores the TN3270 part of the configuration. If you are working from the command line on a Model TX1, the commands for configuring and monitoring TN3270 simply do not appear on the APPN menus.

To change an APPN/TN3270 configuration from the Configuration Program, you make the change, transfer the configuration to the Network Utility, and reboot for it to take effect.

To change an APPN/TN3270 configuration from the command line, you move to talk 6, type `p appn`, and then issue the commands to make the change. You have two options for activating the change:

- Write the configuration to disk and reboot Network Utility to activate it.
- Issue the talk 6 APPN `activate` command to dynamically activate the modified APPN/TN3270 configuration.

  Depending on the configuration items you changed, APPN either makes the change immediately, or restarts APPN (but not the entire Network Utility) to activate the change. For the latter case, if you move to talk 5 and type `p appn` while APPN is restarting, you get the message `APPN is not currently active`. You can poll with talk 5 commands to see when the restart is complete.

  You can recycle the entire TN3270 server function in this way by disabling and enabling it with the `TN3270E config>` command set, and activating each of these configuration changes dynamically.

5.4.3 Implicit and Explicit LU Naming and Mapping

When you configure Network Utility's TN3270 server function, you create a local LU name for every one of the potential concurrent client sessions the Network Utility is intended to support. The LU name you define in the Network Utility need not have any relation to LU names in VTAM.

When a TN3270 client connects to a server over TCP, it can request a specific LU name, or it can place a generic request for any LU of a certain type. If you are configuring a client to request a specific name, you specify one of the local names defined at the server (Network Utility), not a VTAM LU name.

Because a single Network Utility can support thousands of LUs with similar characteristics, it does not require you to individually configure each LU. Rather, you can create a large pool of *implicit* LUs to satisfy clients that do not request a particular LU name. You then add a small number of *explicit* LUs to satisfy the clients that do request a particular name.[6]

As you will see in the example configurations, you define implicit LUs in groups as you define each local PU. You specify an LU name mask and number of LUs, but no NAU address. To configure an explicit LU, you specify an LU name and an NAU address (2-254). When the Network Utility activates the configuration, it reserves the NAU addresses for explicit LUs, and then generates names for the implicit LUs using the group name mask and one of the available NAU addresses.

When a TN3270 client connects in and does not explicitly request an LU name, Network Utility attaches the client to any available implicit LU. At this point, the server function treats all the implicit LUs as being in one large pool without regard to PU boundaries.

Note: The functional PTF of MAS V3.2, available in December 1998, includes significant functional enhancements in the area of LU definition and client mapping. With these enhancements:

- You can define named pools of LUs.
- You can configure mappings between client IP addresses and LU or LU pool names.
- The server can send VTAM a list of dependent LU addresses for each PU, so that VTAM can dynamically create its own LU definitions.
- You can configure multiple local TCP ports for the TN3270 server function.

Refer to the MAS V3.2 *MAS Protocol Configuration and Monitoring Reference Volume 2*, SC30-3885, for more information on configuring these functions.

5.5 Example Configurations

Network Utility as a TN3270E server can be deployed in several configurations. For example, it can be placed either in the remote branch or in the data center. It can attach to the host via a traditional SNA subarea connection or it can use APPN. In the data center, it can be placed in a channel-attached configuration or it can be a stand-alone server that resides on the campus LAN (or ATM cloud) using the channel-attached connection provided by an existing IBM 3745/46 Communication Controller, 2216-400, 3172 Interconnect Controller, an OSA adapter, or an OEM gateway.

One of the most important elements of a TN3270 implementation is scalability. The Network Utility solution can scale to very large configurations while providing high availability and redundancy.

The following scenarios show you how to effectively utilize the Network Utility as a TN3270E server.

[6] The implicit/explicit distinction is solely within the Network Utility. A client can request an implicit LU name, and the Network Utility will satisfy the request if the LU is available. The key point is that the server function will never assign an explicit LU to a client unless the client specifically requests that LU name.

5.5.1 TN3270 Via a Subarea Connection to an NCP

This scenario (shown Figure 3) shows a traditional SNA subarea network with all host access occurring through an IBM 3745/46 Communication Controller with the IBM Network Control Program (NCP). The Network Utility is installed to provide TN3270 server support for downstream workstations both in the local campus and in the remote sites. The Network Utility attaches to the host through the FEP via a normal subarea connection.

Up to 16000 TN3270 sessions can be handled with a single Network Utility installed as shown in Figure 3. As your network grows, the solution can be scaled simply by adding more TN3270E server capacity via additional Network Utilities. You can also set up automatic load balancing among your TN3270E servers by installing a separate IBM router or Network Utility to serve as a Network Dispatcher. (See 5.5.4, “Highly Scalable, Fault-Tolerant TN3270E” on page 77 for an example of how to scale the network.)

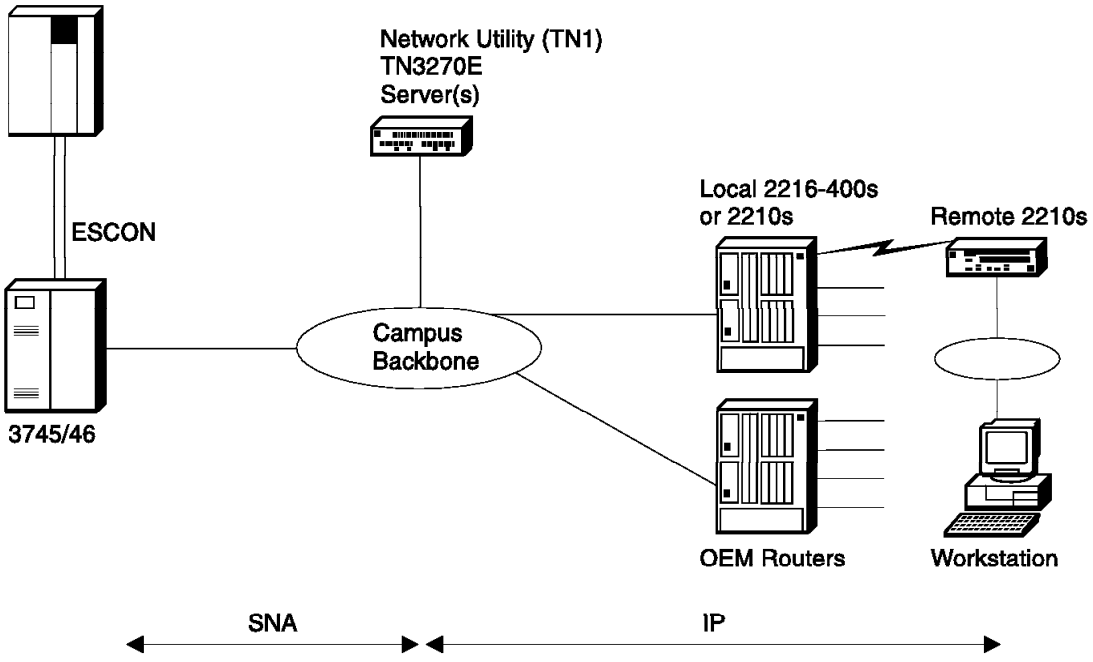

Figure 3. TN3270 Via a Subarea Connection through a 37xx

5.5.1.1 Keys to Configuration

The configuration of the TN3270E server function is very straightforward in this scenario. However, the following points are worth noting:

- There is both an APPN and a subarea implementation of the TN3270E server. Both require APPN support to be installed on the Network Utility and both are configured within the APPN configuration process. This is true even though a pure subarea configuration does not use the APPN function. This is an implementation statement as the TN3270E server function uses the APPN SNA stack for both subarea and APPN connections to the host.

  Note also these additional points relating to APPN and TN3270E server configuration:

  – APPN support must be enabled.

  - You must define a port and one or more link stations to define the connection to VTAM.
  - For subarea configurations, defining a link station and specifying to solicit an SSCP session implicitly defines a PU on the Network Utility. This PU will support up to 253 downstream LUs. If you need more than 253 LUs, then you need to define more than one link station. Each link station needs to use a different service access point (SAP) and a different local node ID (IDNUM).

- When configuring the parameters for the TN3270E server, you can set the IP address of the server to either the internal box IP address or to one of the interface IP addresses. Keep in mind that the address you select for TN3270 may be unavailable for using normal IP Telnet to manage the box.[7]
- The downstream LUs can be defined either as explicit or implicit.
  - Use explicit definitions when you need to ensure that the device will always use the same LU name. (For example, printers would normally use explicit definitions.)
  - Use implicit definitions when you have a large group of devices that can use a common pool of available LUs and do not need to use the same LU name every time.

For a complete look at the configuration parameters needed for this scenario, see Table 2 on page 89.

5.5.2 TN3270 Via a Subarea Connection through a Channel Gateway

This scenario, shown in Figure 4 on page 76, is similar to the previous scenario except that here the Network Utility attaches to the host through a LAN channel gateway such as an IBM 3172, an IBM 2216, an IBM 3746 with the Multiaccess Enclosure (MAE) or an OEM device. These gateways use External Communications Adapter (XCA) pass-through and do not provide the SNA boundary function normally provided by an NCP. With a gateway, this function is provided by VTAM.

If you have an existing gateway with a TN3270 server configured, you can use the Network Utility to offload the existing TN3270 workload or to provide additional TN3270 capacity as your network requirements grow.

An existing 2216 or a 3746 allows you to have multiple channel connections to the host while you can incrementally install Network Utilities for your TN3270E server requirements. The dynamic load balancing features of network dispatcher can be used to optimize efficiency.

---

[7] If you need to use Telnet at this same address, you can configure the TN3270E server to use another port (24 for example) so that telnet can use port number 23. This requires that the TN3270 client workstations be configured to use this same port.

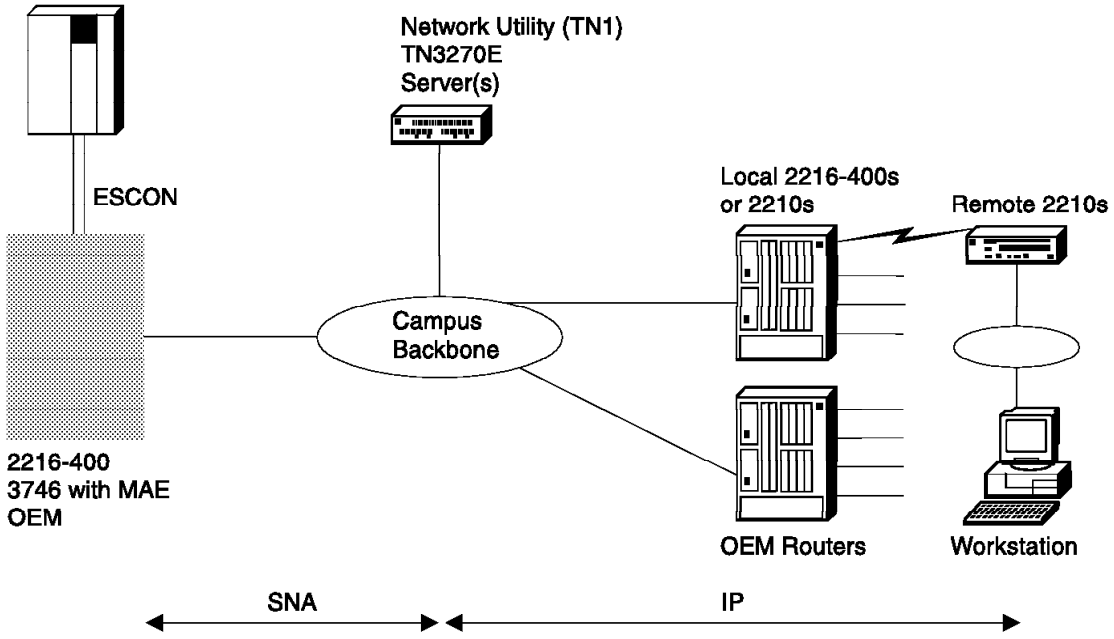

Figure 4. TN3270 Via a Subarea Connection through a LAN Gateway

5.5.2.1 Keys to Configuration

From the Network Utility perspective, the configuration of this scenario is identical to the previous one. The host definitions are also identical. For both scenarios, you just have to define the switched major nodes for the PUs in the TN3270E server.

5.5.3 TN3270 through an OSA Adapter

This scenario is shown in Figure 5 on page 77. Here, the Network Utility attaches to the host through the S/390 Open Systems Adapter (OSA). Like the previous gateway scenario, the SNA boundary function is in the host.

While the TN3270 server function can reside on the host itself, many customers prefer to offload this function externally to another platform. The Network Utility meets this requirement well by providing scalable, cost-effective TN3270E server function without changing your method of host attachment. This allows you to leverage your existing investments.

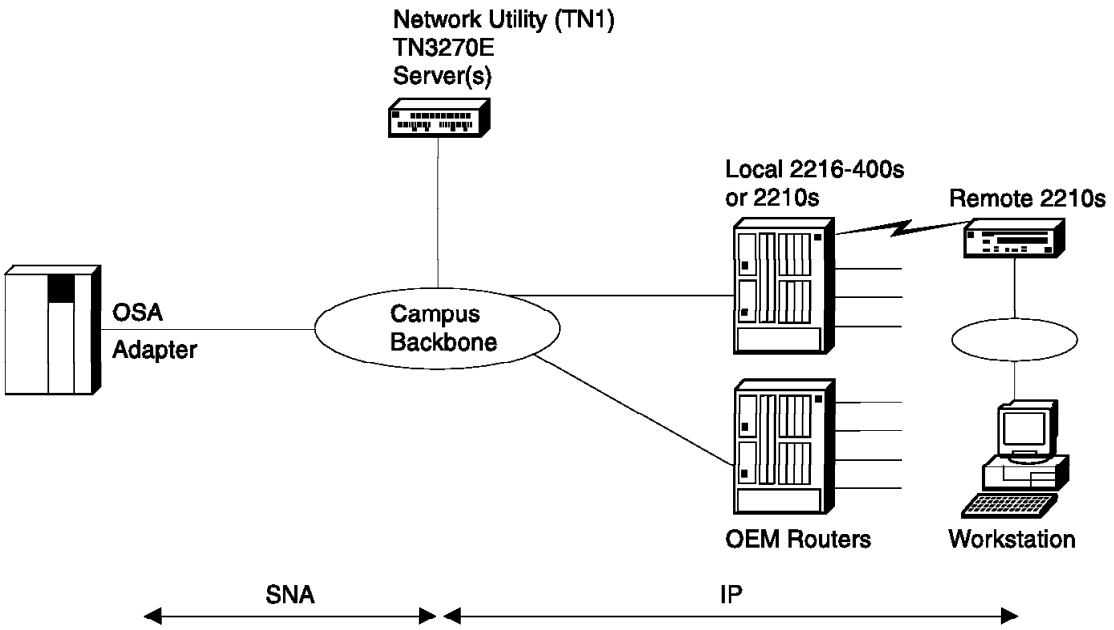

Figure 5. TN3270 Via an OSA Adapter

5.5.3.1 Keys to Configuration

From the Network Utility perspective, the configuration of this scenario is identical to the previous two.

5.5.4 Highly Scalable, Fault-Tolerant TN3270E

This scenario, shown in Figure 6 on page 78, is an extension of the one discussed in 5.5.1, "TN3270 Via a Subarea Connection to an NCP" on page 74. Here, the solution is scaled with multiple Network Utility devices to provide TN3270E server support for large 3270 environments. Also, a separate Network Utility is configured as a Network Dispatcher and deployed to provide load balancing[8] . The new Network Dispatcher Advisor for TN3270 allows the Network Dispatcher to collect load statistics from each Network Utility TN3270E server in real time to achieve the best possible distribution among the TN3270 servers.

The solution provides high availability in the event of a failure in one of the TN3270E servers. The server that the client session is dispatched to is transparent to the user. If a failure occurs, the sessions through that server are lost but the users simply log back on to the host through another Network Utility using the same destination IP address for the TN3270E server.

The Network Dispatcher function can also utilize redundant hardware, with a second Network Utility configured as a Network Dispatcher and serving as a backup to the primary one.

With this configuration, you can scale your TN3270E support to any size simply by adding additional TN3270E server capacity. You can do this incrementally and non-disruptively as your network requirements grow.

---

[8] As of MAS V3.2, the Network Dispatcher function can also dispatch client sessions to the TN3270 server function running in the same Network Utility.

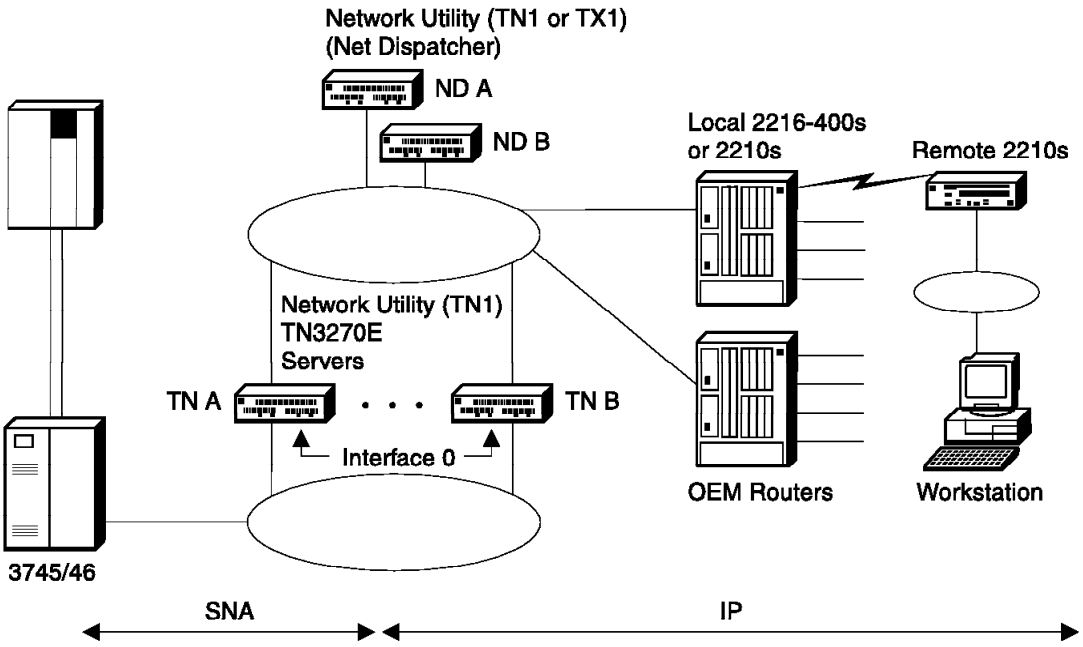

Figure 6. Highly Scalable, Fault-Tolerant TN3270E

5.5.4.1 Keys to Configuration

As far as the TN3270E server is concerned, the configuration is the same whether you have a Network Dispatcher or not. In fact, the TN3270E server is unaware that the traffic from the clients is being dispatched through another machine. See 5.5.1, “TN3270 Via a Subarea Connection to an NCP” on page 74 for the basic configuration points for a TN3270E server. See Table 3 on page 95 for the complete set of configuration parameters for the TN3270E servers for this scenario.

However, the IP addressing needs special attention in this configuration for high availability. In 5.5.1, “TN3270 Via a Subarea Connection to an NCP” on page 74, the TN3270E server was configured with the same address as the router ID (also the same address as the LAN interface). In a Network Dispatcher environment, the IP addressing is somewhat different.

A Network Dispatcher and one or more TN3270E servers form what is called a cluster. An IP address is defined for the cluster and workstations send their TN3270 packets to this IP address. The Network Dispatcher receives these packets and forwards them on to a server in the cluster for processing.

Because the Network Dispatcher does not alter the destination IP address of these packets, each TN3270E server also needs to be configured with this same IP address. However, you have to make sure that the TN3270E servers do not broadcast this address via OSPF or RIP to the network because you do not want these servers to respond to the cluster address. Only the Network Dispatcher should respond to the cluster address.[9]

[9] The cluster address cannot be pinged. The Network Dispatcher does not respond to pings to the cluster address. It only processes TCP and UDP packets.

The router must know the TN3270E server's IP address in order to forward packets to the server function. One way to make this address known to the router is to specify it to an interface as a secondary address. Figure 7 shows an example of this IP addressing scheme for the highly available, fault-tolerant TN3270 configuration depicted in Figure 6 on page 78.

```
TN3270E Server #1 (TNA):
    Internal address  172.128.252.3
    Interface 0       172.128.2.3      (2nd address: 172.128.1.100)
    Interface 1       172.128.1.3
    OSPF Router ID    172.128.1.3
    TN3270E Server    172.128.1.100    (same as cluster address)

TN3270E Server #2 (TNB):
    Internal address  172.128.252.4
    Interface 0       172.128.2.4      (2nd address: 172.128.1.100)
    Interface 1       172.128.1.4
    OSPF Router ID    172.128.1.4
    TN3270E Server    172.128.1.100    (same as cluster address)

Network Dispatcher #1 (NDA):
    Internal address  172.128.252.1
    Interface 0 addrs 172.128.1.1
    OSPF Router ID    172.128.1.1
    Cluster address   172.128.1.100
       Port 23
         Server 1     172.128.1.3
         Server 2     172.128.1.4

Network Dispatcher #2 (NDB):
    Internal address  172.128.252.2
    Interface 0 addrs 172.128.1.2
    OSPF Router ID    172.128.1.2
    Cluster address   172.128.1.100
       Port 23
         Server 1     172.128.1.3
         Server 2     172.128.1.4
```

Figure 7. IP Addressing for the Highly Scalable, Fault-Tolerant TN3270 Scenario

Note that the cluster address is assigned as a second IP address on interface 0 of the Network Utility machines. In this scenario, the LAN segment that interface 0 attaches to does not carry any IP traffic – only the SNA subarea traffic from the TN3270E server to the host.

The configuration of the Network Dispatchers is standard. For the complete set of configuration parameters needed for this scenario, see Table 4 on page 100 for the primary Network Dispatcher. For differences from this configuration for the backup Network Dispatcher, see Table 5 on page 104.

5.5.4.2 Explicit LUs and Network Dispatcher

Special care has to be taken for explicit LU definition in a Network Dispatcher environment. A session request for either an implicit or an explicit LU can be dispatched to any server. This means that the explicit LU has to be defined in each server, because it is not known in advance to which server the session will be dispatched. The explicit LU in this environment (a printer, for example) is represented by two different LUs in VTAM. Each PU in the TN3270E servers that have the LU defined has to have a unique node ID (IDNUM), because VTAM does not allow duplicate PU or LU names to be active at the same time.

When a server has both explicit LUs and implicit LU pools defined, if all the sessions of the pool are used, the server cannot handle any more requests for sessions from this pool. But, the ND still dispatches sessions to this server, because this server does not report 100% load.

One way to handle explicit LUs is to define them on a separate TN3270E server that is outside the Network Dispatcher environment. This is normally acceptable because there are many fewer of them and hence they do not require the dynamic load balancing feature of Network Dispatcher. Also, devices that use explicit definitions (such as printers) often have a much lower acceptable availability requirement even in an otherwise fault-tolerant environment.

5.5.5 TN3270 Via DLUR over APPN

This scenario, shown in Figure 8 on page 81, uses APPN to communicate with the host. The Network Utility uses APPN High Performance Routing (HPR) and establishes a Rapid Transport Protocol (RTP) session with the host. HPR is used all the way from the TN3270E server to VTAM. In case of a failure, this assures nondisruptive session switching to an alternate path if you have parallel gateways. This is especially important in Parallel Sysplex environments.

In addition, HPR is supported over IP through the Enterprise Extender feature of the Network Utility. This is important if you want to place your TN3270E server in a remote location and use IP to transport the APPN traffic back to your data center.

The channel gateway is an APPN network node performing APPN Automatic Network routing (ANR) for the RTP session between the Network Utility and the host.

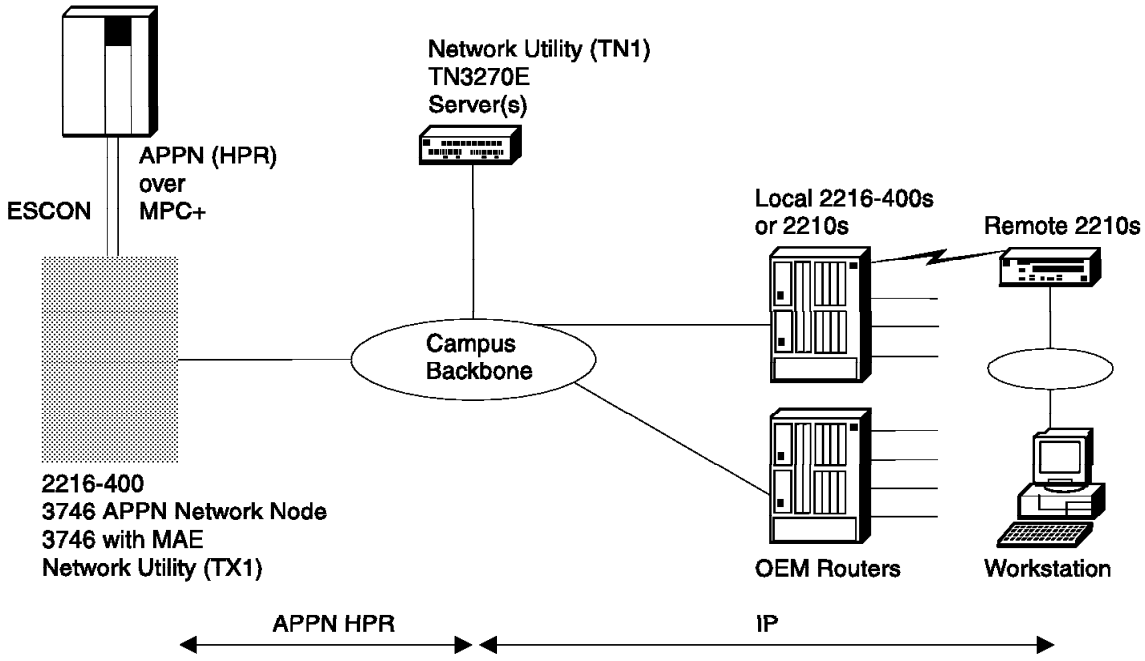

Figure 8. TN3270 Via DLUR over APPN

When connecting a TN3270E server to the host via APPN, you must configure DLUR support on the Network Utility. The DLUR feature extends to APPN nodes the support of T2.0 or T2.1 devices containing dependent LUs. The DLUR function on an APPN network node works in conjunction with a dependent LU server (DLUS). The DLUS function is usually provided by VTAM, although it may reside in any part of a mixed APPN/subarea network.

The dependent LU flows (SSCP-PU and SSCP-LU) are encapsulated in an LU 6.2 pipe (CP-SVR) established between the DLUR APPN node and the DLUS SSCP. The CP-SVR pipe is made up of a pair of LU 6.2 sessions using a new CPSVRMGR mode between the DLUR and the DLUS. This pipe brings the SSCP function (in the DLUS) to the DLUR APPN node where it can be made available to attach T2.0/T2.1 nodes containing dependent LUs.

5.5.5.1 Keys to Configuration

From a downstream workstation perspective, the TN3270E server appears the same whether that server is using SNA subarea or APPN to communicate with the host on the uplink. At the Network Utility, you configure the base TN3270 server parameters the same way as in the SNA subarea scenarios, but the way you configure the local PUs differs. Instead of associating each PU with a subarea link, you configure local PUs without any link association. The DLUR function is responsible for routing traffic on the DLUS-DLUR pipe to and from these local PUs.

APPN requires dependent LU requester (DLUR) support to be configured in the Network Utility. DLUR is quite simple to configure with the only required parameter being the CP name of the dependent LU server (DLUS), which is VTAM.

You have to make some additional host definitions for APPN and DLUR support. See Appendix A, “Sample Host Definitions” on page 167 for an example of these commands.

For a complete look at the configuration parameters needed for this scenario, see Table 6 on page 107.

5.5.6 Distributed TN3270E Server

The previous configurations showed how the Network Utility can be deployed in the data center to centralize the TN3270E server function in your network. This configuration, shown in Figure 9, shows just one example of how the Network Utility can also be placed in a remote location to provide distributed TN3270E server capability.

In this configuration, the Network Utility is providing TN3270E server service to workstations in the remote location. As always with a TN3270 configuration, the workstations are using IP to communicate with the TN3270E server. The TN3270E server is using DLUR over an APPN connection back to the host in the data center.

In this example, the corporate WAN is a public Frame Relay network that carries IP traffic only. Therefore, the Network Utility is configured to use the Enterprise Extender feature which allows the APPN HPR traffic to be carried over the IP-only WAN.

The Enterprise Extender traffic is terminated at the host gateway, which decapsulates the HPR traffic and then passes the APPN traffic through the network node onto the MPC+ path to the host. This is a very fast, low-overhead packet-forwarding function, so a single gateway can handle a large amount of traffic.

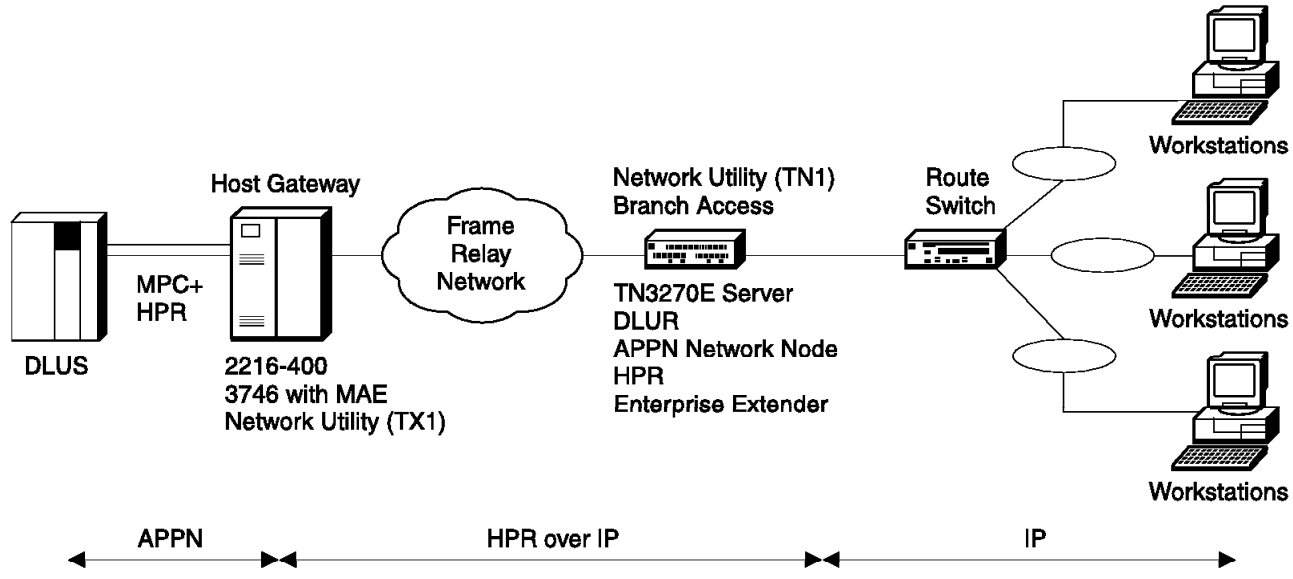

Figure 9. Distributed TN3270E Server

5.5.6.1 Keys to Configuration

From a downstream workstation perspective, the TN3270E server appears the same whether it is in the remote branch or in the data center regardless of whether the upstream connection to the host is using SNA subarea or APPN. Therefore, the TN3270E server function in the Network Utility is configured exactly the same as in the previous scenarios.

APPN and DLUR are configured the same as in 5.5.5, "TN3270 Via DLUR over APPN" on page 80 with one exception, which is the port definition for APPN over the frame relay IP link. When configuring APPN to use HPR over IP (the Enterprise Extender feature), you specify a port type of IP. Then when adding the link station for this port, instead of specifying the MAC address of the adjacent FEP as was done in 5.5.5, "TN3270 Via DLUR over APPN" on page 80, you specify the IP address of the other end of the HPR over IP network, which is the host gateway in this example.[10] The IP network is responsible for the delivery of the traffic to the host gateway over the best path available. You are assured a reliable transport because the connection between the TN3270E Server and the host uses an RTP session.

5.6 Managing the TN3270E Server

This section introduces some of the ways in which you can monitor and manage the TN3270E server function.

Note: The monitoring functions described in this section assume that you are running MAS V3.2 or later operational code. MAS V3.2 introduced several new TN3270 monitoring commands, as well as a TN3270E submenu.

5.6.1 Command-Line Monitoring

To view currently running TN3270 server status from the command line, move first to talk 5 and enter `p appn`. If you get the message `Protocol APPN is available but not configured`, you need to complete your base APPN configuration and reboot Network Utility to activate APPN. As discussed in 5.4.1, "Configuring TN3270 Subarea under the APPN Protocol" on page 71, you need APPN to be active even if you are using only TN3270 subarea connectivity.

Once you have reached the APPN monitoring prompt `APPN >`, type `tn` (short for "TN3270E") to reach the submenu for monitoring TN3270E server status.

The following commands are then available at the `TN3270E >` monitoring prompt:

`list status`

If the system responds `TN3270E is not configured or not active`, you did not enable the TN3270 server function adequately in the currently active APPN configuration. If you get this error and you did configure the function, perhaps the TN3270 server IP address you chose is not active as an interface address or as the internal IP address. Consult the examples of TN3270 configurations in Chapter 6 for other possible reasons, then change your APPN/TN3270 configuration and activate it as described in 5.4.2, "Configuring in the APPN Environment" on page 72.

[10] The host gateway must also be configured with an HPR over IP port in much the same manner as described here.

If the server function is active, this command provides the following information:

- Configuration information currently in use

  TN3270E IP Address
  The server IP address to which clients connect, also the cluster address if you are using Network Dispatcher

  TN3270E Port Number
  The TCP port to which clients connect

  NetDisp Advisor Port Number
  The TCP port to which Network Dispatchers can connect to retrieve load information

  Keepalive type
  Whether and how the server polls clients to see if they are still active. Possible values are:

None	Server does not poll clients, and will discover client absence only when trying to send data.
NOP	Server polls clients at the TCP level, client software need not have capability to respond.
Timing mark	Server polls clients at the TN3270 level, and client software must respond within a certain time window.

  Automatic Logoff
  Whether or not the server disconnects clients after a period of inactivity (with no data flowing in either direction)

- Summary statistics

  Number of connections
  The current number of active TCP connections from TN3270 clients

  Number of connections in SSCP-LU state
  The number of currently active client TCP connections that have an associated LU in this state (received an ACTLU but not yet a BIND)

  Number of connections in LU-LU state
  The number of currently active client TCP connections that have an associated LU in this state (received BIND, fully active)

`list connections`

You can type this command with or without modifiers:

- `list connections`

  Displays all currently active client connections (those with an active TCP connection).

- `list connections` *client ip address*

  Displays all the currently active connections that originated from the specified IP address.

- `list connections` *LU name*

  Displays all the currently active connections that are associated with the specified LU name.

For each of the `list connection` commands, the following information is displayed for each session:

Local LU	The LU name, configured at Network Utility, to which the server function has mapped this client TCP connection
Class	The type of LU, as follows: **IW** Implicit workstation

	EW	Explicit workstation
	IP	Implicit printer
	EP	Explicit printer
Assoc LU	For a workstation LU, the name of any associated printer LU	
Client Addr	The IP address of the client	
Status	Whether the connection is in SSCP-LU state or LU-LU state	
Prim LU	The primary LU name as known to VTAM	
Sec LU	The secondary LU name as known to VTAM	
Idle Min	The number of minutes since this connection carried any user data	

Besides the above list commands, a TN3270 server user needs to be able to query the status of other APPN or SNA resources on which the function depends. The following APPN monitoring commands are of general use:

`aping` - to test connectivity to a remote LU
`li port` - to show interface status
`li link` - to show the status of logical links

If you are using DLUR for your host connection, the following commands are particularly useful:

li appc - to check the status of the DLUS-DLUR pipe
li local - to show the status of internal PUs used by the TN3270 server function

5.6.2 Event Logging Support

In general, APPN/TN3270 ELS messages are intended to capture debug and trace information for IBM support personnel. These functions have extensive logging and trace support, but the ELS messages themselves are tightly packed with low-level information.

Normally, you activate APPN/TN3270 tracing and logging under the direction of IBM support personnel. The general procedure is to enable some of a large list of possible traces as part of your APPN configuration. From the Configuration Program, see the APPN Node Services folder. From talk 6, use the `set trace` command. After you activate this configuration change, the output of these traces flows into a trace table in APPN memory, and also to ELS if you have APPN ELS messages active. Should you have a problem that requires activating traces, IBM support will provide detailed procedures to guide you in capturing debug information.

5.6.3 SNA Management Support

APPN generates SNA alerts for a variety of error conditions, and can forward alerts from other SNA devices. There are no alerts specific to the TN3270 server function, but alerts that a Network Utility itself generates may relate to SNA resources involved with TN3270.

5.6.4 SNMP MIB and Trap Support

Network Utility supports an Internet Draft version of both of the forthcoming standard MIBs for TN3270 server function:

TN3270 Base MIB
TN3270 Response Time MIB

Network Utility support for these MIBs includes the ability to:

- View server configuration, status, and statistics
- Set up client groups for response time collection
- Map LU names from VTAM name to local name to client IP address
- Map client IP addresses to VTAM LU names
- Collect response time data for current client groups

In addition, Network Utility supports the following IETF MIBs relating to APPN and SNA functions:

RFC 2155, APPN
RFC 2051, APPC
RFC 2232, DLUR
RFC 2238, HPR
RFC 1666, SNA NAU
Internet Draft, Extended Border Node

Network Utility supports the following IBM Enterprise-Specific MIBs relating to APPN functions:

APPN Memory
APPN Accounting
APPN HPR NCL
APPN HPR Route Test
APPN Peripheral Access Node (Branch Extender)

These MIBs provide a comprehensive view of APPN and SNA resources within Network Utility, including those being used for TN3270.

5.6.5 Network Management Application Support

The Nways Manager products provide specialized statistical support for TN3270 response time monitoring, as well as the ability to view TN3270 server resources. To initiate response time monitoring, you select a group of one or more clients using an IP address and mask. For each group you define, the manager collects response time statistics into predefined time buckets (less than 1 second, 1 to 2 seconds, and so on). Using the collected information, you can view aggregate historical response time by group, or create custom reports that present the data in different graphical formats.

To view TN3270 resources and their status, you use specific panels that combine information from different tables within the base TN3270 MIB. To view APPN and SNA resources in general, you use specific panels that access information from the APPN MIBs. You can also use integrated browser support to view the information in any of these MIBs.

Nways Manager for AIX provides an APPN-level view of the topology of your network. You can discover participating APPN resources, view them, and view their status as color-coded icons. APPN protocol performance and error events (data and graph) are also provided. This application does not represent Branch Extender or Extended Border Node topologies.

Chapter 6. TN3270E Server Example Configuration Details

This chapter contains diagrams and configuration parameter tables for several of the examples of TN3270E server network configurations in Chapter 5, “TN3270E Server” on page 69. The parameter values shown are from real working test configurations.

For an explanation of the columns and conventions in the configuration parameter tables, see 4.2.2, “Example Configuration Table Conventions” on page 67.

The Network Utility World Wide Web pages contain binary configuration files that match these configuration parameter tables. To access these files, follow the Download link from:

http://www.networking.ibm.com/networkutility

The configurations documented in this chapter are:

Table 1. Cross-Reference of Example Configuration Information

Configuration Description	Parameter Table
5.5.1, “TN3270 Via a Subarea Connection to an NCP” on page 74	Table 2 on page 89
5.5.4, “Highly Scalable, Fault-Tolerant TN3270E” on page 77, for the TN3270 server TN A	Table 3 on page 95
5.5.4, “Highly Scalable, Fault-Tolerant TN3270E” on page 77, for the Network Dispatcher ND A	Table 4 on page 100
5.5.5, “TN3270 Via DLUR over APPN” on page 80	Table 6 on page 107

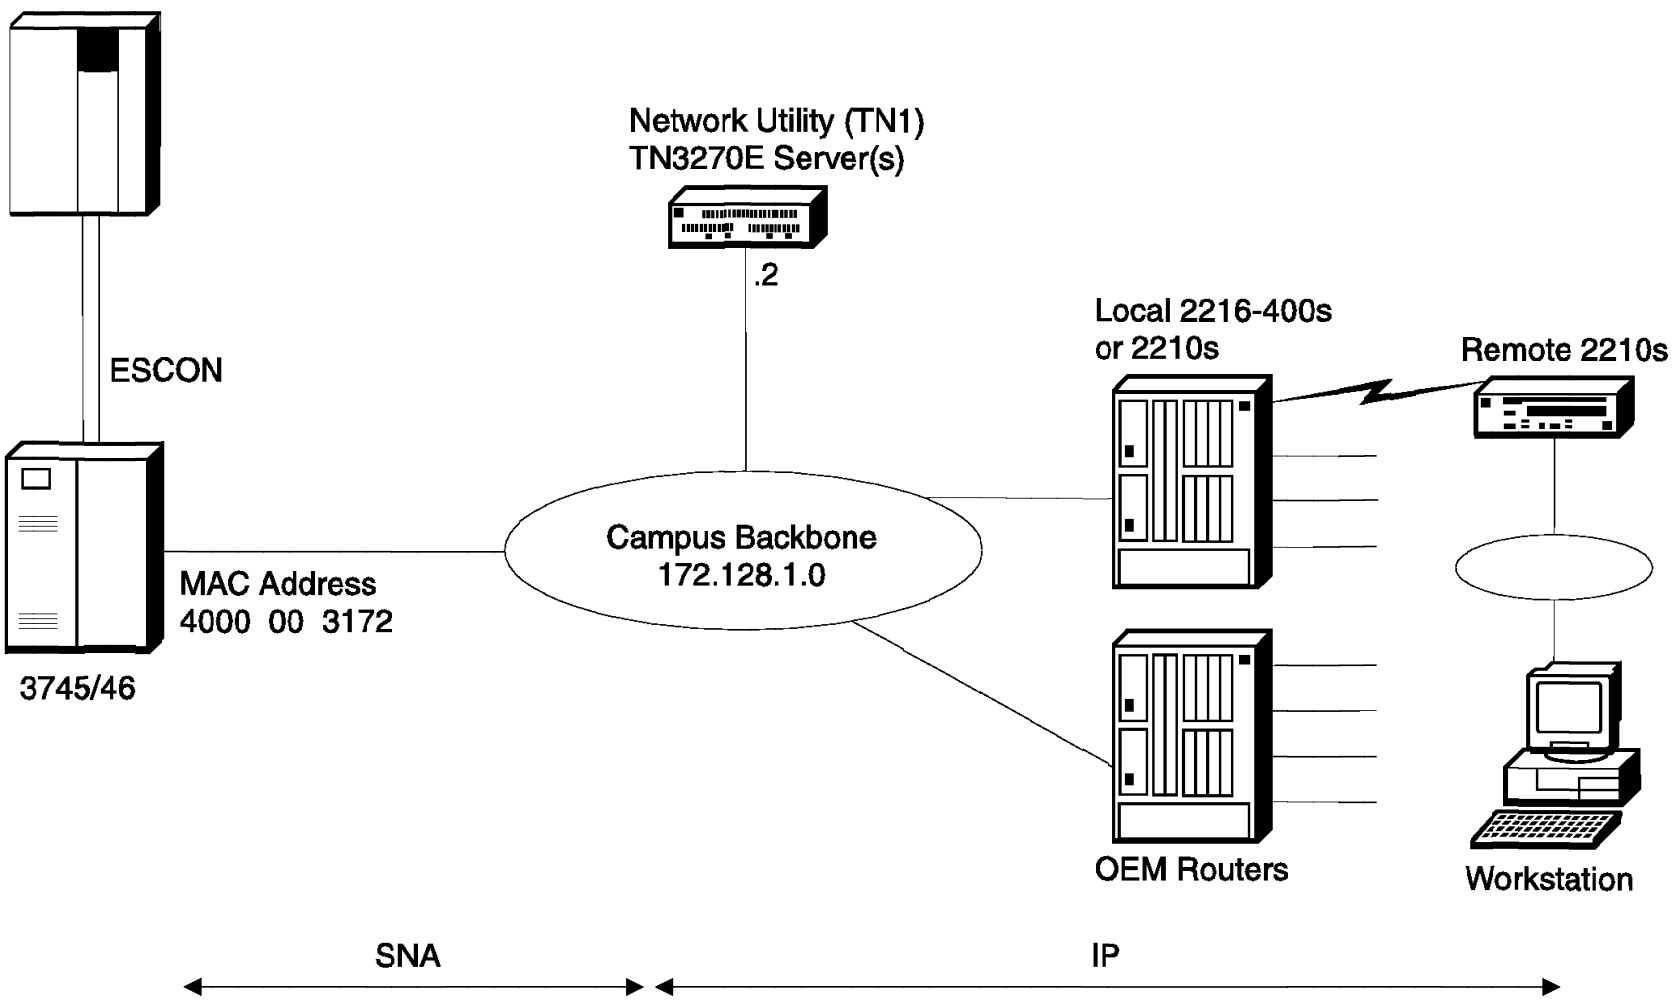

Figure 10. TN3270E Subarea

Table 2 (Page 1 of 4). TN3270E Subarea. See page 74 for a description and 88 for a diagram of this configuration.			
Configuration Program Navigation	**Configuration Program Values**	**Command-Line Commands**	**Notes**
Devices Adapters Slots	Slot1: 2-Port TR	See "add dev" on next row	1
Devices Adapters Ports	Slot 1/Port 1: Interface 0: TR	Config>**add dev tok**	2
Devices Interfaces	Interface 0 MAC Address 400022AA0001	Config>**net 0** TKR Config>**set phy 40:00:22:AA:00:01**	
System General	System name: NU_A Location: XYZ Contact: Administrator	Config>**set host** Config>**set location** Config>**set contact**	
System SNMP Config General	SNMP (checked)	Config>**p snmp** SNMP Config>**enable snmp**	
System SNMP Config Communities General	Community name: admin Access type: Read-write trap Community view: All	SNMP Config>**add community** SNMP Config>**set comm access write**	3
Protocols IP General	Internal address: 172.128.252.2 Router ID: 172.128.1.2	Config>**p ip** IP Config>**set internal 172.128.252.2** IP Config>**set router-id 172.128.1.2**	
Protocols IP Interfaces	Interface 0 (TR slot 1 port 1) IP address: 172.128.1.2 Subnet mask: 255.255.255.0	IP Config>**add address**	
Protocols IP OSPF General	OSPF (checked)	Config>**p ospf** OSPF Config>**enable ospf** (Accept other defaults)	
Protocols IP OSPF Area Configuration General	Area number: 0.0.0.0 Stub area (not checked)	OSPF Config>**set area**	

Table 2 (Page 2 of 4). TN3270E Subarea. See page 74 for a description and 88 for a diagram of this configuration.

Configuration Program Navigation	Configuration Program Values	Command-Line Commands	Notes
Protocols IP OSPF Interfaces	Interface 0 OSPF (checked)	OSPF Config>**set interface** Interface IP address: **172.128.1.2** Attaches to area: **0.0.0.0** (Accept other defaults)	
Protocols APPN General	APPN network node (checked to enable) Network ID: NUBNODE Control point name: CPNU	Config>**p appn** APPN config> **set node** Enable APPN Network ID: **NUBNODE** Control point name: **CPNU** (Accept other defaults)	4
Protocols APPN Interfaces	(highlight Interface 0 Token Ring) (click on the configure tab) Define APPN port (checked to enable) Port name: TR3270 High performance routing (HPR) supported (unchecked to disable) Support multiple PUs (checked to enable)	APPN config>**add port** APPN Port Link Type: **TOKEN RING** Port name: **TR3270** Enable APPN Support multiple PUs High performance routing: **No** (Accept other defaults)	5

Table 2 (Page 3 of 4). TN3270E Subarea. See page 74 for a description and 88 for a diagram of this configuration.

Configuration Program Navigation	Configuration Program Values	Command-Line Commands	Notes
Protocols APPN Interfaces	(highlight Interface 0 Token Ring) (click on the Link stations tab) STAT001 (new definition) General-1 Tab: Link station name: STAT001 Solicit SSCP session (checked) Link support APPN functions (unchecked) General-2 Tab: MAC address of adjacent node: 400000003172 Node ID: 12244 Local SAP address: 04 (click on **Add** to create the Link station) STAT002 (new definition) General-1 Tab: Link station name: STAT002 Solicit SSCP session (checked) Link support APPN functions (unchecked) General-2 Tab: MAC address of adjacent node: 400000003172 Node ID: 12245 Local SAP address: 08 (click on **Add** to create the Link station)	APPN config>**add link** Port name for the link station: **TR3270** Station name: **STAT001** MAC address of adjacent node: **400000003172** Solicit SSCP Session: **Yes** Local Node ID: **12244** Local SAP address: **4** Does link support APPN function?: **No** (Accept other defaults) APPN config>**add link** Port name for the link station: **TR3270** Station name: **STAT002** MAC address of adjacent node:**400000003172** Solicit SSCP Session: **Yes** Local Node ID: **12245** Local SAP address: **8** Does link support APPN function?: **No** (Accept other defaults)	6
Protocols APPN TN3270E Server General	TN3270E (checked to enable) IP address : 172.128.1.2 Automatic logoff (checked to enable)	APPN config>**tn** TN3270E config>**set** Enable TN3270E Server TN3270E Server IP Address: **172.128.1.2** Automatic logoff: **Yes** (Accept other defaults)	7

Table 2 (Page 4 of 4). TN3270E Subarea. See page 74 for a description and 88 for a diagram of this configuration.

Configuration Program Navigation	Configuration Program Values	Command-Line Commands	Notes
Protocols APPN TN3270E Server LUs	Local PU Name: STAT001 (click on **Implicit Pool**) LU name mask: @LU1A Number of implicit workstation definitions: 10 Local PU Name: STAT002 (click on **Implicit Pool**) LU name mask: @LU2A Number of implicit workstation definitions: 10 (click on **LUs** to define explicit LUs) LU name: PC03A NAU address: 5 (click on **Add**)	TN3270E config>**add imp** Station Name: **STAT001** LU name mask: **@LU1A** Number of Implicit LUs in Pool: **10** TN3270E config>**add imp** Station Name: **STAT002** LU name mask: **@LU2A** Number of Implicit LUs in Pool: **10** TN3270E config>**add lu** Station Name: **STAT002** LU name: **PC03A** NAU address: **5**	8

Notes:

1. add dev defines a single port, not an adapter.
2. The configuration program assigns an interface number to all ports of an adapter automatically and you delete the ones you do not want to use. From the command line, you type the add dev command for each port you want to use, and the interface number (also known as ″net number″) is the output of the command.
3. You need a write-capable SNMP community only if you want to download configuration files from the configuration program directly into the router. SNMP is not required to TFTP a configuration file to the router.
4. If you have a pure SNA subarea network with no APPN, then the Network ID can be any value. If you have APPN in your network, then the Network ID should conform to your APPN network naming conventions.
5. APPN must be enabled even though this example uses SNA subarea for the TN3270E server connection to the host. This is because the TN3270E server code uses the APPN SNA stack both for APPN and subarea communications to the host.
6. When you create the link stations, you are also implicitly creating PUs. These PUs are assigned a ″Local Node ID″ here. This must match the ″IDNUM″ in VTAM′s SW Major Node definition. The ID Block is always 077 for a Network Utility. If you need to define multiple link stations (PUs), then each link station has to have a different Local SAP address.

   Setting Solicit SSCP session to yes defines the link as a subarea connection.
7. Beginning with MAS V3.2, TN3270E Server has its own command-line submenu.
8. For implicit LUs, you just have to define the pools. The @LU1A is a template that will be used to create the actual LU names in the pool. In this example, with 10 LUs in the pool, the LU names generated are @LU1A2, @LU1A3, @LU1A4,...@LU1A11 which correspond to LOCADDRs 2-11 for the PU defined in VTAM. Similarly, @LU2A will generate @LU2A2, @LU2A3, @LU2A4. Note that the LU name @LU2A5 is not used because the NAU address of 5 has been reserved for the explicit definition. Therefore, the remaining LUs in the pool are @LU2A6 through @LU2A12.

   For explicit LUs, the LU name given here must match the name defined in the workstation′s 3270 emulation configuration. The NAU address points to the LOCADDR in the appropriate PU definition in the Switched Major node in VTAM.

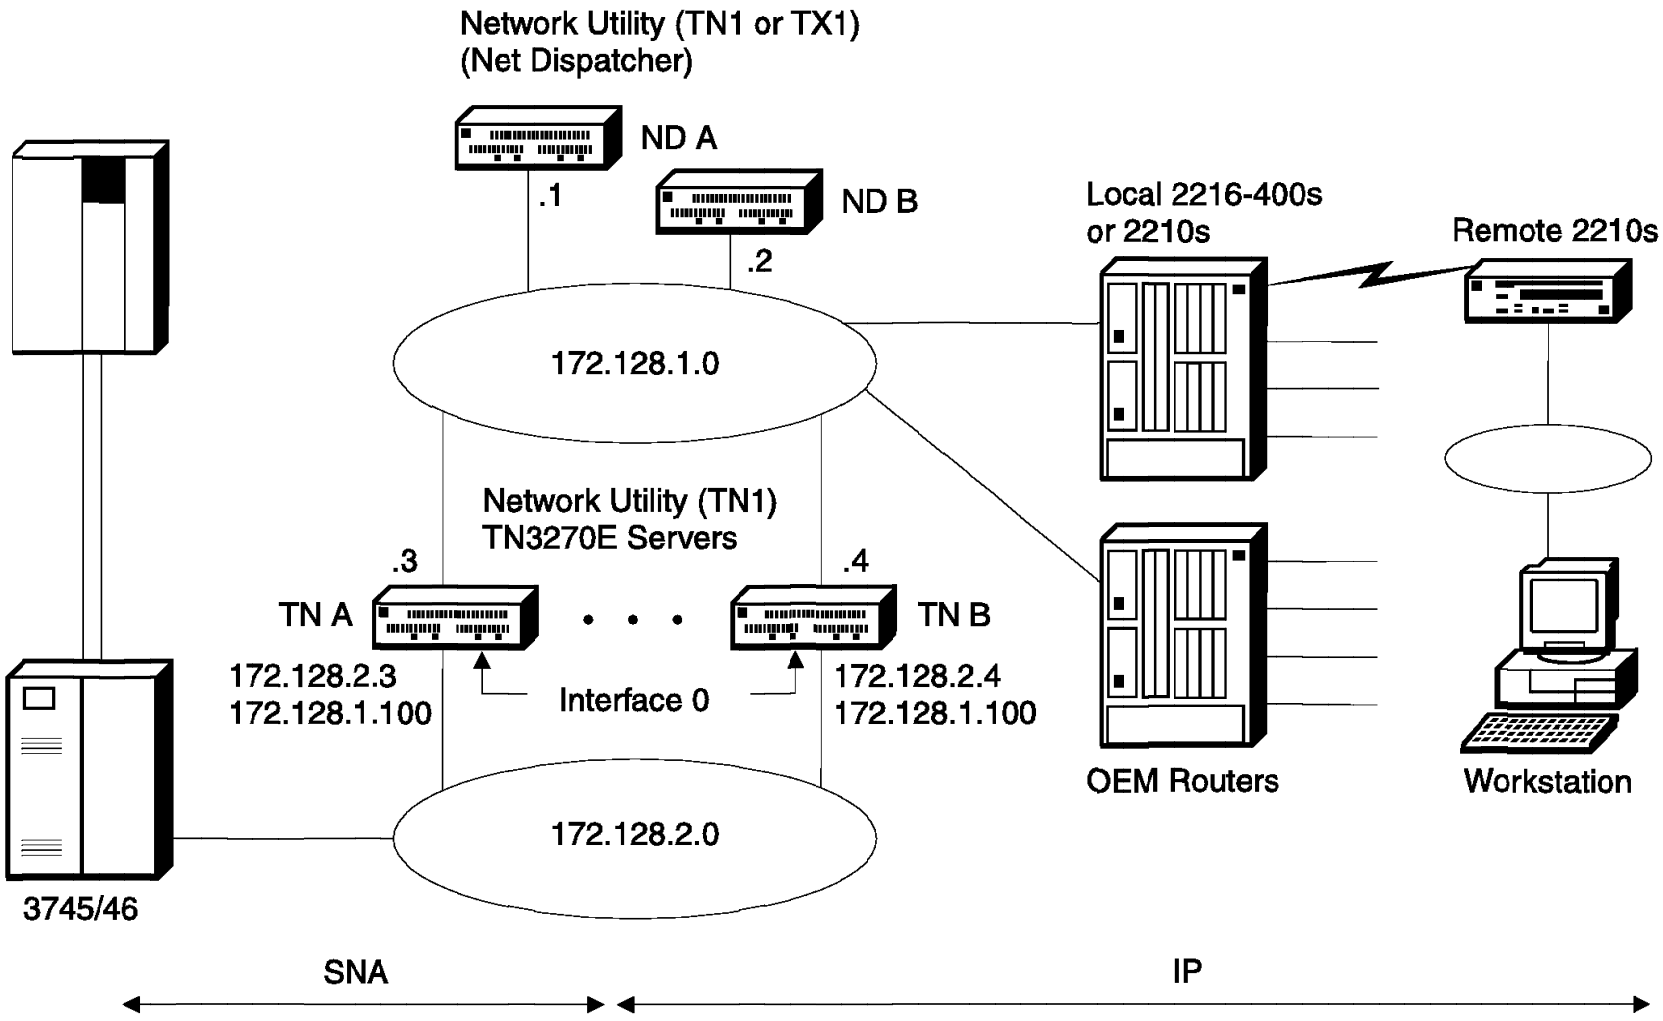

Figure 11. TN3270E Server Config -Highly Scalable, Fault-Tolerant TN3270

Table 3 (Page 1 of 4). TN3270E Server Config -Highly Scalable, Fault-Tolerant TN3270. See page 77 for a description and 94 for a diagram of this configuration.

This table provides the configuration for the TN A server. *See Table 4 on page 100 and Table 5 on page 104 for the configuration of the Network Dispatchers for this example.*

Configuration Program Navigation	Configuration Program Values	Command-Line Commands	Notes
Devices Adapters Slots	Slot1: 2-Port TR	See "add dev" on next row	1
Devices Adapters Ports	Slot 1/Port 1: Interface 0: TR Slot 1/Port 2: Interface 1: TR	Config>**add dev tok** (once for each interface)	2
Devices Interfaces	Interface 0 Mac Address 400022AA0053 Interface 1 Mac Address 400022AA0003	Config>**net 0** TKR config>**set phy 40:00:22:AA:00:53** TKR config>**exit** Config>**net 1** TKR config>**set phy 40:00:22:AA:00:03**	
System General	System name: NU_A Location: XYZ Contact: Administrator	Config>**set host** Config>**set location** Config>**set contact**	
System SNMP Config General	SNMP (checked)	Config>**p snmp** SNMP Config>**enable snmp**	
System SNMP Config Communities General	Community name: admin Access type: Read-write trap Community view: All	SNMP Config>**add community** SNMP Config>**set comm access write**	3
Protocols IP General	Internal address: 172.128.252.3 Router ID: 172.128.1.3 Same Subnet (checked)	Config>**p ip** IP config>**set internal 172.128.252.3** IP config>**set router-id 172.128.1.3** IP config>**enable same-subnet**	4

Table 3 (Page 2 of 4). TN3270E Server Config -Highly Scalable, Fault-Tolerant TN3270. See page 77 for a description and 94 for a diagram of this configuration.

This table provides the configuration for the TN A server. *See Table 4 on page 100 and Table 5 on page 104 for the configuration of the Network Dispatchers for this example.*

Configuration Program Navigation	Configuration Program Values	Command-Line Commands	Notes
Protocols IP Interfaces	Interface 0 (TR slot 1 port 1) IP address: 172.128.2.3 Subnet mask: 255.255.255.0 IP address: 172.128.1.100 Subnet mask: 255.255.255.0 Interface 1 (TR slot 1 port 2) IP address: 172.128.1.3 Subnet mask: 255.255.255.0	IP config>**add address 0 172.128.2.3 255.255.255.0** IP config>**add address 0 172.128.1.100 255.255.255.0** IP config>**add address 1 172.128.1.3 255.255.255.0**	5,6
Protocols IP OSPF General	OSPF (checked)	Config>**p ospf** OSPF Config>**enable ospf**	
Protocols IP OSPF Area Configuration General	Area number: 0.0.0.0 Stub area (not checked)	OSPF Config>**set area**	
Protocols IP OSPF Interfaces	Interface 1 OSPF (checked)	OSPF Config>**set interface** Interface IP address: **172.128.1.3** Attaches to area: **0.0.0.0** (Accept other defaults)	7
Protocols APPN General	APPN network node (checked to enable) Network ID: NUBNODE Control point name: CPNU	Config>**p appn** APPN config> **set node** Enable APPN Network ID: **NUBNODE** Control point name: **CPNU** (Accept other defaults)	8

Table 3 (Page 3 of 4). TN3270E Server Config -Highly Scalable, Fault-Tolerant TN3270. See page 77 for a description and 94 for a diagram of this configuration.

This table provides the configuration for the TN A server. *See Table 4 on page 100 and Table 5 on page 104 for the configuration of the Network Dispatchers for this example.*

Configuration Program Navigation	Configuration Program Values	Command-Line Commands	Notes
Protocols APPN Interfaces	(highlight Interface 0 Token Ring) (click on the **configure** tab) Define APPN port (checked to enable) Port name: TR3270 High performance routing (HPR) supported (unchecked to disable) Support multiple PUs (checked to enable)	APPN config>**add port** APPN Port Link Type**: TOKEN RING** Port name**: TR3270** Enable APPN Support multiple PUs High performance routing**: No** (Accept other defaults)	9
Protocols APPN Interfaces	(highlight Interface 0 Token Ring) (click on the **Link stations** tab) STAT001 (new definition) General-1 Tab: Link station name: STAT001 Solicit SSCP session (checked) Link support APPN functions (unchecked) General-2 Tab: MAC address of adjacent node: 400000003172 Node ID: 12244 Local SAP address: 04 (click on **Add** to create the Link station) STAT002 (new definition) General-1 Tab: Link station name: STAT002 Solicit SSCP session (checked) Link support APPN functions (unchecked) General-2 Tab: MAC address of adjacent node: 400000003172 Node ID: 12245 Local SAP address: 08 (click on **Add** to create the Link station)	APPN config>**add link** Port name for the link station**: TR3270** Station name**: STAT001** MAC address of adjacent node**: 400000003172** Solicit SSCP Session**: Yes** Local Node ID**: 12244** Local SAP address**: 4** Does link support APPN function?**: No** (Accept other defaults) APPN config>**add link** Port name for the link station**: TR3270** Station name**: STAT002** MAC address of adjacent node**:400000003172** Solicit SSCP Session**: Yes** Local Node ID**: 12245** Local SAP address**: 8** Does link support APPN function?**: No** (Accept other defaults)	10

Table 3 (Page 4 of 4). TN3270E Server Config -Highly Scalable, Fault-Tolerant TN3270. See page 77 for a description and 94 for a diagram of this configuration.

This table provides the configuration for the TN A server. *See Table 4 on page 100 and Table 5 on page 104 for the configuration of the Network Dispatchers for this example.*

Configuration Program Navigation	Configuration Program Values	Command-Line Commands	Notes
Protocols APPN TN3270E Server General	TN3270E (checked to enable) IP address : 172.128.1.100 Automatic logoff (checked to enable)	APPN config>**tn** TN3270E config>**set** Enable TN3270E Server TN3270E Server IP Address**: 172.128.1.100** Automatic logoff**: Yes** (Accept other defaults)	11
Protocols APPN TN3270E Server LUs	Local PU Name: STAT001 (click on **Implicit Pool**) LU name mask: @LU1A Number of implicit workstation definitions: 10 Local PU Name: STAT002 (click on **Implicit Pool**) LU name mask: @LU2A Number of implicit workstation definitions: 10	TN3270E config>**add imp** Station Name**: STAT001** LU name mask**: @LU1A** Number of Implicit LUs in Pool**: 10** TN3270E config>**add imp** Station Name**: STAT002** LU name mask**: @LU2A** Number of Implicit LUs in Pool**: 10**	12

Notes:

1. add dev defines a single port, not an adapter.
2. The configuration program assigns an interface number to all ports of an adapter automatically and you delete the ones you do not want to use. From the command line, you type the add dev command for each port you want to use, and the interface number (also known as ″net number″) is the output of the command.
3. You need a write-capable SNMP community only if you want to download configuration files from the configuration program directly into the router. SNMP is not required to TFTP a configuration file to the router.
4. You must enable the ″same-subnet function″ because you are using two interfaces with an IP address within the same subnet. (172.128.1.3 is assigned to TR 1 and 172.128.1.100 (cluster address) is assigned as a second address to TR 0.)
5. Note that Interface 0 has been assigned two IP addresses, one of which is the cluster address used by the Network Dispatcher. The TN3270E Server will be configured for the same address in a subsequent step. All TN3270 traffic will be sent to this address through the Network Dispatcher. In order for this traffic to reach the Network Utility′s internal IP queue, this address needs to be assigned to either an interface address or the internal address. In this example, it has been assigned to an interface as the second address of that interface.
6. Note that Interface 0 is on the LAN segment which is connected to the SNA gateway. This segment carries the LLC traffic from the TN3270 server to the gateway. Depending on the rest of the configuration of your Network Utility, this segment may not have any IP traffic on it. However, since all the TN3270E servers will have the same IP address assigned to the interface on this segment, it has been assigned a subnet address (172.128.2) and all the TN3270E servers will have an address on this subnet also (in this case 172.128.2.3) in order to an IP addressing conflict.
7. It is very important ***not*** to enable OSPF on the Network Dispatcher cluster address. If you do, the cluster address will be broadcast to the network as being on the TN3270E server (in addition to the Network Dispatcher machine).
8. If you have a pure SNA subarea network with no APPN, then the Network ID can be any value. If you have APPN in your network, then the Network ID should conform to your APPN network naming conventions.
9. APPN must be enabled even though the example uses SNA subarea for the TN3270E server connection to the host. This is because the TN3270E server code uses the APPN SNA stack both for APPN and subarea communications to the host.
10. When you create the link stations, you are also implicitly creating PUs. These PUs are assigned a ″Local Node ID″ here. This must match the ″IDNUM″ in VTAM′s SW Major Node definition. The ID Block is always 077 for a Network Utility. If you need to define multiple link stations (PUs), then each link station has to have a different Local SAP address.
11. Beginning with MAS V3.2, TN3270E Server has its own command-line submenu.
12. For implicit LUs, you just have to define the pools. The `@LU1A` is a template that will be used to create the actual LU names in the pool. In this example, with 10 LUs in the pool, the LU names generated are `@LU1A2, @LU1A3,...@LU1A11` which correspond to LOCADDRs 2-11 for the PU defined in VTAM. Similarly, `@LU2A` will generate `@LU2A2, @LU2A3, ... @LU2A11`.

Table 4 (Page 1 of 3). Network Dispatcher Config -Highly Scalable, Fault-Tolerant TN3270. See page 77 for a description and 94 for a diagram of this configuration.

This table provides the configuration for the Primary Network Dispatcher, ND A. *See Table 5 on page 104 for the configuration of the backup Network Dispatcher. See Table 2 on page 89 for the configuration of the TN3270E servers for this example.*

Configuration Program Navigation	Configuration Program Values	Command-Line Commands	Notes
Devices Adapters Slots	Slot 1: 2-Port TR	See ″add device″ on next row	1
Devices Adapters Ports	Slot 1/Port 1: Interface 0: TR	Config>**add dev tok**	2
Devices Interfaces	Interface 0 MAC address: 400022AA0001	Config>**net 0** TKR config>**set phy 40:00:22:AA:00:01**	
System General	System name: NU_A Location: XYZ Contact: Administrator	Config>**set host** Config>**set location** Config>**set contact**	
System SNMP Config General	SNMP (checked)	Config>**p snmp** SNMP Config>**enable snmp**	
System SNMP Config Communities General	Community name: admin Access type: Read-write trap Community view: All	SNMP Config>**add community** SNMP Config>**set comm access write**	3
Protocols IP General	Internal address: 172.128.252.1 Router ID: 172.128.1.1	Config>**p ip** IP config>**set internal 172.128.252.1** IP config>**set router-id 172.128.1.1**	4
Protocols IP Interfaces	Interface 0 (TR slot 1 port 1) IP address: 172.128.1.1 Subnet mask: 255.255.255.0	IP config>**add address**	
Protocols IP OSPF General	OSPF (checked)	Config>**p ospf** OSPF Config>**enable ospf**	

Table 4 (Page 2 of 3). Network Dispatcher Config -Highly Scalable, Fault-Tolerant TN3270. See page 77 for a description and 94 for a diagram of this configuration.

This table provides the configuration for the Primary Network Dispatcher, ND A. *See Table 5 on page 104 for the configuration of the backup Network Dispatcher. See Table 2 on page 89 for the configuration of the TN3270E servers for this example.*

Configuration Program Navigation	Configuration Program Values	Command-Line Commands	Notes
Protocols IP OSPF Area Configuration General	Area number: 0.0.0.0 Stub area (not checked)	OSPF Config>**set area**	
Protocols IP OSPF Interfaces	Interface 0 OSPF (checked)	OSPF Config>**set interface** Interface IP address **172.128.1.1** Attaches to area **0.0.0.0** (Accept other defaults)	
Features Network Dispatcher Router Executor	Executor (checked)	Config>**feat ndr** NDR Config>**enable executor**	
Features Network Dispatcher Router Clusters Detail	Cluster address: 172.128.1.100	NDR Config>**add cluster** Cluster Address**: 172.128.1.100** (Accept other defaults)	
Features Network Dispatcher Router Clusters Ports	Port Number 23	NDR Config>**add port** Cluster Address **172.128.1.100** Port number **23** (Accept other defaults)	
Features Network Dispatcher Router Clusters Servers	Server address: 172.128.1.3 172.128.1.4	NDR Config>**add server** Cluster Address**: 172.128.1.100** Port number**: 23** Server Address**: 172.128.1.3** (Accept other defaults) (Repeat for 172.128.1.4)	

Table 4 (Page 3 of 3). Network Dispatcher Config -Highly Scalable, Fault-Tolerant TN3270. See page 77 for a description and 94 for a diagram of this configuration.

This table provides the configuration for the Primary Network Dispatcher, ND A. *See Table 5 on page 104 for the configuration of the backup Network Dispatcher. See Table 2 on page 89 for the configuration of the TN3270E servers for this example.*

Configuration Program Navigation	Configuration Program Values	Command-Line Commands	Notes
Features Network Dispatcher Router Manager	Manager (checked) Proportion Active: 10 New: 10 Advisor: 80 System: 0	NDR Config>**enable manager** NDR Config>**set manager propor** Active**: 10** New**: 10** Advisor**: 80** System**: 0** (Accept other defaults)	5
Features Network Dispatcher Router Advisors	Advisor (checked) Advisor name: TN3270 Advisor port: 23 Timeout: 10	NDR Config>**add advisor** Advisor name**: 3** (for TN3270) Timeout**: 10** (Accept other defaults) NDR Config>**enable advisor** Advisor name**: 3** (for TN3270) Port number**: 23**	6
Features Network Dispatcher Router Backup	Backup (checked to enable) Backup role: PRIMARY Switch back Strategy: MANUAL	NDR Config>**add backup** Role**: 0**=PRIMARY Switch back strategy**: 1**=MANUAL	7
Features Network Dispatcher Router Reaches	Reach address: (Enter each address and click on **Add**) 172.128.1.3 172.128.1.4	NDR Config>**add reach** Address to reach**: 172.128.1.3** (Repeat for 172.128.1.4)	8
Features Network Dispatcher Router Heart Beats	Source address: 172.128.1.1 Target address: 172.128.1.2 (Enter addresses and click on **Add**)	NDR Config>**add heartbeat** Source Heartbeat address**: 172.128.1.1** Target Heartbeat Address**: 172.128.1.2**	8

Notes:

1. add dev defines a single port, not an adapter.
2. The configuration program assigns an interface number to all ports of an adapter automatically and you delete the ones you do not want to use. From the command line, you type the add dev command for each port you want to use, and the interface number (also known as ″net number″) is the output of the command.
3. You need a write-capable SNMP community only if you want to download configuration files from the configuration program directly into the router. SNMP is not required to TFTP a configuration file to the router.
4. The internal address must be set in order for the advisor and the manager functions to communicate with the executor component of Network Dispatcher.
5. The values for Active, New, Advisor, and System must add up to 100. The Advisor proportion defaults to 0. You need to change this so that the Advisor input can be used to load balance the TN3270 traffic. In this case, it has been set to 80 to give it the a much greater weight than those for active and new connections.
6. The communication port number (defaults to 10008) must match the server′s ″Network Dispatcher advisor port.″
7. Switchback Strategy must be the same for both primary and backup network dispatchers. IBM recommends a manual setting so that you can schedule the switchback at a time when you have the least probability of disrupting your SNA sessions.
8. The reach addresses are the addresses that the Network Dispatcher must be able to reach in order for it to determine that it is functioning correctly. The primary sends this information at regular intervals to the backup. If the backup determines that it has better reachability than the primary, then it will perform a switchover and assume the primary role. Choose at least one host on each subnet that the Network Dispatcher uses. Also, add the addresses for each server in the cluster. In this example, the Network Dispatcher uses only one interface and both servers are on the same subnet as this interface.
9. Here, you are configuring the connection that the primary Network dispatcher will use to send the heart beats to the backup Network Dispatcher. You can define several paths if you have multiple connections between the primary and backup. The heartbeats will be sent over the first path that is available. The most robust solution is to configure a second path between the primary Network Dispatcher and the backup Network Dispatcher using the second slot that is available in each Network Utility.

Table 5. Network Dispatcher Config -Highly Scalable, Fault-Tolerant TN3270. See page 77 for a description and 94 for a diagram of this configuration.

This table provides the configuration differences for the backup Network Dispatcher ND B based on Table 4 on page 100, which gives the configuration for the primary Network Dispatcher. *The definition for the backup Network Dispatcher is the same as for the primary except for the differences that are shown in this table. These differences correspond to the interface addresses and the Network Dispatcher backup functions. The parameters related to the Network Dispatcher that are not shown here must be identical to the values configured on the primary. It is also recommended that the hardware configuration be the same for both the primary and backup Network Dispatchers. See Table 3 on page 95 for the configuration of the TN3270E servers for this example.*

Configuration Program Navigation	Configuration Program Values	Command-Line Commands	Notes
Devices Interfaces	Interface 0 MAC Address 400022AA0002	Config>**net 0** TKR config>**set phy 40:00:22:AA:00:02**	
System General	System name: NU_ND2	Config>**set host**	
Protocols IP General	Internal address: 172.128.252.2 Router ID: 172.128.1.2	Config>**p ip** IP config> **set internal 172.128.252.2** **set router-id 172.128.1.2**	1
Protocols IP Interfaces	Interface 0 (TR slot 1 port 1) IP address: 172.128.1.2 Subnet mask: 255.255.255.0	Config>**p ip** IP config>**add address**	
Protocols IP OSPF Interfaces	Interface 0 OSPF (checked)	Config>**p ospf** OSPF Config>**set interface** Interface IP address**: 172.128.1.2** Attaches to area**: 0.0.0.0** (accept other defaults)	
Features Network Dispatcher Router Backup	Backup (checked to enable) Backup role: BACKUP Switch back Strategy: MANUAL	Config>**feat NDR** NDR Config>**add backup** Role**: 1**=BACKUP Switch back strategy**: 1**=MANUAL	
Features Network Dispatcher Router Heart Beats	Source address: 172.128.1.2 Target address: 172.128.1.1 (Enter addresses and click on **Add**)	NDR Config>**add heartbeat** Source Heartbeat address**: 172.128.1.2** Target Heartbeat Address**: 172.128.1.1**	2

Notes:

1. The internal address must be set in order for the advisor and the manager functions to communicate with the executor component of Network Dispatcher.
2. The backup must be configured with all the same information as the primary Network Dispatcher so that if the primary fails, the backup can assume the full role of primary including the sending of the heartbeats and the reachability information to the primary when it comes back online.

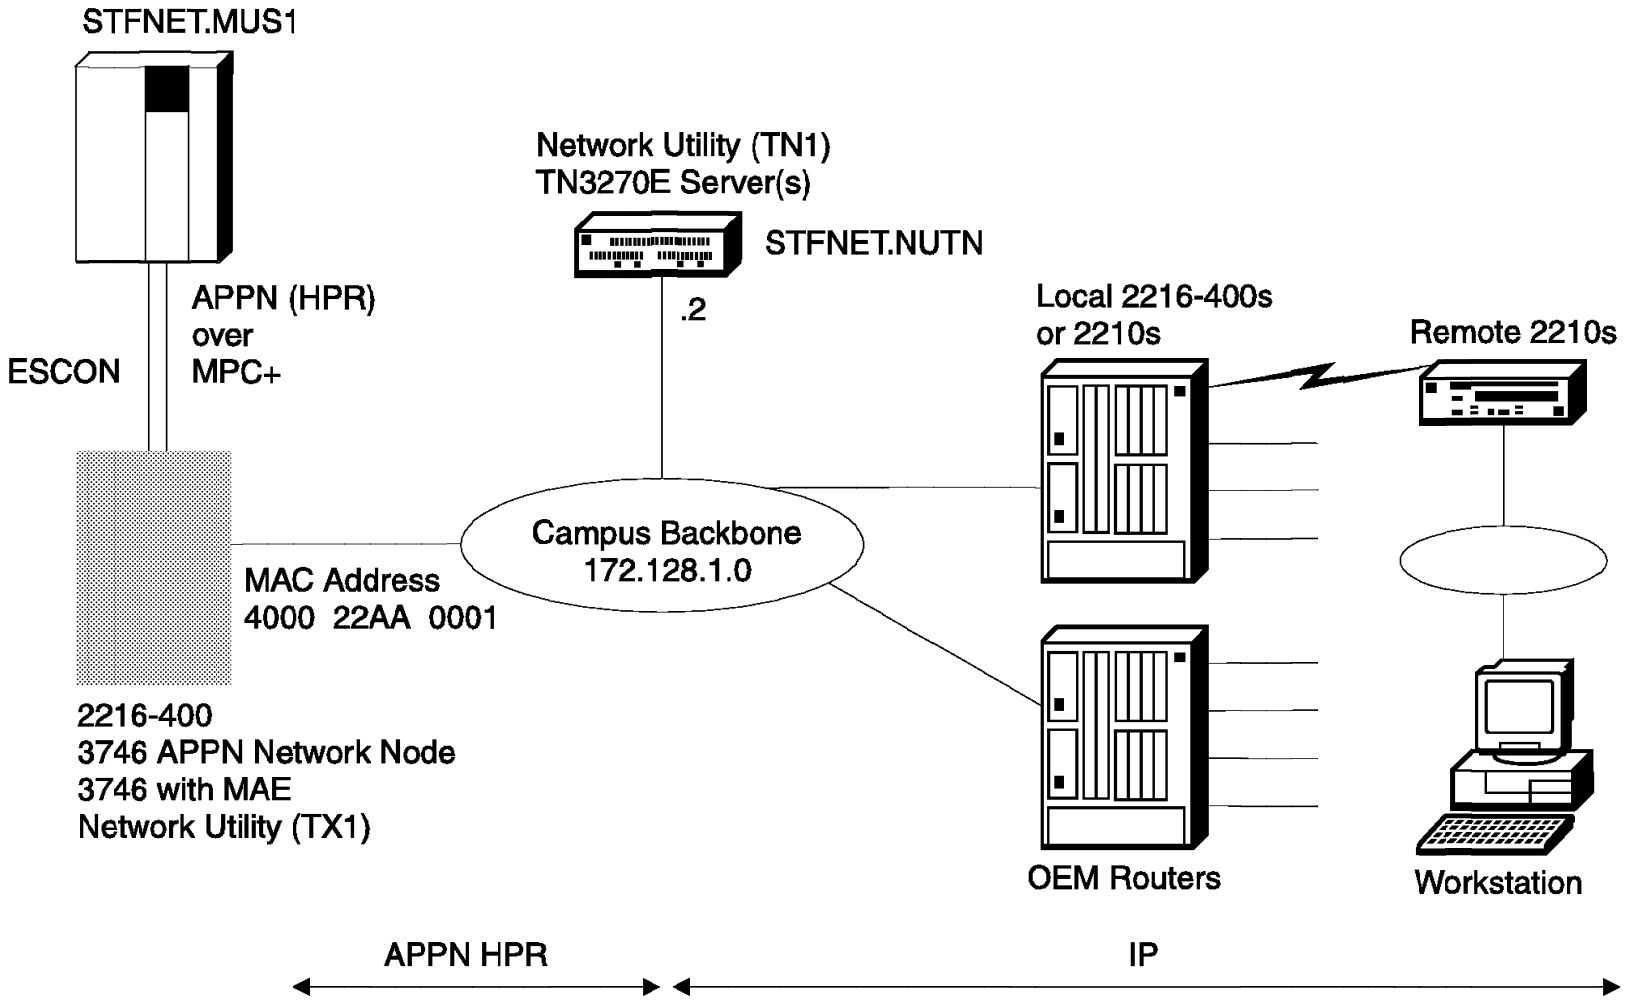

Figure 12. TN3270 Via DLUR over APPN

Table 6 (Page 1 of 3). TN3270 Via DLUR over APPN. See page 80 for a description and 106 for a diagram of this configuration.

Configuration Program Navigation	Configuration Program Values	Command-Line Commands	Notes
Devices Adapters Slots	Slot1: 2-Port TR	See "add dev" on next row	1
Devices Adapters Ports	Slot 1/Port 1: Interface 0: TR	Config>**add dev tok**	2
Devices Interfaces	Interface 0 Mac Address 400022AA0011	Config>**net 0** TKR config>**set phy 40:00:22:AA:00:11**	
System General	System name: NU_A Location: XYZ Contact: Administrator	Config>**set host** Config>**set location** Config>**set contact**	
System SNMP Config General	SNMP (checked)	Config>**p snmp** SNMP Config>**enable snmp**	
System SNMP Config Communities General	Community name: admin Access type: Read-write trap Community view: All	SNMP Config>**add community** SNMP Config>**set comm access write**	3
Protocols IP General	Internal address: 172.128.252.2 Router ID: 172.128.1.2	Config>**p ip** IP config>**set internal 172.128.252.2** IP config>**set router-id 172.128.1.2**	
Protocols IP Interfaces	Interface 0 (TR slot 1 port 1) IP address: 172.128.1.2 Subnet mask: 255.255.255.0	IP config>**add address**	
Protocols IP OSPF General	OSPF (checked)	Config>**p ospf** OSPF Config>**enable ospf**	
Protocols IP OSPF Area Configuration General	Area number: 0.0.0.0 Stub area (not checked)	OSPF Config>**set area**	

Table 6 (Page 2 of 3). TN3270 Via DLUR over APPN. See page 80 for a description and 106 for a diagram of this configuration.

Configuration Program Navigation	Configuration Program Values	Command-Line Commands	Notes
Protocols IP OSPF Interfaces	Interface 0 OSPF (checked)	OSPF Config>**set interface** Interface IP address: **172.128.1.2** Attaches to area: **0.0.0.0** (Accept other defaults)	
Protocols APPN General	APPN network node (checked to enable) Network ID: STFNET Control point name: NUTN	Config>**p appn** APPN config> **set node** Enable APPN Network ID: **STFNET** Control point name: **NUTN** (Accept other defaults)	
Protocols APPN Interfaces	(highlight Interface 0 Token Ring) (click on the configure tab) Define APPN port (checked to enable) Port name: TR001	APPN config>**add port** APPN Port Link Type: **TOKEN RING** Port name: **TR001** Enable APPN (Accept other defaults)	4
Protocols APPN Interfaces	(highlight Interface 0 Token Ring) (click on the Link stations tab) TRTG001 (new definition) General-1 Tab: Link station name: TRTG001 General-2 Tab: MAC address of adjacent node: 400022AA0001 Adjacent Node Type: APPN Network Node (click on **Add** to create the Link station)	APPN config>**add link** Port name for the link station: **TR001** Station name: **TRTG001** MAC address of adjacent node: **400022AA0001** (Accept other defaults)	5
Protocols APPN DLUR	DLUR (checked to enable) Fully qualified CP name of primary DLUS: STFNET.MVS1	APPN config>**set dlur** Enable DLUR Fully qualified CP name of primary DLUS: **STFNET.MVS1** (Accept other defaults)	6
Protocols APPN TN3270E Server General	TN3270E (checked to enable) IP address : 172.128.1.2 Automatic logoff (checked to enable)	APPN config>**tn** TN3270E config>**set** Enable TN3270E Server TN3270E Server IP Address: **172.128.1.2** Automatic logoff: **Yes** (Accept other defaults)	7

Table 6 (Page 3 of 3). TN3270 Via DLUR over APPN. See page 80 for a description and 106 for a diagram of this configuration.

Configuration Program Navigation	Configuration Program Values	Command-Line Commands	Notes
Protocols APPN TN3270E Server Local PUs	Link Station Name: PUPS08T Node ID: 12244 Link Station Name: PUPS18T Node ID: 12245	TN3270E config>**exit** APPN config>**add loc** Station Name**: PUPS08T** Local Node ID**: 12244** (Accept other defaults) APPN config>**add loc** Station Name**: PUPS18T** Local Node ID**: 12245** (Accept other defaults)	8
Protocols APPN TN3270E Server LUs	Local PU Name: PUPS08T (click on **Implicit Pool**) LU name mask: @LU1A Number of implicit workstation definitions: 5 Local PU Name: PUPS18T (click on **Implicit Pool**) LU name mask: @LU2A Number of implicit workstation definitions: 5	APPN config>**tn** TN3270E config>**add imp** Station Name**: PUPS08T** LU name mask**: @LU1A** Number of Implicit LUs in Pool**: 5** TN3270E config>**add imp** Station Name**: PUPS18T** LU name mask**: @LU2A** Number of Implicit LUs in Pool**: 5**	9

Notes:

1. add dev defines a single port, not an adapter.
2. The configuration program assigns an interface number to all ports of an adapter automatically and you delete the ones you do not want to use. From the command line, you type the add dev command for each port you want to use, and the interface number (also known as "net number") is the output of the command.
3. You need a write-capable SNMP community only if you want to download configuration files from the configuration program directly into the router. SNMP is not required to TFTP a configuration file to the router.
4. When using APPN, you can either use High Performance Routing (HPR) or Intermediate Session Routing (ISR). HPR is the default and is what is used in this scenario.
5. The MAC address specified is the MAC address of the APPN host gateway.
6. The CP name of the DLUS is the host VTAM.
7. The Local Node IDs entered for these PUs need to match the IDNUM fields in the PU definitions in the host VTAM.
8. Beginning with MAS V3.2, TN3270E Server has its own command-line submenu.
9. For implicit LUs, you just have to define the pools. The @LU1A is a template that will be used to create the actual LU names in the pool. In this example, with 5 LUs in the pool, the LU names generated are @LU1A2, @LU1A3, @LU1A4, @LU1A5 and @LU1A6 which correspond to LOCADDRs 2-6 for the PU defined in VTAM. Similarly, @LU2A will generate @LU2A2 through @LU2A6.

Chapter 7. Channel Gateway

The Network Utility provides host connectivity through an ESCON channel or parallel channel. It enables the Network Utility to function as a gateway from the host to other networks.

7.1 Configurations Supported

There are three interfaces from host software to a Network Utility.

The first interface is the 8232-compatible support, called LAN Channel Station (LCS). This interface defines a number of commands for direct LAN connection and a blocking and deblocking structure. LAN-ready frames are transmitted from the host to the virtual LAN adapters and conversely. This interface is used by TCP/IP for VM and MVS and AIX/370.

The second interface is the Link Services Architecture (LSA) support, which is accessed in the host through VTAM.

The LSA support is a control interface to allow VTAM to use the Logical Link Control (LLC) portion of the Data Link Control (DLC) layer of the SNA stack. Included is access to LLC Type 1 (connectionless) and LLC Type 2 (connection oriented) data transport. This interface is used by VTAM for both SNA subarea and APPN ISR and HPR data transport.

The third interface is the Multi-Path Channel (MPC+) support, which is accessed in the host through VTAM. The MPC+ support is a protocol layer that allows multiple read and write subchannels to be treated as a single transmission group between the host and channel attached devices. This interface is used by OS/390 for APPN HPR, TCP/IP, and HPDT UDP data transport. Note that the Channel does not support MPC+ subchannel groups which are shared over more than one physical channel interface.

The Network Utility can support 32 ESCON subchannels, in any combination of LCS subchannel pairs, LSA subchannels, and MPC+ groups. This allows a maximum of 16 LCS virtual LAN adapters, or 16 LSA virtual LAN adapters, or 16 MPC+ groups (an MPC+ group must include at least one read subchannel and one write subchannel).

LSA and LCS virtual LAN adapters emulate a token-ring, FDDI, or Ethernet interface for communications with the host. This does not restrict the format of the remote network interface. It is intended only to maintain the existing host interfaces of the 3172 Interconnect Controller, to eliminate host support changes.

Each virtual LAN adapter or MPC+ group can support only one host connection type (LCS/LSA/MPC+). LSA and LCS subchannels can support multiple virtual LAN adapters, for example, one token-ring interface and one Ethernet interface. There is no perceived value to supporting multiple virtual LAN adapters of the same type on a single subchannel or pair, but configuration will not preclude it.

7.2 Host LAN Gateway Function

The host LAN gateway function allows host applications to communicate with LAN-based workstations. The two main host applications supported by the host LAN gateway function are TCP/IP and VTAM. These applications encapsulate LAN frames into channel control words (CCWs) for transport across the channel. This is also referred to as ″blocking″, because a CCW consists of a block of LAN frames sent as a single logical unit. The CCW is then ″deblocked″ by the receiver into individual frames.

Much of the Network Utility LAN gateway function is based on the 3172 Interconnect Controller Program (ICP). Even though there are differences in the 3172 ICP gateway function and the Network Utility Channel function, the hardware and software interfaces between the host and the Network Utility Channel are the same as the interfaces between the host and the 3172 ICP (except for the IP routing support provided within the Network Utility). To preserve the software interface, it is necessary for the Network Utility to create the appearance of a LAN adapter so that the host application still believes it is communicating with a real LAN.

7.3 ESCON Channel Concepts

7.3.1 Subchannels

The ESCON channel interface is divided into 256 logical addresses (inaccurately but consistently referred to as ″subchannels″ for historical reasons). Each host application interface uses one or more subchannels to connect the host application to the Network Utility. Through configuration, each subchannel is assigned a unique relative index, which may or may not match its actual logical address. The ESCON channel may be shared by multiple applications on multiple hosts, but each host application will have dedicated use of its subchannels. (This is not strictly true for MPC+, as explained later, but the statement applies at the MPC+ level; MPC+ subchannels cannot be shared with non-MPC+ applications.) The Network Utility supports up to 32 subchannels at a time.

7.3.2 Channel Protocols

Network Utility supports three channel protocols, corresponding to the three host software interfaces discussed above. Each protocol uses its subchannels differently, and a subchannel can support only one protocol at a time. The channel protocols supported are LAN Channel Station (LCS), Link Services Architecture (LSA) and Multi-Path Channel (MPC+).

7.3.2.1 LAN Channel Station (LCS)

LCS is a channel protocol supported by TCP/IP applications in the host. Each application defines a consecutive pair of subchannels, one for TCP/IP to read from the channel, and one for TCP/IP to write to the channel. The LCS interface allows LAN MAC frames to be transported over the channel, and provides a command interface to activate, deactivate, and query the LAN interfaces. Each MAC frame has a header that identifies the virtual LAN adapter destination of the frame.

7.3.2.2 Link Services Architecture (LSA)

LSA is an interface to support SNA traffic over the channel. Each LSA path is a single bidirectional subchannel between the host application and the Network Utility. The host software (VTAM) issues a read command immediately following each write command to retrieve data from the channel. The Network Utility also issues an Attention command when it has something for the host application to read. LSA has a command interface which allows VTAM to open Service Access Points (SAPs) to communicate with downstream workstations using the IEEE 802.2 logical link control (LLC) interface. The channel blocking/deblocking mechanism for LSA subchannels is the same as for LCS subchannel pairs.

7.3.2.3 Multi-Path Channel (MPC+)

MPC+ is a data link control (DLC) interface for the channel. Each MPC+ path consists of one or more read subchannels and one or more write subchannels, bound together to form a transmission group. MPC+ transmission groups which span more than one physical ESCON channel are not supported in this release. VTAM and the Network Utility exchange XIDs to identify the number and direction of subchannels at initialization, and then each frame has a header to indicate the sending and receiving applications.

7.3.2.4 Blocks

The host channel interface packages control and data frames in blocks of up to 32 KB (36 KB for MPC+). The format of data blocks is different for MPC+ and non-MPC+ host applications. LSA and LCS blocks consist of one or more contiguous frames, each with a header that identifies the destination device by its LAN type and LAN number. MPC+ blocks contain one or more ″discontiguous″ frames, with the first 4 KB of the block containing MPC+ PDU headers and offsets of application data, which is stored in the last 32 KB of the block. MPC+ groups are identified by a ″LAN type″ and ″LAN number″ as well for implementation consistency.

A block of data is transmitted either when it is filled, or when the block delay timer (which determines how long the adapter waits for the block to fill before transmitting) expires. The process of receiving a block of data and forwarding the individual frames to the device driver is called ″deblocking.″

7.3.2.5 Virtual LAN Adapters

First, a little history: the 3172 Interconnect Control Program (on which the Network Utility is partially based) transferred frames from a host channel to one or more LANs. In its configuration, each subchannel was connected to one or more LAN device drivers. Data from the host was received by a deblocker, which would distribute the frames to one of the LAN adapters based on the LAN type and LAN number contained in the frame header. If a host application needed access to multiple LAN adapters, the configuration file would contain one entry for each LAN adapter.

In the Network Utility, instead of each subchannel being connected to one or more real LAN adapters, all of the subchannels are connected to the Base Net Handler, which is in turn connected to one or more virtual net handlers. Each virtual net handler supports one of the three channel protocols (LSA/LCS/MPC+) and sends and receives frames with one of the protocol applications (LLC/IP/APPN), which sends the data to another net handler representing a network connection. There may or may not be any real LAN adapters connected to the Network Utility.

To preserve the existing host interfaces, the Network Utility takes on the appearance of multiple LAN adapters for LSA and LCS connections. Based on configuration parameters, the Virtual net handlers register with the appropriate protocols as either token-ring, Ethernet, or FDDI adapters. The Base Net Handler allows the host to activate and deactivate this ″virtual LAN adapter″ in the same way it controls the 3172′s real LAN adapters. Each virtual LAN adapter has its own MAC address, which allows the Network Utility to appear to the host as one or more LAN adapters on an actual local area network.

A single subchannel (or pair) can be connected to one or more virtual LAN adapters. This is necessary to allow a single host application to communicate with different types of LANs (token-ring, Ethernet, FDDI) over the same subchannel. LAN-bound frames are directed to the correct destination by the LAN type and LAN number in the frame header.

However, the inverse is true only for LSA connections. A single LCS virtual LAN adapter can be connected to only one subchannel. This restriction improves data throughput by allowing host-bound frames to be directed to the correct subchannel by the virtual net handler, without forcing the net handler to examine the MAC address or IP address of each host-bound frame. Multiple VTAMs can share a single LSA net handler if each opens a SAP with a unique number. This cannot be done for the LCS net handler because all TCP/IP traffic uses the multiprotocol SAP number ′AA′x. See Figure 13.

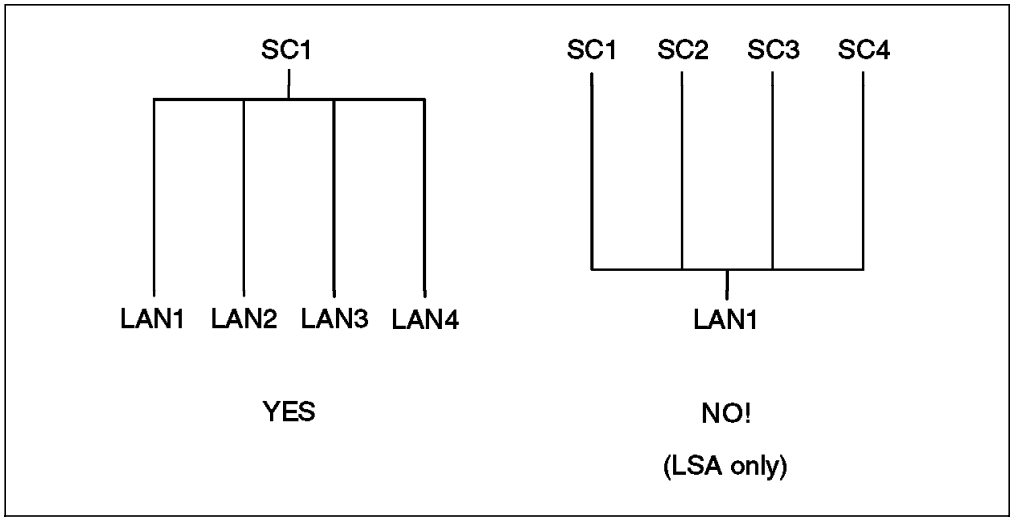

Figure 13. LAN-to-Subchannel Configuration

7.3.2.6 MPC+ Groups

MPC+ does not use the virtual LAN adapter concepts common to both LSA and LCS interfaces, because MPC+ does not support a LAN gateway appearance for the Network Utility. The equivalent interface for MPC+ is the MPC+ group. An MPC+ group is a set of ESCON subchannels configured to act as a single data pipe between the host and Network Utility. An MPC+ group consists of at least one ″read″ subchannel and at least one ″write″ subchannel. Any number of subchannels may be designated as read or write, and multiple MPC+ groups may be defined, subject to a maximum of 32 total subchannels per Network Utility.

Data may be sent over any or all of the active subchannels in an MPC+ group. The MPC+ endpoint is responsible for maintaining data order over a group. The number of subchannels is fixed when the MPC+ group is defined.

MPC+ groups are identified in the microcode using the same ″LAN type″ and ″LAN number″ designation as virtual LAN adapters. As frames are deblocked by the microcode, each frame is given a ″LAN type″ of MPC+ and a ″LAN number″ that corresponds to the MPC+ group associated with the subchannel it was received on. This allows the microcode and net handler to process MPC+ frames in a manner consistent with LSA and LCS frames.

7.3.2.7 LLC Loopback

LLC Loopback is an extension of the virtual LAN adapter concept to allow VTAM connections with the APPN and DLSw functions in the Network Utility. To establish an SNA connection, the LSA interface creates an LLC connection between itself and the remote device across the LAN using IEEE 802.2 frames. See Figure 14.

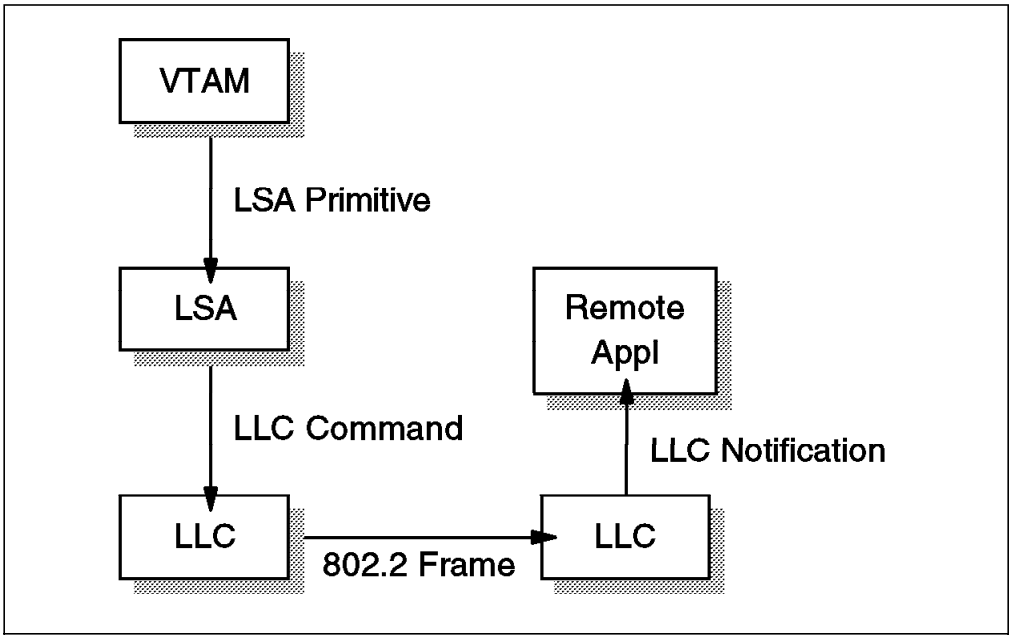

Figure 14. Normal LLC Connection

LLC Loopback allows the Network Utility to communicate directly with other LLC users (APPN and DLSw) in the Network Utility. Instead of turning LLC commands from LSA into 802.2 frames, they are converted into LLC notifications and sent to the appropriate LLC user. See Figure 15 on page 116.

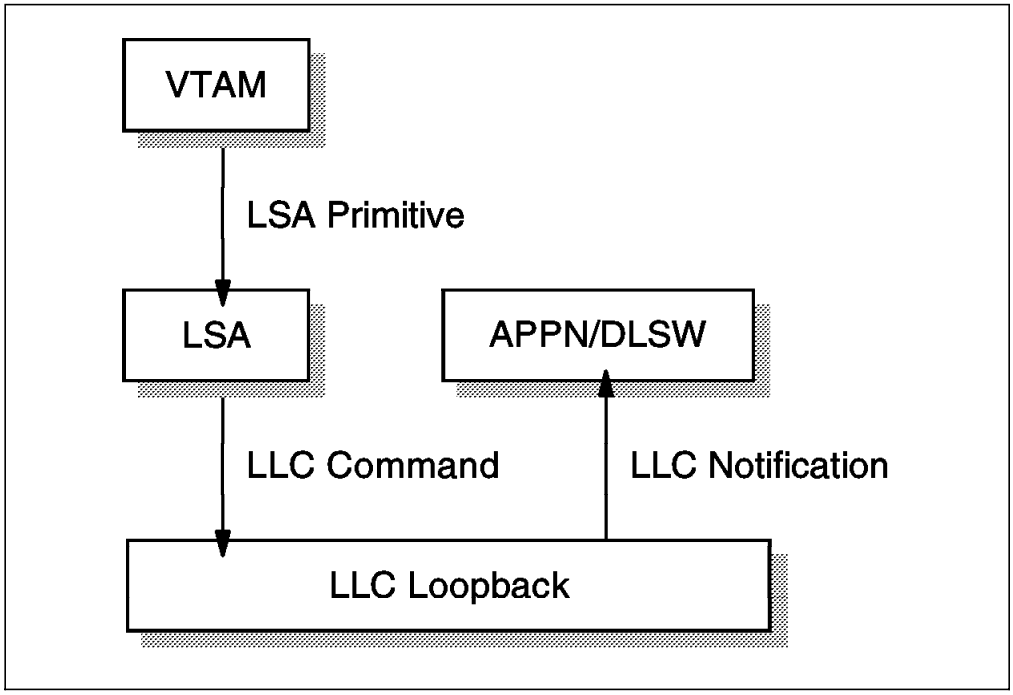

Figure 15. LLC Loopback Connection

LLC Loopback allows the APPN Network Node in Network Utility to act as the adjacent node to VTAM. It also permits VTAM to connect to remote devices and applications using Data Link Switching without requiring changes to VTAM's LSA support, because the loopback connection appears the same as a normal LLC connection to VTAM.

7.4 Example Configurations

This section describes four sample configurations that use the Network Utility as a channel gateway to a mainframe system. Three of the samples show ESCON channel configurations and one shows a parallel channel. These configurations are:

- ESCON Channel Gateway (SNA and IP)
- Parallel Channel Gateway (SNA and IP)
- ESCON Channel Gateway (APPN and IP)
- ESCON Channel Gateway - High Availability

All of these configurations can be built using either the Network Utility Model TN1 or TX1. You do not need the extra function provided by the Model TN1 unless you are planning to configure the TN3270E server function in the same machine.

7.4.1 ESCON Channel Gateway

This scenario is shown in Figure 16 on page 117. The Network Utility is configured to support both SNA and IP traffic into the host from both remote sites and LAN segments at the main site. The ESCON channel adapter is configured with an LSA direct interface to transport the SNA traffic and an LCS interface to perform IP forwarding.

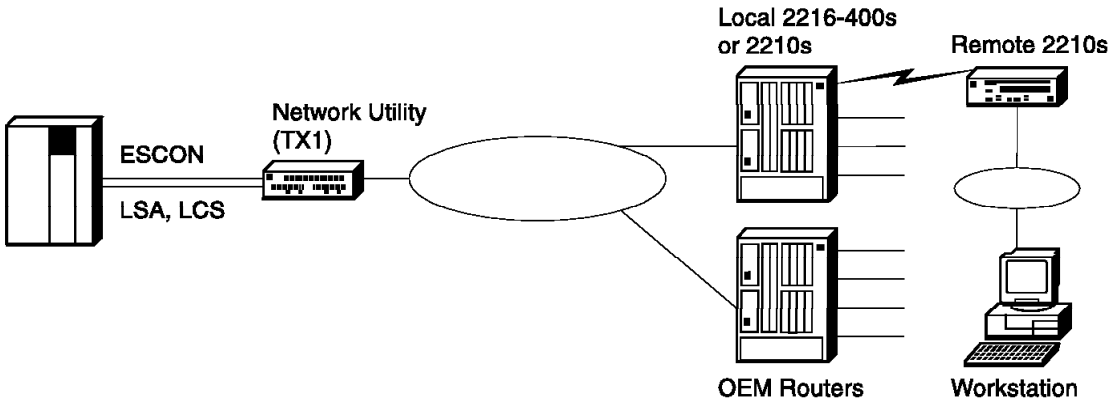

Figure 16. ESCON Channel Gateway

7.4.1.1 Keys to Configuration

The subchannel definitions for both the LCS and the LSA interfaces must match parameters used in the host to define the Network Utility to the host channel subsystem. The key subchannel parameters to configure at the Network Utility are shown in Table 7.

Table 7 (Page 1 of 2). Network Utility Subchannel Configuration Parameters

Command	Description
device	The unit address transmitted on the channel path to select the Network Utility. It is also referred to as subchannel number in S/370 I/O architecture. It is a two-digit hexadecimal value in the range 00 to FF. This value is defined in the host Input/Output Configuration Program (IOCP) by the UNITADD statement on the CNTLUNIT macro instruction for the real device. **Valid Values:** X′00′ to X′FF′ **Default:** None
cu	The Control Unit address defined in the host for the Network Utility. This value is defined in the host IOCP by the CUADD statement on the CNTLUNIT macro instruction. The Control Unit Address must be unique for each logical partition defined on the same host. **Valid Values:** X′0′ to X′F′ **Default:** X′0′
link	This parameter is significant when an IBM 9032 ESCON Director (ESCD) is used between the Network Utility and the host. When an ESCD is used, the link address is the port number of the ESCON Director (ESCD) to which the *host* is attached. If two ESCDs are in the path, it is the host-side port number of the ESCD defined with the dynamic connection. When no ESCD is in the communication path, this value must be set to X′01′. **Valid Values:** X′01′ to X′FE′ **Default:** X′01′

Table 7 (Page 2 of 2). Network Utility Subchannel Configuration Parameters	
Command	**Description**
lpar	Logical partition number. This allows multiple logical host partitions to share one ESCON fiber. This value is defined in the host IOCP by the RESOURCE macro instruction. If the host is not using ESCON Multiple Image Facility (EMIF), use the default of 0 for the LPAR number. **Valid Values:** X′0′ to X′F′ **Default:** X′0′

LPAR and CU Parameters: When defining an LSA, LCS, or MPC+ interface on the Network Utility, you need to specify correct values for the CU and the LPAR parameters.

Notes on the CU Parameter:

The value for the CU needs to be set if you have multiple LPARs or multiple MVS or OS/390 images that need to access the Network Utility. If so, then you will need to create an interface definition (LSA, LCS, or MPC+) for each LPAR and each will use a different value for the CU parameter.

Further, each definition will correspond to a pair of CNTLUNIT and IODEVICE macros in the IOCP definitions. The CUADD parameter in the CNTLUNIT macro will match the CU parameter at the Network Utility for each interface.

Notes on the LPAR Parameter:

The first issue is whether the host is partitioned into multiple logical partitions (LPARs). If it is not, then the LPAR parameter will be zero.

If it is, then you will need a RESOURCE macro in the host Input/Output Configuration Program (IOCP) definitions that specify each partition by name and assign a numeric value to each. This numeric value is used when configuring the Network Utility for the LPAR parameter.

The second issue is whether the channel path identifiers (CHPIDs) are shared between one or more LPARs.[11]

If you are not using shared channels (or you do not have EMIF), then the value for the LPAR parameter will be 0.

Figure 17 on page 119 shows an example where the host is partitioned but the channel paths are not shared between the LPARs.

[11] You need EMIF to share channels between LPARs.

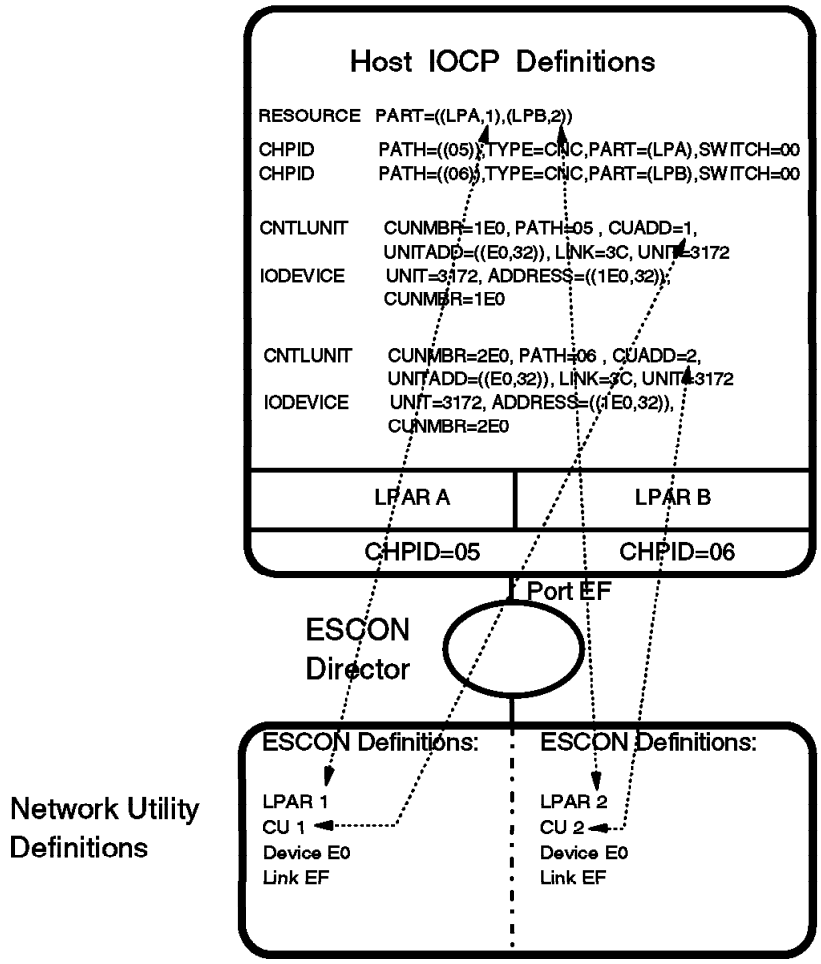

Figure 17. Host/Network Utility Parameter Relationships (Nonshared CHPIDs)

If you are using EMIF on the host, then multiple LPARs can share the same CHPID to the Network Utility. In this case, you will still need two interfaces defined on the Network Utility and each will have a different value specified for the CU parameter. The other parameters can use the same values. Figure 18 on page 120 shows an example where the host is partitioned and EMIF is used to allow both partitions to use the same CHPID.

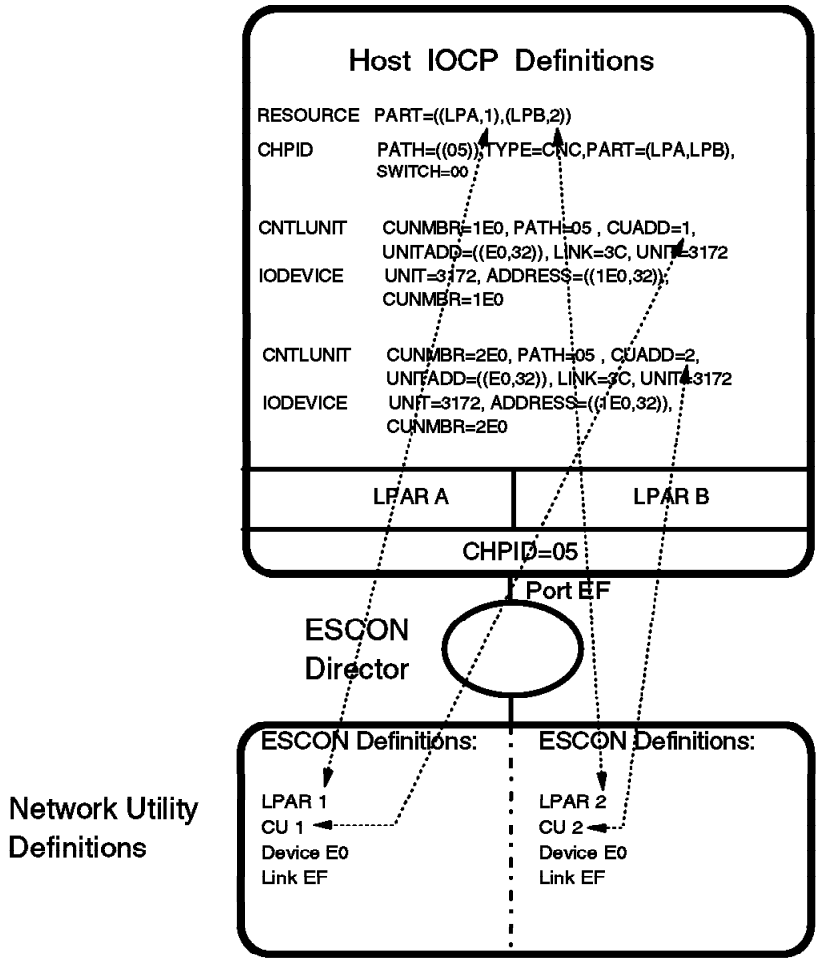

Figure 18. Host/Network Utility Parameter Relationships (Shared Channels)

The LSA Direct Interface: Figure 19 on page 121 shows how the configuration parameters for the Network Utility correlate to the host parameters for an LSA interface definition.

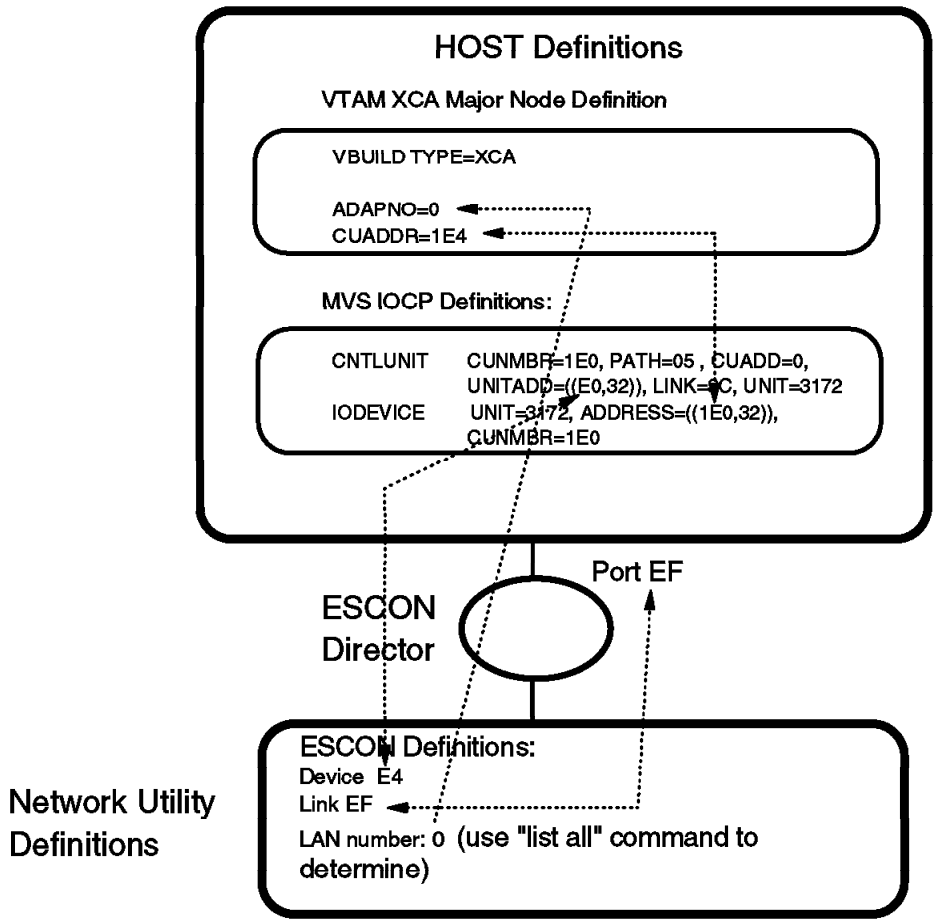

Figure 19. Host/Network Utility Parameter Relationships - LSA

Notes:

1. LSA uses a single bidirectional subchannel between the host and the Network Utility. VTAM issues a read command immediately following each write command to retrieve data from the channel.
2. The device address specified in the Network Utility LSA interface definition must be within the range specified in the UNITADD parameter in the CNTLUNIT macro from the IOCP. For example, the UNITADD parameter in Figure 19 shows that 32 (decimal) device addresses starting at E0 (hex) are being reserved for the Network Utility definition. A device address E4 has been specified for the Network Utility LSA interface. Because E4 is in the range between E0 and FF hex, this is OK as long as no other device (or interface on this Network Utility) tries to use that subchannel.
3. The value specified in the CUADDR parameter in the VTAM XCA major node definition must be within the range specified in the ADDRESS parameter in the IODEVICE macro from the IOCP. For example, the CUADDR parameter in the XCA major node definition in Figure 19 is 1E4 hex, which is in the range 1E0 to 1FF that the ADDRESS parameter in the IODEVICE statement specifies.
4. The values specified for the ADDRESS parameter in the IODEVICE macro and the UNITADD parameter in the CNTLUNIT macro are related *by convention only*. In this example, the value for the ADDRESS parameter has been determined from the value for the UNITADD parameter by prepending a *logical channel identifier* (1 in this case) to the UNITADD value. This will

often be the case. However, when defining the device address on the Network Utility LSA definition, use the UNITADD parameter and not the ADDRESS parameter to determine the valid range of values.

5. When you define an LSA direct interface on the Network Utility, you associate the interface with one of the LAN interfaces on the Network Utility. In effect, this puts the LSA direct interface on this same LAN segment. Every frame with a destination address of the MAC address of the Network Utility adapter on this LAN segment automatically gets forwarded over the channel to the host.

See Appendix A, “Sample Host Definitions” on page 167 for more explanation and samples of host definitions for this interface type.

For a complete look at the configuration parameters needed for this scenario, see Table 2 on page 89.

The LCS Interface: Defining an LCS interface creates a virtual LAN inside the Network Utility. There are two IP stations on this LAN: the Network Utility and the host. This LAN must be a unique IP subnet in the network. A MAC address is also needed for the LCS interface. After you create the LCS interface, do not forget to assign the IP address to this interface.

Important Note

The LCS support described above and documented in the example configurations is the initial 2216 LCS support released in MAS V1R1.1. This type of LCS support can be called ″LCS routing″, because it passes host IP traffic to the IP routing function within the Network Utility. If you are replacing a 3172 with a Network Utility configured with this type of LCS support, you need to configure an additional IP subnet for the virtual LAN segment inside the Network Utility.

MAS V3.2 introduces ″LCS Bridging″ (officially called ″TCP/IP Passthru″), to enable 3172 replacement with no changes to the IP topology of the network. In this mode, the Network Utility simply bridges IP traffic between an LCS bridge port and other configured bridge ports. No IP routing is performed as frames are transferred from one port to another. To enable this mode, you do not specify an IP address for the LCS interface, but you do define a MAC address and enable bridging on it. See the MAS V3.2 *Software User's Guide* for more information on configuring this function.

The functional PTF of MAS V3.2, available in December 1998, provides a third type of LCS support, which can be called ″LCS Passthrough″ or ″3172 Emulation″. This LCS mode mirrors 3172 behavior exactly by mapping an LCS virtual interface to a single LAN interface. Unlike LCS Bridging, where multiple paths exist between the various bridge-enabled interfaces, LCS Passthrough sets up independent fixed paths between specific subchannels and specific LAN adapters. Traffic on one path cannot be seen anywhere else. To enable this mode, you flag the LCS interface for this mode, you do not specify an IP address, and you reference a specific LAN adapter instead of defining an LCS MAC address.

Figure 20 on page 123 shows how the parameters correlate between the host and the Network Utility for an LCS interface definition.

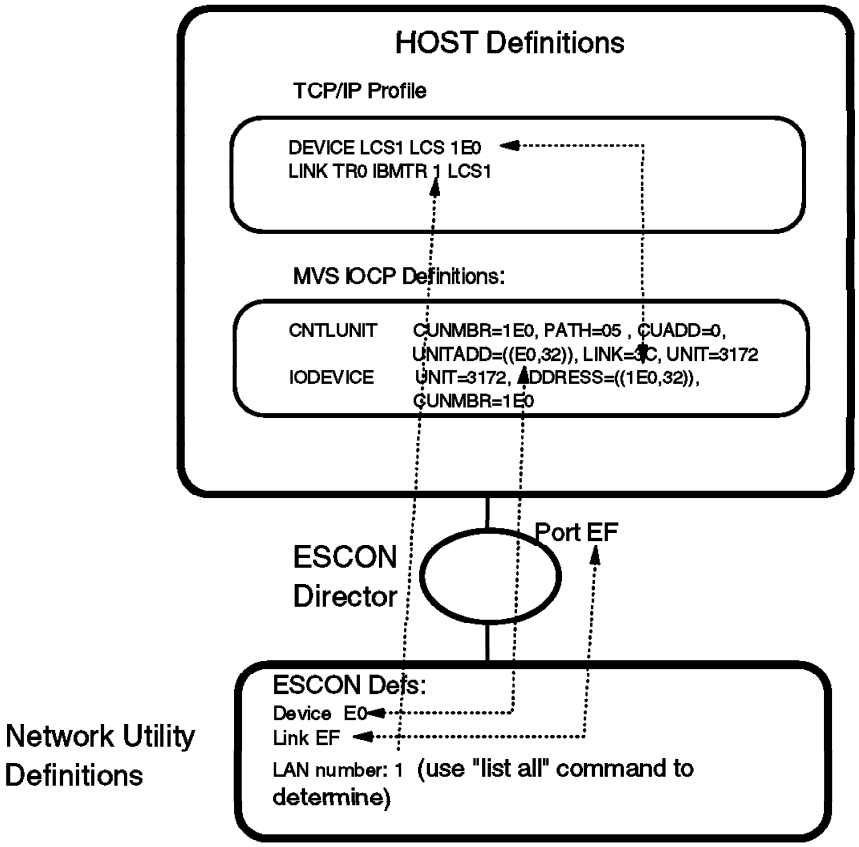

Figure 20. Host/Network Utility Parameter Relationships - LCS

Notes:

1. LCS uses a pair of subchannels, one for reading and one for writing. When configuring the subchannels used by the LCS interface, you actually need to specify only one subchannel address. LCS automatically assigns two adjacent subchannels for the LCS connection, one for the read (device address is odd) and one for the write (device address is even).
2. The device address specified in the Network Utility LCS interface definition must be within the range specified in the UNITADD parameter in the CNTLUNIT macro from the IOCP. For example, the UNITADD parameter in Figure 20 shows that 32 (decimal) device addresses starting at E0 (hex) are being reserved for the Network Utility definition. A device address of E0 has been specified for the Network Utility LCS interface. The Network Utility will automatically allocate E1 also. Since E0 and E1 are in the range E0 to FF hex, this is OK as long as no other device (or interface on this Network Utility) tries to use these same subchannels.
3. The value specified in the DEVICE statement in the host TCP/IP profile must be within the range specified in the ADDRESS parameter in the IODEVICE macro from the IOCP. For example, the DEVICE statement in the host TCP/IP profile in Figure 20 is 1E0 hex, which is in the range 1E0 to 1FF that the ADDRESS parameter in the IODEVICE statement specifies.

4. The values specified for the ADDRESS parameter in the IODEVICE macro and the UNITADD parameter in the CNTLUNIT macro are related *by convention only*. In this example, the value for the ADDRESS parameter has been determined from the value for the UNITADD parameter by prepending a *logical channel identifier* (1 in this case) to the UNITADD value. This will often be the case. However, when defining the device address on the Network Utility LCS definition, use the UNITADD parameter and not the ADDRESS parameter to determine the valid range of values.

See Appendix A, “Sample Host Definitions” on page 167 for more explanation and samples of host definitions for this interface type.

For a complete look at the configuration parameters needed for this scenario, see Table 2 on page 89.

7.4.2 Parallel Channel Gateway

This scenario is shown in Figure 21. It is identical to the ESCON channel gateway except that the connection to the host is via a S/370 Bus and Tag (Parallel Channel) Adapter instead of an ESCON channel. Like the ESCON gateway, this configuration uses an LSA direct connection for the SNA traffic and an LCS interface for the IP traffic.

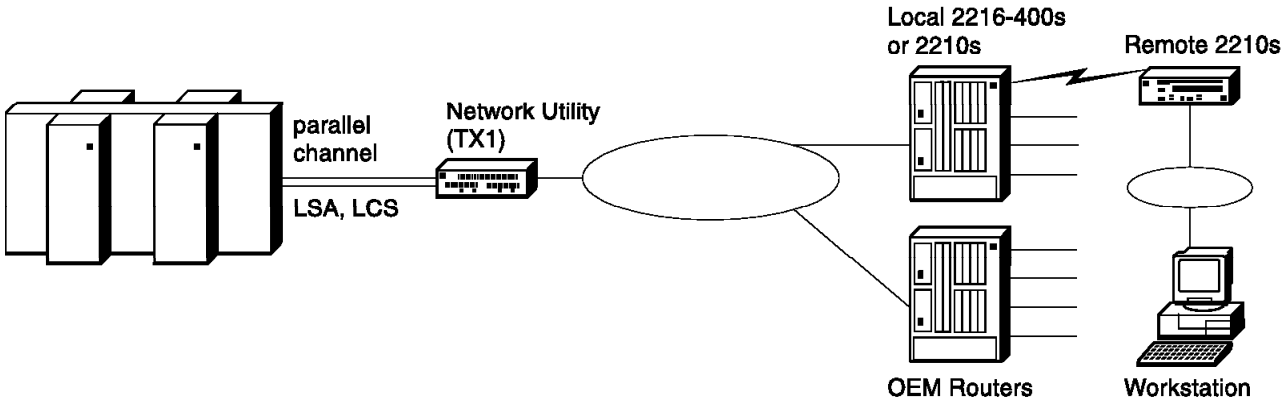

Figure 21. Parallel Channel Gateway

7.4.2.1 Keys to Configuration

The configuration for this scenario is very similar to that for the ESCON gateway (see 7.4.1, “ESCON Channel Gateway” on page 116). The configuration of the LSA and LCS interfaces require fewer parameters because no LPAR, Link Address, or Control Unit values are required for a bus and tag connection. The device address is still required to identify the Network Utility on the channel.

For a complete look at the configuration parameters needed for this scenario, see Figure 5 on page 77. Also, Appendix A, “Sample Host Definitions” on page 167 contains a sample of the host IOCP definition for a Network Utility with a Parallel Channel Adapter.

7.4.3 Channel Gateway (APPN and IP over MPC+)

This scenario is shown in Figure 22. Here, a Multi-Path Channel (MPC+) Group is used to transport both IP and APPN traffic between the Network Utility and the host. MPC+ uses a group of ESCON subchannels to maximize data transfer performance.

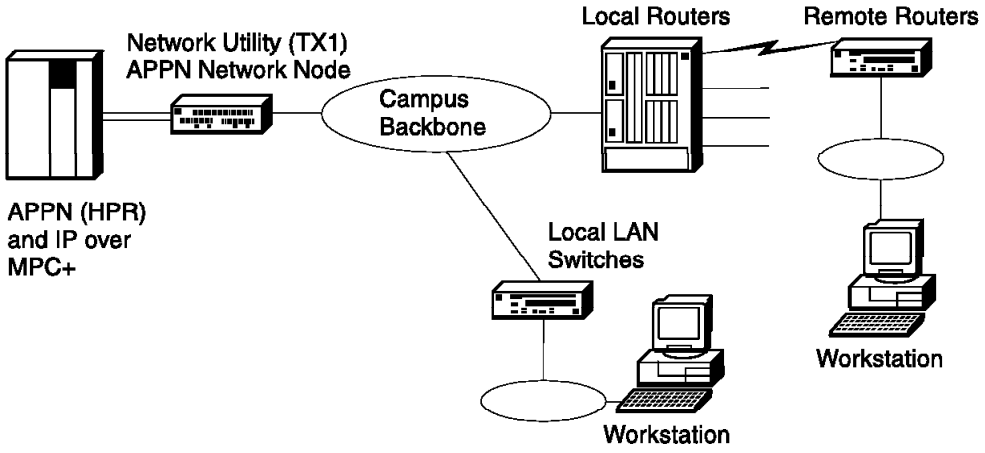

Figure 22. Channel Gateway (APPN and IP)

The APPN traffic coming through the Network Utility is comprised of several different types from the routers in the remote branches:

- TN3270 traffic from TN3270E servers in the branches that are configured with an APPN connection to the host. (See 5.5.6, “Distributed TN3270E Server” on page 82 for an example of this type of configuration.)
- DLUR traffic from the routers in the branches that are providing support for PU 2.0 (dependent) devices.
- APPN host-to-host traffic from distributed processors (such as AS/400 processors) communicating with the mainframe at the central site.

In each of the above cases, the Network Utility is providing ANR forwarding only of the APPN traffic.[12] However, in addition to providing the ANR function, the Network Utility in this scenario could also be configured for TN3270E server support and DLUR support. The DLUR support could provide PU 2.0 devices on the local campus with access to the host and the TN3270E server could provide TN3270 support for workstations and printers on the local campus or for branches that do not have a distributed TN3270E server.

7.4.3.1 Keys to Configuration

Note the following when configuring the Network Utility for this scenario:

- You can either define a separate MPC+ group for your APPN and TCP/IP traffic or you can define a single group that is shared between APPN and TCP/IP.
- An MPC+ group can have as many as 32 subchannels in it. It must have at least one read and one write subchannel defined. From the talk 6 command line (from the `ESCON Add Virtual` prompt), the `sub addr` command is used to

[12] The RTP sessions are between the APPN nodes at each end of the conversations.

add a read subchannel while the `sub addw` command is used to add a write subchannel.

- TCP/IP is configured on an MPC+ interface the same way it is for other interfaces. Specifically, configuring an IP address for the MPC+ virtual net handler enables TCP/IP over the MPC+ interface.
- APPN is configured over the MPC+ connection the same way that it is configured for other interfaces. When you use the `add port` command, specify a port type of `M` for MPC+.
- To run APPN / HPR traffic over a MPC+ Channel, two VTAM definitions need to be created:
  - A Transport Resource List (TRL) element that defines the line control, the subchannels, the number of buffers, and the channel programs to be used
  - A local SNA major node with a local PU definition
- Like the LSA and LCS definitions, the subchannel parameters must match parameters used in the host definitions when defining the Network Utility to the host channel subsystem. See Table 7 on page 117 for a description of the subchannel parameters and Figure 23 for a diagram of how these parameters correlate to the host parameters for an MPC+ definition.

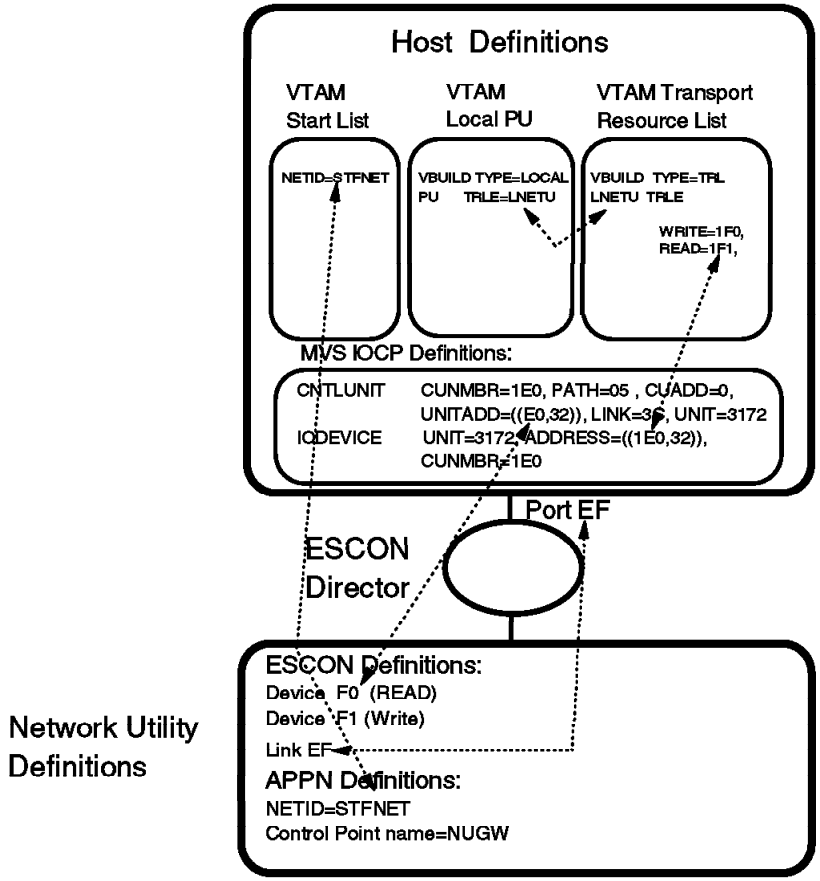

Figure 23. Host/Network Utility Parameter Relationships - MPC+

Notes:

1. The device addresses specified in the Network Utility MPC+ interface definition must be within the range specified in the UNITADD parameter in the CNTLUNIT macro from the IOCP. For example, the UNITADD parameter in Figure 23 shows that 32 (decimal) device addresses starting at E0 (hex) are being reserved for the Network Utility definition. Device addresses F0 and F1 have been specified for the Network Utility MPC+ interface. Because F0 and F1 are in the range E0 to FF hex, this is OK as long as no other device (or interface on this Network Utility) tries to use these same subchannels.
2. The values specified in the VTAM TRL major node definition must be within the range specified in the ADDRESS parameter in the IODEVICE macro from the IOCP. For example, the TRL major node definition in Figure 23 on page 126 specifies 1F0 and 1F1, which are in the range 1E0 to 1FF that the ADDRESS parameter in the IODEVICE statement specifies.
3. The values specified for the ADDRESS parameter in the IODEVICE macro and the UNITADD parameter in the CNTLUNIT macro are related *by convention only*. In this example, the value for the ADDRESS parameter has been determined from the value for the UNITADD parameter by prepending a *logical channel identifier* (1 in this case) to the UNITADD value. This will often be the case. However, when defining device addresses on the Network Utility MPC+ definition, use the UNITADD parameter and not the ADDRESS parameter to determine the valid range of values.

See Appendix A, “Sample Host Definitions” on page 167 for examples of these host definitions.

7.4.3.2 Dynamic Routing Protocols on the ESCON Interface

In a single host environment it is not necessary to run a routing protocol (RIP, for example) on the ESCON subnet. In this case, it is sufficient to add the Network Utility as the default gateway in the host TCP/IP profile.

However, if there are multiple hosts or multiple Network Utility gateways, you should consider running RIP on the ESCON interface. Running a dynamic routing protocol in this environment allows you to route around network failures if an alternate path exists.

Network Utility supports both RIP V1 and V2. RIP V2 offers variable length subnets and other advanced features that RIP V1 does not, and is the recommended choice.

7.4.3.3 Importing the ESCON Subnet into OSPF

If you are running OSPF on your network, then you should import the ESCON subnet into OSPF (unless your host TCP/IP supports OSPF). If this is not done, only workstations connected directly to an interface on the Network Utility will be able to access the TCP/IP host on the ESCON interface.

For a complete look at the configuration parameters needed for this scenario, see Figure 8 on page 81.

7.4.4 ESCON Channel Gateway - High Availability

This scenario is shown in Figure 24. It utilizes redundant Network Utilities, each with an ESCON channel connection to the host. Also, the campus backbones have been duplexed and each Network Utility attaches to a different backbone.

With this configuration, you can still access the host even if you have a failure in one of the campus backbones or a Network Utility. The traffic coming in from the 2216s will still have a valid path to the host through one campus backbone and Network Utility. This is true for both IP and SNA traffic.

The ESCON Director (ESCD) is important in this configuration, especially in Parallel Sysplex environments, because it allows you to fully mesh the connections between the gateways and the LPARs in the sysplex. This provides the highest level of fault tolerance for host access.

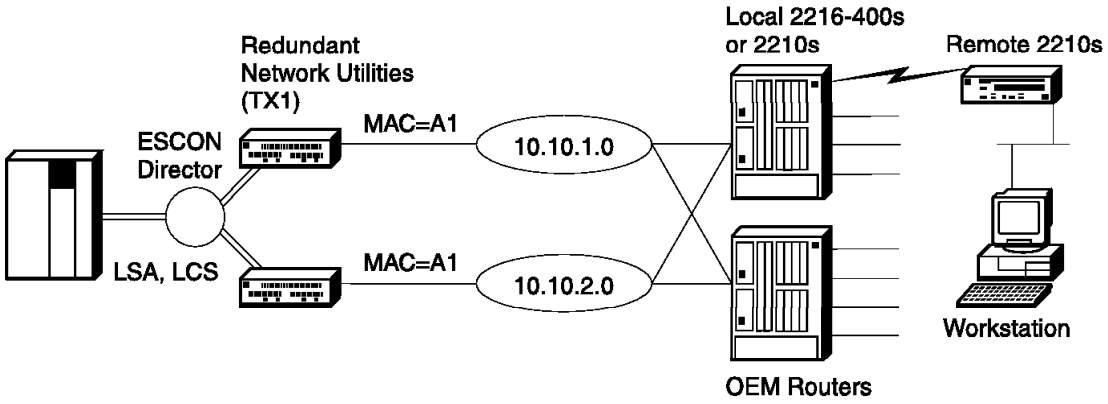

Figure 24. ESCON Channel Gateway - High Availability

7.4.4.1 Keys to Configuration

The configuration for this scenario is very much like the one in 7.4.1, “ESCON Channel Gateway” on page 116. Each Network Utility is configured as a LAN Channel Gateway with a separate LSA and LCS interface defined on each. See Table 2 on page 89 for the parameters needed for configuring a Network Utility as a LAN Channel gateway.

Because each Network Utility is on a different token-ring, the same MAC address can be used for the token-ring interface in each. The IP address used for each interface, however, must be different because each interface is on a different subnet.

Note: While this example shows the use of LSA and LCS connections on the ESCON channel, the use of MPC+ is equally effective in the high-availability environment.

7.5 Managing the Gateway Function

The configuration examples in this chapter and in 9.3.2, “DLSw LAN Channel Gateway” on page 149 show several different uses of channel DLCs:

- A direct LSA interface maps to a LAN interface with no involvement from DLSw or APPN in forwarding frames.

- An LCS Routing or MPC+ virtual interface appears to the IP routing code as another interface, and IP performs its normal routing function to forward frames to other interfaces.

  An LCS Bridging interface appears to the bridge code as another LAN bridge port, and bridging performs its normal function to forward frames to other ports.

- The loopback LSA virtual interface appears as a link to either DLSw or APPN.
- An MPC+ virtual interface can appear as a link to APPN.

To manage the complete range of Network Utility gateway function, you need to manage IP, bridging, DLSw, and APPN as appropriate. This section does not cover these upper-layer functions, but focuses instead on the ways you can monitor and manage channel physical and virtual interfaces.

7.5.1 Command-Line Monitoring

You access the talk 5 commands that show the status of channel resources hierarchically as follows:

1. From the * prompt, type `talk 5` and press Enter to reach the + prompt.
2. From the + prompt, type `int` and press Enter and note the logical interface number for the physical ESCON or PCA interface you are interested in.

   The physical interface is commonly called the *base net*, and may have a number of LSA, LCS, or MPC+ virtual interfaces defined on top of it. The base net and all virtual interfaces each have a different logical interface number.

3. From the + prompt, type `net` *base n number* and press Enter to reach the ESCON or PCA Console subprocess. The command prompt changes to `ESCON>` or `PCA>` as appropriate.

   At these prompts, you can use the `li nets` command to see the current state of every (LSA, LCS, MPC+) virtual interface using this base net. You can also type `li sub` to view the currently running subchannel configuration for this base net.

4. From the base net `ESCON>` or `PCA>` prompt, type `net` *virtual net number* and press Enter to see more detail on a particular virtual interface that uses this base net. The command prompt changes to `LSA>`, `LCS>`, or `MPC+>`, depending on the type of the virtual interface you select.

   Each of these prompts supports a `list` command, to show configuration and current status information relevant to the virtual interface type.

5. To back out from any of these nested levels, type `exit` and press `Ctrl+P` to go back to the * prompt.

For examples and a detailed explanation of the output of these commands, see the chapter ″Configuring and Monitoring the ESCON and Parallel Channel Adapters″ in the *MAS Software User's Guide*.

7.5.2 Event Logging Support

Events occurring within the channel functions are covered by the following ELS subsystems:

ESC	Low-layer ESCON events
PCA	Low-layer parallel channel events
LSA	Events related to LSA virtual interfaces
LCS	Events related to LCS virtual interfaces
MPC+	Events related to MPC+ virtual interfaces

To enable event logging, type `event` from talk 5 or talk 6 to reach the ELS Console or Config subprocess. If you want the logging output to go to talk 2, type `disp sub` *subsystem name* and press Enter to enable normal error reporting, or type `disp sub` *subsystem name* `all` to enable all messages. To get the greatest visibility to a problem, you might enable both one of the ESCON or PCA subsystems, and one of the virtual interface subsystems. If you use these commands from talk 5, you can immediately move to talk 2 and monitor events as they occur.

You can get a feel for the events reported by each of these subsystems using the command `li sub` *subsystem name* from either the talk 5 or talk 6 ELS subprocess.

7.5.3 SNA Management Support

The channel function itself does not send SNA alerts. It does not send traps that can be converted to alerts, but you can enable traps for channel ELS messages and convert those traps to alerts.

7.5.4 SNMP MIB and Trap Support

Network Utility supports an IBM enterprise-specific MIB for ESCON. This MIB provides access to the following information:

- A list of physical interfaces and the fiber signal status of each
- A list of channel links and the host connection status of each
- A list of channel stations with both configuration and normal/error traffic statistics for each.

The ESCON MIB does not define any traps. Parallel channel functions have no MIB support.

Both ESCON and parallel channel interfaces are represented in the Interfaces MIB (RFC 1573), so a management station can access their status and basic per-interface traffic statistics. Network Utility allows a management station to control interface state, and can send traps to report when the interfaces go up or down.

7.5.5 Network Management Application Support

The Network Utility Java-based application provides integrated support for both the ESCON MIB and the Interfaces MIB. You can see color-coded interface status as well as specific panels that present key information from these MIBs. You can also use integrated browser support to view the information in either of these MIBs.

You can disable or enable the emission of interface up/down traps from the Nways Manager products.

Chapter 8. Channel Gateway Example Configuration Details

This chapter contains diagrams and configuration parameter tables for several of the example channel gateway network configurations in Chapter 7, “Channel Gateway” on page 111. The parameter values shown are from real working test configurations.

For an explanation of the columns and conventions in the configuration parameter tables, see 4.2.2, “Example Configuration Table Conventions” on page 67.

The Network Utility World Wide Web pages contain binary configuration files that match these configuration parameter tables. To access these files, follow the Download link from:

```
http://www.networking.ibm.com/networkutility
```

The configurations documented in this chapter are:

Table 8. Cross-Reference of Example Configuration Information

Configuration Description	**Parameter Table**
7.4.1, “ESCON Channel Gateway” on page 116	Table 9 on page 133
7.4.2, “Parallel Channel Gateway” on page 124	Table 10 on page 137
7.4.3, “Channel Gateway (APPN and IP over MPC+)” on page 125	Table 11 on page 141

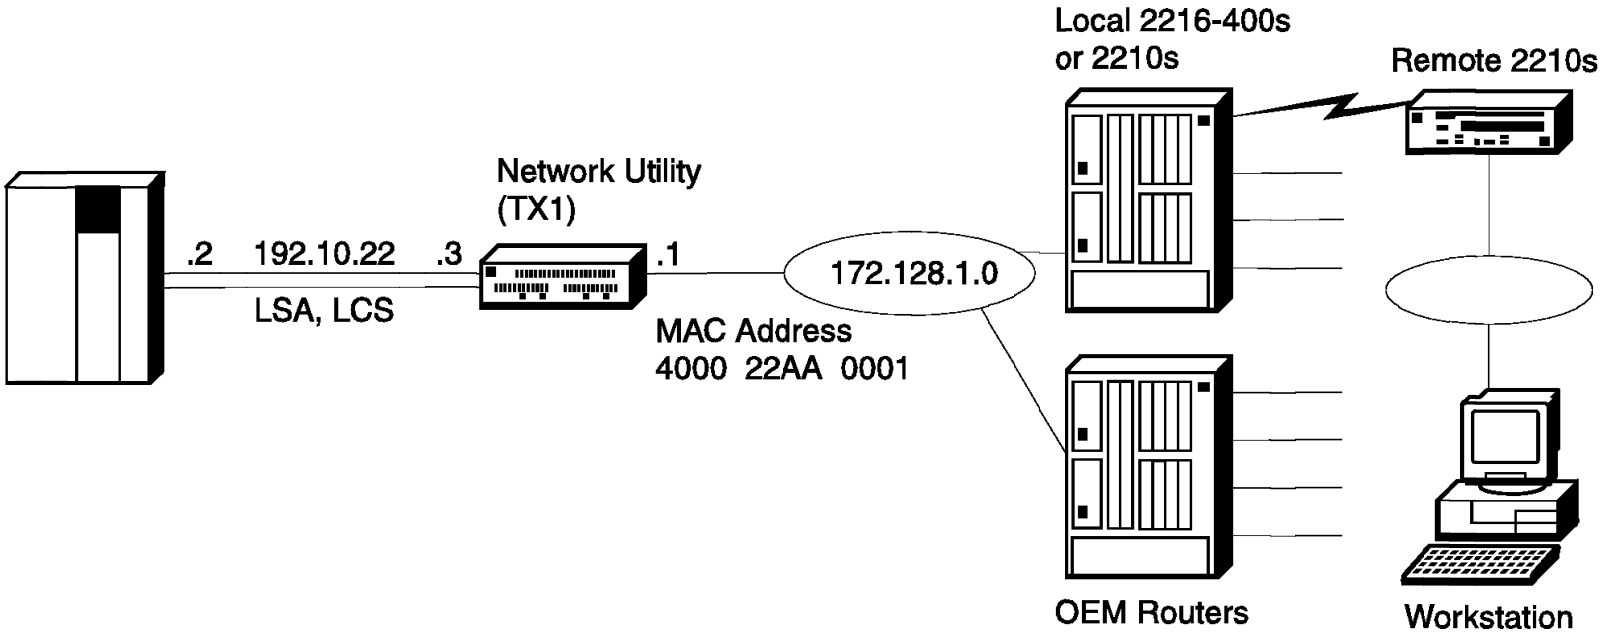

Figure 25. ESCON Channel Gateway

Table 9 (Page 1 of 2). ESCON Channel Gateway. See page 116 for a description and 132 for a diagram of this configuration.

Configuration Program Navigation	Configuration Program Values	Command-Line Commands	Notes
Devices Adapters Slots	Slot 1: 2-Port TR Slot 2: ESCON	See "add device" on next row	1
Devices Adapters Ports	Slot 1 Port 1: Interface 0: TR Slot 2 Port 1: Interface 1: ESCON	Config>**add dev tok** Config>**add dev esc**	2
Devices Interfaces	Interface 0 MAC address: 400022AA0001	Config>**net 0** TKR config>**set phy 40:00:22:AA:00:01**	
Devices Channel Adapters ESCON Interfaces ESCON Interfaces	Interface 2 (new definition) Base Network Number: 1 Protocol Type: LSA Maximum Data Frame: 2052 LAN Net Number: 0 (click on **Add** to create interface 2)	Config>**net 1** ESCON Config>**add lsa** (added as interface 2) ESCON Add Virtual>**maxdata 2052** ESCON Add Virtual>**net 0** (continue in same session with next row)	3,4,5
Devices Channel Adapters ESCON Interfaces ESCON Subchannels	Interface 2 (highlight LSA interface) Device Address: E4 Link Address: EF (click on **Add**)	ESCON Add Virtual>**subchannel add** ESCON Add LSA Subchannel>**device E4** ESCON Add LSA Subchannel>**link EF** (type **exit** twice and then **list all**)	6
Devices Channel Adapters ESCON Interfaces ESCON Interfaces	Interface 3 (new definition) Base Network Number: 1 Protocol Type: LCS LAN Type: Token Ring Maximum Data Frame: 2052 MAC Address: 400022AA0009 (click on **Add** to create interface 3)	Config>**net 1** ESCON Config>**add lcs** (added as interface 3) ESCON Add Virtual>**lantype token** ESCON Add Virtual>**Maxdata 2052** ESCON Add Virtual>**mac 40:00:22:AA:00:09** (continue in same session with next row)	
Devices Channel Adapters ESCON Interfaces ESCON Subchannels	Interface 3 (highlight LCS interface) Device Address: E0 Link Address: EF (click on **Add**)	ESCON Add Virtual>**subchannel add** ESCON Config LCS Subchannel>**device E0** ESCON Config LCS Subchannel>**link EF** (type **exit** twice and then **list all**)	7
System General	System name: NU_A Location: XYZ Contact: Administrator	Config>**set host** Config>**set location** Config>**set contact**	

Table 9 (Page 2 of 2). ESCON Channel Gateway. See page 116 for a description and 132 for a diagram of this configuration.

Configuration Program Navigation	Configuration Program Values	Command-Line Commands	Notes
System SNMP Config General	SNMP (checked)	Config>**p snmp** SNMP Config>**enable snmp**	
System SNMP Config Communities General	Community name: admin Access type: Read-write trap Community view: All	SNMP Config>**add community** SNMP Config>**set comm access write**	
Protocols IP General	Internal address: 172.128.252.1 Router ID: 172.128.1.1	Config>**p ip** IP config>**set internal 172.128.252.1** IP config>**set router-id 172.128.1.1**	
Protocols IP Interfaces	Interface 0 (TR slot 1 port 1) IP address: 172.128.1.1 Subnet mask: 255.255.255.0 Interface 3 (LCS interface) IP address: 192.10.22.3 Subnet mask: 255.255.255.0	IP config>**add address** (once per i/f)	8
Protocols IP OSPF General	OSPF (checked)	Config>**p ospf** OSPF Config>**enable ospf**	8
Protocols IP OSPF Area Configuration General	Area number: 0.0.0.0 Stub area (not checked)	OSPF Config>**set area**	
Protocols IP OSPF AS Boundary Routing	AS Boundary Routing (checked to enable) Import direct routes (checked to enable)	OSPF config>**enable as** Import direct routes (Accept other defaults)	9
Protocols IP OSPF Interfaces	Interface 0 OSPF (checked)	OSPF Config>**set interface** Interface IP address: **172.128.1.1** Attaches to area: **0.0.0.0** (Accept other defaults)	

Notes:

1. add dev defines a single port, not an adapter.
2. The configuration program assigns an interface number to all ports of an adapter automatically and you delete the ones you do not want to use. From the command line, you type the add dev command for each port you want to use, and the interface number (also known as ″net number″) is the output of the command.
3. When you select an interface of type LSA, the “LAN type” field is disabled (gets grayed out) and the “LAN net number” and “loopback” checkbox appears.
4. The “LAN number” field is disabled because a value is assigned by the router automatically. This value must be configured in the host definition for “ADAPTNO.”
5. When you “Add” the interface, a new interface will be generated and it will be assigned the next available interface number.
6. The values that you enter when configuring the subchannels must match values configured at the host. See Appendix A, “Sample Host Definitions” on page 167 for examples of how to match these values.
7. When you add subchannels for an LCS virtual interface, it is only necessary to define one subchannel although LCS requires two. LCS automatically uses the next subchannel in addition to the one defined here. LCS uses the even device address (E0 in this case) as the write subchannel and the odd address (E1) as the read subchannel.
8. You can also use RIP in place of OSPF.
9. You need to import direct routes into OSPF from the ESCON interface because OSPF is not enabled on the ESCON interface. Instead, the subnet on the ESCON interface is imported into OSPF in the Network Utility and then propagated to the network. This is necessary to prevent error messages from occurring at the host if the Network Utility sends the OSPF updates over the LCS connection. TCP/IP at the host does not (yet) support Link State Advertisements from an OSPF router.

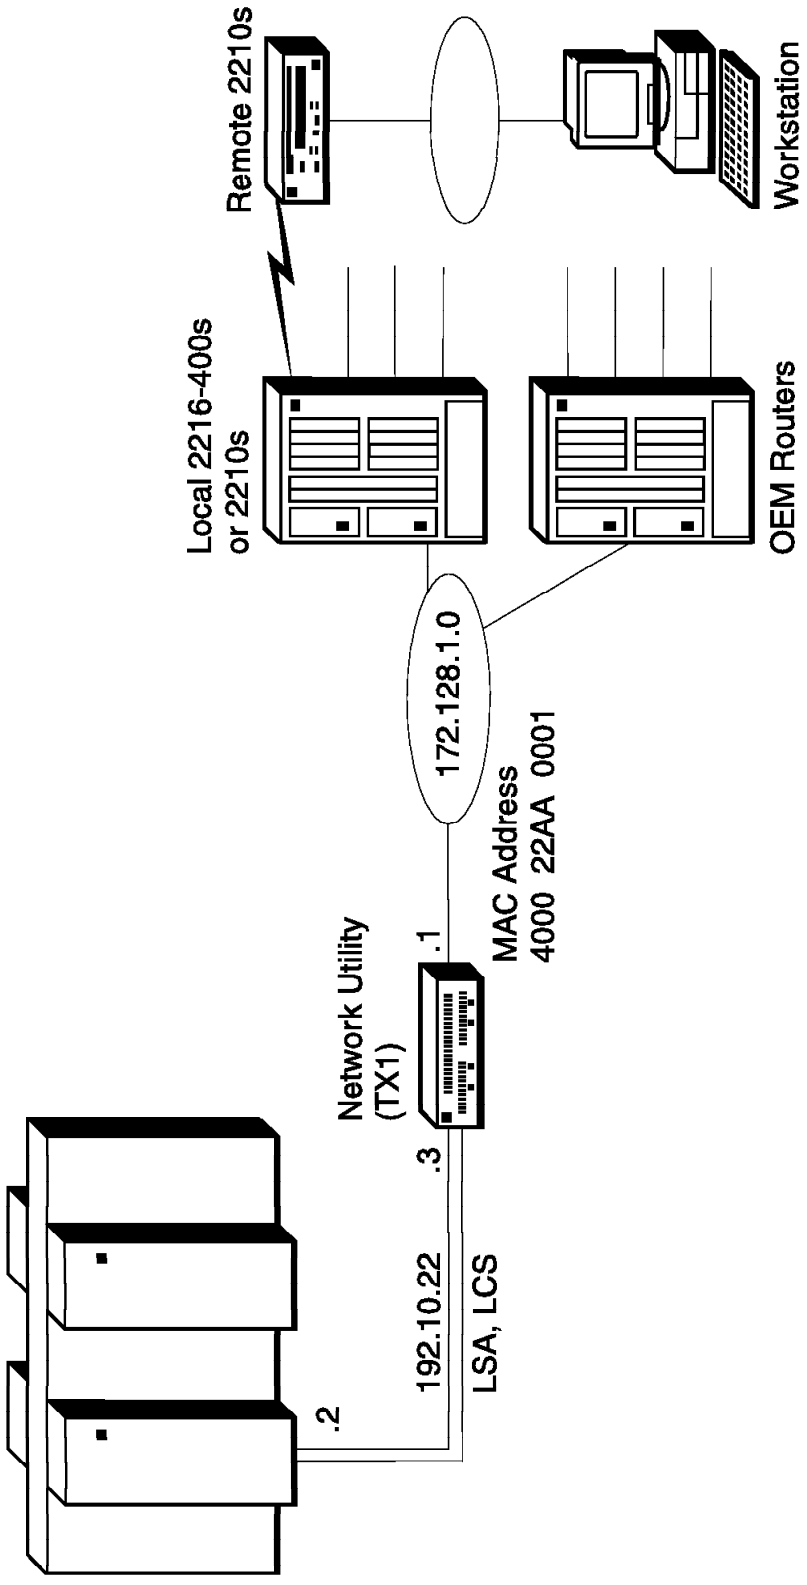

Figure 26. Parallel Channel Gateway

Table 10 (Page 1 of 2). Parallel Channel Gateway. See page 124 for a description and 136 for a diagram of this configuration.

Configuration Program Navigation	Configuration Program Values	Command-Line Commands	Notes
Devices Adapters Slots	Slot 1: 2-Port TR Slot 2: Parallel Channel Adapter (PCA)	See "add device" on next row	1
Devices Adapters Ports	Slot 1/Port 1: Interface 0: TR Slot 2/Port 1: Interface 1: PCA	Config>**add dev tok** Config>**add dev PCA**	2
Devices Interfaces	Interface 0 MAC address: 400022AA0001	Config>**net 0** TKR config>**set phy 40:00:22:AA:00:01**	
Devices Channel Adapters PCA Interfaces PCA Interfaces	Interface 2 (new definition) Base Network Number: 1 Protocol Type: LSA LAN Net Number: 0 (click on **Add** to create interface 2)	Config>**net 1** PCA Config>**add lsa** (added as interface 2) PCA Add Virtual>**net 0** (continue in same session with next row)	3,4,5
Devices Channel Adapters PCA Interfaces PCA Subchannels	Interface 2 (highlight LSA interface) Device Address: 00 Subchannel type: read/write (click on **Add**)	PCA Add Virtual>**subchannel add** PCA Add LSA Subchannel>**device 00** (Type **exit** twice and then **list all**)	6
Devices Channel Adapters PCA Interfaces PCA Interfaces	Interface 3 (new definition) Base Network Number: 1 Protocol Type: LCS MAC Address: 400022AA0009 (click on **Add** to create interface 3)	Config>**net 1** PCA Config>**add lcs** (added as interface 3): PCA Add Virtual>**mac 40:00:22:AA:00:09** (continue in same session with next row)	
Devices Channel Adapters PCA Interfaces PCA Subchannels	Interface 3 (highlight LCS interface) Device Address: 02 Subchannel type: write (click on **Add**)	PCA Add Virtual>**subchannel add** PCA Add LCS Subchannel>**device 02** (Type **exit** twice and then **list all**)	7
System General	System name: NU_A Location: XYZ Contact: Administrator	Config>**set host** Config>**set location** Config>**set contact**	
System SNMP Config General	SNMP (checked)	Config>**p snmp** SNMP Config>**enable snmp**	

Table 10 (Page 2 of 2). Parallel Channel Gateway. See page 124 for a description and 136 for a diagram of this configuration.

Configuration Program Navigation	Configuration Program Values	Command-Line Commands	Notes
System SNMP Config Communities General	Community name: admin Access type: Read-write trap Community view: All	SNMP Config>**add community** SNMP Config>**set comm access write**	
Protocols IP General	Internal address: 172.128.252.1 Router ID: 172.128.1.1	Config>**p ip** IP config>**set internal 172.128.252.1** IP config>**set router-id 172.128.1.1**	
Protocols IP Interfaces	Interface 0 (TR slot 1 port 1) IP address: 172.128.1.1 Subnet mask: 255.255.255.0 Interface 3 (LCS interface) IP address: 192.10.22.3 Subnet mask: 255.255.255.0	IP config>**add address** (once per i/f)	8
Protocols IP OSPF General	OSPF (checked)	Config>**p ospf** OSPF Config>**enable ospf**	8
Protocols IP OSPF Area Configuration General	Area number: 0.0.0.0 Stub area (not checked)	OSPF Config>**set area**	
Protocols IP OSPF AS Boundary Routing	AS Boundary Routing (checked to enable) Import direct routes (checked to enable)	OSPF Config>**enable as** Import direct routes (Accept other defaults)	9
Protocols IP OSPF Interfaces	Interface 0 OSPF (checked)	OSPF Config>**set interface** Interface IP address: **172.128.1.1** Attaches to area: **0.0.0.0** (Accept other defaults)	

Notes:

1. add dev defines a single port, not an adapter.
2. The configuration program assigns an interface number to all ports of an adapter automatically and you delete the ones you do not want to use. From the command line, you type the add dev command for each port you want to use, and the interface number (also known as ″net number″) is the output of the command.
3. When you select an interface of type LSA, the “LAN type” field is disabled (gets grayed out) and the “LAN net number” and “loopback” checkbox appears.
4. The “LAN number” field is disabled because a value is assigned by the router automatically. This value must be configured in the host definition for “ADAPTNO.”
5. When you “Add” the interface, a new interface will be generated and it will be assigned the next available interface number.
6. The values that you enter when configuring the subchannels must match values configured at the host. See Appendix A, “Sample Host Definitions” on page 167 for examples of how to match these values.
7. When you add subchannels for an LCS virtual interface, it is only necessary to define one subchannel although LCS requires two. LCS automatically uses the next subchannel in addition to the one defined here. LCS uses the even device address (02 in this case) as the write subchannel and the odd address (03) as the read subchannel.
8. You can also use RIP in place of OSPF.
9. You need to import direct routes into OSPF from the PCA interface because OSPF is not enabled on the PCA interface. Instead, the subnet on the PCA interface is imported into OSPF in the Network Utility and then propagated to the network. This is necessary to prevent error messages from occurring at the host if the Network Utility sends the OSPF updates over the LCS connection. TCP/IP at the host does not (yet) support Link State Advertisements from an OSPF router.

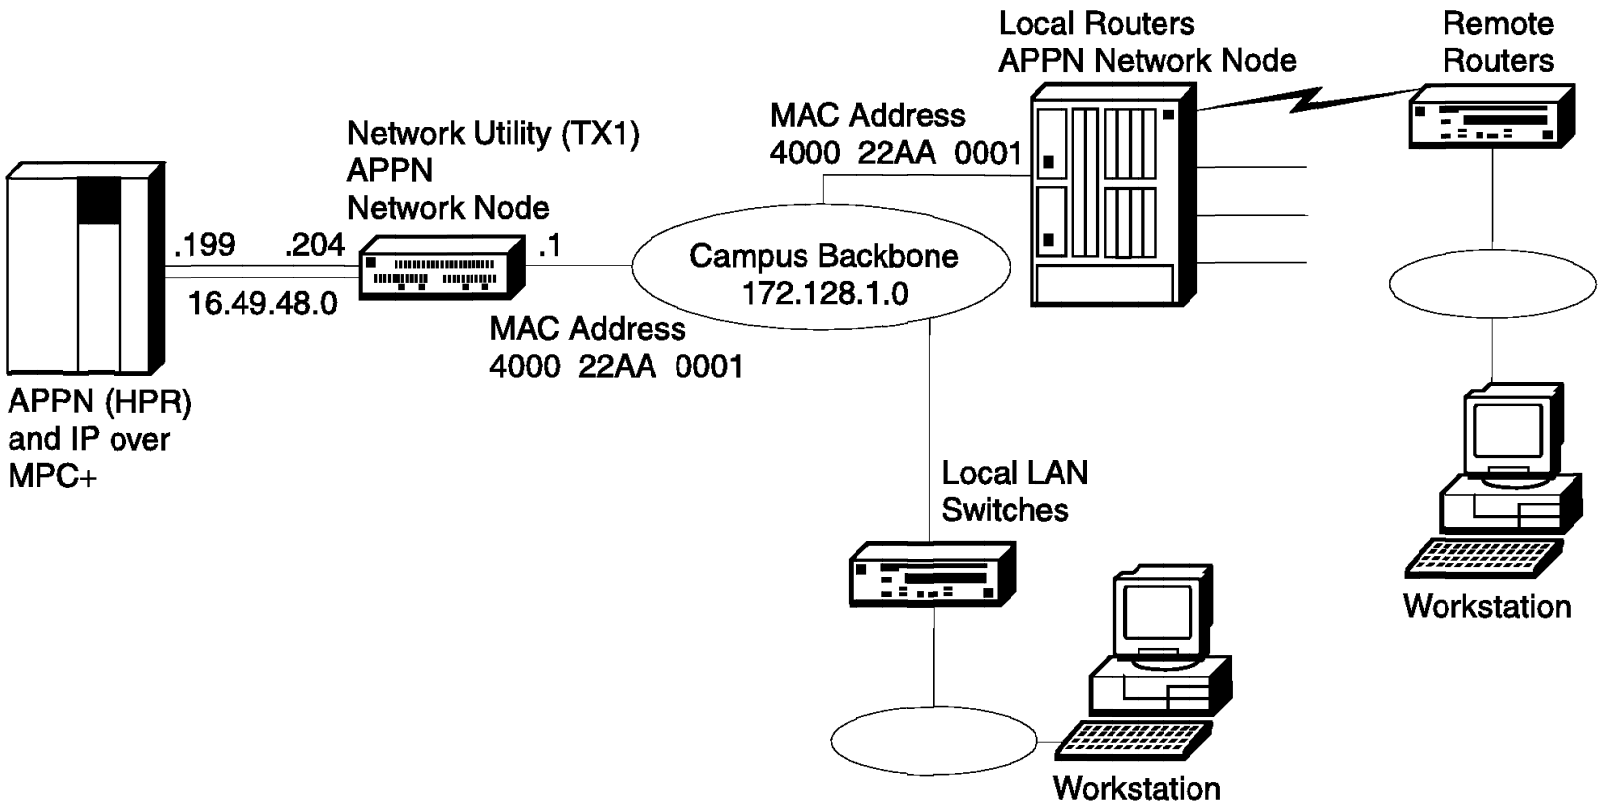

Figure 27. Channel Gateway (APPN & IP over MPC+)

Table 11 (Page 1 of 3). Channel Gateway (APPN & IP over MPC+). See page 125 for a description and 140 for a diagram of this configuration.

Configuration Program Navigation	Configuration Program Values	Command-Line Commands	Notes
Devices Adapters Slots	Slot 1: 2-Port TR Slot 2: ESCON	See "add device" on next row	1
Devices Adapters Ports	Slot 1/Port 1: Interface 0: TR Slot 2/Port 1: Interface 1: ESCON	Config>**add dev tok** Config>**add dev esc**	2
Devices Interfaces	Interface 0 MAC address: 400022AA0001	Config>**net 0** TKR config>**set phy 40:00:22:AA:00:01**	
Devices Channel Adapters ESCON Interfaces ESCON Interfaces	Interface 2 (new definition) Base Network Number: 1 Protocol Type: MPC+ (click on **Add** to create interface 2)	Config>**net 1** ESCON Config>**add mpc** (added as interface 2) ESCON Add Virtual> (continue in same session with next row)	3
Devices Channel Adapters ESCON Interfaces ESCON Subchannels	(highlight interface 2) Device Address: F0 Link Address: EF Subchannel type: Read (click on **Add** to define subchannel) Device Address: F1 Link Address: EF Subchannel type: Write (click on **Add** to define subchannel)	ESCON Add Virtual>**sub addr** ESCON Add MPC+ Read Subchannel>**dev f0** ESCON Add MPC+ Read Subchannel>**link ef** ESCON Add MPC+ Read Subchannel>**exit** ESCON Add Virtual>**sub addw** ESCON Add MPC+ Write Subchannel>**dev f1** ESCON Add MPC+ Write Subchannel>**link ef** (type **exit** twice and then **list all**)	4
System General	System name: NU_A Location: XYZ Contact: Administrator	Config>**set host** Config>**set location** Config>**set contact**	
System SNMP Config General	SNMP (checked)	Config>**p snmp** SNMP Config>**enable snmp**	
System SNMP Config Communities General	Community name: admin Access type: Read-write trap Community view: All	SNMP Config>**add community** SNMP Config>**set comm access write**	5

Table 11 (Page 2 of 3). Channel Gateway (APPN & IP over MPC+). See page 125 for a description and 140 for a diagram of this configuration.

Configuration Program Navigation	Configuration Program Values	Command-Line Commands	Notes
Protocols IP General	Internal address: 172.128.252.1 Router ID: 172.128.1.1	Config>**p ip** IP config>**set internal 172.128.252.1** IP config>**set router-id 172.128.1.1**	
Protocols IP Interfaces	Interface 0 (TR slot 1 port 1) IP address: 172.128.1.1 Subnet mask: 255.255.255.0 Interface 2 (MPC+ interface) IP address: 16.49.48.204 Subnet mask: 255.255.255.0	IP config>**add address** (once per i/f)	
Protocols IP OSPF General	OSPF (checked)	Config>**p ospf** OSPF Config>**enable ospf**	6
Protocols IP OSPF Area Configuration General	Area number: 0.0.0.0 Stub area (not checked)	OSPF Config>**set area**	
Protocols IP OSPF AS Boundary Routing	AS Boundary Routing (checked to enable) Import direct routes (checked to enable)	OSPF Config>**enable as** Import direct routes (Accept other defaults)	7
Protocols IP OSPF Interfaces	Interface 0 OSPF (checked)	OSPF Config>**set interface** Interface IP address: **172.128.1.1** Attaches to area: **0.0.0.0** (Accept other defaults)	
Protocols APPN General	APPN network node (checked to enable) Network ID: STFNET Control point name: NUGW	Config>**p appn** APPN config> **set node** Enable APPN Network ID: **STFNET** Control point name: **NUGW** (Accept other defaults)	

Table 11 (Page 3 of 3). Channel Gateway (APPN & IP over MPC+). See page 125 for a description and 140 for a diagram of this configuration.

Configuration Program Navigation	Configuration Program Values	Command-Line Commands	Notes
Protocols APPN Interfaces	(highlight Interface 0 Token Ring) (click on the configure tab) Define APPN port (checked to enable) Port name: TR001	APPN config>**add port** APPN Port Link Type**: TOKEN RING** Port name**: TR001** Enable APPN (Accept other defaults)	
Protocols APPN Interfaces	(highlight Interface 0 Token Ring) (click on the Link stations tab) TRTG001 (new definition) General-1 Tab: Link station name: TRTG001 General-2 Tab: MAC address of adjacent node: 400022AA0011 Adjacent Node Type: APPN Network Node (click on **Add** to create the Link station)	APPN config>**add link** Port name for the link station**: TR001** Station name**: TRTG001** MAC address of adjacent node**: 400022AA0011** (Accept other defaults)	8
Protocols APPN Interfaces	(highlight Interface 2 ESCON-MPC+) (click on the configure tab) Define APPN port (checked to enable) Port name: MPC001	APPN config>**add port** APPN Port Link Type**: MPC** Interface Number**: 2** Port name**: MPC001** Enable APPN (Accept other defaults)	
Protocols APPN Interfaces	(highlight Interface 2 ESCON-MPC+) (click on the Link stations tab) MPCTG001 (new definition) General-1 Tab: Link station name: MPCTG001 General-2 Tab: Adjacent Node Type: APPN Network Node (click on **Add** to create the Link station)	APPN config>**add link** Port name for the link station**: MPC001** Station name**: MPCTG001** Adjacent Node Type**: 0** = APPN Network Node (Accept other defaults)	

Notes:

1. add dev defines a single port, not an adapter.
2. The configuration program assigns an interface number to all ports of an adapter automatically and you delete the ones you do not want to use. From the command line, you type the add dev command for each port you want to use, and the interface number (also known as ″net number″) is the output of the command.
3. When you “Add” the interface, a new interface will be generated and it will be assigned the next available interface number.
4. The values that you enter when configuring the subchannels must match values configured at the host. See Appendix A, “Sample Host Definitions” on page 167 for examples of how to match these values.
5. You need a write-capable SNMP community only if you want to download configuration files from the configuration program directly into the router. SNMP is not required to TFTP a configuration file to the router.
6. You can also use RIP in place of OSPF.
7. You need to import direct routes into OSPF from the ESCON interface because OSPF is not enabled on the ESCON interface. Instead, the subnet on the ESCON interface is imported into OSPF in the Network Utility and then propagated to the network. This is necessary to prevent error messages from occurring at the host if the Network Utility sends the OSPF updates over the MPC+ connection. TCP/IP at the host does not (yet) support Link State Advertisements from an OSPF router.
8. The destination MAC address in this example is the local router on the right-hand side of the campus backbone in Figure 27 on page 140. This router is also configured to be an APPN network node.

Chapter 9. Data Link Switching

This chapter introduces Data Link Switching (DLSw) and summarizes the DLSw function implemented in Network Utility.

9.1 What Is DLSw?

DLSw is an IBM-invented standard technology for transporting connection-oriented protocols, mainly SNA and NetBIOS, across IP backbone networks. DLSw routers on the edges of an IP network field link establishment requests from native SNA and NetBIOS endstations, search among peer DLSw routers for one serving the target endstation, and then set up a path and relay application data between the endstations through the peer router.

The protocol that flows between DLSw routers is documented in RFC 1795, ″Data Link Switching: Switch to Switch Protocol.″ Clarifications about this protocol and multicast IP-based scalability enhancements are documented in RFC 2166, ″DLSw v2.0 Enhancements″.

Many DLSw implementations provide a *local DLSw* function that connects two links within a single router, as opposed to connecting them across an IP network to another DLSw router. Depending on the DLC types involved, this function may be equivalent to that of a FRAD or X.25 PAD.

9.2 Network Utility DLSw Function

The Network Utility DLSw implementation is nearly identical in function to that of the IBM 2210 and 2216 routers. It can handle the following endstation protocols:

- SNA
  - PU 4/5 to PU 2.0 (and IBM 5394 on SDLC)
  - T2.1 to T2.1
  - PU 4/5 to PU 4/5
- NetBIOS
  - Point-to-point sessions
  - Broadcast datagram traffic
- LAN Network Manager
  - LNM to bridge servers (for example, LBS, CRS, REM)
  - LNM to 8235 intelligent hub
  - LNM to LAN Station Manager

Network Utility DLSw can communicate with endstations across the following data link control (DLC) types:

- 802.2 LLC

  LLC can be carried over any of these interface types:

  - Token-ring
  - Ethernet (10 Mbps or 10/100 Mbps adapters)
  - FDDI
  - PPP links enabled for remote bridging

  - Frame relay PVCs and SVCs enabled for remote bridging (RFC 1490/2427 bridged frame formats)
  - ATM LAN emulation
  - ATM native bridging (RFC 1483 bridged frame formats)

- SDLC

  DLSw can represent the primary station on a multipoint line, multiple secondary stations, or a single fully negotiable station on a point-to-point line.

- QLLC

  DLSw supports any combination of QLLC PVCs and SVCs on a single X.25 interface. It can handle parallel virtual circuits to the same remote DTE address, as well as incoming calls from non-configured SVCs.

- APPN

  You can configure APPN to attach to the DLSw function residing in the same Network Utility. This allows APPN to have links with any PU2.0 or T2.1 SNA end station in the DLSw network, without requiring APPN to be present in remote (especially branch office) routers.

- Channel-LSA

  DLSw supports an internal interface to the ESCON and parallel channel LSA function residing in the same Network Utility. This allow the host to have links with any SNA endstation in the DLSw network, without requiring separate channel gateway and central-site DLSw router products.

With remote DLSw (across IP to another router), Network Utility DLSw supports conversion from TCP DLSw frames to any of the supported DLC types. Local DLSw is supported only for specific combinations of DLC types, as shown here:

	LLC	SDLC	QLLC	APPN	Channel-LSA
LLC	(1)	x	x	(2)	x
SDLC	x	x	x		x
QLLC	x	x	x		x
APPN					
CHANNEL	x	x	x		

Notes:
1 - You should use bridging for local LLC-to-LLC connectivity. The only exception supported by local DLSW is LLC to a frame relay bridge port that is configured as a Boundary Access Node (BAN) port.
2 - APPN has native support for LLC, SDLC, and QLLC, so DLSw does not allow APPN to reach local DLCs of these types.

The following list summarizes some of the other capabilities and features of IBM Network Utility DLSw.

- Dynamic compatibility to all DLSw protocol standards

  IBM DLSw supports RFC 1434+, RFC 1795 (DLSw Version 1), and RFC 2166 (DLSw Version 2). It dynamically detects the protocol level of each partner router with no pre-configuration, and can simultaneously handle partners at different protocol levels.

- Dynamic and on-demand partners

IBM DLSw supports bringing up TCP connections to configured partners only when required, as well as discovering endstations served by non-configured partners, and bringing up those TCP connections on demand.

- Multicast IP discovery

  With the simple configuration of multicast IP addresses or groups, IBM DLSw can perform multicast searches for both endstations and partners. IBM DLSw provides a number of dynamic extensions to the DLSw Version 2 standard, including resource registration and simplified group configuration.

- Traffic prioritization

  There are configuration options allowing you to control not only SNA versus NetBIOS prioritization, but also individual circuit priorities. This is in addition to the Bandwidth Reservation System′s (BRS) extensive support for interface-level traffic prioritization.

- Advanced filtering and static cache entries

  IBM DLSw includes extensive support for MAC address and NetBIOS name lists and static caching, allowing you to control what links are used for searching for resources as well as which remote partners are preferred.

- Load balancing and fault tolerance

  IBM DLSw can cache multiple remote partners and select among them on the basis of neighbor priority, largest frame size support, or first to reply. You can also use the neighbor priority feature to ensure that one central-site router serves only as a backup for another.

  For configurations involving duplicate MAC addresses, you can disable the neighbor priority feature, or set cache parameters to control the paths used to reach those MAC addresses.

9.3 Example Configurations

This section describes three sample configurations that use the Data Link Switching feature of the Network Utility. These configurations are:

- DLSw LAN Catcher
- DLSw LAN Channel Gateway
- DLSw X.25 Channel Gateway

9.3.1 DLSw LAN Catcher

This scenario is shown in Figure 28 on page 148. In this scenario, the SNA traffic in the remote sites uses DLSw to get back to the data center.

The Network Utility is in the data center on the backbone LAN segment. It is a DLSw partner with each remote router and as such requires a TCP session with each. The advantage to this approach is that all of the CPU cycles needed to manage these TCP sessions and to terminate the DLSw connections are concentrated in the Network Utility. Without the Network Utility, the local routers or the host gateway (if DLSw-capable) could be consumed by this workload.

From the host perspective, the SNA LLC2 traffic is bridged into the Network Utility from the host gateway. The host gateway is either an IBM 3745/46, an IBM 3746 with the Multiaccess Enclosure (MAE), or an IBM 2216.

You can take advantage of the 2-port token-ring adapter in the Network Utility by bringing in the IP-encapsulated SNA traffic on one port and delivering LLC2 SNA traffic onto the other token-ring port. Thus, you have twice the bandwidth available with an additional benefit of separating the IP and SNA traffic onto separate rings. Because the Network Utility provides LLC local acknowledgements (spoofing) to the host for each LLC connection, this removes a considerable amount of traffic from the campus backbone in large network environments.

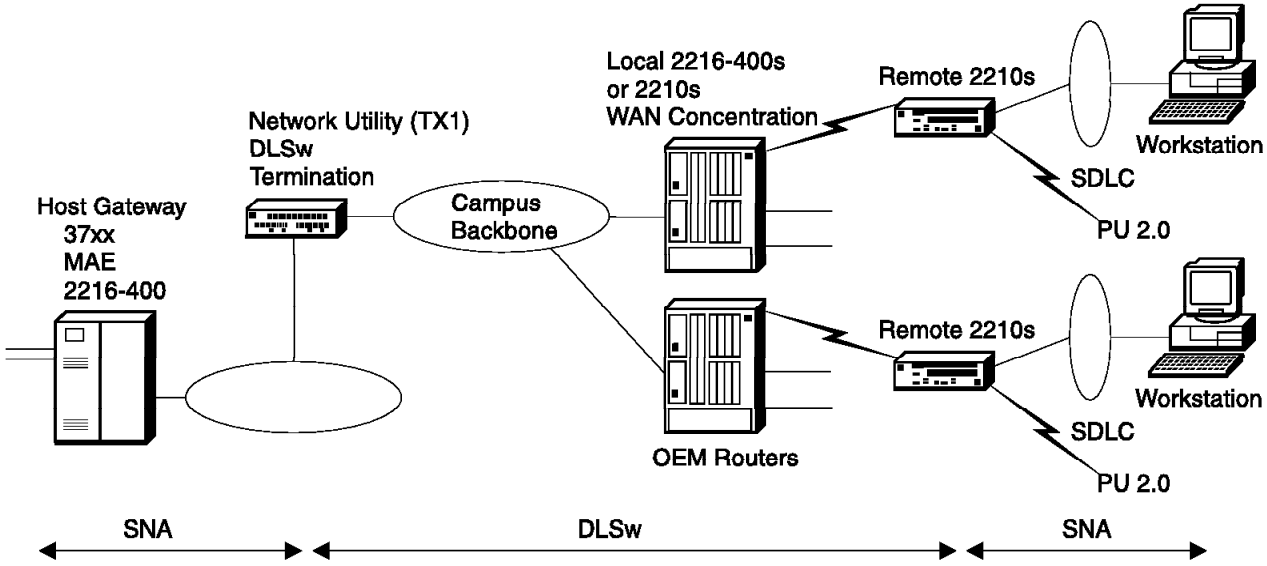

Figure 28. DLSw LAN Catcher

9.3.1.1 Keys to Configuration

For the most part, this is a standard DLSw configuration. However, you should be aware of the following points when configuring the Network Utility as a DLSw LAN Catcher:

- For this scenario, you should configure the Network Utility to allow TCP sessions from any of the remote routers. This is called DLSw dynamic neighbors. This keeps you from having to define the IP address of each DLSw partner on the Network Utility. The default value for dynamic neighbors is ″Enabled.″
- The Network Utility introduces a new parameter for IBM DLSw implementations that allows you to specify how explorer frames are forwarded. This is especially important in the outbound direction from the central site. The parameter is called *enable/disable forwarding explorers* and it gives you the flexibility to specify any of the following options:
  - Disable forwarding of explorer frames

    This option completely disables forwarding of explorer frames.
  - Forward explorer frames to the local TCP connection only

    If you want to block explorer frames from going out on WAN links, then you can specify this option. This is the default value for the Network Utility.
  - Forward explorer frames to all DLSw partners

With this option, explorer frames are sent out to all DLSw partners.

For a complete look at the configuration parameters needed for the DLSw LAN Catcher scenario, see Table 13 on page 159.

9.3.2 DLSw LAN Channel Gateway

This scenario is shown in Figure 29. As in the DLSw LAN catcher scenario, the Network Utility terminates the DLSw sessions from the remote routers. However, in this case, there is an ESCON Channel Adapter in the Network Utility. Instead of bridging the traffic from the DLSw function onto the LAN segment, this configuration passes it directly to the channel via an LSA loopback interface configured in the Network Utility.

This configuration also demonstrates the use of the Network Utility to support SNA traffic from the local campus to the host. This traffic is bridged off the campus backbone through the LSA loopback interface. All SNA devices in the network are configured with the same host destination MAC address which is the MAC address of the LSA loopback interface. This includes the devices at the main site as well as the devices in the remote sites.

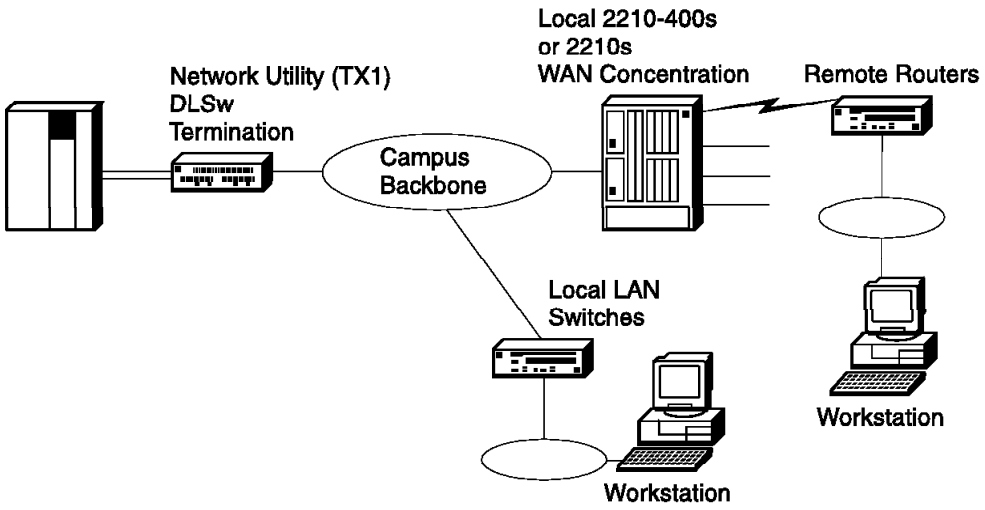

Figure 29. DLSw LAN Channel Gateway

Note: This example illustrates the use of the Network Utility as a channel gateway for DLSw traffic only. However, many of the functions illustrated in the Channel Gateway example configurations on page 116 could be combined with DLSw termination in a valid channel gateway configuration.

9.3.2.1 Keys to Configuration

Note the following points when configuring the Network Utility as a DLSw LAN Channel Gateway:

- An LSA interface must be configured and loopback must be enabled on this interface. Enabling loopback creates a virtual LAN inside the Network Utility. The only two devices on this LAN are the host and the DLSw termination point. A MAC address is defined on the LSA interface that represents the host on the channel. This is the destination MAC address that is configured in the downstream devices.

Note: You can also define an LSA direct connection for the traffic to be bridged in from the local LAN segments. If you do this, then the devices on these segments will have a different destination MAC address from the remote devices because the LSA direct interface will have a different MAC address from the LSA loopback interface.

- When configuring DLSw, you need to open SNA SAPs for the LSA interface as well as the token-ring interface.
- The subchannel configuration for the LSA interface must match parameters configured in the host. See Table 7 on page 117 for a description of the subchannel parameters and Appendix A, “Sample Host Definitions” on page 167 for example host definitions. This information will help you see how these parameters correlate.
- You need to configure a *local TCP connection*. This is done by defining a DLSw partner whose IP address is the internal address of the Network Utility. This is used for the traffic that is bridged from the local LAN segments into the host. This traffic gets bridged into the Network Utility into DLSw where the local TCP connection passes the traffic to the LSA loopback interface.
- The Network Utility currently supports a maximum of 2048 link stations per MAC address/SAP pair (for example, a destination MAC address of 400022AA0099 with SAP 04). If you need more than 2048 workstations, you have to define another LSA interface with a different SAP or a different MAC address. Remember that each LSA interface requires one subchannel of the 32 available on one ESCON channel adapter. You must also define the corresponding XCA major node to support each LSA interface.

9.3.3 X.25 Channel Gateway

This scenario is shown in Figure 30 on page 151. It uses Local DLSw in the Network Utility to map between X.25 addresses and MAC address/SAP pairs. The transport across the WAN is native Qualified Logical Link Control (QLLC), a protocol that allows SNA devices to communicate over X.25 networks. In the Network Utility, local DLSw performs protocol conversion between QLLC and LLC2 frames.

From the remote device perspective, there are two cases to consider:

1. A device on a LAN segment attached to the branch router

   On the workstation, the SNA application generates an LLC frame that it wants to send to the host. If the branch router is an IBM 2210, this LLC frame gets bridged into the 2210 DLSw function, which does three things:

   a. Protocol conversion from the LLC frame to a QLLC frame

   b. Maps the destination MAC address/SAP pair into the appropriate X.25 LCN (PVC) or DTE address (SVC)

   c. Passes the QLLC frame to X.25

   The X.25 PAD function in the branch router creates the LAPB link layer packets and sends them over the PVC (or SVC).

   If some product other than the IBM 2210 plays the role of branch router, it needs to perform these same functions but may do so without using local DLSw.

2. A device directly on the X.25 network (for example, an IBM 3174 Control Unit or an eNetwork Communications Server gateway machine attached via a Wide Area Connector Adapter)

   On these devices, SNA uses QLLC as a native DLC type. It generates a QLLC frame and sends it out over the configured PVC (or SVC).

In each of these cases, at the Network Utility, the LAPB packets are received over the X.25 circuit and passed to QLLC and then on to DLSw. DLSw does two things:

1. Protocol conversion from QLLC into an LLC2 frame
2. Mapping of the X.25 LCN (PVC) or DTE address (SVC) into the MAC address/SAP for the LSA local loopback interface

The traffic is then passed to the LSA loopback interface for transport across the ESCON channel.

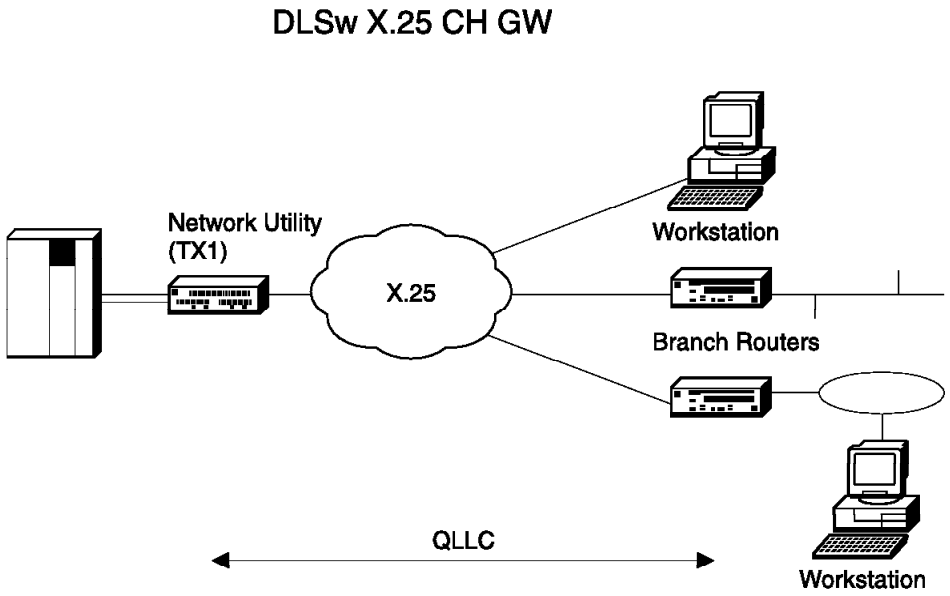

Figure 30. DLSw X.25 Channel Gateway

9.3.3.1 Keys to Configuration

The following list summarizes general configuration tasks you need to perform for this scenario. Please refer to other DLSw and LSA loopback configurations for details. The LSA loopback interface is configured the same as in 9.3.2, “DLSw LAN Channel Gateway” on page 149.

- Add and configure the ESCON and LSA interfaces.
- Add and configure the X.25 interface. From the command line, use the `net` command in talk 6 to enter the X.25 Config subprocess, then use the following commands:
  - `set address` (to set the local DTE address)
  - `add protocol dlsw` (to add DLSw as an X.25 protocol)
  - `add pvc` or `add svc` (to add the individual PVCs or range of SVCs)
- Configure the IP internal address as in other examples.
- Configure DLSw
  - Configure general DLSw (enable, SRB segment, forward explorers locally).

  - Configure the DLSw local TCP connection.
  - Configure DLSw for LSA loopback (open SAPs on the LSA interface).

In addition to these general tasks, you need to configure Network Utility DLSw to map X.25 addresses to the LSA loopback MAC address. There are three ways to do this:

- Configure the X.25 stations individually at DLSw, each with its own destination MAC address. This option applies to both PVCs and SVCs.
- Configure a list of connection IDs, each of which has its own destination MAC address. Some X.25 stations can send a connection ID when they place a call, and Network Utility matches this value to the configured list. This option applies to SVCs only.
- Configure a default destination MAC address for incoming calls that do not contain a connection ID. This option applies to SVCs only.

The remainder of this section describes how to configure each of these three address mapping methods.

If the number of remote X.25 stations is relatively small, then you can configure each remote X.25 device in DLSw to be mapped to the LSA loopback MAC address. To do this using the command line, enter `talk 6` at the * prompt and type the following:

- `protocol dls`
  - `add qllc station` (once for each remote X.25 station). The system prompts you for:
    - Interface number (X.25 interface)
    - PVC or SVC
    - Logical channel number (for PVCs) or DTE address (for SVCs)
    - Source MAC and SAP (can be generated by DLSw)
    - Destination MAC and SAP (enter the LSA loopback MAC address)
    - PU type
    - XID block/num (if the PU type is 2)

To do this using the Configuration Program, do the following:

- Protocols/DLSw/Interfaces/Serial-X25/QLLC Stations
  - Add a QLLC station (enter the same information as above)

If your remote X.25 stations can be configured to send a connection ID when they place a call, [13] you can configure DLSw to map connection ID values to destination MAC addresses. To do this using the command line, enter `talk 6` at the * prompt and type the following:

- `protocol dls`
  - `Add qllc destination` (once for each valid connection ID). The system prompts you for:
    - Connection ID
    - Destination MAC and SAP (enter the LSA loopback MAC address)

To do this using the Configuration Program, do the following:

- Protocols/DLSw/QLLC Destinations
  - Add a QLLC destination (enter the same information as above)

[13] QLLC products frequently present this parameter as a connection password.

Finally, if it is not feasible to configure each remote X.25 station or to use a connection ID, you can use the DLSw ANYCALL feature to accept any incoming X.25 call and map it to the LSA loopback MAC address. To do this using the command line, enter `talk 6` at the * prompt and type the following:

- `protocol dls`
  - `add qllc destination` (once, plus you can add specific connection IDs if you wish). The system prompts you for:
    - Connection ID (use the word 'ANYCALL')
    - Destination MAC and SAP (enter the LSA loopback MAC address)

To do this using the configuration tool, do the following:

- Protocols/DLSw/QLLC Destinations
  - Add a QLLC destination (enter the same information as above)

9.4 Managing DLSw

This section introduces some of the ways in which you can monitor and manage the DLSw function.

9.4.1 Command-Line Monitoring

DLSw supports an extensive set of commands to display status, dynamically modify configuration parameters, and actively control the state of connections. These commands are described in detail in *MAS Protocol Configuration and Monitoring Reference Volume 1*, in the chapter "Configuring and Monitoring DLSw." To access them, enter `talk 5` at the * prompt and `protocol dls` at the + prompt.

Some particularly useful commands for monitoring status are:

`list tcp sess`

Shows the status of all known TCP connections to partner routers. You can see the state of the TCP connections as they come up and go down, as well as the level of the DLSw protocol in use, and summary statistics on the number of DLSw circuits using each connection. If you configure DLSw to accept TCP connections only from dynamic (not configured) partners, this command displays the status of connections as initiated by remote routers. There will be no status if the remote routers are not actively bringing up TCP connections.

If you configure a "local TCP connection" to enable local DLSw function, this connection is flagged as such on the command output so that it can be distinguished from remote partner connections.

`list dls sess all`

Shows the status of all active DLSw sessions. A session, also called a circuit, is defined by a MAC and SAP address 4-tuple and corresponds to an SNA link, not an SNA LU-LU session. Sessions are normally driven up and down by SNA endstations, so the output of this command is dynamic. For every session, you see its identifying MAC and SAP addresses, state, which partner the session is connected through, and an identifier that you can use with the `list dls sess detail` command to get more information. Local DLSw sessions (those that involve only this router) show as two lines of output from this command.

Because a Network Utility may easily have hundreds or thousands of active sessions, you can use different variations of the `list dls session` command

to display only a subset of them. Instead of the keyword ″all,″ you use different keywords to show only those circuits through a given partner, or only those in a given state, and so on. There are roughly 10 keywords defined to select sessions. The output of all these commands pauses when the screen fills, waiting for a keystroke from you to continue or quit. Press the space bar to view the next screen of output.

`list dls mem`
: Shows the status of various pools of DLSw memory, as well as the memory congestion status for all active sessions.

`list llc sess all`
: Shows 802.2 LLC-specific status information for all DLSw sessions that use LLC as the protocol between the router and the endstation. These include sessions running over LAN, channel, ATM, and remotely bridged WAN interfaces. The command output includes more state information as well as the source route to the endstation, if applicable.

`list sdlc sess all`
: Shows SDLC-specific status information for all DLSw sessions that use SDLC as the protocol between the router and the endstation. The command output includes SDLC addressing information as well as state information. If you are working with SDLC devices, this command is more useful than the generic `list dls sess`.

`list qllc sess`
: Shows QLLC-specific status information for all DLSw sessions that use QLLC over X.25 as the protocol between the router and the endstation. The command output includes QLLC addressing information as well as detailed state information. Because the router supports incoming dynamic SVCs, this command is essential to see the status of both configured and dynamic QLLC PVCs and SVCs.

DLSw supports dynamic modification under talk 5 of the vast majority of parameters you can configure under talk 6. DLSw follows the standard model where changes made under talk 5 have an immediate effect but do not survive a box reboot, while changes made under talk 6 take effect only after a box reboot. The talk 5 list commands show the values that are currently active in the running product.

The talk 5 commands `delete` and `disable` give you the power to tear down an existing DLSw connection. For example, you can use `delete dls` *session number* to clean up a hung session and allow the endstations to redrive it. `Delete`/`add` and `disable`/`enable` sequences are powerful methods to recycle configured TCP, SDLC, and QLLC connections.

9.4.2 Event Logging Support

DLSw has several hundred ELS messages defined, ranging from informational messages about normal events, to warnings of serious error conditions. Here are some of the types of DLSw events that can generate ELS messages:

- Initialization and configuration errors
- Partner TCP connection and capabilities frames sent or received
- Explorer frames sent or received for a particular MAC address or NetBIOS name
- Circuit setup/takedown frames sent or received
- DLC link setup/takedown frames sent or received

- Data frames sent or received on active circuits
- Pacing window changes on active circuits
- Memory allocation errors
- Unexpected protocol flows, frames discarded
- Frame flows do not match configuration

Although these messages are used primarily by software engineers to resolve problems, a user with a basic knowledge of the DLSw protocol and DLC link activation flows should be able to make sense of them and debug simple configuration mistakes. By activating these ELS messages and watching the output via talk 2, you should be able at least to answer the question ″Is anything happening?″

″DLS″ is one of the named subsystems within ELS. To activate the standard set of error messages, type `disp sub dls` from the event menu under either talk 6 or talk 5. To activate all DLSw messages, enter `disp sub dls all`. The corresponding commands to deactivate messages begin with `nodisp`. For general information on controlling and viewing ELS messages, please see the *Nways Event Logging System Messages Guide*, SC30-3682.

If you are trying to trace a link activation attempt, DLSw messages alone may not show the complete picture. You can activate the ELS messages for the underlying DLC type as follows:

LLC	`disp sub llc all`
SDLC	`disp sub sdlc all`
QLLC	`disp sub qllc all`
	disp sub x253 all (X.25 layer 3, the packet layer)
Channel-LSA	`disp sub lsa all`

Refer to the *Event Logging System Messages Guide* (on CD-ROM and the 2216 Web Page) for a full list of individual messages and their meaning.

9.4.3 SNA Management Support

Unlike APPN, Network Utility DLSw does not send SNA alerts. It does send traps (described in the following section) and trigger ELS messages that can generate traps. You can convert those traps to alerts.

9.4.4 SNMP MIB and Trap Support

Network Utility DLSw provides full read-only and partial read-write support for the IETF standard DLSw MIB documented in RFC 2024. This large MIB gives visibility to most of the important configuration, status, and accounting information that products implementing RFC 1795 and 2166 should have. This information includes:

- Configuration
  - Node characteristics, for example, dynamic partners are enabled
  - Configured partner information
  - Configured directory/cache entries
- Status
  - Node up or down, for how long
  - Active TCP connections, for how long, dynamic partner information
  - Dynamic directory/cache information
  - Active circuits, for how long, DLC information

- Statistics and Accounting
  - Counts of TCP connections up and down (normal and error)
  - Data and control frames counts per partner
  - Counts of circuits up and down
  - Indices to underlying DLC MIBs for per-circuit frame counts
  - Pacing counts for active circuits

Network Utility DLSw supports all the traps defined in RFC 2024, reporting the following events:

- A TCP connection is terminated due to capabilities exchange failure or a DLSw protocol violation
- A TCP connection comes up or goes down
- A circuit comes up or goes down

DLSw all supports trap control data items so a management station can set the conditions under which a trap is generated.

In addition to RFC 2024, Network Utility DLSw supports IBM-specific DLSw MIB extensions for multicast IP-based groups and for QLLC stations.

9.4.5 Network Management Application Support

The Network Utility Java-based application implemented in the Nways Manager products provides integrated support for the standard DLSw MIB and the IBM-specific DLSw MIB extensions.

To view DLSw resources and their status using these products, you bring up specific panels that present key information from the DLSw MIB and from its underlying DLC-layer MIBs (LLC, SDLC, or X.25). You can also use integrated browser support to view the information in any of these MIBs.

You can control the emission of DLSw traps from the Nways Manager products, so that a given trap is generated always, never, or only under certain conditions.

Nways Manager for AIX can show you a DLSw topology view of your network, including DLSw connectivity, resources, and color-coded status. The topology is refreshed as new nodes are discovered. This application does not present the topology of DLSw IP multicast groups.

Chapter 10. DLSw Example Configuration Details

This chapter contains diagrams and configuration parameter tables for several of the example DLSw network configurations in Chapter 9, “Data Link Switching” on page 145. The parameter values shown are from real working test configurations.

For an explanation of the columns and conventions in the configuration parameter tables, see 4.2.2, “Example Configuration Table Conventions” on page 67.

The Network Utility World Wide Web pages contain binary configuration files that match these configuration parameter tables. To access these files, follow the Download link from:

```
http://www.networking.ibm.com/networkutility
```

The configurations documented in this chapter are:

Table 12. Cross-Reference of Example Configuration Information

Configuration Description	**Parameter Table**
9.3.1, “DLSw LAN Catcher” on page 147	Table 13 on page 159
9.3.2, “DLSw LAN Channel Gateway” on page 149	Table 14 on page 163

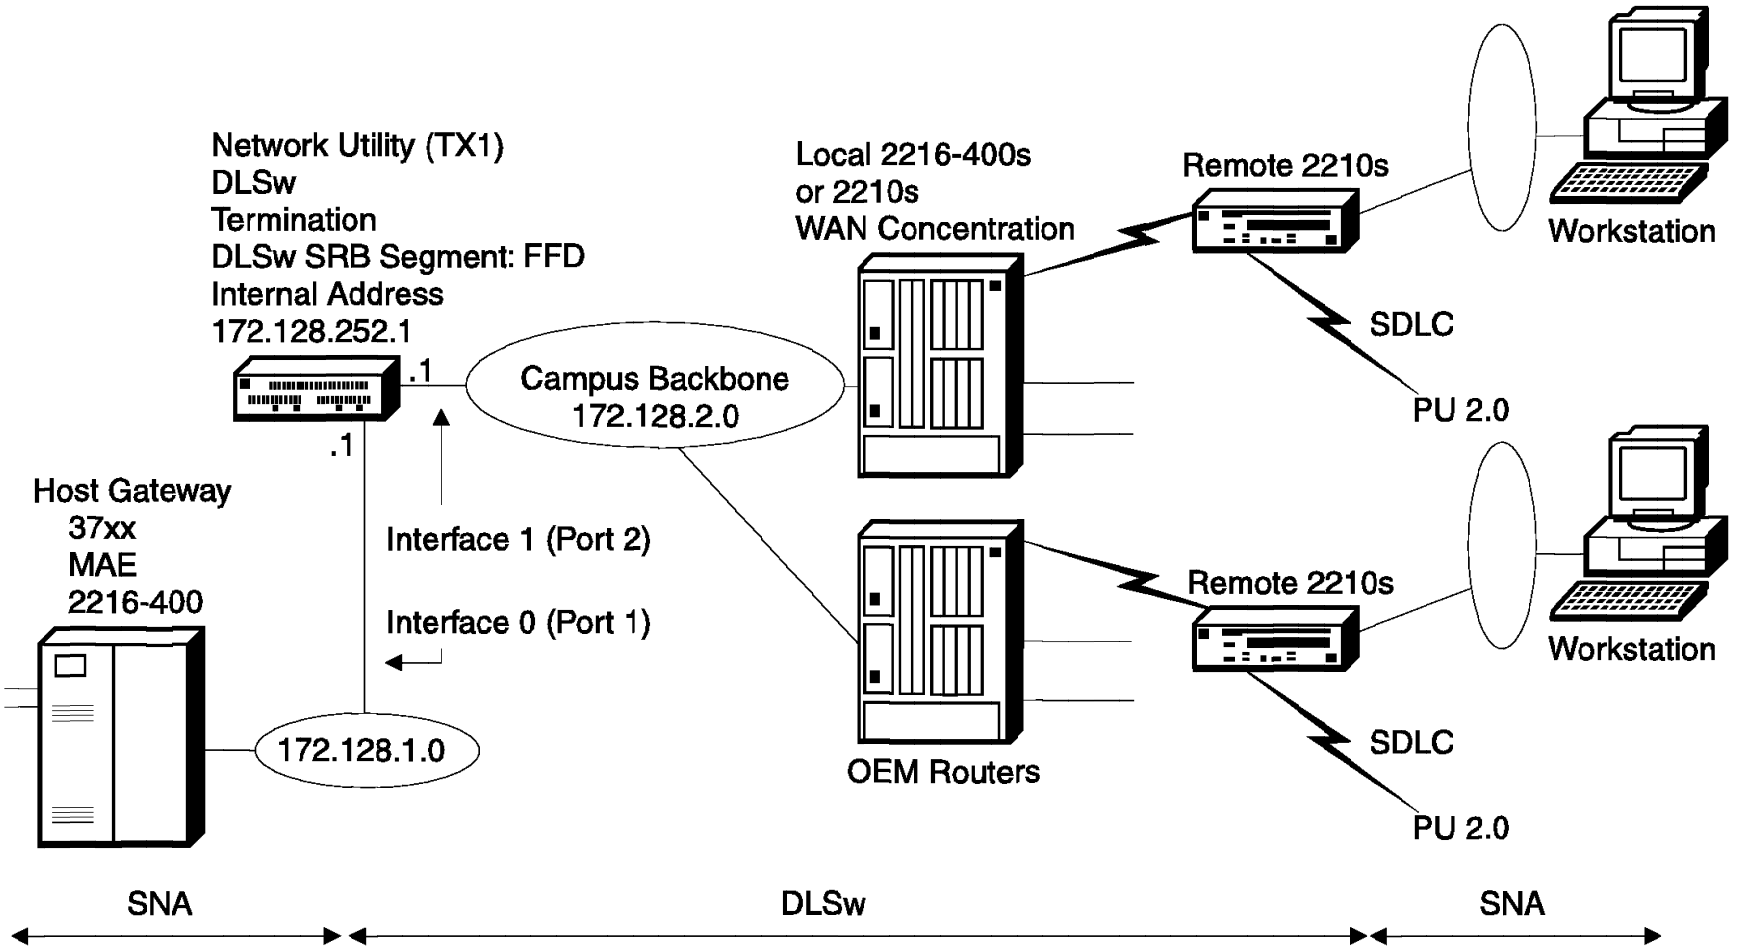

Figure 31. DLSw LAN Catcher

Table 13 (Page 1 of 2). DLSw LAN Catcher. See page 147 for a description and 158 for a diagram of this configuration.

Configuration Program Navigation	Configuration Program Values	Command-Line Commands	Notes
Devices Adapters Slots	Slot 1: 2-Port TR	See ″add device″ on next row	1
Devices Adapters Ports	Slot 1/Port 1: Interface 0: TR Slot 1/Port 2: Interface 1: TR	Config>**add dev tok** (once per port)	2
Devices Interfaces	Interface 0 MAC address: 400022AA0001 Packet size: 4399 Interface 1 MAC address: 400022AA0002 Packet size: 4399	Config>**net 0** TKR config>**set phy 40:00:22:AA:00:01** TKR config>**packet 4399** TKR config>**exit** Config>**net 1** TKR config>**set phy 40:00:22:AA:00:02** TKR config>**packet 4399**	
System General	System name: NU_A Location: XYZ Contact: Administrator	Config>**set host** Config>**set location** Config>**set contact**	
System SNMP Config General	SNMP (checked)	Config>**p snmp** SNMP Config>**enable snmp**	
System SNMP Config Communities General	Community name: admin Access type: Read-write trap Community view: All	SNMP Config>**add community** SNMP Config>**set comm access write**	3
Protocols IP General	Internal address: 172.128.252.1 Router ID: 172.128.1.1	Config>**p ip** IP config>**set internal** IP config>**set router-id**	
Protocols IP Interfaces	Interface 0 (TR slot 1 port 1) IP address: 172.128.1.1 Subnet mask: 255.255.255.0 Interface 1 (TR slot 1 port 2) IP address: 172.128.2.1 Subnet mask: 255.255.255.0	IP config>**add address** (once per i/f)	4
Protocols IP OSPF General	OSPF (checked)	Config>**p ospf** OSPF Config>**enable ospf**	5

Table 13 (Page 2 of 2). DLSw LAN Catcher. See page 147 for a description and 158 for a diagram of this configuration.

Configuration Program Navigation	Configuration Program Values	Command-Line Commands	Notes
Protocols IP OSPF Area Configuration General	Area number: 0.0.0.0 Stub area (not checked)	OSPF Config>**set area**	
Protocols IP OSPF Interfaces	Interface 0 OSPF (checked) Interface 1 OSPF (checked)	OSPF Config> **set interface** Interface IP address **172.128.1.1** Attaches to area **0.0.0.0** (Accept other defaults) OSPF Config>**set interface** Interface IP address **172.128.2.1** Attaches to area **0.0.0.0** (Accept other defaults)	
Protocols DLSw General General	DLSw (checked) SRB segment: FFD Forward explorers: disabled	Config>**p dls** DLSw Config>**enable dls** DLSw Config>**set srb** DLSw Config>**disable forward all**	6
Protocols DLSw General Dynamic Neighbors	Dynamic neighbors (checked)	DLSw Config>**enable dynamic**	7
Protocols DLSw Interfaces	Interface 0 (TR slot 1 port 1) SAP type: SNA (SAPs 0,4,8,C)	DLSw Config>**open 0 sna**	8
Protocols Bridging General	Bridging (checked) DLSw (checked)	Config>**p asrt** ASRT config>**enable br** ASRT config>**enable dls**	9
Protocols Bridging Interfaces	Interface 0 (TR slot 1 port 1) Bridging port (checked) Interface supports: SRB Segment number: 001 MTU size: 4399	('enable br' assumed) ASRT config>**disable transp 1** ASRT config>**enable source 1** ASRT config>**delete port 2**	10

Notes:

1. add dev defines a single port, not an adapter.
2. The configuration program assigns an interface number to all ports of an adapter automatically and you delete the ones you do not want to use. From the command line, you type the add dev command for each port you want to use, and the interface number (also known as ″net number″) is the output of the command.
3. You need a write-capable SNMP community only if you want to download configuration files from the configuration program directly into the router. SNMP is not required to TFTP a configuration file to the router.
4. Only Interface 1 needs to be configured for IP for DLSw to function correctly in this example. Interface 0 is configured here for IP, solely for box management purposes.
5. You can also use RIP in place of OSPF.
6. We disable the forwarding of remote explorers as a general filter, to prevent backbone LAN traffic from generating DLSw search messages on the WAN links to remote sites. This means that all circuits must be initiated by the remote routers. If your network requires the host to be able to initiate connections out to the remote sites, change this parameter to ″forward to all DLSw peers″.

   If the remote routers are IBM routers, you can configure them individually to control which search messages they want to receive, using MAC address and NetBIOS name lists. You can also configure whether each will bring up its TCP connection to Network Utility all the time or drop it when unused, using the *connectivity setup type* parameter.
7. Having dynamic neighbors enabled is the default value, so you do not have to change this panel or issue this command. We show it here to point out that this is the parameter that allows remote DLSw partners (neighbors) to establish TCP connections to this Network Utility without you having to define their IP addresses here. Each remote router needs to be configured with this Network Utility's internal IP address (172.128.252.1) as its partner address.
8. SAPs do not need to be opened on Interface 1 since that interface is only carrying IP traffic and not LLC traffic.
9. ″enable br″ automatically creates TB bridge ports for both token-ring interfaces. Bridge port numbers are 1 and 2, and are independent of adapter port numbers.
10. The disable and enable commands change bridge port 1 from TB to SRB. The ″delete port″ command turns off bridging on interface 1 (bridge port 2). Bridging would be required on this interface if we needed to support local end station traffic bridging from the Campus Backbone to the host.

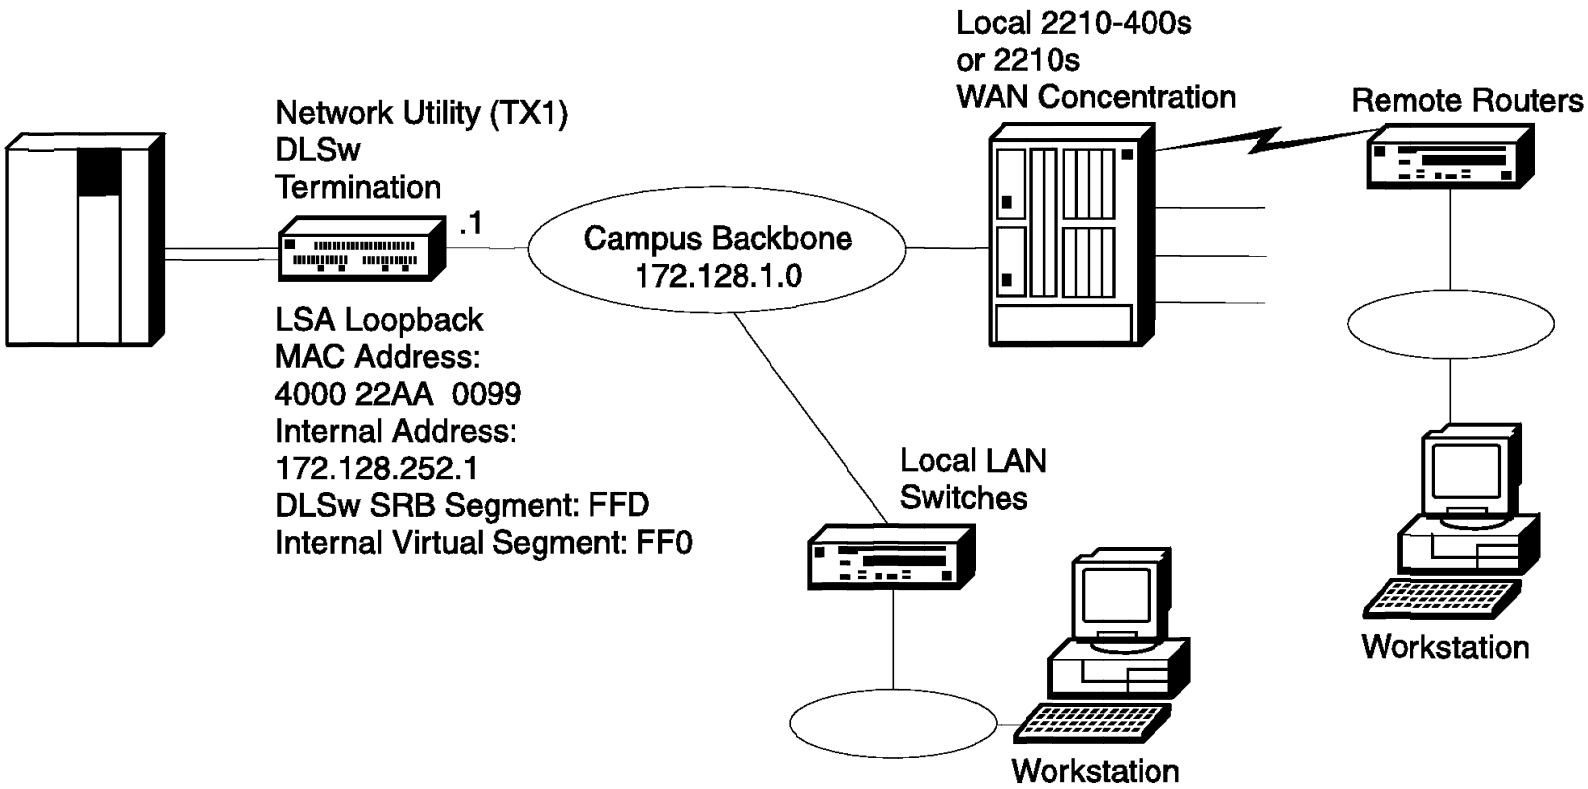

Figure 32. DLSw LAN Gateway

Table 14 (Page 1 of 3). DLSw LAN Gateway. See page 149 for a description and 162 for a diagram of this configuration.

Configuration Program Navigation	Configuration Program Values	Command-Line Commands	Notes
Devices Adapters Slots	Slot 1: 2-Port TR Slot 2: ESCON	See ″add device″ on next row	1
Devices Adapters Ports	Slot 1/Port 1: Interface 0: TR Slot 2/Port 1: Interface 1: ESCON	Config>**add dev tok** Config>**add dev escon**	2
Devices Interfaces	Interface 0 MAC address: 400022AA0001	Config>**net 0** TKR config>**set phy 40:00:22:AA:00:01**	
Devices Channel Adapters ESCON Interfaces ESCON Interfaces	Base network number: 1 Protocol type: LSA (do this first) Loopback (checked - do this second) LAN type: Token Ring Maximum data frame: 2052 MAC address: 400022AA0099	Config>**net 1** ESCON Config>**add lsa** (added as interface 2) ESCON Add Virtual>**enable loopback** ESCON Add Virtual>**mac 40:00:22:AA:00:99** ESCON Add Virtual>**lan tok** ESCON Add Virtual>**maxdata 2052** (continue in same session with next row)	3
Devices Channel Adapters ESCON Interfaces ESCON Subchannels	Interface 2, Base net 1, Protocol LSA Device address: E4 Subchannel type: read/write Link address: EF	ESCON Add Virtual>**subchannel add** (cont'd) ESCON Add LSA Subchannel>**device E4** ESCON Add LSA Subchannel>**link EF** (Type **exit** twice and then **list all**)	
System General	System name: NUA_SC1C Location: XYZ Contact: Admin	Config>**set host** Config>**set location** Config>**set contact**	
System SNMP Config General	SNMP (checked)	Config>**p snmp** SNMP Config>**enable snmp**	
System SNMP Config Communities General	Community name: admin Access type: Read-write trap Community view: All	SNMP Config>**add community** SNMP Config>**set comm access write**	4
Protocols IP General	Internal address: 172.128.252.1 Router ID: 172.128.1.1	Config>**p ip** IP config>**set internal** IP config>**set router-id**	

Table 14 (Page 2 of 3). DLSw LAN Gateway. See page 149 for a description and 162 for a diagram of this configuration.

Configuration Program Navigation	Configuration Program Values	Command-Line Commands	Notes
Protocols IP Interfaces	Interface 0 (TR slot 1 port 1) IP address: 172.128.1.1 Subnet mask: 255.255.255.0	IP config>**add address**	
Protocols IP OSPF General	OSPF (checked)	Config>**p ospf** OSPF Config>**enable ospf**	5
Protocols IP OSPF Area Configuration General	Area number: 0.0.0.0 Stub area (not checked)	OSPF Config>**set area**	
Protocols IP OSPF Interfaces	Interface 0 OSPF (checked)	OSPF Config>**set interface** Interface IP address **172.128.1.1** Attaches to area **0.0.0.0** (Accept other defaults)	
Protocols DLSw General General	DLSw (checked) SRB segment: FFD Forward explorers: local TCP connection only	Config>**p dls** DLSw Config>**enable dls** DLSw Config>**set srb** DLSw Config>**enable forward local**	6
Protocols DLSw General Dynamic Neighbors	Dynamic neighbors (checked)	DLSw Config>**enable dynamic**	7
Protocols DLSw TCP Connections	(add) Neighbor IP address: 172.128.252.1 (this is the router internal IP address)	DLSw Config>**add tcp** DLSw neighbor IP address**: 172.128.252.1** (Accept other defaults)	8
Protocols DLSw Interfaces	Interface 0 (TR slot 1 port 1) SAP type: SNA (SAPs 0,4,8,C) Interface 2 (ESCON-LSA) SAP type: SNA (SAPs 0,4,8,C)	DLSw Config>**open 0 sna** DLSw Config>**open 2 sna**	9
Protocols Bridging General	General Tab: Bridging (checked) DLSw (checked) SRB Tab: Internal Virtual Segment: FF0	Config>**p asrt** ASRT config>**enable br** ASRT config>**enable dls**	10

Table 14 (Page 3 of 3). DLSw LAN Gateway. See page 149 for a description and 162 for a diagram of this configuration.			
Configuration Program Navigation	**Configuration Program Values**	**Command-Line Commands**	**Notes**
Protocols Bridging Interfaces	Interface 0 (TR slot 1 port 1) Bridging port (checked) Interface supports: SRB Segment number: 001 MTU size: 2052	('enable br' assumed) ASRT config>**disable transp 1** ASRT config>**enable source 1**	11

Notes:

1. add dev defines a single port, not an adapter.
2. The configuration program assigns an interface number to all ports of an adapter automatically and you delete the ones you do not want to use. From the command line, you type the add dev command for each port you want to use, and the interface number (also known as ″net number″) is the output of the command.
3. The MAC address representing this LSA loopback interface is the target MAC address that all endstations in the DLSw network use to reach the host through this Network Utility.
4. You need a write-capable SNMP community only if you want to download configuration files from the configuration program directly into the router. SNMP is not required to TFTP a configuration file to the router.
5. You could choose to use RIP in place of OSPF.
6. We enable local forwarding to allow endstations on the local campus to reach the host. We disable the forwarding of remote explorers as a general filter, to prevent backbone LAN traffic from generating DLSw search messages on the WAN links to remote sites. This means that all remote circuits must be initiated by the remote routers. If your network requires the host to be able to initiate connections out to the remote sites, change this parameter to ″forward to all DLSw peers″.

   If the remote routers are IBM routers, you can configure them individually to control which search messages they want to receive, using MAC address and NetBIOS name lists. You can also configure whether each will bring up its TCP connection to Network Utility all the time or drop it when unused, using the *connectivity setup type* parameter.
7. Having dynamic neighbors enabled is the default value, so you do not have to change this panel or issue this command. We show it here to point out that this is the parameter that allows remote DLSw partners (neighbors) to establish TCP connections to this Network Utility without you having to define their IP addresses here. Each remote router needs to be configured with this Network Utility's internal IP address (172.128.252.1) as its partner address.
8. Adding the internal IP address as a neighbor is required to enable DLSw to carry traffic from the ESCON/LSA interface to the backbone LAN.
9. SAPs are opened on Interface 0 to enable the LLC flow to the local LAN switches, and are not required for remote DLSw to work.
10. ″enable br″ automatically creates a TB bridge port for the token-ring interface. The bridge port number is 1, and is independent of adapter port numbers and box interface numbers.
11. The disable and enable commands change bridge port 1 from TB to SRB. Bridging is required on this interface to support local endstation traffic looping through DLSw from the Campus Backbone to the host.

Appendix A. Sample Host Definitions

This appendix contains examples of host definitions for the Network Utility in the configurations used in this manual.

Specifically, definitions for the following environments are presented:

- LSA
- LCS
- MPC+

Additionally, the differences between a Network Utility with an ESCON channel adapter and a Parallel Channel Adapter are highlighted.

For more information on defining the Network Utility to the host, refer to the *Nways Multiprotocol Access Services Software User's Guide Version 3.2*, SC30-3886.

A.1 Overview

There are three steps to define a channel-attached Network Utility to the host:

1. Define the Network Utility to the host channel subsystem

   This will be done either from the I/O configuration program (IOCP) or the Hardware Configuration Definition (HCD), depending on your MVS version. (HCD requires MVS/ESA SP Version 4.2 or later with APAR# OY67361.)

   The definition statements are slightly different for an ESCON channel-attached device than for a parallel channel-attached device. An example of these definitions is given in A.2.1, "Sample Host IOCP Definitions" on page 168.

2. Define the Network Utility as a control unit to the host operating system

   For most systems, the definitions are the same for an ESCON adapter as they are for a Parallel Channel Adapter. Obviously, they depend on the operating system being used. An example of these definitions is given in A.3, "Defining the Network Utility in the Operating System" on page 171.

3. Define the Network Utility to the host TCP/IP or VTAM

   These definitions depend on whether you are defining an LSA (SNA), an LCS (TCP/IP), or an MPC+ (SNA and/or TCP/IP) interface on the Network Utility. Section A.4, "VTAM Definitions" on page 172 shows examples of the required VTAM definitions. Section A.5, "Host IP Definitions" on page 181 shows examples of the required TCP/IP definitions.

A.2 Definitions at the Channel Subsystem Level

You make definitions at this level via the IOCP or with HCD. If HCD is available, you will probably want to use it. HCD offers an improved method of defining system hardware configuration. With HCD several complex steps required for entering hardware configuration data can be accomplished using an interactive dialog. This chapter presents only the IOCP macros that would be generated from HCD.

A.2.1 Sample Host IOCP Definitions

An example of the definitions required in the host I/O Configuration Program (IOCP) for a Network Utility configured with an ESCON adapter is as follows:

```
CHPID           PATH=((05)),TYPE=CNC
CNTLUNIT        CUNUMBR=1E0,PATH=05,CUADD=0,
                UNITADD=((E0,32)),LINK=3C,UNIT=3172
IODEVICE        UNIT=3172,ADDRESS=((1E0,32)),
                CUNUMBR=1E0
```

The following sections describe the IOCP macros that you need for defining the Network Utility at the host.

A.2.1.1 RESOURCE Statement

This identifies the host logical partitions (LPARs) by name and number. This statement is not present if the host is not partitioned *as is the case in the example above*.

- PART=((name1,x),(name2,y)...(nameX,z))

  The name identifies the LPAR and is used in the rest of the channel path definition. The number is the corresponding LPAR number. The LPAR number is used in defining the subchannel on the Network Utility. If the host is not partitioned, the LPAR number is always 0.

A.2.1.2 Channel Path ID (CHPID) Statement

The CHPID identifies the type of channel connection and who uses it.

- PATH=x

  This uniquely identifies the channel path. This value is often called the ″CHPID number″.

- TYPE=CNC

  This indicates that the channel is an ESCON channel. The channel type is CNC for ESCON and BL for block multiplexor (Parallel Channel Adapter).

- SWITCH=x

  This identifies which ESCON Director is in this path. If no director is being used, this parameter is omitted.

- SHARED

  This indicates that the CHPID can be used by multiple LPARs simultaneously. If not present, only one LPAR can use the CHPID at a time.

- PARTITION=(name1,name2,...,nameX)

  This is one form of the PARTITION parameter and it contains an access list of LPARS that indicates which partitions have access to this channel. The names must be included in the RESOURCE statement.

- PARTITION=((name1,...,nameX),(name2,...,nameY))

  This is the other form of the PARTITION parameter. In this form, the first grouping of names is the access list of LPARs, as above. The second grouping is the list of candidate LPARs that an operator could configure to have access to the channel. The second grouping will have at least the same LPARs as the first grouping and it may specify additional LPARs also.

A.2.1.3 Control Unit (CNTLUNIT) Statement

This statement, along with the IODEVICE statement, defines the path from the host to the Network Utility. The CNTLUNIT and IODEVICE statements occur in pairs. If multiple LPARs are being defined to use a single CHPID, there must be a CNTLUNIT and IODEVICE statement for each LPAR.

- CUNUMBR=x

  This is an identifier for the control unit definition.

- PATH=x

  This number identifies the CHPID being used.

- UNIT=3172

  This identifies the type of control unit at the other end of the channel. The value is always 3172 when talking to a Network Utility. The IBM 3172 was the predecessor of the Network Utility ESCON channel function.

- CUADD=x

  This value identifies the control unit address of the Network Utility. The default is 0. For the Network Utility, each LPAR on a given CHPID must have a unique CUADD value. Usually (but not always) the CUADD value will be chosen to match the LPAR number.

- UNITADD=((*addr*,*number*))

  This defines the range of addresses reserved for this control unit, where:

  addr is the hex address of the first subchannel assigned to this control unit

  number is the decimal number of subchannels being assigned to this control unit

  The example above defines a maximum of 32 control unit addresses, or subchannels, starting from E0 hex and going upwards. The device addresses specified on the Network Utility LCS, LSA, or MPC+ interface definition must be from within this range. The Network Utility can use a maximum of 32 subchannels.

- LINK=xx

  The value for the LINK parameter should be set to the port of the ESCON Director (ESCD) that the Network Utility is attached to. Because the ESCD is a switch, you can think of the link parameter as the phone number that the host will use to reach the Network Utility through the switch.

A.2.1.4 IODEVICE Statement

This statement, along with the CNTLUNIT statement, identifies the Network Utility connection to the host.

- ADDRESS=(*addr*,*number*)

  This parameter identifies the range of addresses to the rest of the host, where:

  addr is the hex address being assigned to the first address reserved

  number is the decimal number of subchannels reserved

  This address is different from the UNITADD. It is used in the TCP/IP profile (for LCS), the VTAM XCA Major Node Definition (for LSA), and the VTAM TRL (for MPC+) to identify the subchannels being used.

- CUNUMBR=x

  This identifies the corresponding CNTLUNIT statement to this IODEVICE statement. While the value for this parameter has to be the same for both the CNTLUNIT and the IODEVICE macros, it does not have to relate to any other parameter. It is a good idea, however, to make it the same value specified in the ADDRESS parameter in the IODEVICE macro. The value for CUNUMBR has no significance outside the Channel Path Definition.

- UNIT=3172

  This identifies the type of device that is downstream. It should always be 3172 if the control unit is a Network Utility. The IOCP software in the host does not look at this field. If you are migrating from an IBM 3172 to the Network Utility, you might have a value of UNIT=SCTC in the existing IOCP statement. This is should be changed to 3172 for the Network Utility.

- PARTITION=(name)

  This is the device candidate list and it contains a list of one or more LPARs that have access to the device. This list is a subset of the list of LPARs specified in the CHPID statement and it is used to restrict which LPARs in the channel candidate list are allowed to use these devices. If the host is not partitioned, this field will not be present.

The following is an example of the IOCP statements for defining a Network Utility with a Parallel Channel Adapter (PCA):

```
CHPID                PATH=((05)),TYPE=BL
CNTLUNIT             CUNUMBR=640,PATH=05
                     PROTOCL=S4,UNIT=3172
                     SHARED=N,UNITADD=((40,32))
IODEVICE             UNIT=3172,ADDRESS=((640,32))
                     STADET=N,CUNUMBR=640,TIMEOUT=Y
```

Note the following points concerning the IOCP statements for a Network Utility with a PCA.

- The TYPE is BL for Block Multiplexer
- PROTOCL parameter can be set to the following values, depending on the device capability:

  D Direct-Coupled Interlock (DCI) mode

  S Maximum 3.0 Mbps data streaming speed

  S4 Maximum 4.5 Mbps data streaming speed

  For the Network Utility, set the value to S4. The transfer mode and channel parameter must conform with the PCA setting for transfer mode and channel transfer speed.

- The UNIT parameter on the CNTLUNIT and IODEVICE statements must be set to 3172.
- When an ESCON Converter is the channel path, the CHPID TYPE parameter must be set to CVC; otherwise it is set to BL.

A.3 Defining the Network Utility in the Operating System

The following sections describe the definitions needed for various host operating systems.

A.3.1 Network Utility Definition for VM/SP

You must define the Network Utility to a VM/SP operating system by updating the real I/O configuration file (DMKRIO) with entries for the Network Utility in the RDEVICE and the RCTLUNIT macros. In the following example, 640 is the base unit address and the size of the address range is 32.

```
RDEVICE ADDRESS=(640,32),DEVTYPE=3088
RCTLUNIT ADDRESS=640,CUTYPE=3088,FEATURE=32-DEVICE
```

A.3.2 Network Utility Definition for VM/XA and VM/ESA

You must define the Network Utility to a VM/Extended Architecture (VM/XA or VM/ESA) operating system by updating the real I/O configuration file (HCPRIO) with an entry for the Network Utility in the RDEVICE macro. In the following examples, 640 and 2A0 are base control unit addresses. The address range size, as defined in the UCW or IOCP, is 8 in both examples.

The following example is a VM/XA HCPRIO definition:

```
RDEVICE ADDRESS=(640,8),DEVTYPE=CTCA
```

The following example is a VM/ESA HCPRIO definition:

```
RDEVICE ADDRESS=(2A0,8),DEVTYPE=CTCA
```

A.3.3 Network Utility Definition for MVS/XA and MVS/ESA without HCD

You must define the Network Utility to an IBM Multiple Virtual Storage/Extended Architecture (MVS/XA) or MVS/ESA operating system by updating the MVS Control Program with an entry for the Network Utility in the IODEVICE macro.

For ESCON channels, an example IODEVICE macro is:

```
IODEVICE UNIT=3172,ADDRESS(540,8)
```

For parallel channels, an example IODEVICE macro is:

```
IODEVICE UNIT=CTC,ADDRESS(640,8)
```

The base control unit addresses are 640 and 540. The address range size, as defined in the UCW or IOCP, is 8 in both examples.

A.3.4 Network Utility Definition for MVS/ESA with HCD

The hardware configuration definition (HCD) component of MVS/ESA SP Version 4.2 and 4.3 with APAR #OY67361 offers an improved method of defining system hardware configuration for Network Utility. You can accomplish the several complex steps required for entering hardware configuration data by using an interactive dialog with HCD.

The required configuration data for the Network Utility is:

- When using HCD, with APAR #OY67361, define the Network Utility as (UNIT=3172). For example,

  ```
  IODEVICE UNIT=3172,ADDRESS(740,8)
  ```

- Without HCD, define the Network Utility for:
  - Parallel channels as a 3088 device (UNIT = 3088 or CTC)

    ```
    IODEVICE        UNIT=CTC,ADDRESS(840,8)
    ```

  - ESCON channels as a serial CTC device (UNIT = SCTC)

    ```
    IODEVICE        UNIT=SCTC,ADDRESS(A40,8)
    ```

Notes:

1. If you are using HCD for MVS Version 4 to define your ESCON host connection, you will need APAR # OY67361 to obtain the UIM support for the device definition (UNIT=3172).
2. When you are migrating your IOCP definition and operating system definitions to the HCD environment, it is important that you change all Network Utility device statements to device type (UNIT=3172).

A.3.5 Network Utility Definition for VSE/ESA

You must define the Network Utility to a VSE/ESA operating system by supplying an ADD statement for each channel unit address at initial program load (IPL) time. Code the device type on the ADD statement as CTCA,EML as shown in the following example:

```
ADD 640,CTCA,EML
```

The base control unit address is 640 in the example. For the number of channel unit addresses added, increment the IOTAB storage macro by this count.

A.4 VTAM Definitions

This section gives sample VTAM definitions for an XCA major node, an MPC+ local PU and Transport Resource List (TRL) major node, and an example of defining VTAM for APPN and DLUR support. It also shows an example of a switched major node for a PU in a TN3270 server. This section is not meant to be a complete reference on the subject. For more information on configuring VTAM, refer to the *OS/390 eNetwork Communications Server SNA Resource Definition Reference*, SC31-8565.

A.4.1 VTAM XCA Major Node Definition

When defining a channel gateway using LSA to VTAM, a definition for an External Communications Adapter (XCA) is required. This definition is the same as that used for an IBM 3172. An example is shown in Figure 33 on page 173.

```
*********************************************************************
RAINETU VBUILD TYPE=XCA      1

**
**
RANETUP  PORT  ADAPNO=0,     2                      * X
               CUADDR=285,     3                   * X
               MEDIUM=RING,       4                * X
               SAPADDR=4,            5             * X
               TIMER=60
**
*********************************************************************
RANETUG1 GROUP DIAL=YES,CALL=INOUT,DYNPU=YES
*
RANETUL1 LINE ANSWER=ON,ISTATUS=ACTIVE
RANETUP1 PU   ISTATUS=ACTIVE
RANETUL2 LINE ANSWER=ON,ISTATUS=ACTIVE
RANETUP2 PU   ISTATUS=ACTIVE
RANETUL3 LINE ANSWER=ON,ISTATUS=ACTIVE
RANETUP3 PU   ISTATUS=ACTIVE
RANETUL4 LINE ANSWER=ON,ISTATUS=ACTIVE
RANETUP4 PU   ISTATUS=ACTIVE
RANETUL5 LINE ANSWER=ON,ISTATUS=ACTIVE
RANETUP5 PU   ISTATUS=ACTIVE
RANETUL6 LINE ANSWER=ON,ISTATUS=ACTIVE
RANETUP6 PU   ISTATUS=ACTIVE
RANETUL7 LINE ANSWER=ON,ISTATUS=ACTIVE
RANETUP7 PU   ISTATUS=ACTIVE
RANETUL8 LINE ANSWER=ON,ISTATUS=ACTIVE
RANETUP8 PU   ISTATUS=ACTIVE
RANETUL9 LINE ANSWER=ON,ISTATUS=ACTIVE
RANETUP9 PU   ISTATUS=ACTIVE
```

Figure 33. XCA Major Node Definition Sample for LSA Direct Connection

Notes:

1 TYPE must be XCA.

2 ADAPNO is the LAN number for the Network Utility interface. This value is assigned to the Network Utility's LSA interface when it is created. The value can be obtained from the Network Utility by listing the configuration of the interface from the talk 6 menus or it can be retrieved by entering the `list nets` command from the ESCON console in Talk 5. Note that a wrong value for this parameter is the single most common error in LSA configuration.

3 CUADDR specifies the subchannel to be used to communicate with the Network Utility. This value must be within the range of values specified in the IODEVICE statement in the IOCP definition.

4 This specifies the physical LAN topology to which the LSA interface is attached. This corresponds to the value specified for LANtype for the Network Utility interface. Valid values are MEDIUM=RING for token-ring, MEDIUM=CSMACD for Ethernet, and MEDIUM=FDDI for a Fiber Distributed Data Interface (FDDI) network.

5 SAPADDR is the Service Access Point number VTAM wishes to open on the Network Utility. Note that it is the SOURCE SAP, not the DESTINATION SAP. When more than one active XCA major node refers to the same LAN, all the XCA major nodes have to use different SAPs.

A.4.1.1 LINE Statement

The CALL field can be one of the following:

- IN means only remote devices may establish connections.
- OUT means only VTAM can initiate connections.
- INOUT connections may be initiated at either end.

If VTAM is going to dial out, the Switched Major Node definition must specify a destination with a PATH statement.

An asterisk in the first column indicates a statement has been commented out, and should be ignored. A character in the last column indicates the next line is a continuation of this line.

A.4.2 VTAM Definitions for an MPC+ Connection

An MPC+ connection requires entries in two VTAM control blocks:

- The Local Major Node
- The Transport Resource List (TRL) Major Node

The following is a sample definition for a local SNA major node for a Network Utility MPC+ connection. This is the local PU that resides in VTAM that supports the channel connection defined in the TRL. The connection type must be APPN and you also need to enable HPR.

```
LOCNETU  VBUILD TYPE=LOCAL
MPCNETUP PU     TRLE=MPCNETU,
                XID=YES,
                CONNTYPE=APPN,
                CPCP=YES,
                HPR=YES
```

Notes:

1. TYPE must equal LOCAL on the VBUILD statement.
2. TRLE identifies the TRL being used. The name must match the name of an existing TRL.
3. XID indicates whether XIDs will be exchanged. It must be XID=YES.
4. CONNTYPE must be set to CONNTYPE=APPN since APPN is the only protocol that VTAM uses with an MPC+ connection.
5. CPCP specifies that CP-CP connections with APPN can be established over this MPC+ connection. This could be either set to YES or NO, depending upon your APPN topology.
6. HPR specifies that APPN HPR traffic can flow over this MPC+ connection. HPR is normally used by default, but setting this value to YES ensures it. This is important because an MPC+ connection requires RTP (and HPR).

Next, you need a transport resource list for the MPC+ connection from the Network Utility. An example definition is as follows:

```
        VBUILD TYPE=TRL
MPCNETU TRLE   LNCTL=MPC,
               MAXBFRU=9,
               READ=280,
               WRITE=281,
               MPCLEVEL=HPDT,
               REPLYTO=3.0
```

Notes:

1. TYPE must be TRL.
2. MPCNETU is the name that identifies the TRL. It must match what is specified in the TRLE= field in the local major node definition.
3. LNCTL identifies the connection type. It must be LCNTL=MPC.
4. MAXBFRU is the number of 4K pages per read subchannel.
5. READ/WRITE specifies the subchannels in the MPC+ group, and indicates their direction. The subchannel numbers must be in the range of addresses specified in the IODEVICE statement in the IOCP definition. There can be multiple READ and WRITE parameters in the TRLE statement but there must be at least one of each.

   Note: The designations READ and WRITE here are from the HOST perspective. In the Network Utility MPC+ definition, the designations are from the Network Utility perspective. Therefore, subchannels designated as READ on the host must be designated as WRITE on the Network Utility, and vice versa.
6. REPLYTO is the reply timeout value in seconds.

A.4.3 VTAM Definitions for APPN

If VTAM is configured for DLUS, then it must be an APPN network node. Configuring VTAM as an APPN network node requires certain parameters to be specified in the VTAM startup parameters. These are shown in Figure 34 on page 176. Set the CONNTYPE to APPN and the NODETYPE to a Network Node (NN).

```
ASYDE=TERM,IOPURGE=5M,
CONFIG=IO,
CONNTYPE=APPN,
CPCP=YES,
CSALIMIT=0,
DYNADJCP=YES,
ENCRYPTN=NO,
GWSSCP=YES,
HOSTPU=ISTPUS18,
HOSTSA=18,
HPR=RTP,
NETID=USIBMRA,
NODETYPE=NN,
NOTRACE,TYPE=VTAM,IOINT=0
PPOLOG=YES
SORDER=APPN,
SSCPDYN=YES,
SSCPID=18,
SSCPNAME=RAI,
SSCPORD=PRIORITY,
SUPP=NOSUP,
TNSTAT,CNSL,
VRTG=YES
OSITOPO=LLINES,
OSIMGMT=YES
XNETALS=YES
```

Figure 34. VTAM Start-up Parameters

A.4.4 VTAM Static Definition of TN3270 Resources

VTAM definitions are required for the PUs used by the TN3270E Server. You need a switched major node definition for each PU in the TN3270E server. For example, each PU in the TN3270E server can support up to 253 LUs. If you need 500 3270 sessions, then you will need two PUs in the router and two PU definitions in VTAM.

Figure 35 shows an example of a VTAM switched major node definition for a TN3270E server PU that is connected via DLUR and APPN.

```
LOCNETU  VBUILD TYPE=SWNET
MNETUA   PU     ADDR=01,ISTATUS=ACTIVE,VPACING=0,                      *
                DISCNT=NO,PUTYPE=2,SSCPFM=USSSCS,USSTAB=US327X,        *
                IDBLK=077,IDNUM=02216,IRETRY=YES,MAXDATA=521,          *
                MAXOUT=7,MAXPATH=8,PASSLIM=7,PACING=0,ANS=CONTINUE
*********************************************************************
PNETUA   PATH   PID=1,DLCADDR=(1,C,INTPU),DLCADDR=(2,X,07702216),      *
                DLURNAME=MNETUA
*********************************************************************
JC7LU2   LU     LOCADDR=2
JC7LU3   LU     LOCADDR=3
JC7LU4   LU     LOCADDR=4
```

Figure 35. VTAM Definitions for a TN3270E Server PU (DLUR/APPN)

Figure 36 on page 177 shows an example of a VTAM switched major node definition for a TN3270E server PU that uses a subarea connection to the host.

```
LSAPO8T VBUILD TYPE=SWNET
PUPSO8T PU ADDR=01,IDBLK=077,IDNUM=12244,MAXOUT=7,PACING=0,VPACING=0,
               DLOGMOD=B22NNE,PUTYPE=ANY,
               SSCPFM=USSSCS,MAXDATA=2000,MODETAB=LMT3270
PTO8LU2 LU LOCADDR=02,LOGAPPL=TSO
PTO8LU3 LU LOCADDR=03,LOGAPPL=TSO
PTO8LU4 LU LOCADDR=04,LOGAPPL=TSO
PTO8LU5 LU LOCADDR=05,LOGAPPL=TSO
PTO8LU6 LU LOCADDR=06,LOGAPPL=TSO
```

Figure 36. VTAM Definitions for a TN3270E Server PU (Subarea)

The following sections provide an overview of the statements in the Switched Major Node Definition.

A.4.4.1 VBUILD Statement

The TYPE field must be TYPE=SWNET.

A.4.4.2 PU Statement

This statement defines the type of data flow and the destination. The pertinent parameters are:

- ADDR is an identifier.
- MAXDATA is the maximum packet size VTAM will support over this interface. This value will be negotiated down with the Network Utility during the XID exchange.
- IDBLK/IDNUM identify the remote device when VTAM is communicating with PU 2.0 (dependent) devices.

A.4.4.3 LU Statement

These statements define the logical units (LUs) that can be contacted through this PU. The name on the left of each statement is the name that the host uses to address each LU. The LOCADDR is used by the Network Utility to identify the correct LU in VTAM.

A.4.4.4 PATH Statement

If VTAM is going to dial out, the Switched Major Node definition must specify a destination with a PATH statement. The path statement will be different depending on whether the TN3270E server attaches via a Subarea or a DLUR/APPN connection.

For a subarea connection, the format is:

```
PATH DIALNO=xxyyzzzzzzzzzzzz
```

where:

- xx is a place holder
- yy is the destination SAP number
- zz is the destination MAC address

The example in Figure 36 does not have a PATH statement because in this example, the downstream PU will contact VTAM instead of VTAM dialing out to the device.

The example in Figure 35 on page 176 shows a PATH statement for a TN3270E server PU that is using DLUR to connect to the host. Here, the PATH statement identifies the CP name of the Network Utility (MNETUA) via the DLURNAME parameter. This is needed in order for the LU6.2 conversation between the DLUR and DLUS to be established. Once this session has been established, the SSCP-PU session between VTAM and the TN3270E server PU will be established using the IDBLK/IDNUM value specified by DLCADDR=(2,X,07702216).

A.4.5 VTAM Dynamic Definition of TN3270 Resources

VTAM contains support for several facilities that reduce the amount of user coding required to define its resources such as PUs and LUs. When implemented in Network Utility, TN3270 PUs and LUs appear as switched resources to the VTAM host and require corresponding definitions in VTAM. When large TN3270E environments are being implemented, the definition of these resources could be a very labor intensive task.

VTAM provides a facility that allows switched resources to be defined dynamically. TN3270E can take advantage of this facility to reduce the amount of VTAM definitions the user has to provide. This facility is called *VTAM Dynamic Dial-In* support. This function should not be confused with a similar VTAM function call Dynamic Definition of Dependent LUs (DDDLUs), which requires corequisite function to be present in the TN3270 server. Network Utility does not currently have this corequisite function.

The details for Dynamic Dial-In Support are in the *VTAM Customization Manual* for the release level of VTAM that is installed on your VTAM host. A brief description of this function and its potential use in a TN3270E environment follows.

A.4.5.1 General Overview

Dynamic Dial-In Support makes use of a VTAM exit called the Configuration Services Exit (ISTEXCCS), and a set of model PU/LU definitions that you must define. Each time VTAM receives a connection request from a PU that is not defined to VTAM, the Configuration Services exit is driven and a set of PU and LU definitions are dynamically generated based on the model definitions. These definitions are associated with the PU requesting connection. This set of definitions may be tailored down to the specific PU level, with matching based on information contained in the dial-in PU's XID. This process is repeated each time a connection request is received from a PU that is not defined to VTAM.

Figure 37 on page 179 contains a VTAM definition for a set of Model definitions that could be used to implement Dynamic Dial-In support. Notice that the definition is created in a VTAM member with a VBUILD TYPE= MODEL. This example contains two PU models and two models for LUs. They are the prototypes from which VTAM will generate its dynamic resource definitions. From a practical point of view, if all of the PUs and LUs configured for TN3270E have similar characteristics such as logmode and pacing values, then a single PU and LU definition in the model definition would be sufficient.

The VTAM Configuration Services exit mentioned above could be used to select the appropriate model definition based on XID parameters such as CPNAME and

IDNUM/IDBLK. The corresponding values stored in the VTAM data sets CPNDEF and NIDDEF indicate which model should be used. If these data sets are not defined, the exit has an internal algorithm to select the model and generate LU resource names. This exit routine can be used as is or modified to fit user needs. See “Resource Name Generation” on page 180 for a description of the default name generation algorithm and how you can control what resource names are generated.

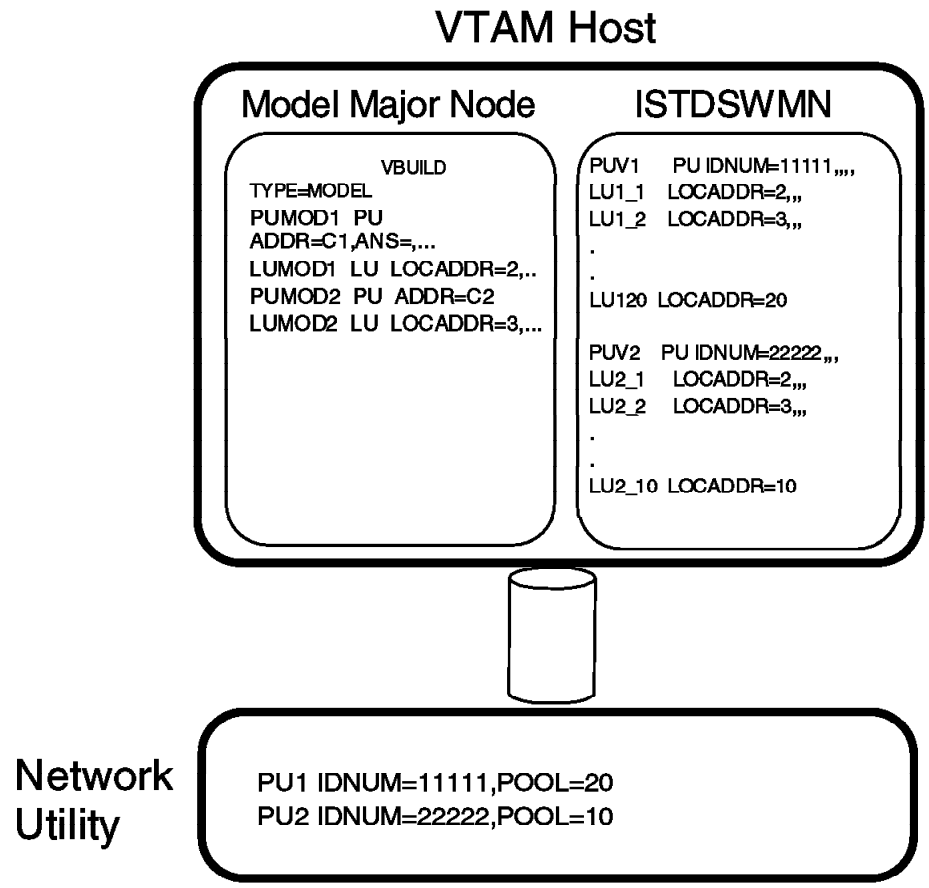

Figure 37. VTAM Dynamic Resource Definition

In Figure 37, the Network Utility has two PUs defined. PU1 has 20 pooled LU resources defined and PU2 has 10 pooled LU resources defined. In the VTAM host a model major node that contains two PU and two LUs is defined. These model definitions will contain all of the normal PU and LU parameters such as polling address, logmode tables, MAXDATA values, and PUTYPE that would normally be defined for specific PUs and LUs. As mentioned earlier, if all of the characteristics of the connecting devices are similar a single PU and LU definition in the model would be sufficient.

In this example, VTAM would dynamically generate the definitions shown in ISTDSWMN, a dynamically created major node. ISTDSWMN contains definitions for a PUV1 and 20 LUs (LU1_1 through LU1_20) and for a PUV2 with 10 LUs (LU2_1 through LU2_10). Notice that the PU names in the Network Utility

configuration do not have to match the names that VTAM generates. The PUs in the Network Utility are correlated to the VTAM PUs based on matching IDNUM values. Notice also that the only definitions required to implement the dynamic definitions are done on the VTAM host. The Network Utility does not know that VTAM is performing the dynamic definition process.

A.4.5.2 Dynamic Dial-In Exit Overview

The Configuration Services Installation EXIT (ISTEXCCS) is shipped with VTAM and can be found in the data set SYS1.SAMPLIB. This exit is described in an appendix of the *VTAM Customization Manual* for the release of VTAM installed on your host.

If the exit is installed in the VTAMLIB data set at VTAM initialization time, it will be loaded and started during VTAM initialization. VTAM calls the exit whenever it receives a REQCONT RU from a real switched PU, or a REQACTPU from a DLUR node representing a PU. These RUs are generated by the receipt of an XID from the connecting device. They contain XID information such as NETID, CPNAME and IDNUM/IDBLK. The exit uses this information to construct a *build* vector. The build vector contains the model names for PU and LU definitions contained in an active Model major node, and the names to be used for the PU and LU definitions that VTAM is to create. VTAM then builds the definitions for the connecting PU in the dynamic Switched Major Node ISTDSWMN. From this point on, VTAM treats these resources just as if they had been predefined by you.

A.4.5.3 Implementing Dynamic Definitions

Resource Name Generation: There are two ways for you to influence what resource names it will associate with the resources that VTAM dynamically generates. The first permits complete control by allowing you to supply specific LU and PU names for the resources definitions that VTAM generates. This support is implemented by coding the appropriate information in the VTAM data sets CPNDEF and/or NIDDEF. These data sets provide the information that the Configuration Services exit uses to construct the Build vector that VTAM uses to generate the dynamic resources. While this approach requires more user VTAM definitions, it is much less definition intensive than manually defining the required VTAM definitions.

At the other extreme, you can allow VTAM to generate the resource names based on an internal algorithm. Implementations not requiring unique naming conventions would be candidates for this approach, which requires minimal VTAM definitions.

The name generated for the PU is of the form cnnnnnss, where:

c can be user-specified in a name definition table

nnnnn is the IDNUM extracted from the received XID

ss is the station address (if specified), otherwise it consists of two blank characters

The name generated for the LU is of the form cnnnnnll where:

c is a user-specified name prefix

nnnnn is the IDNUM extracted from the received XID

ll is the local address of the LU

The details for both of these approaches to resource name generation can be found in the *VTAM Customization Manual*.

Operational Considerations: The dynamically created definitions in the dynamic Switched Major Node ISTDSWMN are kept as long as there are active LU sessions on the LU. If you specify DISCNT=YES on the Model PU definition, all of the dynamic resources associated with the PU will be deleted from ISTDSWMN when all of the sessions on the PU have ended. If you specify DISCNT=NO, these definitions will not be deleted as long as the PU remains active to VTAM.

Security Considerations: When the Configuration Services Exit dynamically defines resources, VTAM has no predefined IDNUM/IDBLK or CPNAME to validate the identity of connecting devices. If this is considered a security exposure, the exit can be modified to consult a list of acceptable IDNUM/IDBLKs or CPNAMEs and compare these values against those contained in the XID from the connecting device. Note that the XID from the connecting device is passed to the exit.

Sources for Additional Information: The details for this facility and additional information can be found in the following VTAM library documents:

- *VTAM Customization Manual*
- *VTAM Resource Definition Reference*
- *VTAM Network Implementation Guide*

You should use the manuals corresponding to the release level of the VTAM installed on your VTAM host.

A.5 Host IP Definitions

Defining the Network Utility to the host for a TCP/IP connection requires you to make changes to the host TCP/IP profile. This section gives an overview of the relevant statements that need to be changed.

A.5.1 DEVICE Statement

This statement defines the subchannel pair being used by TCP/IP. The format is:

DEVICE *name LCS subchannel*

where:

- *name* identifies the subchannel path being used. It has local significance only, and can be anything.
- *subchannel* identifies the even subchannel being used for this connection. This value comes from the IODEVICE statement in the IOCP definition. When specified, that subchannel and the next one are both being used.

A TCP/IP profile must contain one DEVICE statement for each subchannel pair being used.

A.5.2 LINK Statement

This statement identifies which LCS interfaces on the Network Utility are being used on a given subchannel pair. The format is:

LINK *name lantype lannumber devicename*

where:

- *name* identifies the LCS interface. It has local significance only, and can be anything.
- *lantype* identifies the type of LAN interface that the Network Utility LCS interface is emulating. The allowable values are:
  - IBMTR for token-ring
  - ETHERNET for Ethernet V2
  - 802.3 for Ethernet (IEEE 802.3)
  - ETHEROr802.3 for either Ethernet format accepted
  - FDDI for FDDI
- *lannumber* identifies which LCS interface on the Network Utility is being used. The *lannumber* is generated sequentially for each *lantype* on the Network Utility when you add an LCS interface. The *lannumber* can be found by entering `list nets` from the ESCON console in Talk 5. Note that the *lannumber* is not the net number. Having the wrong *lannumber* is the single most common configuration error for an LCS interface.
- *devicename* correlates the LCS interface to a subchannel pair. It must match a previously defined DEVICE statement.

There can be multiple LINK statements associated with a single DEVICE statement. There must be an LCS interface on the Network Utility for each LINK statement.

A.5.3 HOME Statement

This statement specifies the IP address(es) of the host TCP/IP stack. The format is:

```
HOME     ipaddress1     link1
         ipaddress2     link2
```

where:

- *IpaddressX* specifies an IP address on the host.
- *LinkX* specifies which link is associated with this IP address.

There must be only one HOME address for each LINK statement. The HOME address must be in the same IP subnet as the IP address of the LCS interface in the Network Utility, but they must be different addresses.

A.5.4 GATEWAY Statement

This statement identifies the IP routing information for the host. It is divided into three sections:

- Direct Routes are routes directly connected to the host. The subnet containing the Network Utility LCS interface is a direct route.
- Indirect Routes are routes that are accessible via routers. The subnets of the LANs on the Network Utility are indirect routes, for example.

- Default Route is the route to be used if the host doesn't have a direct or indirect route to an IP address.

A.5.4.1 Direct Routes

The format for Direct Routes is:

network firsthop linkname pktsize submask subvalue

where:

- *network* is the non-subnetted part of the IP address.
- *firsthop* indicates the IP address of the next hop in the IP network. For Direct Routes, this should be an equal sign (=).
- *linkname* identifies which link the host should use to get to the addresses on this route. For routes accessible via the Network Utility, this should be the name from the LINK statement associated with the LCS interface on this subnet.
- *pktsize* is the maximum frame size to be used on the interface. It should be less than or equal to the packet size defined in the LCS configuration on the Network Utility. A value of DEFAULTSIZE indicates the default packet size will be used.
- *submask* specifies the subnet mask used on this link. The subnet mask should correspond to the subnet mask defined for the LCS interface in the IP configuration on the Network Utility. This field may also be set to HOST to identify a point-to-point connection. In this case, the network field should contain the full IP address of the LCS interface.
- *subvalue* specifies the subnetted part of the IP address, and together with the network field, should fully specify the IP subnet associated with this LCS interface.

A.5.4.2 Indirect Routes

The format for Indirect Routes is:

network firsthop linkname pktsize submask subvalue

where:

- *network* is the full address of the IP subnet.
- *firsthop* indicates the IP address of the next hop in the IP network. For Indirect Routes accessible via the Network Utility, this should be the IP address of the Network Utility LCS interface.
- *linkname* identifies which link the host should use to get to the addresses on this route. For routes accessible via the Network Utility, this should be the name from the LINK statement associated with the LCS interface on this subnet.
- *pktsize* is the same value as for Direct Routes.
- *submask* should either be 0 or blank if the network field contains the full subnet address.
- *subvalue* should be left blank if there is no subnet mask specified.

A.5.4.3 Default Routes

The format for Default Routes is:

network firsthop linkname pktsize submask subvalue

where:

- *network* should be DEFAULTNET.
- *firsthop* indicates the IP address of the next hop in the IP network. For Default Routes to the Network Utility, this should be the IP address of the Network Utility LCS interface.
- *linkname* identifies which link the host should use to get to the addresses on this route. For routes accessible via the Network Utility, this should be the name from the LINK statement associated with the LCS interface on this subnet.
- *pktsize* is the same value as for Direct Routes.
- *submask* should either be 0 or blank.
- *subvalue* should be blank.

A.5.4.4 START Statement

This statement causes the specified subchannels to be started. The format is:

START *devicename*

where *devicename* is the name on the DEVICE statement above.

There must be a START statement for every DEVICE statement if the customer wishes to activate the devices when TCP/IP is started. If the START statement is not here, the devices can be started using the OBEY file. Note that the name here is the one from the DEVICE statement, not the LINK statement. Note also that the Network Utility LCS interface will remain in the DOWN state until the START has been issued from TCP/IP.

A.5.5 Host TCP/IP Definitions for LCS

This section gives you examples of the above statements required if you are defining an LCS connection.

1. DEVICE statement:

   DEVICE LCS1 LCS 210

   where LCS1 is the device name being defined, LCS is the type of device, and 210 is the host read (Network Utility write) subchannel used for this definition.

2. LINK statement

   LINK ETHLCS1 802.3 0 LCS1

   where ETHLCS1 is the link name, 802.3 is the LAN type to which the LCS interface attaches on the Network Utility, 0 is the LAN number assigned by the Network Utility, and LCS1 is the name of the device (from the device statement above).

   Note: Remember that the LAN number is automatically assigned by the Network Utility when you define the LCS interface. You can obtain it by issuing a `list all` command from the ESCON Config> prompt in the talk 6 process on the Network Utility console.

3. HOME command

```
HOME 9.24.106.72 ETHLCS1
```

where 9.24.106.72 is the IP address of this LCS interface and ETHLCS1 is the name of the link.

4. GATEWAY command

```
GATEWAY 9.24.106   9.24.106.1   ETHLCS1   4096  0
```

where 9.24.106 is the IP address for the network, 9.24.106.1 is the IP address of the default router, ETHLCS1 is the link name defined by the LINK statement above, 4096 is the MTU size, 0 is the subnet mask, and the subnet value has been left blank.

5. Activate the TCP/IP profile

To activate the device defined in 1 on page 184, issue the following command:

```
start lcs1
```

A.5.6 Host TCP/IP Definitions for MPC+

The steps for configuring TCP/IP in the host for an MPC+ connection are the same as for an LCS connection. However, the command syntax for the device and link commands is slightly different. For an MPC+ connection, the syntax for the device command is:

```
DEVICE IPTRL1 MPCPTP
```

where IPTRL1 is the name of the TRL that this connection will use and MPCPTP specifies an MPC point-to-point link.

To define the link, the syntax is:

```
LINK LINK1 MPCPTP IPTRL1
```

where LINK1 is the link name and the other two parameters are the same as those used in the device statement.

Appendix B. Supported RFCs and Other Standards

Nways Multiprotocol Access Services is based on a large set of standards. These standards, draft standards, and proposed standards are issued by the Internet Activities Board and distributed by the Network Information Center. They are referred to by RFC (Request for Comments) numbers:

The following list represents the RFCs supported by Nways Multiprotocol Access Services. This list is subject to future update/correction as additional information is available.

- RFC 768 User Datagram Protocol (UDP)
- RFC 783 Trivial File Transfer Protocol
- RFC 791 Internet Protocol (IP)
- RFC 792 Internet Control Message Protocol (ICMP)
- RFC 793 Transmission Control Protocol (TCP)
- RFC 826 Ethernet Address Resolution Protocol (ARP)
- RFC 854 TELNET Protocol
- RFC 877 IP over DDN X.25 (obsoleted by RFC 1356)
- RFC 888 ″STUB″ Exterior Gateway Protocol
- RFC 894 Transmission of IP Datagrams over Ethernet
- RFC 919 Broadcasting Internet Datagrams
- RFC 922 Broadcasting Internet Datagrams in the Presence of Subnets
- RFC 925 Multi-LAN Address Resolution
- RFC 950 Internet Standard Subnetting Procedure
- RFC 951 Bootstrap Protocol (BOOTP)
- RFC 1009 Requirements for Internet Gateways
- RFC 1027 Using ARP to Implement Transparent Subnet Gateways (Proxy-ARP)
- RFC 1042 Transmission of IP Datagrams over IEEE 802 Networks
- RFC 1058 Routing Information Protocol (RIP)
- RFC 1075 Distance Vector Multicast Routing Protocol (DVMRP)
- RFC 1112 Host Extensions for IP Multicasting
- RFC 1122 Requirements for Internet Hosts - Communications Layers (missing some support but being addressed)
- RFC 1123 Requirements for Internet Hosts
- RFC 1144 Compressing TCP/IP headers for Low-Speed Serial Links
- RFC 1155 Structure and Identification of MIBs
- RFC 1156 MIB-I (obsoleted by RFC 1213)
- RFC 1157 SNMP
- RFC 1171 The Point-to-Point Protocol (PPP) (obsoleted by RFC 1331 and 1548)
- RFC 1172 PPP Initial Configuration Options (obsoleted by RFC 1332)

- RFC 1191 Path MTU Discovery
- RFC 1212 Concise MIB Definitions
- RFC 1213 MIB-II
- RFC 1220 PPP Extensions for Bridging
- RFC 1231 Token-Ring MIB (obsoleted by RFC 1748)
- RFC 1236 IP to X.121 Address Mapping for DDN
- RFC 1247 OSPF Version 2—TOS not supported (obsoleted by RFC 1583)
- RFC 1253 OSPF Version 2 MIB
- RFC 1256 ICMP Router Discovery Messages
- RFC 1284 Ethernet MIB (obsoleted by RFC 1398)
- RFC 1286 Bridge MIB (obsoleted by RFC 1493, 1525)
- RFC 1293 Inverse Address Resolution Protocol (InARP)
- RFC 1294 Multiprotocol Interconnect over Frame Relay (obsoleted by RFC 1490)
- RFC 1315 MIB for Frame Relay DTEs
- RFC 1317 MIB for RS-232-like Hardware Devices
- RFC 1321 The MD5 Message-Digest Algorithm
- RFC 1331 Point-to-Point Protocol (PPP) (obsoleted by RFC 1548)
- RFC 1332 PPP Internet Protocol Control Protocol (IPCP)
- RFC 1334 PPP Authentication Protocols (obsoleted by RFC 1994)
- RFC 1338 Supernetting (CIDR) (obsoleted by RFC 1519)
- RFC 1350 TFTP Version 2
- RFC 1356 Multiprotocol Interconnect over X.25 and ISDN
- RFC 1390 Transmission of IP and ARP over FDDI Networks
- RFC 1398 Ethernet MIB (obsoleted by RFC 1623)
- RFC 1434 DLSw
- RFC 1471 PPP Link Control Protocol MIB
- RFC 1483 Multiprotocol Encapsulation over ATM Adaptation Layer 5
- RFC 1490 Multiprotocol Interconnect over Frame Relay (obsoletes RFC 1294)
- RFC 1492 Access Control Protocol
- RFC 1493 Definitions of Managed Objects for Bridges (obsoletes RFC 1286)
- RFC 1512 FDDI Management Information Base
- RFC 1519 Classless Inter-Domain Routing (CIDR)
- RFC 1525 Definitions of Managed Objects for Source Routing Bridges (obsoletes RFC 1286)
- RFC 1541 Dynamic Host Configuration Protocol (DHCP) (same relay-agent support as for BootP)
- RFC 1542 Clarifications and Extensions for the Bootstrap Protocol (obsoletes RFC 1532)

- RFC 1548 The Point-to-Pont Protocol
- RFC 1549 PPP in HDLC Framing
- RFC 1573 Evolution of Interfaces Group of MIB-II
- RFC 1576 TN3270 Current Practices
- RFC 1577 Classical IP and ARP over ATM
- RFC 1583 OSPF Version 2—TOS not supported (obsoletes RFC 1247)
- RFC 1584 Multicast Extensions to OSPF (MOSPF)
- RFC 1593 SNA APPN Node MIB
- RFC 1623 Ethernet MIB (obsoleted by RFC 1650)
- RFC 1626 Default IP MTU for use over ATM AAL5
- RFC 1638 Bridging Control Protocol (BCP)
- RFC 1646 TN3270 Extensions for LU Name and Printer Selection
- RFC 1647 TN3270 Enhancements
- RFC 1650 Definitions of Managed Objects for the Ethernet-like Interface Types
- RFC 1654 Border Gateway Protocol (BGP) 4
- RFC 1657 Definitions of Managed Objects for BGP-4
- RFC 1661 The Point-to-Point Protocol (PPP) (Obsoletes RFC 1548)
- RFC 1662 PPP in HDLC-like Framing (Obsoletes RFC 1549) (LQM option not supported)
- RFC 1695 Definitions of Managed Objects for ATM Management Version 8.0
- RFC 1723 RIP Version 2
- RFC 1747 Definitions of Managed Objects for SNA Data Link Control: SDLC
- RFC 1748 IEEE 802.5 MIB
- RFC 1755 ATM Signalling Support for IP over ATM
- RFC 1771 A Border Gateway Protocol 4 (BGP-4)
- RFC 1793 Extending OSPF to Support Demand Circuit
- RFC 1795 AIW Version 1 Data Link Switching (DLSw)
- RFC 1812 Requirements for IP Version 4 Routers (obsoletes RFC 1716)
- RFC 1858 Security Considerations for IP fragment filtering
- RFC 1883 Internet Protocol, Version 6 (IPv6) Specification
- RFC 1884 IP Version 6 (IPv6) Addressing Architecture
- RFC 1885 Internet Control Message Protocol (ICMPv6) Version 6 (IPv6) Specification
- RFC 1933 Transition Mechanisms for IPv6 Hosts and Routers
- RFC 1962 The PPP Compression Control Protocol (CCP)
- RFC 1970 Neighbor Discovery for IP Version 6 (IPv6)
- RFC 1972 A Method for the Transmission of IPv6 Packets over Ethernet Networks

- RFC 1974 PPP Stac LZS Compression Protocol
- RFC 1981 Path MTU Discovery for IP version 6
- RFC 1994 CHAP - PPP Challenge Handshake Authentication Protocol (obsoletes RFC 1334)
- RFC 2019 A Method for the Transmission of IPv6 Packets over FDDI Networks
- RFC 2024 Definitions of Managed Objects for Data Link Switching
- RFC 2043 PPP SNA Control Protocol (SNACP)
- RFC 2051 Definitions of Managed Objects for APPC
- RFC 2058 Remote Authentication Dial In User Service (RADIUS)
- RFC 2063 IP Version 6 over PPP
- RFC 2080 RIPng for IPv6
- RFC 2097 The PPP NetBIOS Frames Control Protocol (NBFCP)
- RFC 2118 Microsoft Point-To-Point Compression (MPPC) Protocol
- RFC 2138 Remote Authentication Dial In User Service (RADIUS)
- RFC 2139 RADIUS Accounting
- RFC 2155 Definitions of Managed Objects for APPN
- RFC 2166 AIW DLSw Version 2 Standard
- RFC 2178 OSPF Version 2 (obsoletes RFC 1583)
- RFC 2205 Resource ReSerVation Protocol (RSVP)—V 1 Functional Specification
- RFC 2206 RSVP Management Information Base (MIB) using SMIv2
- RFC 2208 RSVP Version 1 Applicability Statement - Some Guidelines on Deployment
- RFC 2210 The Use of RSVP with IETF Integrated Services
- RFC 2211 Specification of the Controlled-Load Network Element Service
- RFC 2213 Integrated Services Management Information Base (MIB) using SMIv2
- RFC 2215 General Characterization Parameters for Integrated Service Network Elements
- RFC 2232 Definitions of Managed Objects for DLUR using SMIv2
- RFC 2238 Definitions of Managed Objects for HPR using SMIv2
- RFC 2320 Managed Objects for Classical IP and ARP Over ATM Using SMIv2 (IPOA-MIB)
- RFC 2338 VRRP - Virtual Router Redundancy Protocol

The bridging is also based on:

- ISO 10038 - ANSI/IEEE Std 802.1D: Media access control (MAC) bridges

The ATM support for bridging, routing and LAN Emulation is also based on:

- LAN Emulation Over ATM: Version 1.0 Specification, ATM Forum
- User-Network Interface Specification - Version 3.0, ATM Forum

- User-Network Interface Specification - Version 3.1, ATM Forum
- Q.2110 (Service-specific connection-oriented protocol), ITU-T
- Q.2130 (Service-specific coordination function), ITU-T
- Q.2931 (Signalling messages), ITU-T
- I.363 (AAL Type 5 Common Part Protocol), ITU-T
- ATM_Forum/94-0737R4, ″LAN Emulation Client Management: Version 1.0 Specification″, May, 1995.

Appendix C. Special Notices

This publication is intended to help networking professionals quickly understand the functions of the Network Utility Models TN1 and TX1. The information in this publication is not intended as the specification of any programming interfaces that are provided by the Network Utility Models TN1 or TX1. See the PUBLICATIONS section of the IBM Hardware Announcement for the Network Utility 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, North Castle Drive, Armonk, NY 10504-1785.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

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

Any pointers in this publication to external Web sites are provided for convenience only and do not in any manner serve as an endorsement of these Web sites.

Any performance data contained in this document was determined in a controlled environment, and therefore, the results that may be obtained in other

operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.

The following terms are trademarks of the International Business Machines Corporation in the United States and/or other countries:

ACF/VTAM
APPN
DB2
ESCON
IBM
MVS/XA
Nways
OS/390
S/370
SP
VSE/ESA
400
AIX
AS/400
eNetwork
EtherJet
MVS/ESA
NetView
OS/2
Parallel Sysplex
S/390
VM/ESA
VTAM

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc.

Java and HotJava are trademarks of Sun Microsystems, Incorporated.

Microsoft, Windows, Windows NT, and the Windows 95 logo are trademarks or registered trademarks of Microsoft Corporation.

PC Direct is a trademark of Ziff Communications Company and is used by IBM Corporation under license.

Pentium, MMX, ProShare, LANDesk, and ActionMedia are trademarks or registered trademarks of Intel Corporation in the U.S. and other countries.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Other company, product, and service names may be trademarks or service marks of others.

Appendix D. Related Publications

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

D.1 International Technical Support Organization Publications

For information on ordering these ITSO publications see “How to Get ITSO Redbooks” on page 197.

- *IBM 2216 Nways Multiaccess Connector Description and Configuration Scenarios - Volume I*, SG24-4957
- *IBM 2210 Nways Multiprotocol Router IBM 2216 Nways Multiaccess Connector Description and Configuration Scenarios - Volume II*, SG24-4956
- *IBM 2210 Nways Multiprotocol Router Description and Configuration Scenarios - Volume I*, SG24-4446
- *IBM 2216 Multiaccess Connector ESCON Solutions*, SG24-2137
- *Inside APPN - The Essential Guide to the Next-Generation SNA*, SG24-3669
- *TCP/IP Tutorial and Technical Overview*, GG24-3376

D.2 Redbooks on CD-ROMs

Redbooks are also available on CD-ROMs. **Order a subscription** and receive updates 2-4 times a year at significant savings.

CD-ROM Title	Subscription Number	Collection Kit Number
System/390 Redbooks Collection	SBOF-7201	SK2T-2177
Networking and Systems Management Redbooks Collection	SBOF-7370	SK2T-6022
Transaction Processing and Data Management Redbook	SBOF-7240	SK2T-8038
Lotus Redbooks Collection	SBOF-6899	SK2T-8039
Tivoli Redbooks Collection	SBOF-6898	SK2T-8044
AS/400 Redbooks Collection	SBOF-7270	SK2T-2849
RS/6000 Redbooks Collection (HTML, BkMgr)	SBOF-7230	SK2T-8040
RS/6000 Redbooks Collection (PostScript)	SBOF-7205	SK2T-8041
RS/6000 Redbooks Collection (PDF Format)	SBOF-8700	SK2T-8043
Application Development Redbooks Collection	SBOF-7290	SK2T-8037

D.3 Other Publications

These publications are also relevant as further information sources:

- *Network Utility Models TN1 and TX1 Installation and Initial Configuration Guide*, GA27-4167
- *IBM 2216 Nways Multiaccess Connector and Network Utility Introduction and Planning Guide*, GA27-4105
- *IBM 2216 Nways Multiaccess Connector and Network Utility Service and Maintenance Manual*, SY27-0350
- *IBM 2216 Nways Multiaccess Connector Installation and Initial Configuration Guide*, GA27-4106
- *Nways Multiprotocol Access Services Software User's Guide Version 3.2*, SC30-3886

- *Nways Multiprotocol Access Services Using and Configuring Features Version 3.2*, SC30-3993
- *Nways Multiprotocol Access Services Protocol Configuration and Monitoring Reference Vol. 1 Version 3.2*, SC30-3884
- *Nways Multiprotocol Access Services Protocol Configuration and Monitoring Reference Vol. 2 Version 3.2*, SC30-3885
- *Nways Event Logging System Messages Guide*, SC30-3682-10 (available at http://www.networking.ibm.com)
- *AIS/MAS/MRS/MSS/MSSC Configuration Program User's Guide for Nways Multiprotocol Access, Routing and Switched Services*, GC30-3830
- *VTAM Network Implementation Guide V4R4 for MVS/ESA*, SC31-8370
- *VTAM Resource Definition Reference V4R4 for MVS/ESA*, SC31-8377
- *Fiber Optic Link Planning*, GA23-0367
- *SNA APPN Architecture Reference*, SC30-3422
- *OS/390 eNetwork Communications Server SNA Resource Definition Reference*, SC31-8565

How to Get ITSO Redbooks

This section explains how both customers and IBM employees can find out about ITSO redbooks, redpieces, and CD-ROMs. A form for ordering books and CD-ROMs by fax or e-mail is also provided.

- **Redbooks Web Site** http://www.redbooks.ibm.com/

  Search for, view, download or order hardcopy/CD-ROMs redbooks from the redbooks Web site. Also read redpieces and download additional materials (code samples or diskette/CD-ROM images) from this redbooks site.

  Redpieces are redbooks in progress; not all redbooks become redpieces and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

- **E-mail Orders**

  Send orders via e-mail including information from the redbook order form to:

	IBMMAIL	**Internet**
In United States:	usib6fpl at ibmmail	usib6fpl@ibmmail.com
In Canada:	caibmbkz at ibmmail	lmannix@vnet.ibm.com
Outside North America:	dkibmbsh at ibmmail	bookshop@dk.ibm.com

- **Telephone Orders**

United States (toll free)	1-800-879-2755	
Canada (toll free)	1-800-IBM-4YOU	
Outside North America	(long distance charges apply)	
(+45) 4810-1320 - Danish	(+45) 4810-1220 - French	(+45) 4810-1270 - Norwegian
(+45) 4810-1420 - Dutch	(+45) 4810-1020 - German	(+45) 4810-1120 - Spanish
(+45) 4810-1540 - English	(+45) 4810-1620 - Italian	(+45) 4810-1170 - Swedish
(+45) 4810-1670 - Finnish		

This information was current at the time of publication, but is continually subject to change. The latest information for customers may be found at http://www.redbooks.ibm.com/ and for IBM employees at http://w3.itso.ibm.com/.

IBM Intranet for Employees

IBM employees may register for information on workshops, residencies, and redbooks by accessing the IBM Intranet Web site at http://w3.itso.ibm.com/ and clicking the ITSO Mailing List button. Look in the Materials repository for workshops, presentations, papers, and Web pages developed and written by the ITSO technical professionals; click the Additional Materials button. Employees may also view redbook, residency and workshop announcements at http://inews.ibm.com/.

IBM Redbook Fax Order Form

Fax your redbook orders to:

United States (toll free)	1-800-445-9269
Canada	1-403-267-4455
Outside North America	(+45) 48 14 2207 (long distance charge)

Please send me the following:

Title	Order Number	Quantity

First name Last name

Company

Address

City Postal code Country

Telephone number Telefax number VAT number

- Invoice to customer number
- Credit card number

Credit card expiration date Card issued to Signature

We accept American Express, Diners, Eurocard, Master Card, and Visa. Payment by credit card not available in all countries. Signature mandatory for credit card payment.

Index

Numerics

A

B

C

D

E

M

N

O

P

Q

R

S

T

U

V

W

X

Y

ITSO Redbook Evaluation

IBM Network Utility Description and Configuration Scenarios
SG24-5289-00

Your feedback is very important to help us maintain the quality of ITSO redbooks. **Please complete this questionnaire and fax it to: USA International Access Code + 1 914 432 8264 or:**

- Use the online evaluation form found at http://www.redbooks.ibm.com
- Send your comments in an Internet note to redbook@us.ibm.com

Which of the following best describes you?
__Customer __Business Partner __Solution Developer __IBM employee
__None of the above

Please rate your overall satisfaction with this book using the scale:
(1 = very good, 2 = good, 3 = average, 4 = poor, 5 = very poor)

Overall Satisfaction _____

Please answer the following questions:

Was this redbook published in time for your needs? Yes____ No____

If no, please explain:

_____

_____

_____

_____

What other redbooks would you like to see published?

_____

_____

_____

Comments/Suggestions: (THANK YOU FOR YOUR FEEDBACK!)

_____

_____

_____

_____

_____

SG24-5289-00
Printed in the U.S.A.