

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



Planning Series:

Protocol Descriptions

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



Planning Series:

Protocol Descriptions

Note

Before using this information and the product it supports, be sure to read the general information under “Notices” on page xiii.

Third Edition (October 2001)

This edition applies to the 3745 Communication Controller Models A and 3746 Nways® Multiprotocol Controller Models 900 and 950.

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

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

Department CGFA
Design & Information Development
IBM Corporation
PO Box 12195
Research Triangle Park NC 27709
U.S.A.

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

© Copyright International Business Machines Corporation 1988, 2001. All rights reserved.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures	ix
Tables	xi
Notices	xiii
Electronic Emission Notices	xiv
Industry Canada Class A Emission Compliance Statement	xiv
Avis de conformité aux normes d'Industrie Canada	xiv
European Union (EU) Mark of Conformity Statement	xiv
Japanese Voluntary Control Council for Interference (VCCI) Statement	xv
Korean Communications Statement	xvi
Taiwanese Class A Warning Statement	xvi
New Zealand Radiocommunications (Radio) Regulations	xvi
Trademarks	xvi
Year 2000 Statement	xvii
Safety	xviii
Notice to Users	xviii
What Is New in This Edition	xviii
About This Guide	xix
Who Should Use the 3745/3746 Planning Series	xix
How the 3745/46 Planning Series Is Organized	xix
Where to Find More Information	xxii
Additional Information on the Web	xxiii
CD-ROM	xxiii
Accessing CD-ROM Information	xxiv
How to Use the 3745/3746 Planning Series	xxv
Your Responsibility as a Customer	xxv
Finding Your Way Around in the New Planning Series	xxviii
Chapter 1. APPN / HPR Overview	1
Introduction	1
Node Types	2
APPN Network Node	2
APPN End Node	3
LEN End Node	3
Enhanced Support for LEN End Nodes	4
Other Node Types	5
Boundary and Peripheral Node	6
Composite Node	6
Interchange Node	6
Virtual Routing Node	7
Border Node	8
Extended Border Node	8
Branch Extender	8
APPN Node Structure	10
Node Operator	10
Node Operator Facility	10
Application Transaction Program	11
Control Point	11

Intermediate Session Routing	12
Logical Unit	12
Path Control	13
Data Link Control	13
Names	14
Network-Accessible Units	15
Network Identifiers	15
Network Names	15
Network-Qualified Names	15
Multiple LEN Nodes with the Same CPName	15
Addresses	16
Domains	17
High-Performance Routing	17
HPR Base and Towers	18
HPR Link Support	19
Automatic Network Routing	20
Rapid Transport Protocol	21
Class of Service	21
ARB Flow Control and Congestion Control	22
RTP Alive Timer	26
Nondisruptive Path Switch	27
Multilink Transmission Groups	27
HPR MLTG Requirements	27
HPR MLTG Overview	28
Priority	30
Route Calculation	30
HPR-Only Route for Path Switch	31
Timers	31
Alive Timer	31
Short_req Timer	31
Path Switch Timer	31
Migration	32
Interoperation with Existing APPN Nodes	32
No Configuration Restrictions	32
HPR Uses APPN Control Point Protocols and Algorithms	32
Shared Topology	32
Migration Planning	32
Dependent LU Requester/Server	33
3746-9x0 APPN/HPR Network Node Implementation	34
Terminology and Implementation Specifics	34
Session Establishment and Routing	35
3746 Network Node Processor Backup	38
Network Node Connectivity	39
DLC Support of Error Recovery	39
Token-Ring	39
Frame Relay	39
SDLC	40
X.25	40
Dependent LU Requester Implementation	40
VTAM DLCADDR Keyword in PATH Statement	42
Migration of SNI Connections to 3746 Models 900 and 950	43
EBN/SSE Usage Rules	46
Networks with DLURs: PLU, DLUS, and DLUR in Three Different Subnetworks	46

Networks with DLURs: DLUS and DLUR in Same Subnetwork	47
Networks with DLURs: PLU and DLUS in Same Subnetwork	48
Networks with DLURs: PLU and DLUR in Same Subnetwork	49
Networks with DLURs: PLU in Subarea Network, DLUS and DLUR in Same APPN Network	50
APPN Networks: 3746-9X0 Used as NNS of VTAM ENs	51
APPN Networks: 3746-9X0 Connected to a VTAM EBN	52
Mixed APPN and Subarea Networks: 3746-9X0 between EBN and NCP/Subarea (with SLU)	53
Mixed APPN and Subarea Networks: 3746-9X0 between EBN (with SLU) and NCP/Subarea	54
Internal APPN Connection between a 3745 and a 3746-900	55
Example of an Internal APPN Link	55
3746-900 Configuration Requirements	55
Definitions Required for the Internal APPN Link	56
Examples of 3746 Network Nodes in APPN/HPR Networks	58
External APPN Connection between a 3745-Composite Network Node and a 3746-900	61
HPR MLTGs in the 3746	63
User-Defined Parameters	64
Operation of 3746 Network Nodes in SNA/APPN/HPR Networks	66
APPN/HPR Network Node	66
Dependent Logical Unit Requester (DLUR)	67
High-Performance Routing	68
HPR Multilink Transmission Group	69
3746 Model 900 as a Mixed SNA and APPN/HPR Node	69
X.25 Network Connectivity	69
Frame-Relay Networking	70
Physical Media	70
LU-LU Sessions with Session Pacing	70
Link Station Name Included in XID3 Frame	70
VTAM DeactPU/DeactLU Support by NNP	71
3746-9x0 MAE APPN/HPR Implementation	71
MAE Implementation Specifics	71
APPN over DLSw	71
Supported Traffic Types	72
Topology Safestore	73
Route Test	73
No User Application Programming Interfaces	74
Limited Resource Link Stations	74
Session-Level Security	74
Parallel TGs	74
DLUR Restrictions	74
Connection Network Restrictions	74
APPN over DLSw Restrictions	75
Summary of Implemented APPN Functions	76
Summary of Supported DLCs and APPN Functions	81
 Chapter 2. Internet Protocol (IP) Overview	 83
The Need for Transparency in Communication	83
IP Addressing	84
The IP Address	85
IP Address Classes	85
Host and Network Numbers	86

IP Addresses for Routers	86
Subnets	86
Subdividing the Network	87
The Subnet Mask	87
Example Subnet Mask	87
Building Good Masks	88
Multiple Destinations	88
Broadcasting	88
Multicasting	89
Domains	89
The Hierarchical Name Space	90
Fully Qualified Domain Names (FQDNs)	90
Generic Domains	91
Country Domains	91
3746-9x0 Router Node Implementation	91
Overview of 3746 IP Routing	92
IP General Functions	93
Broadcasting	93
Supernetting (Route Aggregation)	94
Multiple IP Addresses for a Network Interface	95
Multiple Interfaces to the Same Subnet	98
User Datagram Protocol (UDP)	98
Transmission Control Protocol (TCP)	98
Internet Control Message Protocol (ICMP)	98
Address Resolution Protocol (ARP)	100
ARP Caching	100
Proxy-ARP	101
Routing Protocols	101
Routing Information Protocol (RIP) Version 1	101
Routing Information Protocol (RIP) Version 2	102
RIP Metrics	102
Open Shortest Path First Protocol (OSPF)	103
Not So Stubby Area	105
Border Gateway Protocol Version 4 (BGP)	107
Interoperability of Routing Protocols	107
OSPF Specifics	108
Internal Applications	109
Ping	110
Traceroute	111
Bootstrap Protocol (BOOTP)	112
Security	112
3746 IP Filters	113
3746 Access Control	114
Filters Versus Access Controls	115
Filters	115
IP over 3746 Networks	116
IP Interfaces	117
ESCON Attachment	118
ESCON: General Configuration	119
ESCON: Port Sharing	120
ESCON: IP Addresses and Subnet Addresses Rules	121
ESCON: IP Configuration	124
Token-Ring Attachment	125
Token-Ring: General Configuration	125

Token-Ring: Port Sharing	125
Token-Ring: IP Configuration	126
External IP Connection Between 3745 and 3746-900	127
Internal IP Connection Definitions	128
Internal IP Connection Between 3745 and 3746-900	132
Telnet Operations through the 3746 Network Node	133
Basic Operation	134
Network Virtual Terminal	135
NVT Printer	135
Full-Screen Capability	135
Command Structure	135
Host/Client Implementations	135
Mandatory Access Control Entry	136
List of Abbreviations	137
Glossary	141
Bibliography	145
Customer Documentation for the 3745 (All Models) and 3746 (Model 900)	145
Additional Customer Documentation for the 3745 Models 130, 150, 160, and 170	151
Customer Documentation for the 3746 Model 950	152
Required Documentation	156
Related Documentation	156
Index	159

|
|

Figures

1.	Composite Network Node Acting As an Interchange Node	7
2.	Virtual Routing Node	8
3.	Structure of an APPN or LEN Node	14
4.	Session with Several Session Stages	16
5.	HPR Base and Towers	18
6.	Basic Multiplexer System	24
7.	Burst Timing and the Three Parameters	25
8.	Multilink and Parallel TGs	29
9.	Dependent LU Requester/Server	33
10.	3746 NN Structure	35
11.	CP-CP Sessions	36
12.	Intermediate Session Routing	37
13.	3746 NN Intermediate Session Routing	38
14.	3746 Network Node DLUR Connectivity	40
15.	Example Network: Adding APPN Before SNI-to-APPN Migration	44
16.	Example Network: SNI-to-APPN Migration	44
17.	Example Network: Upgrading 3746-900 to a 3746-950	45
18.	Networks with DLURs: PLU, DLUS, and DLUR in Three Different Subnetworks	46
19.	PLU and DLUR in Same Subnetwork	47
20.	PLU and DLUS in Same Subnetwork	48
21.	PLU and DLUR in Same Subnetwork	49
22.	PLU in Subarea Network, DLUS and DLUR in Same APPN Network	50
23.	3746-9X0 Used as NNS of VTAM ENs	51
24.	3746-9X0 Connected to a VTAM EBN	52
25.	3746-9X0 between EBN and NCP/Subarea (with SLU)	53
26.	3746-9X0 between EBN (with SLU) and NCP/Subarea	54
27.	Internal APPN Link (3746-900)	56
28.	3746 as an ANR Node	58
29.	3746 as an APPN/RTP Node (with Boundary Function)	59
30.	3746 as a DLUR/RTP Node (with Boundary Function)	59
31.	3746 as an APPN/RTP Node with SNI Network	60
32.	External Connection of 3745-CNN and 3746-NN	61
33.	Defining MLTGs in CCM	63
34.	Mode Configuration Panel	64
35.	TG Row Configuration Panel	65
36.	CoS/TG User-Defined Parameters Panel	65
37.	Station Configuration - TG Characteristics Panel	66
38.	Dependent LU Support	67
39.	APPN over DLSw	72
40.	Communicating across Multiple Networks	83
41.	The Internet Protocol (IP) Overlays Networks	84
42.	3746 Network Nodes with the IP Feature (FC 5033)	84
43.	Classes of IP Addresses	85
44.	Routers Need Two Addresses	86
45.	Hierarchical Structure of Domains	90
46.	IP Address Support Example	92
47.	Three Networks and Three Routers Before Consolidation of the Routers	95
48.	Three Networks and One 3746 IP Router After Consolidation	96
49.	A Single 3746 IP Router Serving Four IP Networks through a Single Port	97

50.	Packet Internet Groper (Ping)	110
51.	Traceroute	111
52.	IP Filters	113
53.	IP Access Control	114
54.	3746-9X0 IP Base	116
55.	IP Interfaces	118
56.	ESCON Port, Host Link, and Link Station	118
57.	Multiple 3746 IP ESCON Links	119
58.	ESCON Host Links and Link Stations	120
59.	ESCON Port Sharing	121
60.	Four ESCP Stations with the Same IP Addresses	122
61.	Five ESCP Stations with the Same IP Addresses	123
62.	ESCP Stations on Separate Subnets	124
63.	Token-Ring Encapsulation	125
64.	Token-Ring Port Sharing	126
65.	Internal IP PtP Connection	127
66.	Example Configuration	128
67.	Port Configuration	130
68.	IP over Token-Ring Parameters	131
69.	OSPF/RIP Parameters per IP Address	131
70.	RIP Parameters per IP Address	132
71.	Direct NCP-IP 3746 Routing	133
72.	NCP-IP 3746 Routing Over ESCON	133
73.	Telnet Connection via a 3746 Network Node	134
74.	Telnet Command Structure	135

Tables

1.	Customer Tasks	xxv
2.	Location of Old Planning Guide Chapters in New Planning Guides	xxviii
3.	Networks with DLURs: PLU, DLUS, and DLUR in Three Different Subnetworks	46
4.	PLU and DLUR in Same Subnetwork	47
5.	PLU and DLUS in Same Subnetwork	48
6.	PLU and DLUR in Same Subnetwork	49
7.	PLU in Subarea Network, DLUS and DLUR in Same APPN Network . . .	50
8.	3746-9X0 Used as NNS of VTAM ENs	51
9.	3746-9X0 Connected to a VTAM EBN	52
10.	3746-9X0 between EBN and NCP/Subarea (with SLU)	53
11.	3746-9X0 between EBN (with SLU) and NCP/Subarea	54
12.	Summary of Implemented Base Functions (APPN Version 2)	76
13.	Summary of Implemented Optional Functions (APPN Version 2)	79
14.	Summary of Supported DLCs and APPN Functions	81
15.	Use of IP Address Classes	85
16.	IP Address Structure	86
17.	IP Addresses for a Router	86
18.	The Generic Top-Level Domains	91
19.	IP Multiple Access Parameters	114
20.	Customer Documentation for the 3745 Models X10 and X1A, and 3746 Model 900	145
21.	Additional Customer Documentation for the 3745 Models 130 to 170 . .	151
22.	Customer Documentation for the 3746 Model 950	152

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area.

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

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

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

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

Electronic Emission Notices

Federal Communications Commission (FCC) Statement

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

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

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Industry Canada Class A Emission Compliance Statement

This Class A digital apparatus complies with Canadian ICES-003.

Avis de conformité aux normes d'Industrie Canada

Cet appareil numérique de la classe A est conforme à la norme NMB-003 du Canada.

European Union (EU) Mark of Conformity Statement

This product is in conformity with the protection requirements of EU Council Directive 89/336/EEC on the approximation of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a non-recommended modification of the product, including the fitting of non-IBM option cards.

This product has been tested and found to comply with the limits for Class A Information Technology Equipment according to CISPR 22/European Standard EN 55022. The limits for Class A equipment were derived for commercial and industrial environments to provide reasonable protection against interference with licensed communication equipment.

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

Zulassungsbescheinigung laut dem Deutschen Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) vom 30. August 1995 (bzw. der EMC EG Richtlinie 89/336).

Dieses Gerät ist berechtigt, in Übereinstimmung mit dem Deutschen EMVG das EG-Konformitätszeichen - CE - zu führen.

Verantwortlich für die Konformitätserklärung nach Paragraph 5 des EMVG ist die IBM Deutschland Informationssysteme GmbH, 70548 Stuttgart.

Informationen in Hinsicht EMVG Paragraph 3 Abs. (2) 2:

Das Gerät erfüllt die Schutzanforderungen nach EN 50082-1 und EN 55022 Klasse A.

EN 55022 Klasse A Geräte müssen mit folgendem Warnhinweis versehen werden:

“Warnung: Dies ist eine Einrichtung der Klasse A. Diese Einrichtung kann im Wohnbereich Funkstörungen verursachen; in diesem Fall kann vom Betreiber verlangt werden, angemessene Maßnahmen durchzuführen und dafür aufzukommen.”

EN 50082-1 Hinweis:

“Wird dieses Gerät in einer industriellen Umgebung betrieben (wie in EN 50082-2 festgelegt), dann kann es dabei eventuell gestört werden. In solch einem Fall ist der Abstand bzw. die Abschirmung zu der industriellen Störquelle zu vergrößern.”

Anmerkung:

Um die Einhaltung des EMVG sicherzustellen, sind die Geräte, wie in den IBM Handbüchern angegeben, zu installieren und zu betreiben.

Japanese Voluntary Control Council for Interference (VCCI) Statement

This product is a Class A Information Technology Equipment and conforms to the standards set by the Voluntary Control Council for Interference by Technology Equipment (VCCI). In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

この装置は、情報処理装置等電波障害自主規制協議会（VCCI）の基準に基づくクラスA情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

Korean Communications Statement

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

A급 기기(업무용)

이 기기는 업무용으로 전자파적합등록을 받은 기기이오니
판매자 또는 이용자는 이점을 주의하시기 바라며, 만약
구입하였을 때에는 구입한 곳에서 가정용으로 교환하시기
바랍니다.

Taiwanese Class A Warning Statement

警告使用者：
這是甲類的資訊產品，在
居住的環境中使用時，可
能會造成射頻干擾，在這種
情況下，使用者會被要求
採取某些適當的對策。

New Zealand Radiocommunications (Radio) Regulations

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

Trademarks

The following are trademarks or registered trademarks of International Business Machines Corporation or Tivoli Systems Inc. in the United States, or other countries, or both:

AIX	NetView
ACF/VTAM	Nways
Advanced Peer-to-Peer Networking	Operating System/2
APPN	OS/2
AS/400	OS/390
CICS	Processor Resource/Systems Manager
DB2	PS/2
Enterprise Systems Connection	RETAIN
Architecture	RS/6000
Extended Services	S/370
ESCON	S/390
ESCON XDF	S/390 Parallel Enterprise Server
ES/3090	System/36
ES/9000	System/370
IBM	System/390
the IBM logo	SystemView
LPDA	Tivoli
Multiprise	TME
MVS	VM/ESA
MVS/ESA	VSE/ESA
MVS/XA	VTAM

Freelance is a trademark of Lotus Development Corporation in the United States, or other countries, or both.

Java, all Java-based trademarks and logos, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and/or other countries.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks or registered trademarks of Microsoft Corporation.

Intel and Pentium are registered trademarks of Intel Corporation in the U.S. and other countries.

UNIX is a registered trademark of the Open Group in the United States and other countries.

Other company, product, and service names may be trademarks or service marks of others.

Year 2000 Statement

This product is Year 2000 ready. When used in accordance with its associated documentation, it is capable of correctly processing, providing, and/or receiving date data within and between the 20th and 21st centuries, provided all other products (for example, software, hardware, and firmware) used with the product properly exchange accurate date data with it.

The 3745 and 3746 controllers require a certain level of microcode to be Year 2000 ready. For more detailed information, go to one of the following Web sites and click **Support**.

www.ibm.com/networking/3745
www.ibm.com/networking/3746

Safety

First see the following manual: *3745 Communication Controller All Models, 3746 Nways Multiprotocol Controller Models 900 and 950 Safety Information*, GA33-0400.

Notice to Users

The IBM 3746 Expansion Unit Model 900 and IBM 3746 Nways Multiprotocol Controller Model 950 are manufactured according to the International Safety Standard IEC 60950.

Active Remote Couplers (ARCs) and the X.21 Interface, housed within the IBM 3746 Expansion Unit Model 900 and IBM 3746 Nways Multiprotocol Controller Model 950, and supplied by IBM, do not use or contain excessive voltages. An excessive voltage is one that exceeds 42.4 V peak ac or 60 V dc. They interface with the IBM 3746 Expansion Unit Model 900 and IBM 3746 Nways Multiprotocol Controller Model 950 using Safety Extra Low Voltages (SELV) only.

What Is New in This Edition

This book has been revised to include the following changes and enhancements:

- Multiple IP interfaces to the same subnet
- The not so stubby area (NSSA) function
- Nexthop awareness over Frame Relay
- Nexthop awareness over ESCON
- Multiple LEN nodes with the same CPName
- HPR/RTP adaptive rate-based (ARB) responsive mode (ARB2)
- LU-LU sessions with session pacing
- Link station name included in XID3 frame
- VTAM DeactPU/DeactLU support by NNP

The technical changes and additions are indicated by a vertical line (|) to the left of the change.

About This Guide

The *3745/3746 Planning Series* is designed to help you plan the installation and configuration of the IBM 3745 Communication Controller Models A and IBM 3746 Nways® Multiprotocol Controller Models 900 and 950. The *Planning Series* also describes the information you must gather to install and integrate 3746 Controllers into Advanced Peer-to-Peer Networking®/High-Performance Routing (APPN®/HPR) and Internet Protocol (IP) environments.

The *3745/3746 Planning Series* consists of a set of Planning Guides that replace, update and obsolete the *Planning Guide*.

Who Should Use the 3745/3746 Planning Series

The *3745/3746 Planning Series* is intended for network planners, network specialists, and system programmers responsible for collecting the information required for the installation and network integration of 3745 Communication Controller Models A and 3746 Expansion Unit Model 900 in an SNA environment, as well as the 3746-950 and 3746-900 as APPN/HPR network nodes and IP routers.

How the 3745/46 Planning Series Is Organized

Important:

1. If you already use the existing *Planning Guide*, IBM recommends that you read the new *Planning Series* to learn about new features and to become familiar with the new structure in which planning information is presented.
2. When planning the installation and configuration of 3746 controllers you must use the *IBM 3745 Communication Controller Models A, IBM 3746 Nways Multiprotocol Controller, Models 900 and 950: Overview* along with the *3745/3746 Planning Series* to have all required information.
3. The 3745/3746 documentation is updated periodically in response to your needs and to reflect product evolutions. Because of the time delay necessary to update hard media (books that are printed and available on CD-ROM), it is highly recommended that you periodically check the IBM 3745/3746 documentation on the Web for the latest versions of the documents (see “Additional Information on the Web” on page xxiii).

Refer to the appropriate Planning Guide for the parameters to be customized for the installation and operation of:

- 3745 Communication Controller Models A
- 3746 Nways Multiprotocol Controller Models 900 and 950
- Network Node Processor (NNP)
- Multiaccess Enclosure (MAE)
- Service Processor
- Distributed Console Access Facility (DCAF) and TME® 10 remote consoles
- Java™ Console
- Network management

When you define 3746 resources controlled by NCP, record the information in the worksheets provided for the Controller Configuration and Management application.

The *3745/3746 Planning Series* consists of the following planning guides:

Overview, Installation, and Integration, GA27-4234

Starts with a general overview of 3746 planning and then explains the various 3745 and 3746 installation and upgrade scenarios.

The guide also explains the options available for the basic integration of the controller and its service processor into your network. There are MOSS-E worksheets for these options, which are to be filled out for the IBM service representative who does the actual controller installation or upgrade. The appendixes:

- Show the panels of the MOSS-E service processor customization function
- Describe the support offered by each level of the 3746 Licensed Internal Code

ESCON Channels, GA27-4237

After an overview of ESCON® architecture and the adapters, the publication explains the configuration and tuning. This can be done with either the ESCON Generation Assistant (EGA) tool or the Controller Configuration Management (CCM) tool.

The publication also includes examples of various types of ESCON configurations.

Note: For information about using ESCON adapters on the MAE, refer to the *Multiaccess Enclosure Planning* guide.

Token Ring and Ethernet, GA27-4236

Helps you with the configuration and definitions of your 3746 Network Node token-ring adapters (TRAs) for APPN/HPR-, IP-, and NCP-controlled traffic.

There are MOSS-E worksheets for the token-ring information needed by the IBM service representative to install or update your machine.

Although no longer available from IBM, the guide explains 3746 Ethernet support and Ethernet adapter configuration.

The token-ring (IEEE 802.5) and Ethernet (IEEE 802.3) standards are discussed in the appendixes.

Note: For Multiaccess Enclosure Ethernet information, refer to the *Multiaccess Enclosure Planning* guide.

Serial Line Adapters, GA27-4235

Provides an overview of the serial line adapters, and describes the support for X.25, frame relay, PPP, and SDLC.

The two ways that the 3746 supports ISDN (LIC16 adapter¹ and terminal adapters) are explained, including how ISDN lines can be used as backups for other types of lines.

¹ No longer available

An appendix describes the frame-relay support in each NCP level since frame relay was introduced in NCP Version 6.

Note: For Multiaccess Enclosure ISDN information, refer to *Multiaccess Enclosure Planning*.

Physical Planning, GA27-4238

Gives information to help you plan the physical site used by the 3745/3746 frames, Service Processor, and Network Node Processor: the physical dimensions, electrical characteristics, and so on. It also gives this information for the various components of the 3745/3646, such as the Multiaccess Enclosure, Controller Extension, LICs, LCBs, ARCs, and so on.

The cable descriptions include feature codes (FCs) and part numbers used when ordering them.

The guide includes and explains the controller installation sheets, which show what IBM has installed on your machines.

Plugging sheets for keeping track of your installed LICs, ARCs, and cables are provided along with examples and explanations of their use.

Note: This type of information for the Multiaccess Enclosure is in the *Multiaccess Enclosure Planning* guide.

Management Planning Guide, GA27-4239

Starts with a management overview covering:

- The Tivoli® NetView® program
- Performance management
- Service Processor
- Network Node Processor
- APPN Topology Integrator

Then there are chapters about:

- APPN/HPR Network Node management
- NetView Performance Monitor
- Remote console support
- IBM Remote Support Facility
- 3746 IP router management
- Multiaccess Enclosure APPN/HPR Network Node management
- X.25 network

There are MOSS-E worksheets for the network management parameters needed by the IBM service representative to install or upgrade your machine.

The guide also explains MOSS-E Service Processor Customization.

There is an example of ESCON Management Information Base (MIB) definitions.

Note: For Multiaccess Enclosure management information, refer to the *Multiaccess Enclosure Planning* guide.

Multiaccess Enclosure Planning, GA27-4240

Provides information about the Multiaccess Enclosure and its adapters (ATM, ESCON, and so on) and how to configure them.

For information about:

- Multiaccess Enclosure APPN/HPR Network Node management, refer to the *3745/3746 Planning Series: Management Planning Guide*
- Physical site planning and the cables, refer to the *3745/3746 Planning Series: Physical Planning*

Protocol Descriptions, GA27-4241

Is an in-depth description of these protocols used by the 3746:

- APPN/HPR
- IP

The detailed discussions of how the 3746 and Multiaccess Enclosure support these protocols help you understand the purpose of the protocol parameter definitions and what types of information are needed for the most efficient operation of your 3745/3746-connected networks.

CCM Planning Worksheets (online)

These example worksheets for the 3746 and MAE can be used to plan the actual definitions of the many CCM parameters you need to configure for your 3746.

These worksheets are available in PDF format at:

ibm.com/networking/did/3746bks.html#Customer

Where to Find More Information

While planning a migration, you must use the following documents in addition to the *3745/3746 Planning Series* guides:

- *IBM 3745 Communication Controller Models A and 170, 3746 Nways Multiprotocol Controller Models 900 and 950: Overview*, GA33-0180
- *IBM 3745 Communication Controller All Models, 3746 Nways Multiprotocol Controller Model 900: Console Setup Guide*, SA33-0158 (This guide contains information about remote console access to 3745/3746-900s via an SNA/subarea, APPN, or TCP/IP path and using a modem.)

Also, you may need to use the following additional documents:

- *IBM 3746 Nways Multiprotocol Controller Model 900 and 950: Controller Configuration and Management: User's Guide*, SH11-3081 (IBM recommends that you prepare controller definitions before installing a 3746. To obtain a stand-alone version of the Controller Configuration and Management that runs on an OS/2® workstation, contact your IBM marketing representative.)
- *3746 Nways Multiprotocol Controller Model 950: User's Guide*, SA33-0356 (This guide contains information about routine operations, installing and testing the communication line adapters, service processor, and remote consoles.)
- *Planning for Integrated Networks*

Be sure to use the latest editions of these documents. This will ensure that you have up-to-date and complete information about the 3746 controllers.

The following *IBM International Technical Support Organization* redbooks provide useful information about 3746 implementation:

- *APPN Architecture and Product Implementations Tutorial*, GG24-3669

- *IBM 3746 Nways Multiprotocol Controller Model 950 and IBM Model 900: APPN Implementation Guide*, GG24-2536
- *Subarea Network to APPN Network Migration Guide*, SG24-4656
- *IBM 3746 Nways Multiprotocol Controller Model 950 and Model 900: IP Implementation Guide*, SG24-4845

Be sure to see the other relevant documents listed in the bibliography at the back of this guide.

Additional Information on the Web

You can access the latest news and information about IBM network products, customer service and support, and information about microcode upgrades at:

www.ibm.com/

The latest versions of the *Planning Series* and other 3745/3746 documentation are available in PDF format at:

www.ibm.com/networking/did/3746bks.html#Customer

CD-ROM

Starting with engineering change F12380, the Licensed Internal Code (LIC) is shipped on a CD-ROM. The complete 3745/3746 documentation set is also included on the CD-ROM.

Examples: 3745 Models A and 3746 *Planning Series*, 3746 NNP and Service Processor Installation and Maintenance Guides, CCM *User's Guide*, 3746-950 *User's Guide*, and others. See the bibliography for the complete name and form number of the books.

3745/3746 documentation is in PDF format. Acrobat Reader for OS/2® is included on the CD-ROM to allow you to read the PDF files and print all or part of a book.

Accessing CD-ROM Information

To access the CD-ROM from a service processor equipped with a CD-ROM drive, use the following procedure:

1. Install the CD-ROM in the service processor CD-ROM drive.
2. In the MOSS-E main panel, open the **View** menu and select **Information**.
3. Double-click **CD-ROM documentation**. Your browser automatically opens and displays the documentation home page.
4. Click any highlighted text (blue and underlined) to go to the material that interests you:

- a. Click **Documentation** to access 3745/3746 books.
- b. Click the icon marked PDF that corresponds to the item that interests you.

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

Press **Esc** to display the Reader menus that allow you to print all or part of the file.

When you close the Acrobat Reader, you return to the browser.

When you close the browser, you return to the MOSS-E Documentation menu.

Each document file has one or more of the following identifiers:

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

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

How to Use the 3745/3746 Planning Series

Your Responsibility as a Customer

You are responsible for performing the tasks listed in Table 1. These tasks are not performed by IBM personnel as part of the machine installation and basic operations. They can, however, be performed by IBM on a fee basis.

Table 1 (Page 1 of 3). Customer Tasks	
Task	Where to Find Information
Network design:	<p>Network design is not covered in this book. Refer to the following IBM books for SNA, APPN/HPR, and IP network planning guidance:</p> <ul style="list-style-type: none">• <i>Planning for Integrated Networks</i>• IBM redbooks:<ul style="list-style-type: none">– <i>Subarea Network to APPN Network Migration Guide</i>– <i>IBM 3746 Nways Multiprotocol Controller Model 950 and IBM Model 900: APPN Implementation Guide</i>– <i>IBM 3746 Nways Multiprotocol Controller Model 950 and IBM Model 900: IP Implementation Guide</i>– <i>IBM Nways 2216 Multiaccess Connector Description</i>– <i>IBM 2216 Multiaccess Connector ESCON Solutions</i>
<p>Physical planning:</p> <p>Before the IBM service representative arrives to install your controller, make sure that you have met the necessary requirements for the following:</p> <ul style="list-style-type: none">• Electric power• Floor space with service clearances• Space for the cables• The RSF switched line• The Controller Expansion (FC 5023)• Other components (such as the service processor).	<p>“Physical Planning Details” chapter in the <i>3745/3746 Planning Series: Physical Planning</i></p>
<p>Controller hardware configuration definitions:</p> <p>Decide what type of attachments (lines) and how many of each type you need.</p>	<p>This input is necessary for the IBM ordering system (CF3745). For more information, refer to the <i>3745/3746 Planning Series: Physical Planning</i>.</p>

Table 1 (Page 2 of 3). Customer Tasks

Task	Where to Find Information
<p>Software definitions and tuning:</p> <ul style="list-style-type: none"> • ESCON port, host link, and station definitions; ESCON resource, TCP/IP, and VTAM® tuning • Token-ring port and station definitions; PU and LU maximum limits; port sharing with NCP-controlled traffic; duplicate addresses; token-ring APPN, IP, and/or NCP resource tuning and VTAM tuning • Serial line (SDLC, PPP, frame-relay, and X.25) port and station definitions; location of CLPs, LICs, LCBs, and ARCs; maximum CLA line connectivity; CLP backups • Multiaccess Enclosure: hardware planning and configuration; software configuration and tuning • Use of the Controller Configuration and Management (CCM) application. 	<p>Refer to:</p> <ul style="list-style-type: none"> • “ESCON Adapters” chapter in the <i>3745/3746 Planning Series: ESCON Channels</i> • “ESCON Channel Adapter” chapter in the <i>3745/3746 Planning Series: Multiaccess Enclosure Planning</i> • “ESCON Configuration Examples” chapter in the <i>3745/3746 Planning Series: ESCON Channels</i> • “Token-Ring Adapters” chapter in the <i>3745/3746 Planning Series: Token Ring and Ethernet</i> • “Serial Line Adapters” chapter in the <i>3745/3746 Planning Series: Serial Line Adapters</i> • “3746 SDLC Support” chapter in the <i>3745/3746 Planning Series: Serial Line Adapters</i> • <i>3745/3746 Planning Series: Multiaccess Enclosure Planning</i> • <i>3745/3746 Planning Series: Physical Planning</i> • <i>IBM Controller Configuration and Management User's Guide</i>, SH11-3081. <p>Also refer to:</p> <ul style="list-style-type: none"> • <i>IBM 3746 Nways Multiprotocol Controller Model 950 and IBM Model 900: APPN Implementation Guide</i> (an IBM redbook) • <i>IBM 3746 Nways Multiprotocol Controller Model 950 and IBM Model 900: IP Implementation Guide</i> (an IBM redbook).
<p>Filling out:</p> <ul style="list-style-type: none"> • 3746 plugging sheets To keep a record of the processors and couplers (and their addresses) installed in the 3746 frame. • <i>CCM User's Guide</i>, SH11-3081 worksheets To plan the 3746 and MAE logical resource definitions. They can then be used when configuring the 3746 and MAE using the CCM. 	<p>Refer to:</p> <ul style="list-style-type: none"> • “Plugging Sheets for 3745 and 3746” chapter in the <i>3745/3746 Planning Series: Physical Planning</i> • <i>3745/3746 Planning Series: CCM Planning Worksheets</i>

Table 1 (Page 3 of 3). Customer Tasks

Task	Where to Find Information
<p>NetView definitions in VTAM, the MOSS-E, NPM, CCM, NetView/360, and Tivoli NetView (formerly NetView for AIX) for:</p> <ul style="list-style-type: none"> • APPN traffic • IP traffic • NetView alert path 	<p>Refer to:</p> <ul style="list-style-type: none"> • “3746 Management Overview” chapter in the <i>3745/3746 Planning Series: Management Planning Guide</i> • “3746 APPN/HPR Network Node Management” chapter in the <i>3745/3746 Planning Series: Management Planning Guide</i> • “3746 IP Router Management” chapter in the <i>3745/3746 Planning Series: Management Planning Guide</i>.
<p>Controller, service processor, and network node processor definitions. For example:</p> <ul style="list-style-type: none"> • Link IPL port information • Password management • NetView alert reporting path definitions • DCAF LU definitions • Ethernet port definitions for SNMP • Service processor token-ring and IP LAN addresses 	<p>Refer to “Controller and Service Processor Integration” chapter in the <i>3745/3746 Planning Series: Overview, Installation, and Integration</i>.</p> <p>Fill out the worksheets in the various <i>Planning Series</i> guides. These worksheets are used by the IBM service representative during installation.</p>
<p>Remote console definitions (using DCAF):</p> <ul style="list-style-type: none"> • Ensure that the necessary hardware and software is available for the type of console attachment chosen • Service processor definitions for DCAF • DCAF installation and configuration on the remote console 	<p>Refer to:</p> <ul style="list-style-type: none"> • “Remote Customer Consoles” chapter in the <i>3745/3746 Planning Series: Management Planning Guide</i> • For the 3746-900, refer to the <i>3745 Console Setup Guide</i> • For the 3746-950, refer to the <i>IBM 3746 Nways Multiprotocol Controller Model 950 User’s Guide</i>
<p>Connection to the IBM remote support facility (RSF):</p> <ul style="list-style-type: none"> • Service processor connection (modem) definitions • Customer definitions for RSF records. 	<p>Refer to the “Connecting to the IBM Remote Support Facility” chapter in the <i>3745/3746 Planning Series: Management Planning Guide</i></p>
<p>Problem determination through the MOSS-E and NetView</p>	<p>For the 3746-900, refer to:</p> <ul style="list-style-type: none"> • <i>Problem Analysis Guide</i> accessed online from the MOSS-E • <i>3745 Models A: Alert Reference Guide</i> • <i>3745 All Models: Advanced Operators Guide</i>

Finding Your Way Around in the New Planning Series

If you are familiar with the layout of the old *3745 Communication Controller Models A and 3746 Models 900 and 950: Planning Guide*, GA33-0457, Table 2 should help you find which of the eight new books of the planning series contains the information that you need.

Note: Some of the chapters in the *Planning Guide* have been split into two or more new chapters in one or more new guides.

Table 2 (Page 1 of 2). Location of Old Planning Guide Chapters in New Planning Guides

Old Planning Guide		New Planning Series Book	
Chapter	Chapter Name	Chapters	Guide Name
1	3745 and 3746 General Information	--	Not included in the new guides
2	APPN/HPR Overview	1	<i>Protocol Descriptions</i>
3	Internet Protocol (IP) Overview	2	<i>Protocol Descriptions</i>
4	3746 ATM Support	4	<i>Multiaccess Enclosure Planning</i>
5	Token-Ring/802.5	B	<i>Token-Ring and Ethernet</i>
6	Ethernet Overview	C	<i>Token-Ring and Ethernet</i>
7	Frame Relay Overview	4, 5	<i>Serial Line Adapters</i>
8	Point-to-Point Protocol (PPP) Overview	4	<i>Serial Line Adapters</i>
9	X.25 Overview	2, 3, 5, 7	<i>Serial Line Adapters</i> <i>Management Planning</i>
10	ISDN Adapters	8	<i>Serial Line Adapters</i>
11	ESCON Overview	1	<i>ESCON Channels</i>
12	3745 and 3746 Installation and Upgrade Scenarios	2	<i>Overview, Installation, and Integration</i>
13	Configuration Scenarios	6	<i>Multiaccess Enclosure Planning</i>
14	3746 Planning Overview	1	<i>Overview, Installation, and Integration</i>
15	ESCON Adapters	1, 2, 3	<i>ESCON Channels</i>
16	Token-Ring Adapters	1, 2, 3	<i>Token-Ring and Ethernet</i>
17	Ethernet Adapters	4, 5	<i>Token-Ring and Ethernet</i>
18	Serial Line Adapters	1	<i>Serial Line Adapters</i>
19	3746 SDLC Support	3, 4	<i>Serial Line Adapters</i>
20	Multiaccess Enclosure	1	<i>Multiaccess Enclosure Planning</i>
21	Multiaccess Enclosure Adapters Overview	2	<i>Multiaccess Enclosure Planning</i>
22	ESCON Channel Adapter	8	<i>Multiaccess Enclosure Planning</i>
23	Multiaccess Enclosure ISDN Support	5	<i>Multiaccess Enclosure Planning</i>
24	3746 Configuration Overview	--	Not included in the new guides
25	Welcome to the CCM	--	Not included in the new guides
26	Multiaccess Enclosure Configuration	7	<i>Multiaccess Enclosure Planning</i>
27	3746 Base Frame ESCON Configuration Examples	1	<i>ESCON Channels</i>
28	Configuring the MAE ESCON Channel Adapter	8	<i>Multiaccess Enclosure Planning</i>

<i>Table 2 (Page 2 of 2). Location of Old Planning Guide Chapters in New Planning Guides</i>			
Old Planning Guide		New Planning Series Book	
Chapter	Chapter Name	Chapters	Guide Name
29	3746 Management Overview	1	<i>Management Planning</i>
30	3746 APPN/HPR Network Node Management	2	<i>Management Planning</i>
31	3746 IP Router Management	6	<i>Management Planning</i>
32	MAE APPN/HPR Network Node Management	2	<i>Management Planning</i>
33	MAE IP Router Management	6	<i>Management Planning</i>
34	Controller and Service Processor	3	<i>Overview, Installation, and Integration</i>
35	Customer Consoles and DCAF	4 1 1	<i>Management Planning</i> <i>Overview, Installation, and Integration</i> <i>Token-Ring and Ethernet</i>
36	Connecting to the IBM Remote Support Facility	5	<i>Management Planning</i>
37	Performance Management with NetView Performance Monitor	3	<i>Management Planning</i>
37	3746 IP Router Management	6	<i>Management Planning</i>
38	MOSS-E Worksheets for Controller Installation (3745)	A A A	<i>Overview, Installation, and Integration</i> <i>Management Planning</i> <i>Token-Ring and Ethernet</i>
39	Parameter Cross-Reference Table	B	<i>Overview, Installation, and Integration</i>
40	CCM Worksheets for Controller Configuration Definitions	1	<i>CCM Planning Worksheets (online)</i>
41	Multiaccess Enclosure Worksheets	2	<i>CCM Planning Worksheets (online)</i>
42	Familiarizing Yourself with the Installation Sheets	2	<i>Physical Planning</i>
43	Plugging Sheets for the 3746 Nways Multiprotocol Controller	3	<i>Physical Planning</i>
44	Physical Planning Details	1	<i>Physical Planning</i>
A	3746-9x0 Microcode Levels (EC)	D	<i>Overview, Installation, and Integration</i>
B	ESCOM MIB	A	<i>Management Planning</i>
C	MOSS-E Service Processor Customization Function	C	<i>Overview, Installation, and Integration</i>

Chapter 1. APPN / HPR Overview

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

For ACF/VTAM® and NCP implementations of HPR, refer to:

- *ACF/VTAM V4R3 HPR Early User Experiences*
- *ACF/NCP V7R3: New Functions*

Introduction

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

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

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

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

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

Advanced Program-to-Program Communication (APPC) is usually provided as system software. The APPC software provides two interfaces. The first, a programming interface "at the top," responds to requests from application programs

that need to communicate. One example of such an interface is the CPI-C. The second interface, “at the bottom”, exchanges data with the communications hardware.

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

The following sections, provide a description of each node type, showing the differences between them.

Node Types

Each node type in the APPN architecture has a specific function set. The following describes how each node type works. Table 12 on page 76 and Table 13 on page 79 summarizes all the option sets defined in the APPN architecture, and shows their implementation in the 3746.

APPN Network Node

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

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

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

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

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

APPN End Node

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

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

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

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

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

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

LEN End Node

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

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

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

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

Enhanced Support for LEN End Nodes

After migrating from a VTAM® and NCP composite node configuration and replacing the NCPs with 3746-900 network nodes (900 NNs), networks experience VTAM-to-branch session failures. Sessions can be activated from the branch to the central VTAM, but not from VTAM to the branch because VTAM can no longer find the branch LUs. The following sections explain the cause of these session failures and the enhanced support solutions.

Before Migration: In the VTAM and NCP composite node configuration, the network has all devices and their LUs predefined in VTAM switched major node decks. Most devices are defined to VTAM as LEN end nodes, and their LUs are defined as independent LUs associated with the particular device. When VTAM activates a connection to one such device, it finds the device predefinition (the PU statement) and the associated LU statements. The PU statement tells VTAM to treat this device as a LEN end node, even though the device actually is configured as an APPN network node. In this way, if any application or other LU wants to establish a session with any of those independent LUs, VTAM knows where to send the session request (BIND) and which device can successfully process it.

After Migration: After migrating to the 900 NN configuration, the branch boxes are no longer connected to VTAM; they are connected to the 900 NN control point. This means that all the switched major node PU and LU definitions in VTAM are now useless (the PUs remain connectable forever, and the LUs remain undefined). Because the branch LU definitions in the VTAM switched major node deck are not usable and the branch LUs are not predefined in the 900 NNs, VTAM can no longer find the branch LUs to set up VTAM-to-branch sessions, with the exception of VTAM-to-branch-initiated sessions, when the adjacent node is a network node with no CP-CP sessions. In those cases, predefinition in the 900 NN of all the LUs residing in the branch network nodes is still mandatory.

Functional Enhancements: The following enhancements have been made to address this migration issue:

1. The first enhancement enables the 900 NN to learn some of the names of the LUs in a LEN end node without predefining them. The 900 NN obtains the PLU names from the BINDs that the LEN end node sends in and then defines a directory entry for each of those LUs. The LU directory entries enable the 900 NN to answer a received Locate, Find (PLU_name_from_BIND) with a

Found, allowing sessions toward the LEN end node LU to succeed after that LU has sent a BIND.

Note: Once the 900 NN creates an LU definition, it keeps the directory entry as long as the connection to the LEN end node is alive. This might result in multiple network nodes having the same definition for a given LEN end node LU. If the LEN end node has multiple links to different network nodes and sends a blind BIND for the same PLU name on each of the links, each of the network nodes will respond Found to a Locate for that LU name. The same situation can occur if the LEN end node is actually an APPN end node that has no network node server (which is the only reason why it would send a BIND with no RSCV) and has some links predefined with CP-CP sessions disabled.

Unfortunately, this solution addresses only a part of the problem. Some branch LUs never send a BIND until being prompted by receipt of a BIND from VTAM.

2. The second enhancement addresses the inability of the 900 NN to respond to Locates on behalf of branch network nodes.

When the branch device connects to the 900 NN, the 900 NN sees this device as another APPN network node that is configured not to support CP-CP sessions on the link. Architecturally, one network node does not respond to Locates for LUs that are located on another network node. So, even if all the branch LUs are predefined in each 900 NN, because the 900 NN will not respond on behalf of the branch device (as a network node), VTAM is still not able to find the branch LUs.

The purpose of this second enhancement is to have the 900 NN respond to Locates on behalf of the branch network nodes. Predefinition in the 900 NN of all the LUs residing in the branch network nodes is **still mandatory** for VTAM-to-branch-initiated sessions, when the adjacent node is a network node with no CP-CP sessions as well as when it is a LEN end node.

A branch-to-VTAM BIND from a LEN end node will always create an LU predefinition, but it will **not** cause an LU predefinition to be created if the BIND is received from a network node unless there is some other DEFINE_PARTNER_LU_LOCATION already defined that tells the 900 NN to respond to Locates on behalf of that network node.

Other Node Types

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

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

Boundary and Peripheral Node

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

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

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

Composite Node

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

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

Interchange Node

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

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

Multiple VTAMs and NCPs may connect using subarea protocols (see for example VTAM1/NCP in Figure 1 on page 7). Session establishment is possible between

any LU in the subarea network and any LU in the APPN network. The VTAM host to which APPN nodes attach, or the VTAM host owning the NCPs to which APPN nodes attach, must have implemented APPN VTAM, because it is responsible (as an “interchange node”) for the conversion of subarea to APPN protocols and conversely; other VTAMs within the subarea network can be back-level VTAMs. From the viewpoint of the APPN nodes, LUs owned by VTAMs (for example, VTAM2 or VTAM3) other than the VTAM providing the interchange function, are considered to reside on APPN end nodes.

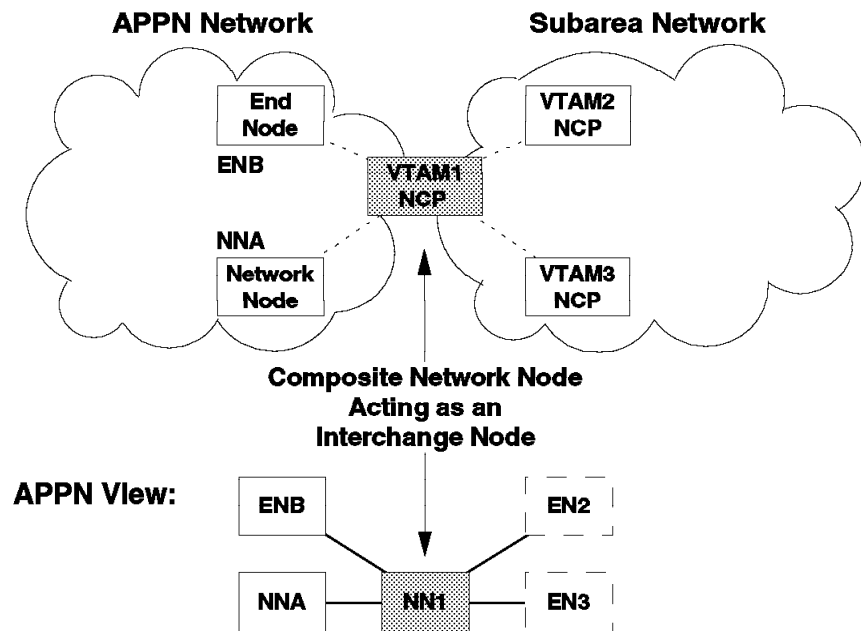


Figure 1. Composite Network Node Acting As an Interchange Node

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

Virtual Routing Node

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

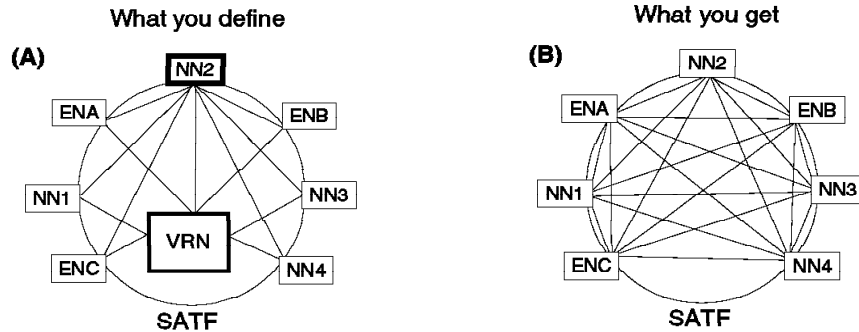


Figure 2. Virtual Routing Node

A virtual routing node (VRN) is not a node, but it is a way to define an APPN node's attachment to a shared-access transport facility. It reduces end node definition requirements by relying on the network node server to discover the common connection and supply necessary link-level signaling information as part of the regular Locate search process; LU-LU session data can then be routed directly, without intermediate node routing, between APPN nodes attached to the SATF.

Border Node

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

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

Extended Border Node

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

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

Branch Extender

The Branch Extender Network Node (BrNN) is designed to optimize the connection of a branch office to an APPN WAN backbone network. The BrNN isolates all the end nodes on one or more branch office LANs from the backbone WAN. The domain of a BrNN might contain only end nodes and cascaded BrNNs. The domain of a BrNN does not contain network nodes or nodes with DLUR. The link on which the BEX is seen as an EN is called the *uplink*. The link on which the BEX is seen as an NN is called the *downlink*.

When configuring a BrNN, configure link stations to the backbone to be uplinks. This causes the BrNN to appear as a conventional end node to the backbone. From the perspective of the backbone, all resources in the domain of the BrNN appear to be owned by the BrNN, hiding the topology of the BrNN's domain from the backbone and reducing the number of broadcast locates in the backbone.

A BrNN presents a conventional network node interface over downlinks. End nodes in the domain of the BrNN register their resources with the BrNN and use the BrNN as a conventional network node server.

A BrNN:

- Reduces of the number of network nodes in a large APPN network
- Hides branch office topology from the WAN and hides WAN topology from the BrNN
- Enables, direct, peer-to-peer communication between defined branches connected to the same connection network
- Reduces CP-CP session traffic on the WAN link

The following items are limitations of Branch Extender:

- Network nodes are allowed to connect only over links that a BrNN defines as uplinks.
- Only end nodes or cascaded BrNNs can be attached to a BrNN downlink.
- Border nodes acting as end nodes and DLUR nodes cannot be attached to a BrNN downlink.
- A node cannot connect to a Branch Extender over an uplink and a downlink at the same time.
- A BrNN can have CP-CP sessions with only one network node at a time.

The following CCM parameters have been added to configure BEX:

APPN Branch Extender

A node level parameter to enable the function.

Permit Search for Unregistered LUs

A branch extender normally registers all its LUs with its network node server, and indicates to the network node server that it should **not** be searched when the network node server is performing a domain broadcast search for an unknown LU.

Setting this parameter to YES causes the BEX to indicate to the network node server that it **should** be searched. You can set it to YES if the BEX node will be the network node server for an EN that does not register its LUs, so that searches can be propagated down to that EN. You can also set it to YES when the BEX node is DLUR and your network node server does not allow registration of DLUS-served dependent LUs.

Branch Uplink

This parameter allows you to define the upstream link or links on which your node shows an EN appearance.

Link to Preferred Network Node

This parameter allows you to define which adjacent node will be your network node server. You can define more than one network node server

If the preferred network node server goes down and there is no other preferred network node server available, CP-CP sessions will be established with any available NN. As soon as a link to a preferred network node server is reestablished, the CP-CP sessions on the non-preferred network node server are forced down and then brought back up with the preferred network node server.

APPN Node Structure

This section briefly describes the structure and components of APPN network nodes, end nodes, and LEN nodes. An overview of each is presented in Figure 3 on page 14.

Node Operator

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

A node operator can be any of the following entities:

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

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

Node Operator Facility

This component allows communication between the node operator and the control point (CP), intermediate session routing (ISR), and LUs. Node Operator Facility (NOF) initializes the CP and ISR components when the node is started. It also performs functions such as the following when requested to do so by the node operator:

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

Application Transaction Program

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

No matter what the verbs are, they perform the following tasks in a transaction program:

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

Control Point

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

The CP consists of the following components:

Configuration Services

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

Topology and Routing Services

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

Directory Services

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

Session Services

Session services (SS) activates and deactivates CP-CP sessions used by the CP components to exchange network information. Also, SS assigns unique session identifiers to sessions, and assists LUs in activating and deactivating sessions.

Address Space Manager

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

Management Services

Management services (MS) monitors and controls resources of a node. It can generate alerts to the network operator.

Intermediate Session Routing

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

The components of the ISR are:

Session Connector Manager

Session connector manager (SCM) interfaces with the ASM to obtain the LFSID and the transmission group in the direction of the destination LU. Perform intermediate BIND processing.

Session Connector

Session connector (SC) connects two stages of a session. The ISR must receive a frame and forward it to another node that can be the destination or another intermediate node. This is performed using a table that associates a source address in an incoming frame with an address on which the frame must be forwarded. These addresses are based on the LFSID of each session and are carried in the TH field of the SNA frame (specifically in the origin and destination address fields).

Logical Unit

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

Path Control

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

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

Data Link Control

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

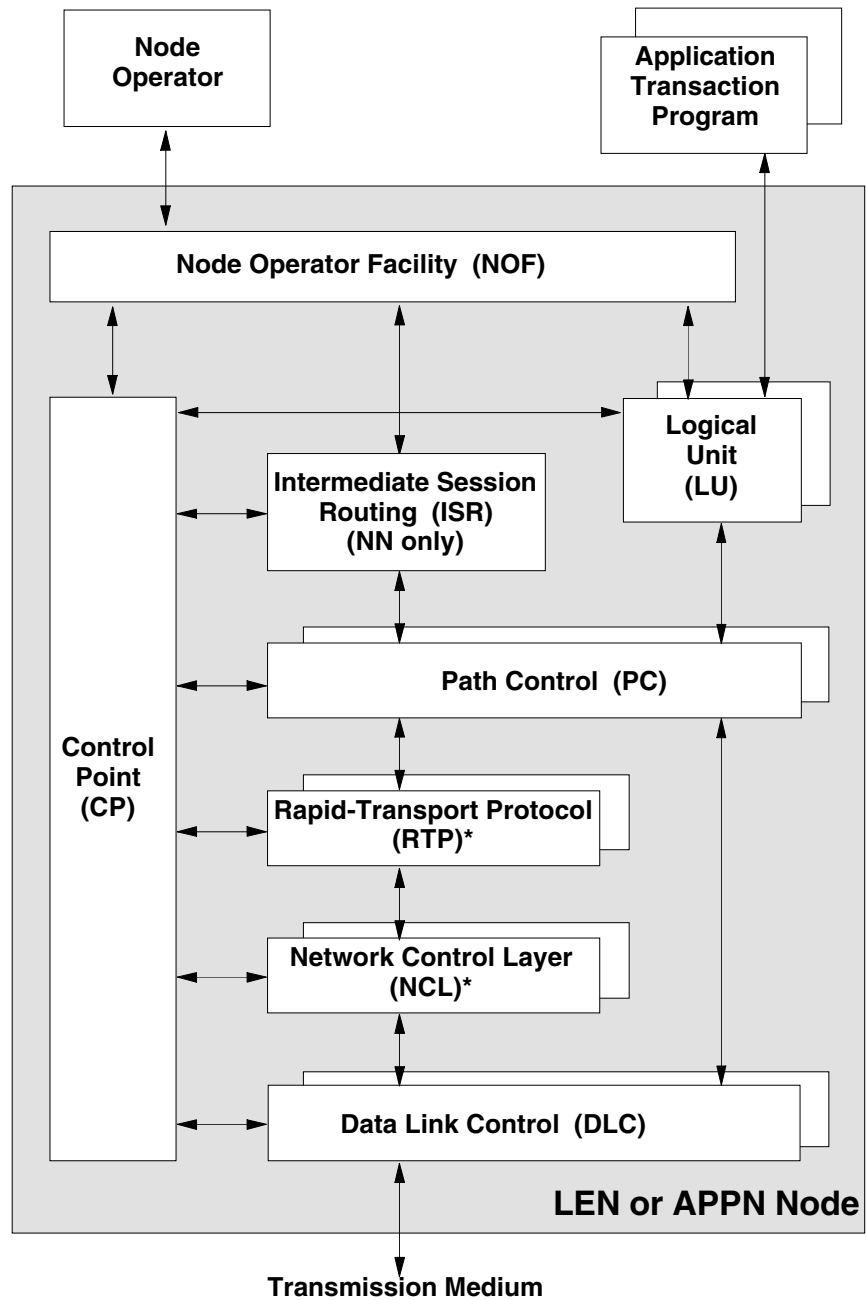


Figure 3. Structure of an APPN or LEN Node

Names

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

Network-Accessible Units

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

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

Network Identifiers

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

Network Names

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

Network-Qualified Names

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

Multiple LEN Nodes with the Same CPName

Concurrently active connections to multiple LEN nodes which are all configured with the same CPName are supported, and which all have dependent LUs which must be served by a DLUS. Suppose the NNP has an active connection to NETA.LEN, which is not an NN and which does not support parallel TGs. Furthermore, suppose an XID is received from another node named NETA.LEN. Without this capability, the NNP will not allow the second XID exchange to complete (it does not break the already active connection). This function allows the exchange to complete, but only if the received XID indicates that parallel TGs are not supported and that ACTPU is required. Mechanically what the NNP does, it ignores the CPName in the received XID, and creates a unique “dummy” CPName for the device for use internally by the NNP. When the XID exchange completes, the NNP passes the whole received XID up to the DLUS, which causes VTAM to search its switched major node for the matching IDBLK/IDNUM and subsequently send the ACTPU and ACTLUs. Furthermore, this function works without any predefinition of the LEN nodes in the NNP. Manually adding all those link station definitions is not a concern.

Addresses

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

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

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

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

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

The session identifier is associated with:

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

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




Figure 4. Session with Several Session Stages

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

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

Domains

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

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

Note: In traditional subarea networking, a domain is the part of the network managed by a VTAM system services control point (SSCP). In this publication, the term *domain*, refers to an APPN domain unless explicitly stated otherwise.

High-Performance Routing

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

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

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

In an HPR network, a new form of routing is used, which is called *automatic network routing (ANR)*. ANR is a source-routing protocol, which means the sender of a packet provides the information about the physical path that the packet will use through the network in the network header. Because HPR provides nondisruptive path switching, the HPR architecture handles route changes in mid-session.

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

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

The process of supporting the end-to-end sessions across the HPR network is known as *Rapid Transport Protocol (RTP)*.

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

HPR Base and Towers

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

Multilink Transmission Group (Option Set 1404)
Dedicated RTP Connections (Option Set 1403)
Control Flows over RTP (Option Set 1402)
RTP Functions for HPR (Option Set 1401)
Base Functions for HPR (Option Set 1400)
APPN End Node or Network Node

Figure 5. HPR Base and Towers

Base Functions (Option set 1400)

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

A new packet format, called a *network layer packet (NLP)*, transports data in the HPR subnet.

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

Prior to establishing an RTP connection, a route setup protocol is used over the desired path. Link and node APPN topology update information indicates the appropriate level of HPR support by means of a new HPR control vector. Base-level nodes participate by adding the appropriate ANR information for the inbound and outbound links. When the route setup messages are exchanged

between two nodes where one or both are base-level nodes, they flow in a FID2 PIU.

RTP Functions for HPR (Option set 1401)

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

- **Rapid Transport Protocol (RTP)**

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

- **Nondisruptive Path Switch**

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

- **APPN/HPR Boundary Function Support**

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

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

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

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

Control Flows Over RTP Tower (Option set 1402)

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

Dedicated RTP Connections (Option set 1403)

See "Rapid Transport Protocol" on page 21.

Multilink Transmission Group (Option set 1404)

See "Multilink Transmission Groups" on page 27.

HPR Link Support

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

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

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

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

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

Automatic Network Routing

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

HPR uses the ANR routing mode to route session traffic, including binds and unbinds, through an HPR network between nodes supporting the High-Performance Routing Transport option.

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

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

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

Rapid Transport Protocol

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

Rapid Transport Protocol provides end-to-end error recovery with selective retransmission, nondisruptive path switching and adaptive rate-based (ARB) flow control. RTP can be implemented on network nodes or end nodes.

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

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

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

Class of Service

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

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

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

ARB Flow Control and Congestion Control

In HPR, sessions with the same CoS are multiplexed over a single RTP connection. The ARB flow control mechanism, used by RTP, ensures the fairness among multiple RTP connections through the network.

To ensure fairness among multiple sessions using one RTP connection, HPR uses not only the ARB mechanism for RTP connections but also the existing session pacing over an RTP connection to prevent one session from monopolizing buffers in the two RTP components. In an APPN sense, the RTP connection can be seen as one stage (hop) on the session path.

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

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

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

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

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

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

The limitations of ARB are:

- Has slow ramp-up, especially for new connections that have some fairness problems
- Does not coexist well with TCP/IP (this is based on its extreme reaction of small losses)
- Does not work well with high-speed links (T3 and beyond), since the ARB mechanism is designed for lower-speed links

ARB2 is an improved version of the original ARB. ARB2 delivers higher performance in several key areas of own-control aspects:

- Faster ramp-up speeds for transmission startup and better support for short transmissions
- Better bandwidth allocation fairness
- Better support of high-speed links

The operational principles behind the ARB2 algorithm are:

- Transmission startup should get fast ramp-up performance
- Transmissions which are short should get BW “priority”
- Transmissions which are sustained should get fair BW allocation
- Steady-state link utilizations should be as high as possible
- The algorithm should be tolerant of a tunable amount of loss
- The algorithm should be programmable to accommodate implementation idiosyncrasies and coexistence with other protocols
- The algorithm should “self-calibrate” as much as necessary to account for system unknowns
- The challenge has been in blending all of these characteristics using all the necessary compromises

The core concept that the ARB2 algorithm implements is that the sender of On/Off traffic patterns needs to be controlled dynamically. Mathematical models of On/Off systems can be employed to provide a framework for understanding the physics of the system and to calculate the control parameters that are embedded in the ARB2 algorithm. The mathematical model provides accurate estimates of the algorithm’s key parameter: the receiver threshold where the sender is instructed to decrease its average transmission rate. The model also provides a method for determining the mathematical connection between the various control parameters, such as the rate increase and decrease amounts.

Subsequent calculations providing the analytic connections between the control settings follows from the basic model (see Figure 6 on page 24). The mathematical model provides a framework in which the dynamics of the queue length trajectory can be understood and thereby controlled, providing an essential tool for *adjusting* the algorithm.




Figure 6. Basic Multiplexer System

Consider a simple system comprised of N HPR/RTP senders sharing a link as shown in Figure 6. All operating *statically* in the sense that their average transmission rates are static. Although the analysis of this system would seem straightforward, in actuality, senders do not submit traffic at constant rates: traffic is submitted in On/Off patterns. See Figure 7 on page 25 for sender traffic patterns that occur with ARB traffic sources. In Figure 7 on page 25 there are three key pattern parameters:

- The peak rate
- The cycle-average rate
- The on time (or burst length)

After looking at many On/Off system studies, it is clear that the key to controlling congestion periods is controlling these three parameters of the sender's traffic.

However, neither the peak rate nor the burst length can be controlled. Therefore it is generally established by the sender's link speed, and the burst length should not become too small due to system efficiency reasons, such as packet overhead in transmission. But, because the cycle-average rate is determined by the on-period spacing, it is possible to change the average rate in order to control burst spacing.

In general, On/Off system losses can be severe even when the average link utilization is small (as little as 5%). This is due to coincidence probabilities in the burst arrival patterns and the subsequent duration of the coincidence relative to the queue capacity. As burst periods become less frequent or more spaced out, coincidence probabilities generally decrease, unless burst lengths increase at the same time.

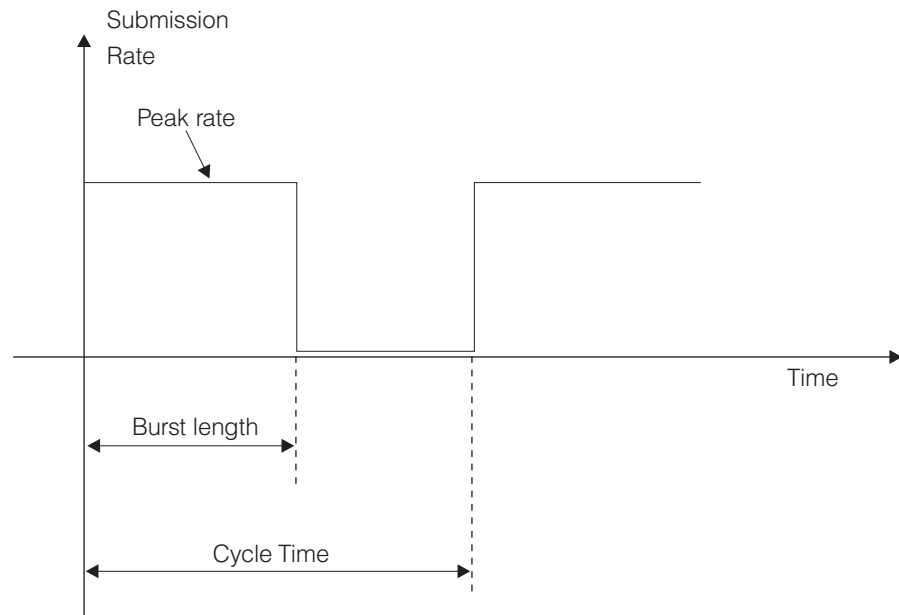


Figure 7. Burst Timing and the Three Parameters

The main principle behind a rate control scheme with ARB senders that have On/Off patterns is that the On/Off patterns must be controlled to provide high-link utilization with limited loss probabilities. As the On periods become more spaced out, the loss probabilities decrease, but so does link utilization. Contrarily, when On periods become too close, loss probabilities increase.

The role of the ARB2 algorithm is to determine dynamically the proper spacing between On periods. Determining the average sending rate is simply the mechanism used to achieve the control and is not necessarily the fundamental goal of the control problem. In fact, there are network scenarios in which the total of all average rates is small when compared to the shared link rate and still losses occur. This is due to the bursty nature of the On/Off patterns. These scenarios show that achieving the proper average rates is not necessarily the fundamental goal. In systems with smooth traffic shaping provided by the senders, the goal is different. The fundamental control goal is determining the sending rate, since this can be achieved for all network scenarios and still have high link utilizations. In that case, there are virtually no coincidence probabilities that cause queues to overflow, as they only overflow when the sum of the sending rates exceeds the shared link rate.

The approach taken in ARB2 is to have senders decrease their On-period spacing until waiting times in the shared queue exceed a certain value. The value analyzes a system with N senders sharing a link running at high utilization (for example, 90%). For a system with N senders, the statistics of the waiting times can be calculated analytically. If the measured waiting time exceeds a percentile (determined analytically), then it is clear that the sender On-period spacing has gotten too small, so the sender increases the spacing.

With this as the guiding principle, the ARB2 receiver simply monitors delays on the connection, and, when the measured waiting time exceeds a prior determined percentile, it sends slowdown messages to the sender. When the sender receives a slowdown message, the On-period spacing is increased by decreasing the average rate. (The other details on fairness and rate increase speeds are not

considered since they are details on achieving various qualities of ARB2 behavior and not fundamental to the principle of ARB2.)

The receiver continues to send slowdown messages for as long as the measured waiting time exceeds the prior determined percentile. This is done because the actions of other senders are not immediately known and serious congestion must be avoided. In addition, when the measured waiting time drops below the percentile, the receiver sends an increase message to the sender, thus causing it to decrease its On-period spacing. It does this rather than waiting for the measured waiting time to decrease to a lower value, because it is necessary to react strongly to the measured congestion (with potentially several slowdown messages) that may have been an over-reaction. The goal again is to find the proper On-period spacing so that the system operates near the proper point on the link utilization curve.

In conclusion, the receiver measures the waiting line experienced by connection traffic and when it exceeds an amount set using an analysis of the On/Off system, it sends a slowdown message to the sender. This causes the system to recover toward the originally intended load point. (The load point is set offline and is subject to all the physics of On/Off systems.) If the operating point was too large, subsequent measurements will drive the system again towards the operating point. The result of this method is to provide a system that attempts to operate near its intended load point. If it is not near the load point, it takes action (in a distributed manner) to drive towards the intended load point. Due to ARB2 algorithm actions, it continues to drive towards the intended load point when it varies from this point. In addition, to assure that new connections are allowed to join the system, the continual small oscillation around the load point provides the opportunity for senders to learn about load changes.

Note: The notion of reaching a state in which no change occurs might not be productive in that it limits the speed at which adaptations to new loads can occur. Small oscillation around the intended load point is a useful element of the algorithm design.

As a result of this behavior, ARB2 can be seen to vary the On-period spacing continually in an attempt to keep system operation near the prior chosen operating point. The operating point is not straightforwardly chosen since this is a system with bursty traffic submission.

RTP Alive Timer

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

The purpose of the Alive timer is as follows:

- It keeps limited resource links active, where limited resource links are automatically deactivated in HPR when no traffic flows over the link for a

specified period of time. When there is no session traffic, RTP sends *liveliness* messages at intervals set by the Alive timer.

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

Nondisruptive Path Switch

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

Multilink Transmission Groups

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

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

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

HPR MLTG Requirements

MLTGs have advantages over single-link TGs and parallel TGs in a number of cases:

Where the traffic demand can exceed existing TG capacity

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

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

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

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

Where individual links are unreliable

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

Where you have a subarea network including multilink transmission groups

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

Additional design objectives of the MLTG architecture include:

- The need to support mixed link types within MLTGs. All supported SNA link types are also supported in HPR MLTGs.
- The need to support mixed link speeds within MLTGs.
- The need to minimize system definition.

HPR MLTG Overview

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

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

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

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

As regards routing and ANR labels, MLTGs are treated the same as single-link TGs. See “Automatic Network Routing” on page 20. An MLTG is assigned one ANR label for each direction.

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

Some functions are not supported in HPR MLTG, including:

- Limited resource
- Connection networks
- Non-activation XID

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




Figure 8. Multilink and Parallel TGs

Priority

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

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

Route Calculation

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

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

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

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

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

To take full advantage of the HPR function:

- Upgrade adjacent APPN nodes. HPR-HPR-APPN routes are better than HPR-APPN-HPR routes.
- Add links with the heaviest traffic, for example backbone links, to the HPR network first.

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

HPR-Only Route for Path Switch

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

Timers

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

- Alive timer
- Short_req timer
- Path switch timer

Each is described in more detail below.

Alive Timer

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

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

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

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

Short_req Timer

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

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

$$\text{Short_Req} = (0.9 \times \text{Previous_Short_Req}) + (1.8 \times \text{RTT}).$$

Path Switch Timer

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

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

Migration

The following features of HPR ease migration from APPN:

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

Each is described in more detail in the following sections:

Interoperation with Existing APPN Nodes

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

No Configuration Restrictions

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

HPR Uses APPN Control Point Protocols and Algorithms

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

Shared Topology

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

Migration Planning

The benefits appear when you have at least one HPR subnet operating, see "Route Calculation" on page 30. You could start by nominating for HPR operation either:

- The longest chain of HPR capability.
 1. List all HPR-capable stations and nodes.
 2. Starting from a VTAM host, work outwards to the furthest HPR-capable unit. That becomes your first HPR subnet, or chain of subnets.
 3. Using *CCM User's Guide*, SH11-3081, define all the required 3746 Network Nodes and stations as HPR-capable.
- The links with the heaviest traffic, either by current or predicted performance, work outwards to adjacent APPN nodes.

Dependent LU Requester/Server

The dependent LU requester (DLUR) and dependent LU server (DLUS) functions allow SSCP-PU and SSCP-LU data to flow through an APPN network (see Figure 9 on page 33). The session establishment data for dependent LUs flows on top of two APPN LU6.2 sessions. (For each DLUR and DLUS, there is one *contention winner* session and one *contention loser* session.) From the DLUR function to the dependent PU, the flow is the same of the NCP or VTAM boundary function to a PU (ACTPU, ACTLU, BIND, and so on).

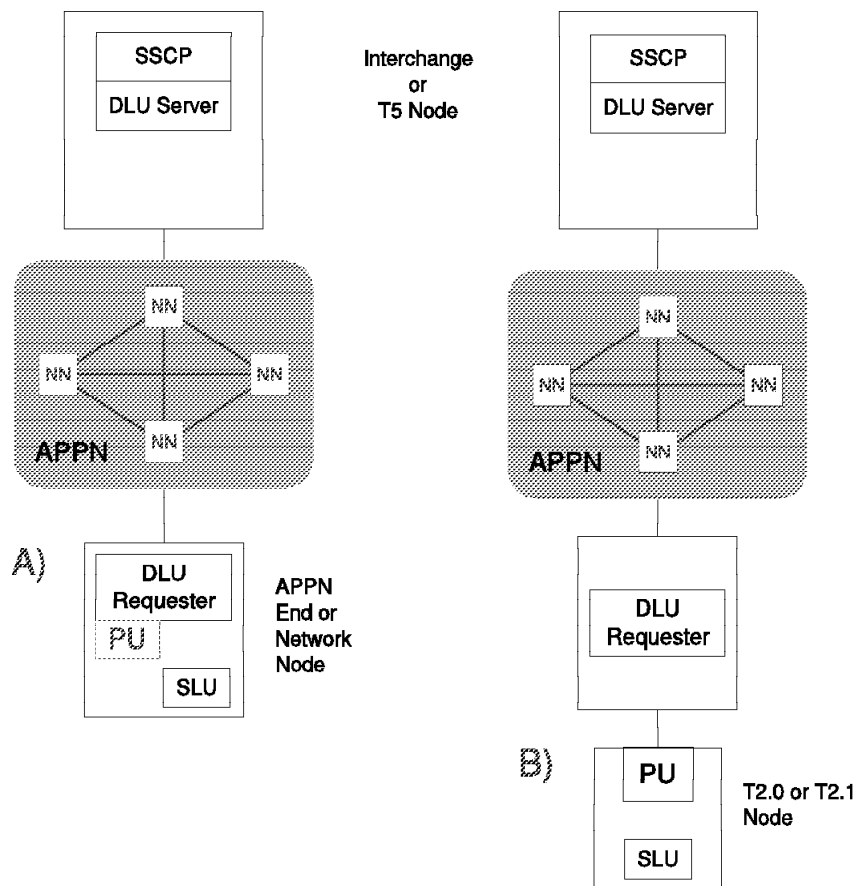


Figure 9. Dependent LU Requester/Server.

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

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

The dependent LU requester function (option set 1067) is an enhancement to an APPN end node or network node. This function is the client side of the dependent

LU server function in which SSCP-PU and SSCP-LU flows to a PU T2.0 or APPN or LEN node attached to the requester are encapsulated within LU 6.2 sessions as mentioned previously.

The requester function provides a remote boundary function for dependent LUs. This option set relieves the restriction that PU T2.0 nodes be directly attached (or bridged, or data link switched, or frame relayed) to the VTAM or NCP boundary function. The dependent LU requester function can reside in the same node as the secondary LU or be provided by a node adjacent to and upstream from the secondary LU (see Figure 9 on page 33).

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

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

3746-9x0 APPN/HPR Network Node Implementation

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

Terminology and Implementation Specifics

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

Figure 10 on page 35 depicts how the APPN functions are split up between the network node processor (NNP) and the adapters within the 3746 frame.

Note: Adapter means the CLP, TRP2, ESCP2, or CBSP2 processor and the associated line interface (LIC), token-ring interface (TIC3), Ethernet interface (TIC3 bridged), or ESCON® (ESCC) couplers.

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




Figure 10. 3746 NN Structure. A full APPN stack is composed of functions performed on the network node processor (NNP) and within the 3746-9x0.

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

- Node operator facility (NOF)
- Topology and routing services (TRS)
- Directory services (DS)
- Configuration services (CS)
- Session services (SS)
- Dependent LU requester (DLUR)

APPN functions performed within the 3746 adapters are:

- Data link control (DLC)
- Path control (PC)
- Address space manager (ASM)
- Session connector*

* SCM = session connector manager
SC = session control (in SNA)

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

Session Establishment and Routing

During APPN session establishment, CP functions on the NNP participate in locating session partners and are responsible for APPN route calculation. Figure 11 on page 36 depicts how CP-CP session data flows between the NNP and the control points of adjacent nodes.

Irrespective of the coupler that the node is connected to, CP-CP data will always traverse:


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

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

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

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

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



Legend:

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

Figure 11. CP-CP Sessions

The BIND, which is the first SNA request unit flowing on the newly calculated route between two session partners, will trigger the address space manager (ASM) function running on the 3746 processors to assign LFSIDs. In addition, a session connector will be generated to enable intermediate session routing on the 3746 NN. The session connector can be an intraprocessor (within the same 3746 processor), or an interprocessor (between two different processors connected through the 3746 connectivity switch [CS]). See Figure 12 on page 37.

Figure 13 on page 38 illustrates the data flows during and after session establishment. End node A (EN A) is token-ring-connected to 3746 NN, while end node C (EN C) is SDLC-connected. In both cases the 3746 NN (NN B) is providing the network node server function, having CP-CP sessions with both end nodes.




Figure 12. Intermediate Session Routing

To locate the session partner and calculate the best session path, APPN functions within the NNP are invoked. Initiated by the BIND, CP functions available on the 3746 processors will assign local form session identifiers (LFSIDs) for this session and generate a session connector (SC). Note that for this session an interprocessor SC applies. If both EN A and EN C were connected to couplers controlled by the same processor, an intraprocessor SC would result.

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




Figure 13. 3746 NN Intermediate Session Routing. All shaded components are involved in session setup. The dark-shaded components are also involved in intermediate session routing.

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

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

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

3746 Network Node Processor Backup

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

Network Node Connectivity

Refer to the *3745 Communication Controller Models A and 170, 3746 Nways Multiprotocol Controller Models 900 and 950: Overview*, GA33-0180 for information about:

- The number of PUs, frame-relay DLCIs, and sessions available on the 3746-9x0
- The total number of PUs, APPN and dependent LU sessions, and lines that a 3746 Network Node can handle (no matter what type 2 adapter configuration is used)

DLC Support of Error Recovery

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

- Using link-level error recovery

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

- Not using link-level error recovery

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

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

Token-Ring

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

Frame Relay

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

SDLC

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

X.25

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

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

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

Dependent LU Requester Implementation




Figure 14. 3746 Network Node DLUR Connectivity

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

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

- For the DLUR to initiate activation, through a call-in (to VTAM), no system definitions are required for existing token-ring LANs. Refer to the “Token-Ring Configuration” chapter in the *3745/3746 Planning Series: Token Ring and Ethernet*.

For SDLC lines, new definitions are needed in VTAM. Even if the SDLC line is leased, it must be referenced using a switched major node definition. Refer to “3746 SDLC Support” chapter in the *3745/3746 Planning Series: Serial Line Adapters*.

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

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

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

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

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

- Multiple DLURs can support the same PU, but only one at a time. The SSCP-PU session is through only one DLUR at any given time.
- Redial occurs as follows for DLUS-supported PUs:
 - When a DLUS initiates the activation of a PU but receives a negative response, VTAM attempts to redial over each valid path statement for the PU until successful, or until all valid paths have been tried. Redial does not occur if the negative response indicates that the PU is already active, or that the fully qualified procedure correlation identification (PCID) is not unique. If the fully qualified PCID is not unique, the DLUS attempts to redial the PU over the same path with a newly generated PCID.
 - Redial does not occur if PU activation was initiated by the DLUR.

- When a protocol violation, topology database update (TDU) error, or CPSVRMGR session outage signal is received for a particular DLUR, VTAM attempts to redial every active or pending-active PU served by that DLUR for which a valid PATH statement is found. If the PU is already active, then VTAM performs the following actions:
 - If the PU was defined with ANS=CONT and the DLUR supports this function, VTAM performs giveback processing prior to attempting redial.
 - If the PU was defined with ANS=STOP or the DLUR does not support ANS=CONT, VTAM deactivates the PU prior to attempting redial.
- When a CP-CP session between a DLUS and DLUR fails, the CPSVRMGR session is deactivated, enabling reactivation of this CP-CP session.
- Information regarding DLUR-attached resources can be processed by the configuration services XID exit. This exit routine can be coded so that VTAM processes or denies requests for contact from known switched devices or so that VTAM processes or denies requests for PU activation from a DLUR. For details, refer to the *VTAM Customization*, LY43-0063. (Available to IBM licensed customers only).

VTAM DLCADDR Keyword in PATH Statement

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

The format of the DLCADDR keyword is:

DLCADDR=(element,data_format,element_value)

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

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

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

For VTAM V4 R2: PTF UW28497
 For VTAM V4 R3: PTF UW28498.

- The third suboperand defines the value of the element.

An example of DLCADDR encoding is given here:

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

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

Migration of SNI Connections to 3746 Models 900 and 950

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

- A peripheral border node (PBN) allows access to adjacent APPN/HPR networks.
- An extended border node (EBN) allows access to non-adjacent APPN/HPR networks.

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

Therefore, if your “partners” do not migrate to APPN but keep using SNI, you must keep a 3745/NCP that will handle your SNI traffic, while the coupled 3746-900 Network Node handles your APPN traffic. As shown in Figure 15 on page 44, the SNI traffic can flow over 3746 adapters to the NCP.

One exception to the above statement is when an SNI network is connected to a 3745 that is part of a CNN. Figure 31 on page 60 shows such a configuration. The 3745/NCP handles the SNI flows and routes these flows through an internal APPN link to the 3746-900 network node. Figure 27 on page 56 shows an example of how an internal APPN link works. The TRP2 that handles the APPN link also contains the HPR/RTP function. RTP converts APPN (ISR) flows from the SNI connection to HPR flows, which can then be routed through the APPN network as any other HPR flows.

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

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

This does not affect the SNI connections on the 3746-900; they stay under control of the NCP running in the 3745 CCU. See Figure 15 on page 44.

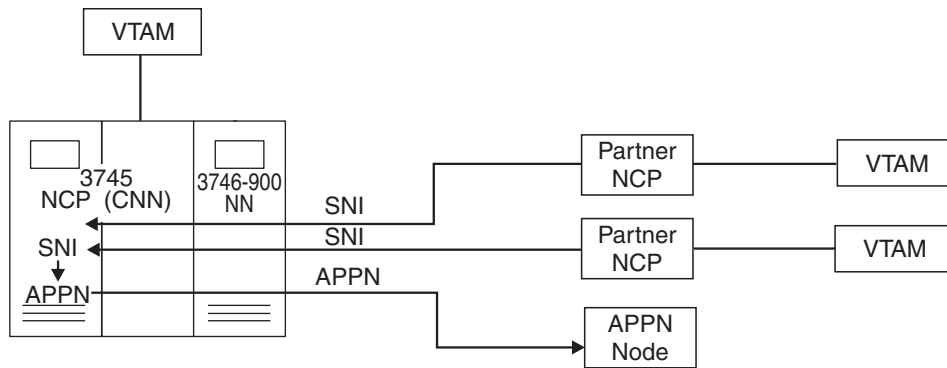


Figure 15. Example Network: Adding APPN Before SNI-to-APPN Migration

If there is a need to route traffic between an SNI port (or any SNA/subarea port of the 3745/3746-900) and an APPN port controlled by the 3746-900 network node, then the NCP must be defined as part of a composite network node (CNN). NCP does the SNI-to-APPN and APPN-to-SNI conversions for the communication with the 3746-900 network node. This communication uses the coupling between the 3745 and the 3746-900, (see “Internal APPN Connection between a 3745 and a 3746-900” on page 55).

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

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

In the partner NCP, the connection is border node, and in the 3746-900, the line will be defined in the 3746-900 network node instead of NCP. See Figure 16.

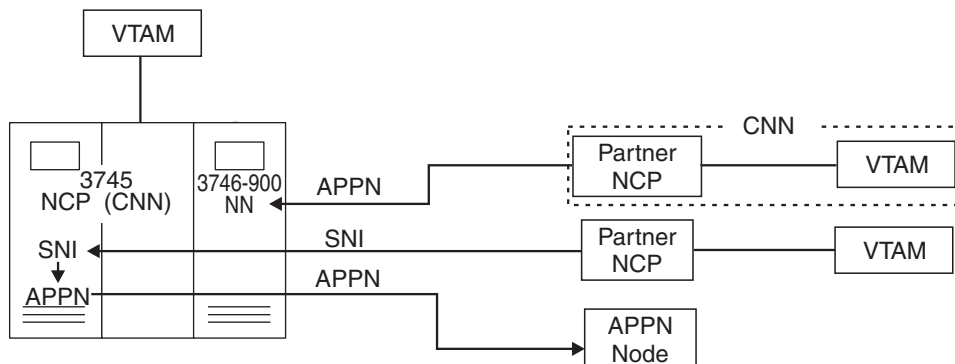


Figure 16. Example Network: SNI-to-APPN Migration

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

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

This can be accomplished when there are no more SNI links or other requirements for NCP. See Figure 17 on page 45.

Note: Without Session Services Extensions (SSE) support in the 3746 NN, this configuration supports independent sessions across the network boundary

between the 3746 NN and a CNN providing the extended border node (EBN) function. Dependent sessions across this boundary are not supported. The support of dependent sessions requires the 3746 NN directly adjacent to the CNN to support SSE. This is supported by the 3746 and VTAM (which has SSE support) acting together as a CNN.

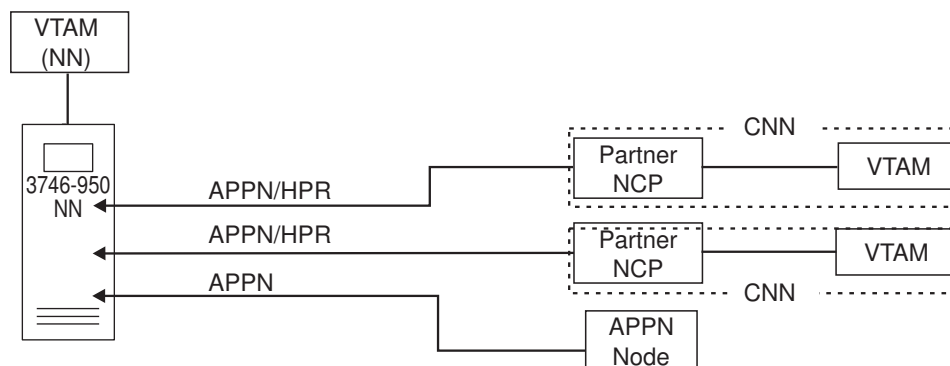


Figure 17. Example Network: Upgrading 3746-900 to a 3746-950

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

EBN/SSE Usage Rules

This section gives examples of independent and dependent connections across subnetwork boundaries using the extended border node (EBN) function.

VTAM, the IBM 2216 Nways Multiaccess Controller, and the IBM 3746 Multiaccess Enclosure support all of EBN.

The 3746 supports only the Session Services Extensions (SSE) part of EBN.

Networks with DLURs: PLU, DLUS, and DLUR in Three Different Subnetworks




Figure 18. Networks with DLURs: PLU, DLUS, and DLUR in Three Different Subnetworks

Table 3. Networks with DLURs: PLU, DLUS, and DLUR in Three Different Subnetworks		
Path	Required Nodes	Comments
DLUS-PLU	1 EBN on <i>each side</i>	
DLUS-DLUR	1 EBN on <i>dependent logical unit server (DLUS) side</i>	SSE not needed on dependent logical unit requester (DLUR) side.
PLU-DLUR	1 EBN on <i>primary logical unit (PLU) side</i>	SSE not needed on DLUR side.
	An <i>extended boundary</i> connection is needed between networks A and C.	More than one extended boundary is allowed.
	A <i>peripheral boundary</i> connection is needed between networks B and C.	Only <i>one</i> peripheral boundary is allowed: when one boundary (peripheral or extended) exists, a second boundary (peripheral or extended) is not allowed between networks A and B.

Networks with DLURs: DLUS and DLUR in Same Subnetwork




Figure 19. PLU and DLUR in Same Subnetwork

Table 4. PLU and DLUR in Same Subnetwork		
Path	Required Nodes	Comments
DLUS-PLU	1 EBN on <i>each side</i>	
DLUS-DLUR		SSE not needed, DLUS and DLUR are both in same network.
PLU-DLUR	1 EBN on <i>PLU side</i>	SSE not needed on DLUR side.
	An <i>extended boundary</i> needed between networks A and B.	A peripheral boundary is not allowed between networks A and B in this case.

Networks with DLURs: PLU and DLUS in Same Subnetwork




Figure 20. PLU and DLUS in Same Subnetwork

Table 5. PLU and DLUS in Same Subnetwork		
Path	Required Nodes	Comments
DLUS-PLU		SSE in DLUS and PLU only needed for: <ul style="list-style-type: none"> • Secondary logical unit (SLU)-initiated sessions • Sessions requiring queuing (initiate-only, initiate-or-queue, and queue-only queuing) • Third-party initiated (TPI) sessions • Auto-logon (initiate-or-notify initiate)
DLUS-DLUR	1 EBN on <i>DLUS side</i>	SSE not needed on DLUR side.
PLU-DLUR	1 EBN on <i>PLU side</i>	SSE not needed on DLUR side.

Networks with DLURs: PLU and DLUR in Same Subnetwork




Figure 21. PLU and DLUR in Same Subnetwork

Table 6. PLU and DLUR in Same Subnetwork		
Path	Required Nodes	Comments
DLUS-PLU	1 EBN on <i>each side</i>	
DLUS-DLUR	1 EBN on <i>DLUS side</i>	SSE not needed on DLUR side.
PLU-DLUR		SSE not needed, both PLU and DLUR are in same network.
	An <i>extended boundary</i> needed between networks A and B.	A peripheral boundary is not allowed between networks A and B in this case.

Networks with DLURs: PLU in Subarea Network, DLUS and DLUR in Same APPN Network




Figure 22. PLU in Subarea Network, DLUS and DLUR in Same APPN Network

Table 7. PLU in Subarea Network, DLUS and DLUR in Same APPN Network		
Path	Required Nodes	Comments
DLUS-PLU PLU-DLUR	The subarea VTAM, interchange node (ICN), and DLUS must support SSE.	The “owning” control point (CP) and network node server (NNS) must both support SSE. In this example, the CP (PLU) is a subarea VTAM and the NNS (PLU) is the ICN. The DLUS is both the CP and NNS for the SLU. The DLUR does not need an SSE.
DLUS-DLUR		SSE not needed because both are in same network.
	No EBNs needed.	There are no APPN subnetwork boundaries.
		Subarea resources appear to other network nodes as if they are APPN end nodes (ENs) served by the ICN in the path to that resource. In other words, the PLU looks as if it is on an EN served by the ICN.

APPN Networks: 3746-9X0 Used as NNS of VTAM ENs




Figure 23. 3746-9X0 Used as NNS of VTAM ENs

Table 8. 3746-9X0 Used as NNS of VTAM ENs		
Path	Required Nodes	Comments
	By default, CP-CP sessions are not set up by VTAM.	VTAM assumes that it has an NNS that supports SSE. To have CP-CP sessions set up, define a NETSERVER list in the VTAM EN and code SLUINIT=OPT for every 3746 in the list.
	Only PLU-initiated sessions can be set up through the 3746 without SSE. SSE required for: <ul style="list-style-type: none"> • Secondary logical unit (SLU)-initiated sessions • Sessions requiring queuing (initiate-only, initiate-or-queue, and queue-only queuing) • Third-party initiated (TPI) sessions • Auto-logon (initiate-or-notify initiate) 	

APPN Networks: 3746-9X0 Connected to a VTAM EBN




Figure 24. 3746-9X0 Connected to a VTAM EBN

Table 9. 3746-9X0 Connected to a VTAM EBN		
Path	Required Nodes	Comments
PLU (subnetwork A)→ SLU (subnetwork B)	If the 3746 is not SSE capable , EBN changes 'Search/Find' to 'Search Only' before sending it to the 3746.	
PLU (subnetwork B)→ SLU (subnetwork A)	If the 3746 is not SSE capable , EBN should change 'Search/Find' to 'Search Only' before sending it to the PLU.	
	<p>Only PLU-initiated sessions can be set up through the 3746 without SSE.</p> <p>SSE required for:</p> <ul style="list-style-type: none"> • Secondary logical unit (SLU)-initiated sessions • Sessions requiring queuing (initiate-only, initiate-or-queue, and queue-only queuing) • Third-party initiated (TPI) sessions • Auto-logon (initiate-or-notify initiate) 	

Mixed APPN and Subarea Networks: 3746-9X0 between EBN and NCP/Subarea (with SLU)




Figure 25. 3746-9X0 between EBN and NCP/Subarea (with SLU)

Table 10. 3746-9X0 between EBN and NCP/Subarea (with SLU)		
Path	Required Nodes	Comments
SLU-SSCP		Subarea path has a NNS that supports SSE.
SSCP-PLU PLU-SLU	<p>1 EBN on <i>either side</i></p> <p>SSE required for (the 3746 is the NNS of DLU for sessions that originate in network A):</p> <ul style="list-style-type: none"> • Secondary logical unit (SLU)-initiated sessions • Sessions requiring queuing (initiate-only, initiate-or-queue, and queue-only queuing) • Third-party initiated (TPI) sessions • Auto-logon (initiate-or-notify initiate) 	

Mixed APPN and Subarea Networks: 3746-9X0 between EBN (with SLU) and NCP/Subarea




Figure 26. 3746-9X0 between EBN (with SLU) and NCP/Subarea

Table 11. 3746-9X0 between EBN (with SLU) and NCP/Subarea		
Path	Required Nodes	Comments
SSCP-PLU PLU-SLU	<p>1 EBN on <i>either side</i></p> <p>SSE required in 3746. The 3746 is the NNS originating logical unit (OLU) for sessions the originate in network A.</p> <p>SSE required for:</p> <ul style="list-style-type: none"> • Secondary logical unit (SLU)-initiated sessions • Sessions requiring queuing (initiate-only, initiate-or-queue, and queue-only queuing) • Third-party initiated (TPI) sessions • Auto-logon (initiate-or-notify initiate) 	

Internal APPN Connection between a 3745 and a 3746-900

On a 3745/3746-900, the controller bus coupler (CBC) provides an internal connection (an APPN transmission group) between the 3746-900 network node and the 3745 (NCP) operated as a composite network node (CNN). This internal APPN link is set up using a 3746-900 token-ring port. Any token-ring port, except the TIC3 port of the CBSP2 (port address 2080), can be used (see Figure 27 on page 56). The TRP2 for this TIC3 performs the APPN routing between the CBC link and the other ports of the 3746 Network Node.

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

Example of an Internal APPN Link

Figure 27 on page 56 shows an example of a 3745/3746-900 configuration using the internal APPN communication between the 3746 Network Node and the CNN. The components are:

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

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

3746-900 Configuration Requirements

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

- 3745 with a single CCU (Models 17A, 21A, and 31A)
One TIC3 on a TRP2 must be available.
- 3745 with two CCUs (Models 41A, and 61A):
 - If only **one** NCP communicates with the 3746 Network Node at a time (for example, in the “3745 twin-standby mode”), one TIC3 on a TRP2 must be available, for example, on the TRP2s connecting CCU B.
 - If **two** NCPs communicate at the same time with the 3746 Network Node, two TIC3s must be available. For load distribution and backup reasons, they should preferably be on two different TRP2s. For example, the internal APPN link to CCU B could use a TIC3 on the TRP2 connecting CCU B and the APPN internal link CCU A could use a TIC3 on another TRP2.

Definitions Required for the Internal APPN Link

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

Important

The CNN SAP value must be different from the NCP SAP (default X'04').

3746 Network Node (CCM Definitions): One token-ring port definition is required for each CCU/3746 Network Node link. Figure 27 is a single link example.

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

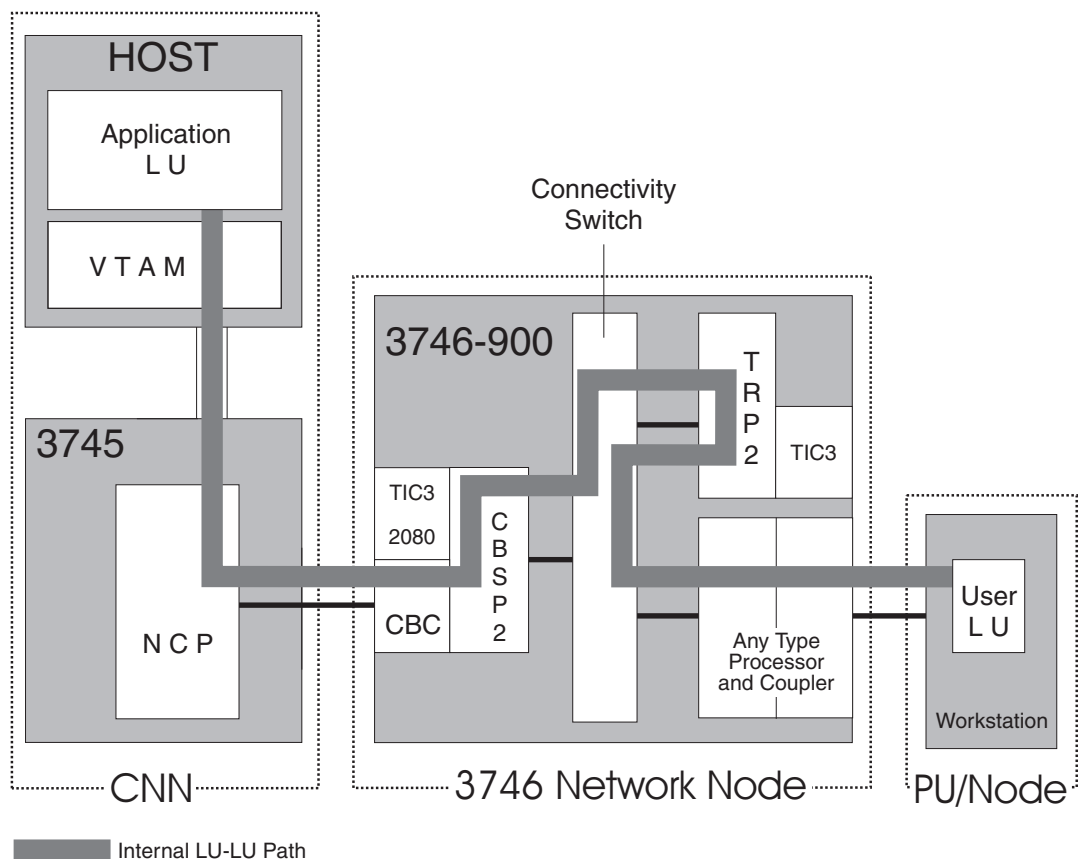


Figure 27. Internal APPN Link (3746-900)

The link station definition depends on what controls the TIC3:

1. If the TIC3 port connects LAN stations controlled **only by the 3746 Network Node, and these stations do not need to be defined in the 3746 Network Node**, then, using the Controller Configuration and Management (CCM User's Guide, SH11-3081) application:
 - Define the TIC3 port with the option "Accept any incoming calls" set to YES.
 - Do not make a link station definition for this TIC3 port.

2. If the TIC3 port connects LAN users controlled **only by NCP in the CNN**, then, using *CCM User's Guide*, SH11-3081:
 - Define the TIC3 port with the option “Accept any incoming calls” set to NO.
 - For this TIC3 port-define NCP as a link station:
 - The locally administered address (LAA) of the station must have the same value as the port LAA.
 - The destination service access point of the station must be set to X'04' (the default value for the NCP service access point).
3. If the TIC3 port connects LAN users controlled **only by the 3746 Network Node and defined in the 3746 Network Node**, then, using *CCM User's Guide*, SH11-3081:
 - Define the TIC3 port with the option “Accept any incoming calls” set to NO.
 - For this TIC3 port define NCP as a link station:
 - The locally administered address (LAA) of the station must have the same value as the port LAA.
 - The destination service access point of the station must be set to X'04' (the default value for the NCP service access point).

CNN (VTAM and NCP Definitions)

- Token-ring (port) definitions in NCP

The token-ring port address is the address of the TIC3 defined in the 3746 Network Node (using the *CCM User's Guide*). The locally administered address must be the same in the NCP keyword LOCADD as in *CCM User's Guide*.

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

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

- Link station definitions in VTAM

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

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

Examples of 3746 Network Nodes in APPN/HPR Networks

This section contains examples of 3746 Network Nodes in APPN/HPR networks as:

- An ANR node, see Figure 28
- An APPN/RTP Node with boundary function (see Figure 29 on page 59)
- A DLUR/RTP Node with boundary function (see Figure 30 on page 59)

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

ADP = adapter
ANR = automatic network routing
DLC = data link control
DLUS = dependent LU server
HPR = High-Performance Routing
ISR = intermediate session routing
RTP = Rapid Transport Protocol.

Example 1: ANR Node




Figure 28. 3746 as an ANR Node

Example 2: APPN/RTP Node with Boundary Function

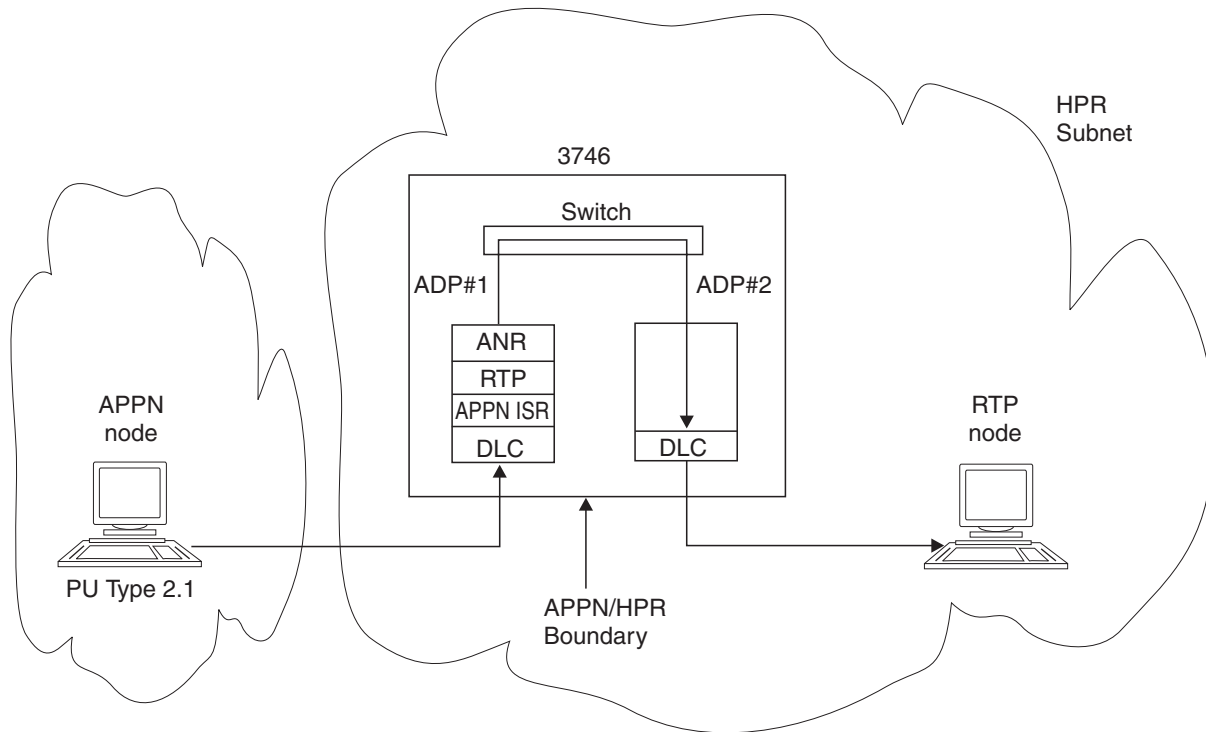


Figure 29. 3746 as an APPN/RTP Node (with Boundary Function)

Example 3: DLUR/RTP Node with Boundary Function




Figure 30. 3746 as a DLUR/RTP Node (with Boundary Function)

Example 4: APPN/RTP Node with SNI Network




Figure 31. 3746 as an APPN/RTP Node with SNI Network

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

A simple way to interconnect a 3745-CNN and the attached 3746-900 Network Node is to use an external connection (token-ring or serial link). Compared both the internal APPN connection (see 55), this connection does not support as much traffic between the 3745-CNN and the 3746-NN, for example traffic between nodes such as ENA and ENB.

For example, in Figure 32, the advantages of a connection network, duplicate MAC addresses, and optimal route computation are achieved by using an “external” token-ring interconnecting the 3745-CNN and the attached 3746-900 NN.

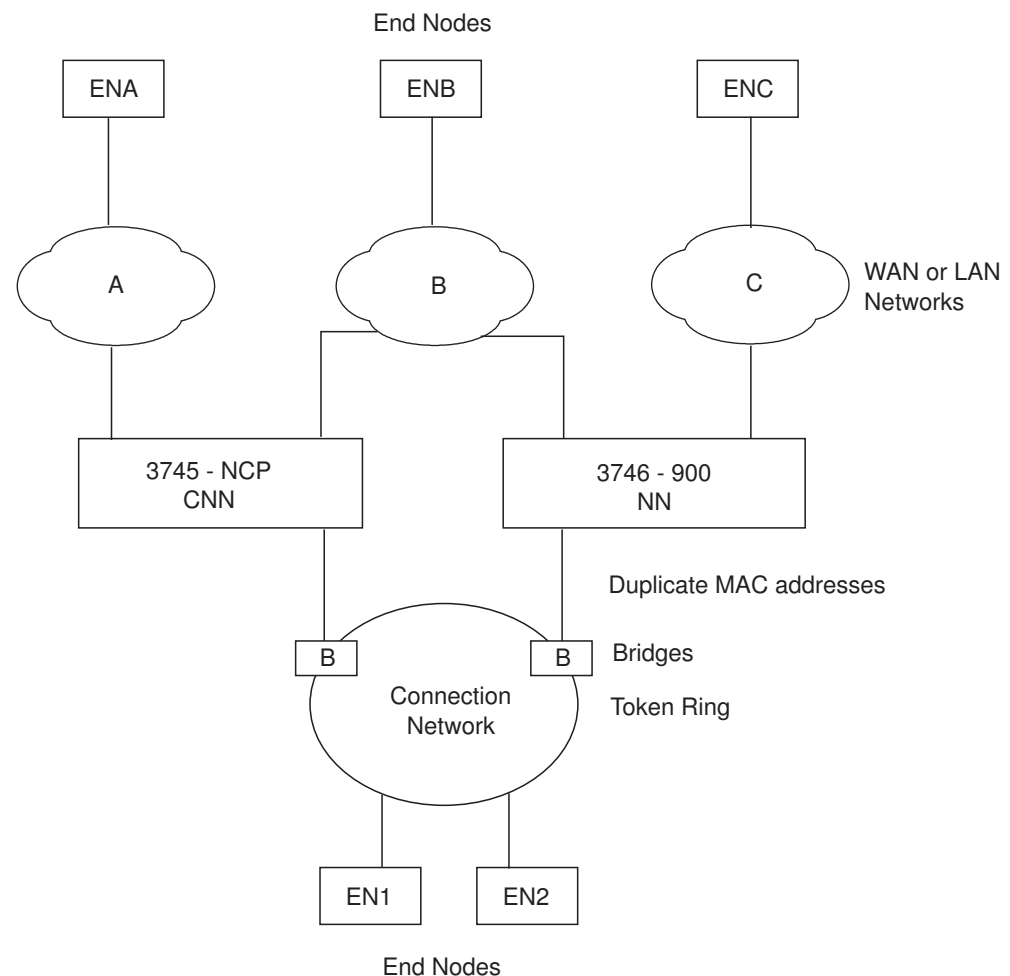


Figure 32. External Connection of 3745-CNN and 3746-NN

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

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

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

Referring to Figure 32, suppose EN1 has CNN as a network node server:

Case 1: EN1 establishes an LU session with EN2

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

Case 2: EN1 establishes an LU session with ENA

As ENA is not reachable by the 3746 NN, the 3745 CNN will send back the RSCV containing the route:

EN1 → CNN → ENA

Here, only CNN is used and the 3746 NN is not involved.

Case 3: EN1 establishes a LU session with ENB

In its route computation, the CNN (as a network node of EN1) will find two possible routes:

EN1 → CNN → ENB

Or

EN1 → 3746 NN → ENB

This is because CNN uses the topology database with all nodes and TGs. CNN will choose the best route that maps the CoS of the required session.

If the best route is:

EN1 → 3746 NN → ENB

EN1 will send the BIND and all data directly to 3746 NN. The CNN is not involved.

Case 4: EN1 establishes a LU session with ENC

This is more or less identical to case 3. CNN finds one route that is:

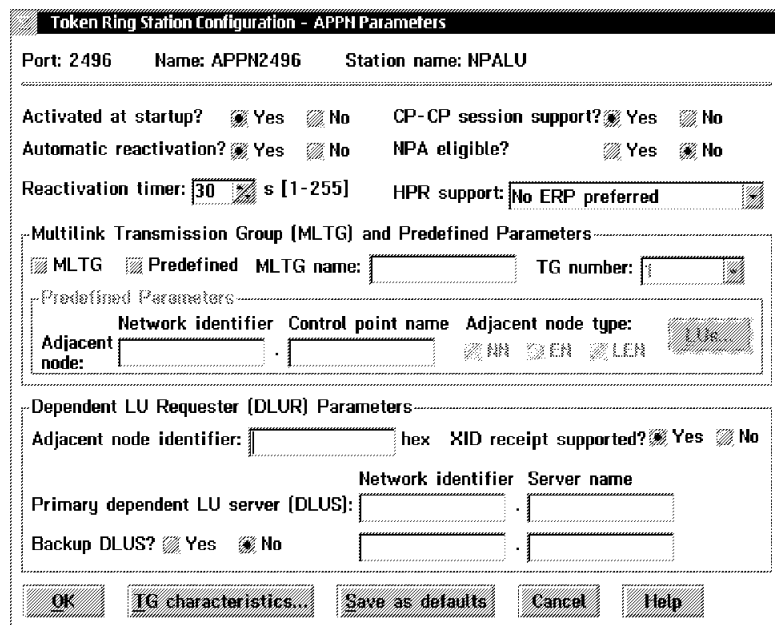
EN1 → 3746 NN → ENC

and thus EN1 will send the BIND and all data directly to 3746 NN. The CNN is then no longer involved.

HPR MLTGs in the 3746

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

HPR MLTGs are defined in CCM from the APPN parameters panel for APPN stations. You can specify the name and TG number (1 - 20) of the MLTG here. Figure 33 shows the MLTG definition panel.



The dialog box is titled "Token Ring Station Configuration - APPN Parameters". It contains the following fields and controls:

- Port: 2496, Name: APPN2496, Station name: NPALU
- Activated at startup? ☒ Yes ☐ No
- Automatic reactivation? ☒ Yes ☐ No
- Reactivation timer: 30 s [1-255]
- CP-CP session support? ☒ Yes ☐ No
- NPA eligible? ☐ Yes ☒ No
- HPR support: No ERP preferred (dropdown menu)
- Multilink Transmission Group (MLTG) and Predefined Parameters:
 - ☒ MLTG ☐ Predefined
 - MLTG name: (text field)
 - TG number: 1 (dropdown menu)
- Predefined Parameters:
 - Adjacent node: (text field)
 - Network identifier: (text field)
 - Control point name: (text field)
 - Adjacent node type: LUa... (dropdown menu)
 - Buttons: NN, EN, LEN
- Dependent LU Requester (DLUR) Parameters:
 - Adjacent node identifier: (text field) hex
 - XID receipt supported? ☒ Yes ☐ No
 - Primary dependent LU server (DLUS): (text field)
 - Backup DLUS? ☐ Yes ☒ No
 - Network identifier: (text field)
 - Server name: (text field)
- Buttons: OK, TG characteristics..., Save as defaults, Cancel, Help

Figure 33. Defining MLTGs in CCM

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

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

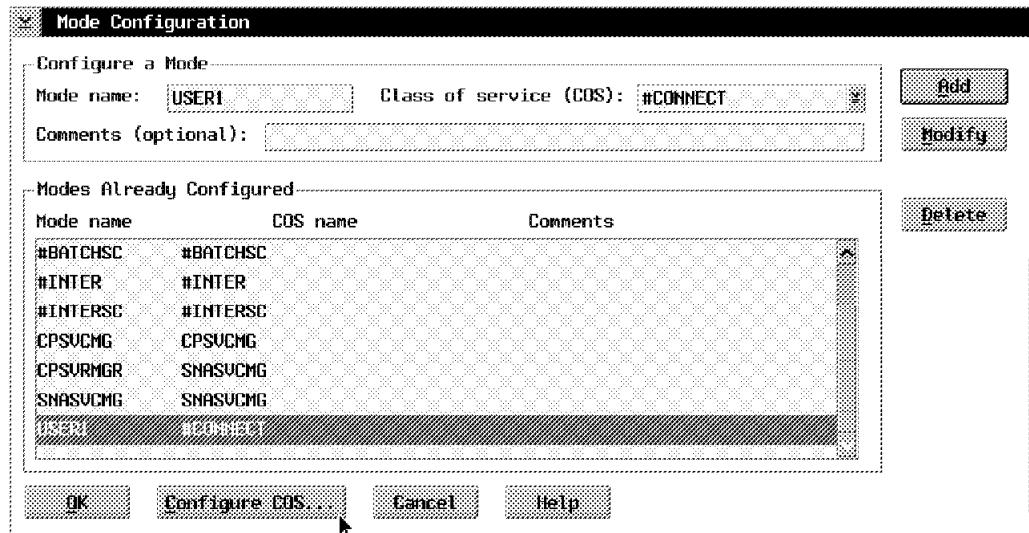
The previous example shows how to define MLTGs at each link station. There is also a central way of defining and checking all MLTG definitions in CCM. From the main CCM panel, select **Configuration** → **APPN** → **MLTG**. This displays a panel where all defined MLTGs and all defined links stations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

User-Defined Parameters

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

1. Class of service and LOGMODE definitions

User-defined parameters are enabled only for user defined class of services (CoS) and logmodes. Changing standard logmodes such as #CONNECT or #INTER would reduce UDP effectiveness. To define your own CoS and LOGMODE, go to the Configuration menu from the CCM main panel, select **APPN** and then select **MODE/CoS** to display the following panel:



The screenshot shows the 'Mode Configuration' dialog box. It has a title bar 'Mode Configuration'. Inside, there's a section 'Configure a Mode' with fields for 'Mode name' (containing 'USER1'), 'Class of service (COS):' (containing '#CONNECT'), and 'Comments (optional):'. To the right of these fields are 'Add', 'Modify', and 'Delete' buttons. Below this is a section 'Modes Already Configured' containing a table with three columns: 'Mode name', 'COS name', and 'Comments'. The table lists several modes: #BATCHSC, #INTER, #INTERSC, CPSUCMG, CPSVRMGR, SNASUCMG, and USER1. The 'USER1' row is highlighted. At the bottom of the dialog are 'OK', 'Configure CoS...', 'Cancel', and 'Help' buttons. A mouse cursor is pointing at the 'Configure CoS...' button.

Mode name	COS name	Comments
#BATCHSC	#BATCHSC	
#INTER	#INTER	
#INTERSC	#INTERSC	
CPSUCMG	CPSUCMG	
CPSVRMGR	SNASUCMG	
SNASUCMG	SNASUCMG	
USER1	USER1	

Figure 34. Mode Configuration Panel

Click **Configure CoS...** to create your own class of services. Enter the CoS Name and click **Add**. Select **TG Row** to specify the different TG weights and associated UDPs ranges. This displays the following panel:

Transmission Group Row Configuration

COS: #USER1

Configure a Transmission Group (TG) Row

Weight: 10 numerical [0-255] Minimum Maximum

Propagation delay: Minimum Maximum

Security: Nonsecure Guarded radiation

Effective capacity (EC): 0 255 numerical [0-255]

Relative cost per byte (RCPB): 0 255 numerical [0-255]

Relative cost per unit of time (RCPUT): 0 255 numerical [0-255]

TG Rows Already Configured

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

Buttons: Add, Modify, Delete, OK, User defined parameters..., Cancel, Help

Figure 35. TG Row Configuration Panel

Enter the TG weight you want to define, and then click **User-defined parameters....** This displays the following panel:

COS/TG User Defined Parameters

COS: #USER1 TG row: 1

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

Buttons: OK, Save as defaults, Cancel, Help

Figure 36. CoS/TG User-Defined Parameters Panel

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

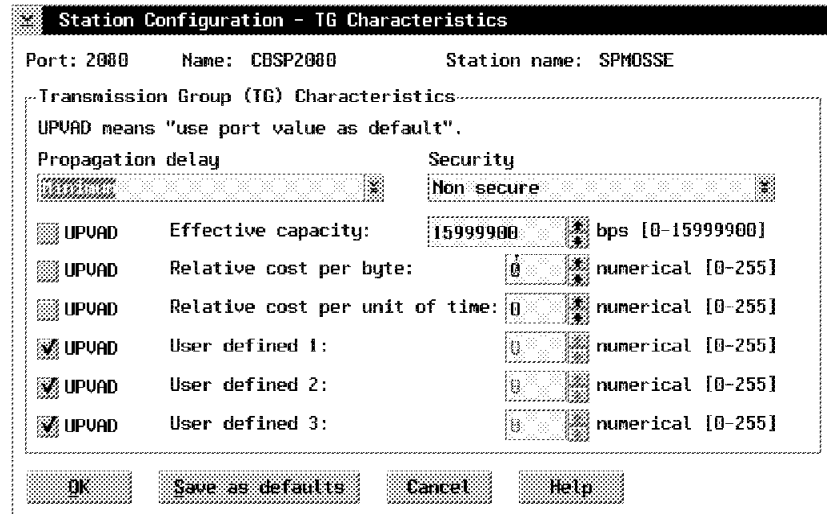
2. UDPs definitions at PORT level

You can define specific UDP values for each APPN port you define from CCM. These values will be compared with the ranges defined in CoS and LOGMODE tables and the corresponding TG weight will be associated with the port. Depending on the UDP values, different weights will be associated with each port, thus allowing APPN to compute the best route.

3. TGs characteristics modification at station level

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

From the Station Definition panel, under CCM, click **APPN parameters...** and then **TG characteristics...** to display the following panel:



The image shows a software window titled "Station Configuration - TG Characteristics". At the top, it displays "Port: 2000", "Name: CBSP2000", and "Station name: SPMOSSE". Below this is a section titled "Transmission Group (TG) Characteristics" with a note: "UPVAD means 'use port value as default'." The panel contains several settings: "Propagation delay" is set to 0; "Security" is set to "Non secure"; "Effective capacity" is 15999900 bps; "Relative cost per byte" is 0; "Relative cost per unit of time" is 0; and three "User defined" fields (1, 2, and 3) are all set to 0. Each of these fields has a checkbox labeled "UPVAD" to its left. The checkboxes for "User defined 1", "User defined 2", and "User defined 3" are checked. At the bottom of the panel are four buttons: "OK", "Save as defaults", "Cancel", and "Help".

Figure 37. Station Configuration - TG Characteristics Panel

Deselect the pushbutton for the parameters you want to alter, and then enter the new parameter values. Click **OK** after you have defined the new values. If **UPVAD** (use port value as default) is selected, then the port parameters are used as default for this link.

Operation of 3746 Network Nodes in SNA/APPN/HPR Networks

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

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

APPN/HPR Network Node

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

The 3746 network node supports the following types of end-node connections (see Figure 38 on page 67):

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

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

Dependent Logical Unit Requester (DLUR)

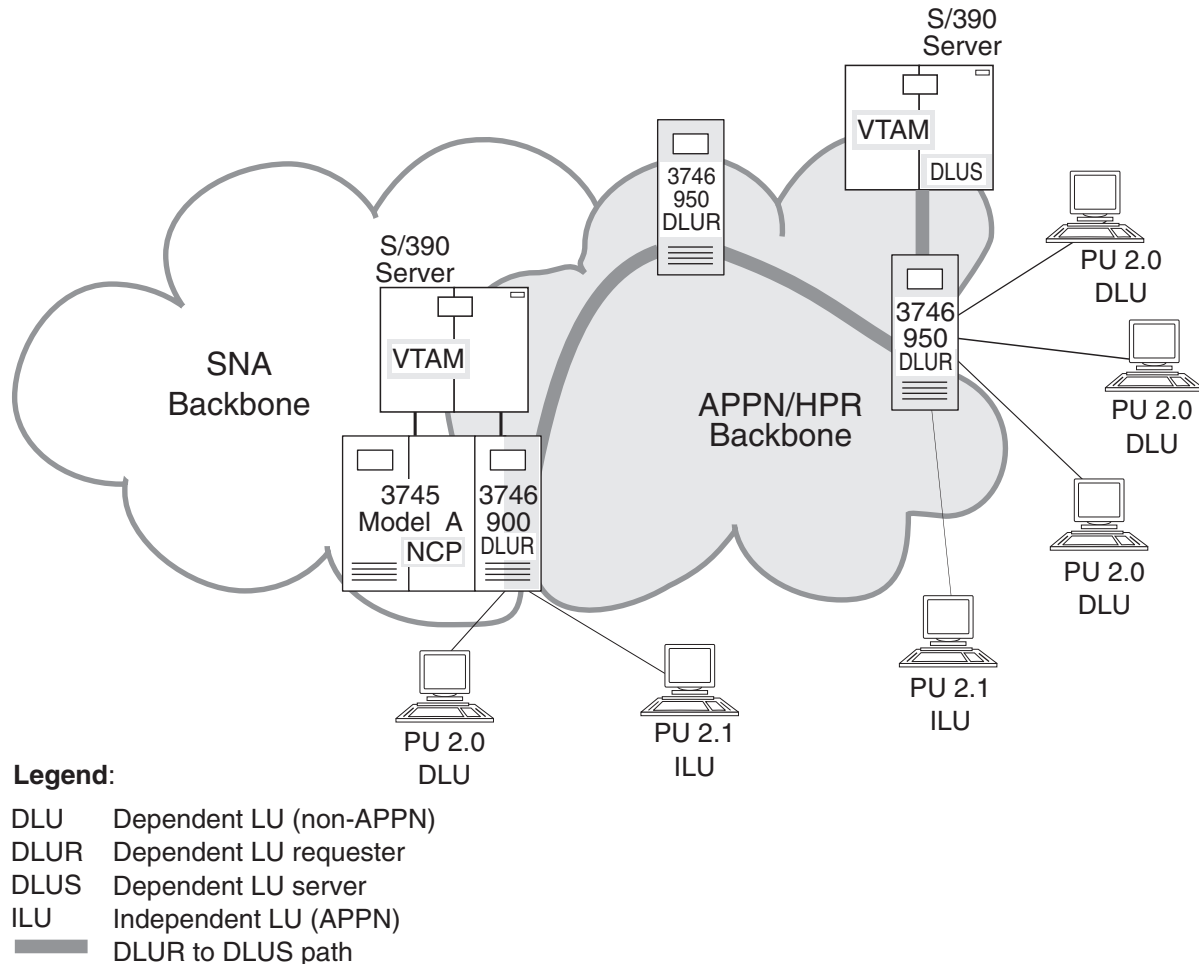


Figure 38. Dependent LU Support

The 3746 network node allows existing host-dependent SNA devices to access S/390 applications across an APPN/HPR backbone network. For example (see Figure 38). A physical unit type 1.0 or type 2.0 attached to the rightmost controller can access applications in any of the two S/390® servers.

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

In an APPN environment, the dependent LUs (DLUs) must reside on, or be owned by an APPN node providing the DLUR function. The DLUR node requests the dependent logical unit server (DLUS) of a VTAM network node to provide the SSCP services for its dependent LUs. To support session establishment for the

dependent LUs, the traditional SSCP-PU or SSCP-LU data flows through two LU 6.2 sessions between the DLUR node and DLUS node.

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

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

High-Performance Routing

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

- Dynamic rerouting around failed nodes and links without session loss
- Better routing performance
- Enhanced congestion control, which improves link efficiency
- A reduction in the amount of storage required in intermediate nodes
- Very high data throughputs between the S/390 servers and the network
- Synergy with the Parallel Sysplex® processor implementation providing end-to-end non-disruptive path switching up to the applications. This includes the S/390 support of multinode persistent session (MNPS).

The HPR architecture is made of two layers:

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

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

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

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

- Transport of APPN and SNA boundary traffic (DLUR)
- Selective transmission, based on class of service
- Flow control and network congestion avoidance (adaptive rate based [ARB])
- End-to-end error recovery and selective retransmission

- Non-disruptive rerouting around network failures

HPR Multilink Transmission Group

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

The MLTG is defined with a single TG number. It is reported and seen as a single TG by topology management, and viewed as a single TG in the route calculation process. Based on link error rates, error recovery can be run on an individual link basis. Links can be dynamically removed from the MLTG when no longer needed, thus resulting in cost savings. HPR MLTG is supported over SDLC, frame-relay and X.25 links, token-ring and Ethernet LANs.

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

In a mixed SNA and APPN/HPR network, as shown in Figure 38 on page 67, the 3746-900 NN can operate as an SNA node (PU type 4) for the network resources owned by the NCP running in the 3745, and as an APPN/HPR network node (PU type 2.1) for the resources owned by the 3746 APPN/HPR control point.

In Figure 38 on page 67, the 3745/3746-900 NN is channel-attached to a VTAM that operates as an interchange node (ICN). This ICN allows SNA devices connected to the SNA backbone to access applications over the APPN/HPR backbone, and SNA/APPN devices connected to the APPN/HPR backbone to access S/390 applications in the SNA backbone, therefore providing any-to-any networking.

X.25 Network Connectivity

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

- Support of qualified logical link control (QLLC) connections for APPN, SNA/DLUR, and HPR traffic
- Routing of mixed APPN, SNA/DLUR, HPR, IP, and SNA/NCP traffic on the same X.25 link
- PVC and SVC connections
- Up to 2.048-Mbps speed
- Direct DTE-to-DTE attachment
- Support of SDLC, frame-relay, ISDN, X.25, X.25 NCP, and X.25 NPSI links attached to a single communication line processor (CLP)

Frame-Relay Networking

The IBM 3746 Models 900 and 950 provide frame-relay networking functions. They are similar to those provided by NCP for 3746-900, with the exception of frame-relay switching functions (refer to *3745/3746 Planning Series: Serial Line Adapters*).

Physical Media

The 3746-900 Network Node and 3746-950 provide the same connectivity. They both support:

- Token-ring and Ethernet LANs
- Leased frame-relay links
- Switched and leased SDLC links
- ESCON channels
- Frame-relay network connections
- X.25 network connections

The 3746 Network Node supports RFC 1490, including:

- Frame-relay BAN
- Frame-relay BNN

For more information about the adapters that support network node connectivity, refer to the *IBM 3745 Communication Controller Models A, IBM 3746 Nways Multiprotocol Controller, Models 900 and 950: Overview*.

LU-LU Sessions with Session Pacing

Without this support, the 3746 does not bind with the DLUR LU when the secondary window size is “0” on a non-negotiable bind. The session is established with the host when connected through a direct-connected link. A non-negotiable LU bind, that does not support adaptive pacing and specifies infinite window size (using receive window size of 0) is not accepted by the 3746 intermediate network node. Without this support, the window size must not be 0 or greater than 63.

This function accepts this specific LU bind in the APPN tower, assuming that the DLC layer controls the pacing through the RNR with a remote PU to avoid network node congestion that is not controlled further by the APPN layer.

Link Station Name Included in XID3 Frame

For APPN ESCON links (ALS) from the 3746-900 to TPF systems, some APPN links are associated with NN in NNP and some with ICN (VTAM/NCP). The focus is on the XID frames that are sent to the TPF APPN EN when these links are activated. Without this function, the XID sent to the links associated with NNP does not contain the link station name, where as the XID sent to the ICN associated links does contain the link station name.

This function adds the optional Control Vector CV 0EF7 (Link Station Name) in the XID3 for the NNP associated links.

VTAM DeactPU/DeactLU Support by NNP

A DLUR does not behave like NCP if a dependent LU does a call in. This LU goes through traffic then the PU is deactivated from VTAM. The PU is forced down, but the DLUR reactivates it immediately.

From a user's perspective, the LU never drops. From the perspective of the node, this is a verb, *define_dlur_defaults* that defines the DLUR characteristics (such as, the default primary, backup DLUS, and various DLUR timeouts).

This function allows new capability with a field called, *drop_dependent_link_when_no_session*.

This field is on the CCM panel that defines the other DLUR information (such as pri DLUS, and so on).

3746-9x0 MAE APPN/HPR Implementation

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

MAE Implementation Specifics

APPN over DLSw

The MAE supports APPN over DLSw for connectivity to nodes through a remote DLSw partner. An example is shown in Figure 39 on page 72. The MAE supports only remote DLSw.

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

When APPN is configured on the MAE to use a data link switching (DLSw) port, DLSw provides a connection-oriented interface (802.2 LLC Type 2) between the APPN component in the MAE and APPN nodes and LEN end nodes attached to a remote DLSw partner.

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

Figure 39 on page 72 shows how TCP/IP and DLSw are used to transport APPN traffic over an IP network.




Figure 39. APPN over DLSw

Supported Traffic Types

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

For token-ring and Ethernet LAN ports

Each protocol that uses a port must have a unique SAP address, with the exception of DLSw (that can use the same SAP address as other protocols because DLSw frames will not be destined for the local MAC address, but rather a DLSw MAC address). A unique SAP address identifies the APPN link station for HPR traffic (local HPR SAP address parameter).

If ISR traffic is destined for a link station, then a different SAP address (local APPN SAP address parameter) must be used. The ISR traffic uses LLC Type 2 LAN frames. The HPR traffic is handled in a similar fashion to LLC Type 1 LAN frames and must have a different SAP address. The default SAP address for HPR traffic is X'C8'. If X'C8' has already been used by another protocol on a port, the default must be overridden. Note that there is only one APPN link station even though APPN ISR and HPR traffic use different SAP addresses.

For frame-relay ports

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

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

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

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

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

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

For PPP ports

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

Topology Safestore

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

The NNP and MAE only save the topology to their hard disks once a day during garbage collection.

Route Test

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

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

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

No User Application Programming Interfaces

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

Limited Resource Link Stations

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

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

Session-Level Security

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

Parallel TGs

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

DLUR Restrictions

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

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

Connection Network Restrictions

The router APPN support has the following connection network restrictions:

- Connection networks defined on the router network node are supported only on token-ring and Ethernet LAN ports.
- The same connection network (virtual routing node [VRN]) can be defined on only one LAN. However, the same VRN can be defined on multiple ports having the same characteristics to the same LAN.

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

APPN over DLSw Restrictions

The following restrictions apply to APPN over DLSw:

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

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

Summary of Implemented APPN Functions

Note: See the notes at the end of the tables (page 78) for abbreviations and explanation of terms used.

Table 12 (Page 1 of 3). Summary of Implemented Base Functions (APPN Version 2)

Function		CS OS/390 (VTAM)	3746	3746
Option Set	Description	V2 R8	NNP	MAE FC 3001
Configuration Services				
001	LEN-level XID3	EN/NN	NN	NN
002	All XID3 States	EN/NN	NN	NN
003	Link Station Role Negotiation	EN/NN	NN	NN
006	CP Name on XID3	EN/NN	NN	NN
007	TG Number Negotiation	EN/NN	NN	NN
008	Multiple TGs	EN/NN	NN	NN
010	Single-Link TG	EN/NN	NN	NN
1001	Secondary-Initiated Non-Activation XID	EN/NN	NN	NN
1004	Adjacent Node Name Change	EN/NN	NN	NN
Intermediate Session Routing				
011	LFSID Addressing	EN/NN	NN	NN
013	Priority Queuing for Transmission	2	NN	NN
Address Space Manager				
020	Extended BIND and UNBIND	EN/NN	NN	NN
021	Adaptive Pacing for Independent LU BINDs	3	NN	NN
023	Bind Segmenting and Reassembly	-	NN	NN
024	Adaptive Pacing for Dependent LU BINDs	-	NN	NN
Session Services				
030	CP-CP Sessions	EN/NN	NN	NN
031	CP-CP Capabilities Exchange	EN/NN	NN	NN
033	FQPCID Generation	EN/NN	NN	NN
034	CD-Initiate	EN/NN	NN	NN
035	Reconstruct CD-Initiate Reply	EN/NN	NN	NN
036	CoS/TPF	EN/NN	NN	NN
037	BIND (ILU=PLU)	EN/NN	NN	NN
038	Limited Resource	EN/NN	NN	NN
039	BIND without RSCV from Any LEN or APPN Node	EN/NN	NN	NN
040	Propagate Unrecognized CVs	NN	NN	NN
041	Session RU Segmenting and Reassembly	EN/NN	NN	NN
042	Interleaved Segments	EN/NN	NN	NN
1015	CP-CP Session Activation Enhancements	EN/NN	NN	NN
Directory Services				
050	Register EN Resources	EN/NN	NN	NN
051	Locate/Find/Found	EN/NN	NN	NN
052	Reconstruct GDS Variables for Locate Reply and CD-Initiate Reply	EN/NN	NN	NN
053	Participate in Network Searches	EN/NN	NN	NN
054	Send Wildcard Reply	7	NN	NN
055	Broadcast and Directed Searches	EN/NN	NN	NN
056	ENCP Search Control	EN/NN	NN	NN
057	Partial Directory Entries	-	NN	NN
059	Accept Unqualified LU Name	EN/NN	NN	NN
060	Locate Chains - Locate(keep)	EN/NN	NN	NN
061	Sending Locate to a Gateway	NN	NN	NN
062	Cache Resource Locations	NN	NN	NN
063	Favor Explicit Replies	NN	NN	NN
064	Network-Qualified LU Names	EN/NN	NN	NN
065	Central Directory Client	NN	NN	NN
066	Abbreviated Resource Hierarchy	NN	NN	NN
068	Authentic Net ID Indicator	EN/NN	NN	NN

Table 12 (Page 2 of 3). Summary of Implemented Base Functions (APPN Version 2)

Function		CS OS/390 (VTAM)	3746	3746
Option Set	Description	V2 R8	NNP	MAE FC 3001
069	DS Support for Domain LEN Resources	EN/NN	NN	NN
1103	Retry Referred Search	NN	NN	NN
1104	Topology-Based Directory Nonverify	5	NN	NN
1105	PCID Modifier	EN/NN	NN	NN
1109	Surrogate Owner	NN	NN	NN
1117	Bypass of Directed Locate Not Allowed	5	-	-
1119	Report Branch Topology to Manager	NN	NN	NN
1120	Branch Awareness	NN	NN	NN
1121	Branch Extender Function	NN	NN	NN
Topology and Routing Services				
070	Process Local Resource Change	EN/NN	NN	NN
073	Initial Topology Exchange	EN/NN	NN	NN
074	Flow Reduction Sequence Numbers	EN/NN	NN	NN
075	Resource Sequence Numbers	EN/NN	NN	NN
076	Topology Broadcast	NN	NN	NN
077	Garbage Collection	NN	NN	NN
078	Topology Isolation at Net ID Boundaries	NN	NN	NN
079	Build RSCV	NN	NN	NN
080	Calculate Route Using Connection Networks	NN	NN	NN
081	Class-of-Service Manager	EN/NN	NN	NN
082	Route Randomization	NN	NN	NN
083	Member of Connection Network	EN/NN	NN	NN
084	Select One-Hop Routes	NN	NN	NN
085	Select Network Routes	NN	NN	NN
086	Topology Awareness of CP-CP Sessions	NN	NN	NN
087	Garbage Collection Enhancements	NN	NN	NN
088	TDU Flow Improvements	NN	-	-
1202	Safe-Store of Topology DB	NN	NN	NN
Node Operator Command Set				
090	Common Node Operator Command Set	EN/NN	NN	-
091	Network Node Operator Command Set	NN	NN	-
Intermediate Session Routing				
100	Extended/Unextended BIND and UNBIND	EN/NN	NN	NN
101	Fixed Session-Level Pacing	EN/NN	NN	NN
102	Adaptive Session-Level Pacing	EN/NN	NN	NN
103	Intermediate Session Segmenting/Reassembly	-	NN	NN
104	Routing BIND and UNBIND	EN/NN	NN	NN
105	Intermediate Session Routing for Dependent LU Sessions	EN/NN	NN	NN
106	Intermediate Session Routing for Type 6.2 LU-LU Sessions	EN/NN	NN	NN
Management Services - Multiple-Domain Support				
150	MDS Common Base	4	NN	NN
151	MDS End Node Support	4	NN	-
152	MDS Network Node Support	4	NN	NN
153	MDS High Performance Option	4	-	-
154	MDS Transport Confirmation Option	-	-	-
Management Services - MS Capabilities				
160	MS_CAPS Base End Node Support	4	NN	-
161	MS_CAPS Have a Backup or Implicit FP	4	NN	1
162	MS_CAPS Be a Sphere of Control End Node	4	-	-
163	MS_CAPS Base Network Node Support	4	NN	NN
164	MS_CAPS Have a Subarea FP	6	-	-
Management Services - Entry Point Alerts				
170	EP Alert Base Subset	4	NN	NN
171	Problem Diagnosis Data in Alert	-	NN	NN

Table 12 (Page 3 of 3). Summary of Implemented Base Functions (APPN Version 2)

Function		CS OS/390 (VTAM)	3746	3746
Option Set	Description	V2 R8	NNP	MAE FC 3001
174	Operator-Initiated Alert	4	-	-
175	Qualified Message Data in Alert	-	-	-
176	Self-Defining Message Text Subvector in Alert	-	-	-
177	LAN Alert	-	-	-
178	SDLC/LAN LLC Alert	-	-	-
179	X.21 Alert	-	-	-
180	Hybrid Alert	-	-	-
181	X.25 Alert	-	-	-
182	Held Alert for CPMS	-	-	NN
183	Resolution Notification Support	-	-	-
184	Operations Management Support in Alert	-	-	-
Miscellaneous				
1013	Interoperability with Peripheral Border Node	NN	NN	NN
Notes: CS OS/390 = IBM eNetwork Communications Server for OS/390®, which includes VTAM EN/NN = End Node and Network Node NN = Network Node - = Not Supported 1 Backup Focal Point only. 2 Supported by NCP's BF only. 3 VTAM will respond to a BIND pacing request received, but will never set the pacing request indicator on a BIND. 4 Function supported through NetView. 5 VTAM does perform topology database lookup (to see if an unknown resource is a NNCP), but does not skip sending the APPN locate. This locate is then sent as a directed search to the NN. Because of this processing, VTAM has implemented option set 1117. 6 VTAM/NetView can be the alert focal point for EP_ALERTs (function sets 170 to 184, except for 174), but does not generate EP_ALERTs itself. 7 In some cases, CS OS/390 sends wildcard replies, but does not allow wildcard definitions.				

Table 13 (Page 1 of 2). Summary of Implemented Optional Functions (APPN Version 2)				
Function		CS OS/390 (VTAM)	3746	3746
Option Set	Description	V2 R8	NNP	MAE FC 3001
Configuration Services				
1002	Adjacent Link Station Name	EN/NN	NN	NN
1003	Short-Hold Mode	EN/NN	-	-
1006	Dynamic Name Change	EN/NN	-	-
1007	Parallel TGs	EN/NN	NN	NN
CP Capabilities				
1011	Multiple Local LUs	EN/NN	NN	-
1012	LU Name = CP Name	-	NN	NN
1014	Peripheral Border Node	-	-	-
1016	Extended Border Node	NN	-	NN
1017	Gateway	-	-	-
1018	Delete EN Resources Before Registering	EN/NN	-	-
Dependent LU Support				
1060	Prerequisite for Session Services Extensions CP Support	EN/NN	NN	NN
1061	Prereqs. for SSE NNS Support	NN	NN	NN
1062	Session Services Ext. CP Support	EN/NN	NN	NN
1063	Session Services Ext. NNS Support	NN	NN	NN
1064	Session Services Ext. PLU Node Support	EN/NN	-	-
1065	Session Services Ext. CP(SLU) (SSCP) Support	EN/NN	-	-
1066	Dependent LU Server	NN	-	-
1067	Dependent LU Requester	-	NN	NN
1071	Generalized ODAI Usage	-	NN	NN
Cryptography Support				
1070	Session Cryptography	EN/NN	-	-
Directory Services				
1100	Safe-Store of Directory Cache	NN	-	-
1101	Preloaded Directory Cache	NN	NN	NN
1102	EN Authorization	-	-	-
1106	Central Directory Server	NN	-	-
1107	Central Resource Registration (of LUs)	EN/NN	NN	NN
1108	Nonverify	2	-	-
1116	DLUS-Served LU Registration NNS Support	NN	NN	NN
1118	EN TG Vector Registration	EN/NN	-	-
Topology and Routing Services				
1200	Tree Caching and TG Caching	NN	NN	NN
1201	Permanent Storage Medium	EN/NN	NN	NN
1203	Detection and Elimination of TDU Wars	3	NN	NN
Intermediate Session Routing				
1300	Tuning Values for ISR	EN/NN	NN	-
1301	Nonpaced Intermediate Session Traffic	EN/NN	-	-
High Performance Routing				
1400	HPR Base (ANR)	NN	NN	NN
1401	Rapid Transport Protocol	EN/NN	NN	NN
1402	Control Flows over RTP	EN/NN	NN	NN
1403	Dedicated RTP Connections	-	-	-
1404	Multilink TG (MLTG)	-	NN	-
1405	HPR Border Node	NN	-	NN
Management Services - File Services				

Table 13 (Page 2 of 2). Summary of Implemented Optional Functions (APPN Version 2)				
Function		CS OS/390 (VTAM)	3746	3746
Option Set	Description	V2 R8	NNP	MAE FC 3001
1500 1501	File Services Support Base Network Operator Support for File Services	- -	- -	- -
Management Services - Change Management				
1510	CM Base	-	-	-
1511	CM Production Only Activate	-	-	-
1512	CM Execution Window Timing	-	-	-
1513	CM Activate Report	-	-	-
1514	CM Alter Active Install	-	-	-
1515	CM Object Disposition Install	-	-	-
1516	CM Initiate Command	-	-	-
1517	CM Cancel Command	-	-	-
1518	CM Activate Last	-	-	-
Management Services - Operations Management				
1520 1521	Common Operations Services Operations Management	- 1	- -	- -
Notes: EN/NN = End node and network node NN = Network node - = Not Supported 1 Function provided by NetView. 2 Only at the NN(OLU). 3 CS OS/390 support is prestandard; most of the option set is implemented.				

Summary of Supported DLCs and APPN Functions

Table 14 gives an overview of the DLC types and APPN functions supported by APPN-capable hardware.

<i>Table 14. Summary of Supported DLCs and APPN Functions</i>		
DLC TYPE	MAE	3746
Token-Ring	IHD	IHD
Ethernet	IHD	IHD
Twinax	n/a	n/a
Frame-Relay BNN	IHD	IHD
Frame-Relay BAN	IHD	IHD
Point-to-Point Protocol	IH	n/a
APPN over DLSw	ID	ID (see note)
HPR over IP	H	H (see note)
SDLC Leased Line	ID	IHD
X.25 PVC and SVC (Previewed)	ID	IHD
APPN over PPP over ISDN	IH	n/a
APPN over Frame Relay over ISDN	IHD	n/a
APPN over ATM Forum Compliant LAN Emulation	IHD	n/a
Native HPR over ATM	H	n/a
APPN over PPP over V.25 bis	IH	n/a
SNA over Asynch	n/a	n/a
ESCON	IH	IH
Parallel channel	IH	n/a
Note: Needs the MAE installed. Legend: I = Intermediate Session Routing H = High-Performance Routing D = Dependent LU Requester, this refers to the port providing the connection to the downstream PU (DSPU).		

Chapter 2. Internet Protocol (IP) Overview

This chapter gives an overview of the Internet Protocol (IP). It is not just one protocol, but a group of protocols that make up an architecture designed for communication between and across diverse network platforms and network architectures.

This chapter describes:

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

The Need for Transparency in Communication

The computer networks of large organizations have evolved through the merger, interconnection, and expansion of previously autonomous networks. Such interconnected networks support diverse protocols, multiple speed requirements, and incompatible topologies. For example, suppose there is a community of users, A, B, C, and D (see Figure 40), who need to reach each other across four networks.




Figure 40. Communicating across Multiple Networks

They would each need to know about intervening networks:

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

This is obviously an impractical and unsatisfactory situation.

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

The Internet Protocol (IP) is just such an architecture.² IP provides a unified appearance to users and applications independently of the actual networks

² Internet also refers to an *interconnected network*, a network consisting of two or more otherwise autonomous networks.

connecting them by the simple expedient of a common addressing scheme. This relationship of IP to networks is shown in Figure 41 on page 84.




Figure 41. The Internet Protocol (IP) Overlays Networks

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




Figure 42. 3746 Network Nodes with the IP Feature (FC 5033). They link the networks and route data between the them.

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

IP Addressing

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

The IP Address

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

myname@ibm.com.

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

128.2.7.9

is a valid Internet address. Its binary form is:

B'10000000 00000010 00001111 00001001'

The numeric form is used by the IP software, which also provides the mapping from symbolic to numeric forms.

IP Address Classes

The address shown above is an example of one class of IP address but is not suitable for all network configurations. The proportion of networks to hosts can be vastly different from one enterprise to another. IP addresses are thus arranged in **classes** to suit different network characteristics. See Figure 43.




Figure 43. Classes of IP Addresses

Each class, having a different proportion of network number to host number, can be assigned to different sizes of network. See Table 15.

Table 15. Use of IP Address Classes		
Class	Maximum Number of Networks	Maximum Number of Hosts per Network
A	126	16 777 214
B	16382	65534
C	2 097 150	254
D	Not applicable	Not applicable
E	Not applicable	Not applicable

Host and Network Numbers

The IP address consists of a pair of numbers:

IP address = <network number><host number>

The network number is administered centrally by the Internet Network Information Center (the InterNIC) and is unique for each network. The host number (also known as host address or host ID), is assigned to each host within a network and belongs to the authority that owns the network. Table 16 shows the components of the example address given in “The IP Address” on page 85.

Table 16. IP Address Structure		
IP address=	<Network number>	<Host number>
128.2.7.9	128.2	7.9

IP Addresses for Routers

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




Figure 44. Routers Need Two Addresses

It must have only one address when seen from each network, but each network has its own network number. Thus, the router must have two addresses, one for each network it connects. Its resulting addresses are shown in Table 17.

Table 17. IP Addresses for a Router		
IP address	Network number	Host number
128.2.7.9	128.2	7.9
131.5.1.2	131.5	1.2

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

Subnets

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

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

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

Subdividing the Network

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

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

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

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

Administrators can assign subnet addresses to represent, for example:

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

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

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

The Subnet Mask

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

- Ones (1's) indicate bit positions assigned to the subnet number.
- Zero (0's) indicate bit positions assigned to the host number.

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

Example Subnet Mask

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

- The first byte is the subnet number, the second the host number. This gives 254 (256 minus 2 with the values 0 and 255 being reserved) possible subnets, each having up to 254 hosts. The subnet mask is 255.255.255.0.
- The first 12 bits are used for the subnet number and the last four for the host number. This gives 4094 possible subnets (4096 minus 2) but only 14 hosts per subnet (16 minus 2). The subnet mask is 255.255.255.240.

There are many other possibilities.

Building Good Masks

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

Multiple Destinations

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

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

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

- Broadcast addresses
- Multicast addresses

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

Broadcasting

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

Limited broadcast address

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

Network-directed broadcast address

If the following conditions are true:

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

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

Subnet-directed broadcast address

If the following conditions are true:

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

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

Multicasting

Broadcasting has a major disadvantage— its lack of selectivity. If an IP datagram is broadcast to a subnet, every host on the subnet will receive it, and have to process it to determine whether the target protocol is active. If it is not, the IP datagram is discarded. Multicasting avoids this overhead by using groups of IP addresses. Each group is represented by a 28-bit number, which is included in a Class D address. See Figure 43 on page 85.

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

There are two kinds of host groups:

Permanent

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

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

A permanent group exists even if it has no members.

Transient

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

Domains

The system of symbolic names (see “The IP Address” on page 85,) is more convenient to use than the IP numeric address structure. However, this requires:

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

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

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

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

The Hierarchical Name Space

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

For example, consider the name `lcs.mit.edu`. The lowest-level domain name, `lcs`, is a subdomain of `mit`, which again is a subdomain of `edu` (education) which is called a top-level domain. We can also represent this naming concept by a hierarchical tree. See Figure 45.

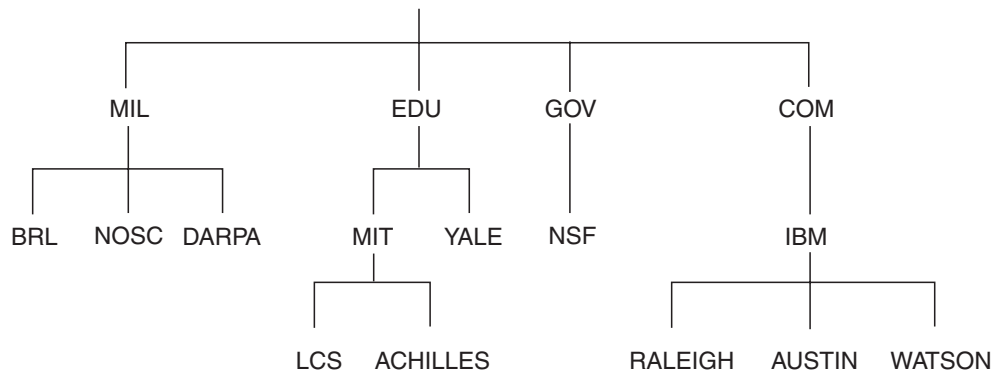


Figure 45. Hierarchical Structure of Domains

Figure 45 shows the chain of authority assigning domain names. This tree is only a tiny fraction of real namespace. Table 18 on page 91 shows some of the top-level domains. The single domain above the top-level domains has no name and is referred to as the *root domain*.

Fully Qualified Domain Names (FQDNs)

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

Generic Domains

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

Table 18. The Generic Top-Level Domains	
Domain Name	Meaning
edu	Educational institutions
gov	Government institutions
com	Commercial organizations
mil	Military groups
net	Networks
int	International organizations
org	Other organizations

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

Country Domains

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

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

3746-9x0 Router Node Implementation

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

- IP routing providing connectivity over:
 - ESCON channels:
To IBM and other equipment manufacturers' (OEM) processors that adhere to the ESCON architecture, possibly through one or two cascaded ESCON Directors (IBM 9032 or 9033).
 - Token-ring LANs
 - Ethernet LANs:
Ethernet access requires the 3746-9x0 Ethernet feature. Refer to the "Ethernet Adapters" chapter in the *3745/3746 Planning Series: Token Ring and Ethernet* and the "Physical Planning Details" chapter in the *3745/3746 Planning Series: Physical Planning* for detailed information on this feature.
 - Leased lines:
Leased lines using either the PPP or frame-relay protocol.

- Frame-relay and X.25 networks
- Dynamic routing protocol support, using:
 - RIP Version 1
 - OSPF Version 2
 - BGP Version 4
- ARP and proxy ARP
- BOOTP relay

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

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

Overview of 3746 IP Routing

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




Figure 46. IP Address Support Example

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

Networks can be:

- Multiaccess broadcast networks

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

Examples of multiaccess broadcast networks are token-ring, Ethernet, and FDDI.³

- Multiaccess nonbroadcast networks

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

- Point-to-point (PtP) networks

These allow direct communication between two adjacent nodes. No broadcast is required on PtP connections. From an IP routing perspective, a PtP connection comprises a separate network, although only two stations connect. Instead of the term PtP network, the term PtP connection is more commonly used. Examples of point-to-point networks are PPP, Serial Line Internet Protocol (SLIP),³ SNAlink, and ESCON connections.

IP General Functions

Broadcasting

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

- Limited or local-wire broadcast

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

- Subnet-directed broadcast

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

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

³ Not supported by the 3746.

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

- Limited or local-wire broadcast

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

- Subnet-directed broadcast

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

- Network-directed broadcast

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

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

Supernetting (Route Aggregation)

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

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

This process of combining multiple networks into a single entry is termed *address aggregation* or *address summarization*. It is also called *supernetting* because routing is based upon network masks that are shorter than the natural network mask of the IP address, in contrast to subnetting where the network masks are longer than the natural mask.

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

Multiple IP Addresses for a Network Interface

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

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

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

Figure 47 shows a situation that requires three routers to handle communication between three IP networks.




Figure 47. Three Networks and Three Routers Before Consolidation of the Routers

Figure 48 shows the situation after consolidation into one 3746 IP Router. Users see the network addresses as before and require no definition changes.




Figure 48. Three Networks and One 3746 IP Router After Consolidation

2. Route aggregation of multiple Class C addresses.

Suppose an enterprise requests 1000 new IP host addresses for its internal network. They are given **four** contiguous ranges of class C addresses, but the four networks are connected to a single port on the 3746 IP Router.

Figure 49 shows that the single port can be assigned four IP addresses, one for each IP network. To each network, the 3746 IP Router appears as a single IP address.




Figure 49. A Single 3746 IP Router Serving Four IP Networks through a Single Port

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

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

Multiple Interfaces to the Same Subnet

By default, the IP addresses assigned to the network interfaces must each be in a different network or subnet. This restriction can be removed by enabling the same subnet option. This option allows the router to have multiple IP ports to the same subnet.

Multiple IP interfaces to the same subnet are useful under only one of the following conditions:

- OSPF Point-to-Multipoint configured on the IP interfaces
- Nexthop awareness enabled on the IP interfaces and static routes defined for the routes that go through the IP interfaces.

For example, for IP over token-ring, two ports (TIC3s) can be configured in the 3746 for the same physical token-ring LAN with only one port being used at a time. By default, the same subnet option is disabled.

User Datagram Protocol (UDP)

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

Transmission Control Protocol (TCP)

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

Internet Control Message Protocol (ICMP)

The 3746 IP Router recognizes the following ICMP types:

0 Echo Reply

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

3 Destination Unreachable

The 3746 can generate the following ICMP destination unreachable messages:

- Network unreachable

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

- Fragmentation needed (the Do Not Fragment [DF] bit is set)

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

- Protocol unreachable

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

- Port unreachable

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

- Source route failed

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

4 Source quench

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

5 Redirect

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

8 Echo Request

Echo request messages are responded to the echo reply message. ICMP Echo request are generated by the Ping application (see “Ping” on page 110) at the user’s request.

11 Time exceeded

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

12 Parameter problem

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

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

13 Timestamp

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

14 Timestamp reply

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

15 Information request

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

16 Information reply

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

17 Address mask request

Address mask request messages are dropped if:

- The destination IP address matches the 3746’s internal IP address.

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

18 Address mask reply

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

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

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

Address Resolution Protocol (ARP)

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

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

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

ARP Caching

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

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

ARP table entries that have not been recently used or updated (refreshed) are discarded. Permanent entries are never discarded. The default time for an entry to

be retained before being discarded is five minutes. User commands exist to delete all (dynamic), or specific ARP (dynamic and permanent) entries. Dynamic ARP table entries are also removed when a network interface is considered down.

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

Proxy-ARP

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

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

Note: An option exists to disable this option.

- The physical transmit interface of the route to the destination IP address is the same interface from which the ARP request is received.
- The route found for the destination protocol address is a filter.

Routing Protocols

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

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

Each is explained in more detail in the following sections.

Routing Information Protocol (RIP) Version 1

RIP Version 1 is implemented in both the 3746 and the MAE. It includes support for *split horizon*.

To speed network convergence, triggered updates are sent.

Note: *Poison reverse* is not supported.

Route Acceptance Policy: For each interface, you can configure, whether or not:

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

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

Route Advertisement Policy: For each interface, you can configure, whether or not:

- RIP will be advertised.
- Network, subnet, or host routes will be advertised.
- Default routes, static routes, or both will be advertised.

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

Routing Information Protocol (RIP) Version 2

RIP Version 2 is implemented in both the 3746 and the MAE. This support of RIP2 enables, for example, the ESCON adapter to reduce loads on the host by using an IP multicast address to broadcast periodic RIP2 messages.

RIP2 has the following features:

Support for Variable Subnet Masks: This is the major enhancement brought by RIP V2. RIP V2 supports subnet masks of variable sizes within the same router, while in RIP V1 all the subnet masks had to be the same for a given router.

Password Authentication: Defines an authentication key that is added to each RIP V2 advertisement packet originating from the interface that is being defined as supporting RIP V2.

Multicasting (Instead of Broadcasting): RIP V1 broadcasts advertisement packets over all local interfaces.

RIP V2 sends RIP2 advertisement packets, from the interfaces where RIP V2 is defined, to the multicast group having the class D address 224.0.0.9.

RIP Metrics

Where you have multiple interfaces, the RIP metrics allow you set the weight of one interface higher than another. For example, for traffic that can flow out over either of two interfaces, it will always go out over the interface with the lower weight (lower metric value).

In-Metric: This metric is added to the RIP routes learned on an interface prior to adding the routes to the routing table.

When an interface receives RIP messages, the metric is used when storing routes in the routing table. Later, when looking up a route in the table, the stored metric is taken into account in the route.

The in-metric range is 1 to 15.

Out-Metric: This metric is added to RIP routes advertised on an interface. The metric is set on RIP messages that are sent to other routers on the interface.

The out-metric range is 0 to 15.

Open Shortest Path First Protocol (OSPF)

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

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

There is no support for type of service (ToS)-based routing; that is, ToS 0 is the only supported ToS.

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

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

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

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

- Numbered point-to-point (PtP)

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

- Unnumbered point-to-point (PtP)

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

- Broadcast

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

- Nonbroadcast multiaccess (NBMA).

Viewing a network, for example a frame-relay network, as an NBMA network can be used when the network is fully meshed (meaning virtual circuits exist between any pair of routers). In routers that are eligible to become designated

routers, neighbors must be configured whether or not the neighbor is eligible to become a designated router. The configurable poll interval defines the interval at which the designated router will attempt to contact the neighboring routers to establish an adjacency. NBMA connections are supported for frame-relay networks. If the network is partially meshed, it is more useful to view the network as a point-to-multipoint network.

- Point-to-multipoint (PtM)

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

- Virtual link

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

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

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

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

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

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

Importing Routes into OSPF: The 3746 IP Router configuration determines whether or not non-OSPF routes can be imported as OSPF AS external routes, and advertised as OSPF ASE LSAs throughout the OSPF routing domain. Imported

routes can be imported as ASE type 1 (router) or ASE type 2 (network) routes. An ASE type 1 route always overrides type 2 routes. It has a single metric that is the sum of the path cost to the AS border router, and the AS border router metric. Conversely, the metric for a type 2 ASE has internal and external components, namely the path cost to the AS border router and the route's external metric. When comparing ASE type 2 routes to the same destination, the one with the lower external metric will be always be preferred, independently of the internal metric. The metric used for both the OSPF ASE type 1 route and the external component of the type 2 route is the metric from the protocol from which it is imported.

Default Route: 3746 IP allows generation of a default routes and advertisement of the routes through an AS external (ASE) LSA. For details see "Interoperability of Routing Protocols" on page 107.

Not So Stubby Area

The not so stubby area (NSSA) function is compatible with RFC 1587.

The OSPF specification defines two general classes of area configuration:

- The first allows type 5 LSAs to be flooded throughout the area. In this configuration, type 5 LSAs can be originated by routers internally to the area or flooded into the area by area border routers. These areas, referred to here as *type 5 capable areas* or just plain *areas* in the OSPF specification, are distinguished by the fact that they can carry transmit traffic. The backbone is always a type-5 capable area.
- The second type of area configuration, called *stub*, allows no Type-5 LSAs to be propagated into or throughout the area and instead depends on default routing to external destinations.

NSSAs are defined in much the same manner as existing stub areas. To support NSSAs, a new option bit, the "N" bit, and a new type of LSA (Type 7) are defined. The N bit ensures that routers belonging to a NSSA agree on its configuration. Similar to the stub area's use of the "E" bit, both NSSA neighbors must agree on the setting of the N bit or the OSPF neighbor adjacency will not form.

Type-7 LSAs provide for carrying external route information within a NSSA. Type-7 AS external LSAs have virtually the same syntax as the Type-5 AS external LSAs with the obvious exception of the link-state type. There are two major semantic differences between Type 5 and Type 7 LSAs:

- Type-7 LSAs can be originated by and advertised throughout a NSSA; as with stub areas, Type-5 LSAs are not flooded into NSSAs and do not originate there, except on NSSA border routers.
- Type-7 LSAs are advertised only within a single NSSA. They are not flooded into the backbone area or any other area by border routers, though the information that they contain can be propagated into the backbone area.

To allow limited exchange of external information across an NSSA border, NSSA border routers translate selected Type-7 LSAs received from the NSSA into Type-5 LSAs. These Type-5 LSAs will be flooded to all Type-5 capable areas. NSSA border routers can be configured with address ranges so that several Type-7 LSAs can be represented by a single Type-5 LSA. The NSSA border routers that perform translation are configurable thus creating efficient forwarding of Type-5 LSAs originating from aggregated Type-7 LSAs. In addition, an NSSA border

router can originate a default Type 7 LSA (IP address of 0.0.0.0) into the NSSA, and must if no NSSA internal Type 7 default route exists.

Default routes are necessary because NSSAs do not receive full routing information and must have a default route in order to route to AS-external destinations. Like stub areas, NSSAs can be connected to the backbone at more than one area border router, but cannot be used as a transit area. Note that the default route originated by an NSSA border router is never translated into a Type-5 LSA. However, a default route originated by an NSSA internal AS boundary router (one that is not also an area border router) can be translated into a Type-5 LSA.

Like stub areas, the importing of OSPF summary routes (Type-3 LSAs) into NSSAs is a configuration option. However particular care should be taken to ensure that OSPF internal routes are always chosen over OSPF external (Type 7) routes. This can happen when other IGPs, like RIP and ISIS, leak routing information between an OSPF NSSA and another OSPF area. In these cases, all OSPF summary routes should be imported into the affected NSSAs. The recommended default behavior is to import OSPF summary routes into NSSAs.

As with standard OSPF stub areas, when summary routes are not imported, a Type-3 summary LSA for the default route should originate from NSSA border routers. The Type-3 summary default route ensures intra-AS connectivity to the rest of the OSPF domain, because it takes precedence over any Type-7 external default route that might originate on a NSSA internal router. This Type-3 summary default route prevents the OSPF domain's internal traffic that is normally forwarded by OSPF summary routes from exiting the AS through any NSSA Type-7 external default route originated by a NSSA internal router. Type-7 external defaults generated on NSSA internal routers and the no-summary option are mutually exclusive. When summary routes are imported into the NSSA, only Type-7 default routes (as opposed to Type 3) originate into the NSSA from NSSA border routers.

NSSA border routers never generate Type-4 summary LSAs for their NSSA ASBRs since NSSA ASBRs generate only Type-7 LSAs that are only flooded throughout the NSSA.

For the IBM 3746-9x0, the primary motivation for implementing NSSAs are:

External Route Aggregation to limit route table size

Presently, there is no mechanism to aggregate external advertisements from many OSPF routers. Nor is there any way to keep external routes from being flooded throughout the entire OSPF routing domain.

Separate default route to firewall per area

In many enterprise networks, it is desirable to have a separate Internet firewall in each location. This can be accomplished using NSSAs without having the area-specific default route advertised throughout the routing domain.

Implementation of ExtraNets

Multiple enterprises could use an OSPF backbone with attached NSSAs (one per enterprise) to implement an extranet (network administered by multiple cooperating parties) and still have flexible routing within the NSSA.

Border Gateway Protocol Version 4 (BGP)

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

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

Interoperability of Routing Protocols

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

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

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

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

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

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

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

OSPF Specifics

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

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

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

Metric type 2

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

Metric type 1

means that the metric is comparable to internal paths.

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

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

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

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

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

- Interior gateway protocol (RIP, OSPF) routes are preferred over BGP routes.
- OSPF routes are always preferred over other type of routes (other than routes to directly attached networks).
 - OSPF intra-area routes are preferred to inter-area routes.
- RIP routes and static routes can override OSPF ASE routes. This can be configured if external routes will be compared to OSPF ASEs as type 1 or type 2 routes. Given the preceding definitions of the configurable OSPF ASE route, static and RIP routes can override OSPF ASE routes in the following situations:
 - The OSPF comparison is configured as type 1, and the OSPF ASE route is an OSPF ASE type 2 route.
 - The OSPF comparison is configured as type 1, and the OSPF ASE route is an OSPF ASE type 1 route but has a higher cost than the RIP or static route.

- The OSPF comparison is configured as type 2, and the OSPF ASE route is an OSPF ASE type 2 route but has a higher cost than the RIP or static route.

Nexthop Up/Down Awareness: Nexthop awareness allows the router to sense whether a neighboring router is active or inactive. When this option is enabled, the router makes a more accurate determination of whether a static route that uses the neighboring router as its next hop will function. It also allows the router to determine over which network interface a static route's next hop can be reached when that next hop is in an IP subnet that is defined on multiple network interfaces. Nexthop awareness is supported on Frame Relay and ESCON ports.

Nexthop Awareness over Frame Relay: When static routes are defined over frame relay ports, it allows the 3746 to update the status, active or not, of the data link connection identifiers (DLCIs) in the OSPF tables. This ensures the use of active DLCIs for the route selection, and enables high-availability network design.

Nexthop Awareness over ESCON: It enables the reporting of Transaction Processing Facility (TPF) status, active or inactive, to the neighboring LAN-attached router. This allows end-user connection requests, especially those user reconnection requests received after a failure, to be automatically routed through the 3746 to an operational TPF. It enables high-availability access to TPF applications, in conjunction with multiple ESCON ports (Same Subnet option). No operator intervention is required in case of TPF failure.

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

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

Internal Applications

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

- Ping
- Traceroute

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

Ping

Packet internet groper (Ping) is a TCP/IP applications that sends one or more messages to a specified destination host requesting a reply and measures the round trip time. Traditionally, if you could Ping a host, then other applications such as Telnet or FTP could reach that host. With the advent of security measures on the Internet, which control access to networks by application protocol, port number, or both, this is no longer strictly true. Nonetheless, the first test to reach a host is still to attempt to ping it.




Figure 50. Packet Internet Groper (Ping)

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

```
ping host
```

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

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

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

Be aware that IP is a connectionless protocol and that a Ping request and its reply may traverse different routes through the network. Being able to ping a node does not necessarily mean that the reply can be returned as well. Ping requests contain a destination address that is set to the target host indicated in the Ping command. The 3746 IP Routers sets the source IP address to its router ID. Therefore, to enable the target host to return the Ping reply, an active route between this host and the 3746 Router ID is required.

Note: For pings to be successful, make sure that the Ping ICMP messages are not filtered along the route between the 3746 IP Router and the destination host.

In normal operation, Ping builds and sends an ICMP echo request. Ping waits up to 1 second to receive an ICMP echo reply to calculate the round-trip time, and display the size of the echo request/reply packet (not including IP header), its source IP address, and the round-trip time in milliseconds. 3746 IP echo request/reply messages are always 64 bytes long (excluding IP header).

Ping keeps count of the echo requests it sends and the echo replies it receives. It keeps a sum of the round-trip times, and calculates maximum and minimum round-trip times. In addition, it keeps count of the ICMP destination unreachable (network or host unreachable). 3746 IP Ping stops when the user presses any key. Ping statistics are then displayed, which include:

- The number of echo request sent
- The number of echo replies received
- The percentage of packets lost (no replies received)
- The number of destination unreachable messages received
- The number of echo request not sent due to lack of buffers

Traceroute

The Ping command described in the previous section can be used to verify connectivity between two hosts. In addition, for example when the Ping command fails, the Traceroute command can be used to determine the route that IP datagrams follow from host to host.

Traceroute relies on ICMP time exceeded and destination unreachable messages. It sends a UDP datagram with a TTL of 1 to the destination host. The first router to see the datagram will decrement the TTL to 0 and return an ICMP Time Exceeded message as well as discarding the datagram. In this way, the first router in the path is identified. This process can be repeated with successively larger TTL values in order to identify the series of routers in the path to the destination host. Traceroute sends UDP datagrams to the destination host that reference a port number outside the normally used range. This enables Traceroute to determine when the destination host has been reached, that is, when an ICMP destination (Port) unreachable message is received.

For each TTL value, the UDP datagram is sent three times, waiting up to three seconds between transmissions to receive a response. Traceroute continues this process until either a TTL value of 33 has been reached or an ICMP destination unreachable message has been received.

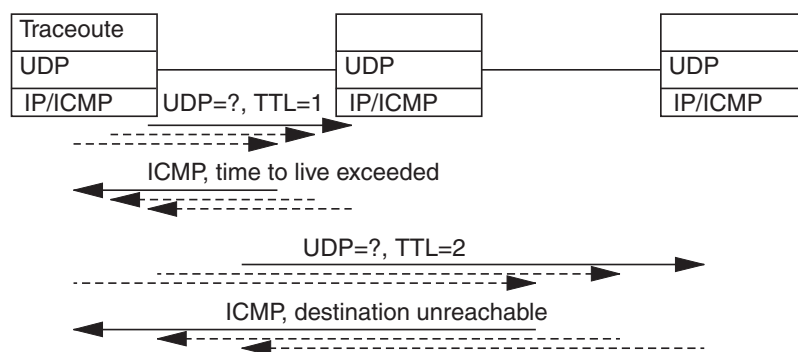


Figure 51. Traceroute

When a route exist between the 3746 IP Router and a destination node, each intermediate router will return an ICMP time exceeded message. However, a situation might exist where somewhere along the route the Traceroute frame cannot be forwarded.

Note: For Traceroute to be successful, make sure that the Traceroute ICMP messages are not filtered along the route between the 3746 IP Router and the

destination host. For each TTL value, Traceroute displays a line with the following information:

- 'BF' each time the 3746 IP is unable to send the UDP frame due to lack of buffers.
- '*' each time no response is received.
- The source IP address of the response, if different from the previous response received.
- The round-trip time in milliseconds.
- '!P' if the response is a destination unreachable (protocol unreachable).
- '!H' if the response is a destination unreachable (host unreachable).

Bootstrap Protocol (BOOTP)

One restriction with the scheme described in the previous section is related to the use of the limited broadcast address for BOOTP requests; it requires that the server is on the same subnet as the recipient.

The 3746 IP Router, however, has a mechanism to forward BOOTP requests to one or multiple preconfigured BOOTP servers. If the BOOTP request's router address is zero, the 3746 BOOTP relay agent sets it to the IP address of its receiving interface.

The 3746 IP BOOTP relay agent discards the BOOTP request message when:

- BOOTP request's "hops" value is greater than the value configured by the user.
- BOOTP request's "hops" value is greater than or equal to 16.
- BOOTP request's "seconds (retry)" value is greater than the value configured by the user.

This allows BOOTP servers on the local network a chance to respond before the BOOTP relay agent has forwarded to remote BOOTP servers.

The BOOTP relay agent forwards a BOOTP reply message to the interface whose IP address matches the BOOTP reply's router IP address. Because the BOOTP client may not know its IP address and therefore may not be able to respond to an ARP request, the 3746 IP BOOTP relay agent installs an entry in the outgoing interface ARP cache table using the client IP and hardware address specified in the BOOTP reply message.

Once a BOOTP server has responded and the BOOTP client has processed the reply, it can proceed with the transfer of the boot file and execute the full boot process. The transfer of the boot file is usually done using the Trivial File Transfer Protocol (TFTP). The full boot process will replace the minimum IP protocol stack used by BOOTP and TFTP by a normal IP protocol stack transferred as part of the boot file and containing the correct customization for the client.

Security

The 3746 IP Router offer two methods to control its IP forwarding functions, namely by defining:

- Filters
- Access controls

3746 IP Filters

The 3746 IP filters are essentially user-defined static routes. Any packet received or originated by the 3746 IP Router in which the destination IP address matches a filter is discarded. No error message will be generated. Each filter, like a static route, consists of an IP address and a mask. When an IP datagram is forwarded, the destination address is ANDed with the filter mask. If the result equals the filter address, then the datagram is discarded.

$\text{destination_address} \& \text{filter_mask} = \text{filter_address} = \text{discard!}$

The 3746 IP Router installs filters in its routing table; therefore the number of filters is limited by the size of the routing table. Routes that match a filter are not advertised when using a dynamic routing protocol.

Example of Routing table w/tn Filters

9.0.0.0	255.0.0.0	Filter
9.100.0.0	255.255.0.0	Direct
9.150.0.0	255.255.0.0	9.200.1.1

FILTERS

- Defined as static routes
- Consists of an IP address + mask
- Saved in IP routing table & caches
- Packets discarded if destination address matches
- Routes matching filters not advertised
- OSPF routes takes precedence on filters
- Can be overridden by a more specific route (best choice)




Figure 52. IP Filters

3746 Access Control

Access control provides a way for the user to control which IP packets are forwarded by the 3746 IP Router based on the packet's source IP address, destination IP address, IP protocol number, and TCP or UDP port number. It is applied to all IP packets that are received or originated by the 3746 IP Router. Access control is an optional function that by default is disabled.

Note: The 3746 implements access control for the whole of the 3746, **not at the port (interface) level**. You cannot select individual ports for access control.

- Inclusive 9.150.2.2
- Exclusive 9.150.0.0 / 255.255.0.0
- Exclusive 9.100.0.0 protocol number = 17 (UDP)

ACCESS CONTROL

- Specifically defined (not routes)
- Inclusive or exclusive
- IP source address (+ mask)
- IP destination address (+ mask)
- Range of protocol numbers
- Range of port numbers
- Based on ordered lists
 - First matching if multiple entries
 - Less efficient than filters.
- Cache of 8 most recently seen packets.




Figure 53. IP Access Control

Multiple access entries can be specified, each having the parameters shown in Table 19.

Table 19 (Page 1 of 2). IP Multiple Access Parameters	
Parameter	Description
Access control type	The access control type can be inclusive or exclusive indicating that if this entry is matched the datagram will be forwarded, or discarded, respectively.
Source IP address	Value compared with the datagram's source IP address.
Source IP address mask	Value ANDed with the datagram's source IP address before comparing it with the access control entry's source IP address value.
Destination IP address	Value compared with the datagram's destination IP address.

Table 19 (Page 2 of 2). IP Multiple Access Parameters	
Parameter	Description
<i>Destination IP address mask</i>	Value ANDed with the datagram's destination IP address before comparing it with the access control entry's source IP address value.
<i>First protocol number</i>	The lower bound of a range of IP protocol numbers; a datagram matches the access control entry only if its IP protocol number is greater than or equal to this value.
<i>Last protocol number</i>	The upper bound of a range of IP protocol numbers; a datagram matches the access control entry only if its IP protocol number is less than or equal to this value.
<i>First port number</i>	The lower bound of a range of TCP/UDP port number; a packet in which the IP protocol number indicates TCP or UDP matches the access control entry only if its TCP or UDP port number is greater than or equal to this value.
<i>Last port number</i>	The upper bound of a range of TCP/UDP port numbers; a packet in which the IP protocol number indicates TCP or UDP matches the access control entry only if its TCP or UDP port number is less than or equal to this value.

The access control entries are organized into a single list that is used for all interfaces. The access list is ordered, meaning that if a packet matches multiple entries in the list, the first matching entry in the list is the one that takes effect. If no matching access control entry is found, the 3746 discards the datagram.

Note: CCM always adds an all-inclusive entry at the end of the list, meaning that datagrams are forwarded unless discarded because of other entries in the list. Use the Telnet operator interface to ensure that this entry is on that list.

Attention

PTR Z237641 needs to be examined to ensure correct IP addressing on the Service LAN. The **exclusive** and **inclusive** statements should be coded with care according to the particular installation.

Filters Versus Access Controls

Filters

A filter on a destination IP address is added to the IP routing table with indication **filter**. The position of a route in the routing table depends on its IP address. Filters are saved into the main routing table maintained on the CBSP2 and the routing table cache on each of the 3746 processors.

A filter is a route like a static route. The primary difference is that filters are not advertised by RIP, or BGP, and if a filter is replaced by another type of route (for example, OSPF) that later goes away, the filter is not automatically restored. OSPF internal routes (intra-area and inter-area) take preference over filter routes. This is to prevent unexpected reachability problems in OSPF networks. Because it is a link state protocol, OSPF does not take filters into account in its advertisements, so filters in an OSPF network can cause reachability problems. To be sure to filter the traffic, access control must be used instead.

Filters have another shortcoming: they can be partially overridden by a more specific route. Assume a filter for 9.0.0.0 with mask 255.0.0.0, and a route for 9.100.0.0 with mask 255.255.0.0 has been installed. Because the routing table search looks for the most specific match, datagrams to 9.100.x.x will follow the route instead of the filter.

Access Control: Access control entries are searched sequentially for each datagram until a matching entry is found. This may lead to considerable overhead when many access control entries have been specified. There is no special search of a routing table for a filter. The routing table is searched once for the best matching route, and if that match turns out to be a filter, the datagram is dropped.

Recommendation: In general, the use of filters is advised rather than using access control. However, you should be aware of the shortcomings of filters mentioned above.

IP over 3746 Networks

Figure 54 depicts how TCP/IP functions have been implemented on the 3746-9x0 components identified.

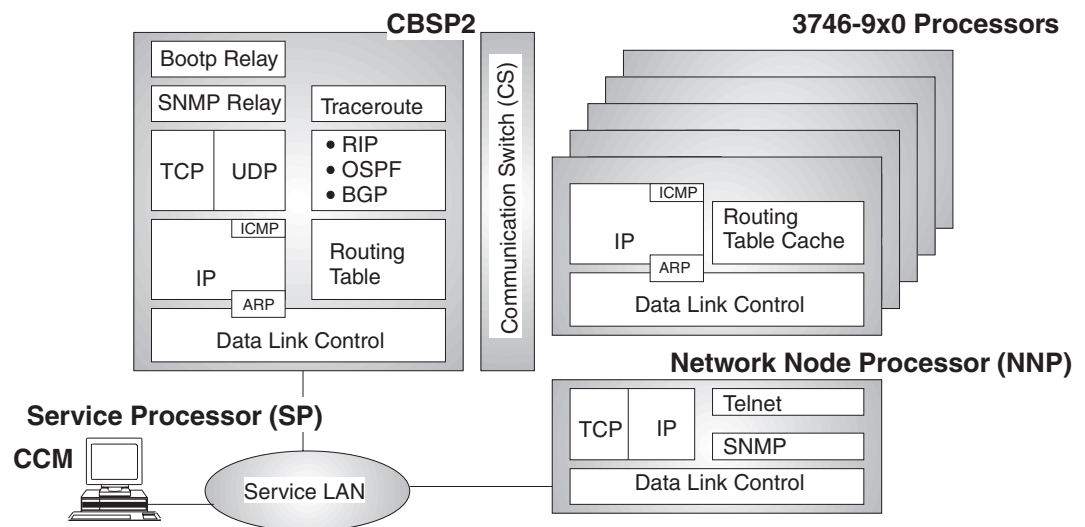


Figure 54. 3746-9X0 IP Base

Each 3746-9x0 processor, including the CBSP2, provides IP routing and data link control (DLC) functions. The DLC functions enable data transport to adjacent nodes. The DLC functions are shared between protocol stacks (IP and APPN/HPR). The IP functions allow IP transport between each of the 3746-9x0 adapters, and to and from externally attached IP equipment. To perform the IP routing, each adapter maintains a routing table cache. The code that performs the IP routing functions is present on all processors. However, it will be activated only on processors on which an IP interface has been configured. The routing cache is built during the transport of IP data, using routing information learned from the CBSP2.

In addition to the IP forwarding and DLC functions, more functions are required. All dynamic IP routing protocols (RIP, OSPF, and BGP) run on the CBSP. The CSBP2 is responsible for processing routing (RIP/OSPF/BGP) packets received

from, and generating routing packets destined for, external equipment. The adapters are instrumental in forwarding these packets between the external equipment and the CBSP2, and conversely.

The CSBP2 builds and maintains a routing table using static routes and IP interfaces defined on the 3746-9x0, and routing information learned from the dynamic routing protocols. To minimize the routing information maintained on each of the adapters, each processor maintains a routing table cache that contains a subset of the information stored in the CSBP2 routing table. When, during IP forwarding, a processor is not capable of making a routing decision based on the information maintained in its routing table cache, the appropriate routing information is retrieved by querying the CBSP2. Routes learned from the CSBP2 are stored within the routing table cache, and used to route successive IP datagrams. The Routing Table Caches contain one entry per destination host and **not** per destination network. This is to prevent less specific routes from masking more specific routes.

For example, if the most specific route that matches a particular destination IP address is the default route, if the default route is saved in the cache, then everything matches it. Thus everything is forwarded on the **default route** and more specific routes are ignored. Caching everything as host routes avoids the complications of dealing with more specific routes.

All cached entries are subject to an aging mechanism, and will be periodically removed.

Attached to the service LAN is the network processor (NNP). The main functions of the NNP are the 3746 IP operator and SNMP functions. In addition, the NNP contains the 3746 IP configuration file. When the 3746 IP code is restarted, each adapter retrieves a copy of the 3746 IP configuration file and configures itself. For details of SNMP functions, refer to “3746 Management Overview” chapter in the *3745/3746 Planning Series: Management Planning Guide*.

Also attached to the service LAN is the service processor (SP). The SP provides a repository of the microcode running on the 3746-9x0 adapters. Each 3746-9x0 adapter loads its microcode from the SP during an initial microcode load (IML). The 3746-9x0 controller configuration and management (CCM) tool can be accessed from the SP.

IP Interfaces

The 3746 IP Router supports attachment of IP equipment using ESCON, token-ring, Ethernet, or serial lines. For serial line-attached equipment, frame-relay, X.25, and Point-to-Point Protocol (PPP) connections can be used (see Figure 55 on page 118). Both switched, for example, using a public switched telephone network (PSTN), and leased-line PPP and frame-relay connections are supported.




Figure 55. IP Interfaces

Each of these interfaces is discussed in more detail in the following section.

ESCON Attachment

ESCON attachment can be used to provide native IP transport, using the channel data link control (CDLC), between the 3746 IP and host systems running TCP/IP for MVS™ (at least Version 3 Release 1). See Figure 56. The host systems can be directly attached to the 3746-9x0, or connected through an ESCON Director (ESCD).

The ESCON implementation on the 3746-9x0 is based on the concept of host links and link stations. A host link is a logical connection between a host system and the 3746-9x0. A link station represents a point-to-point link between the 3746-9x0 and host programs such as TCP/IP for MVS, VTAM, and/or TPF. Host link and link station definitions need to be entered before IP (or APPN/HPR or SNA subarea) communication is possible.




Figure 56. ESCON Port, Host Link, and Link Station

On each 3746-9x0 you can install up to 16 ESCON adapters (15 on a 3746-900 attached to a dual CCU Model 3745). Each ESCON adapter controls a single ESCON coupler. An ESCON coupler provides a single physical connection to a single host system, unless you are using an ESCON Director (ESCD). An ESCD provides a physical connection to multiple host systems (See Figure 57).

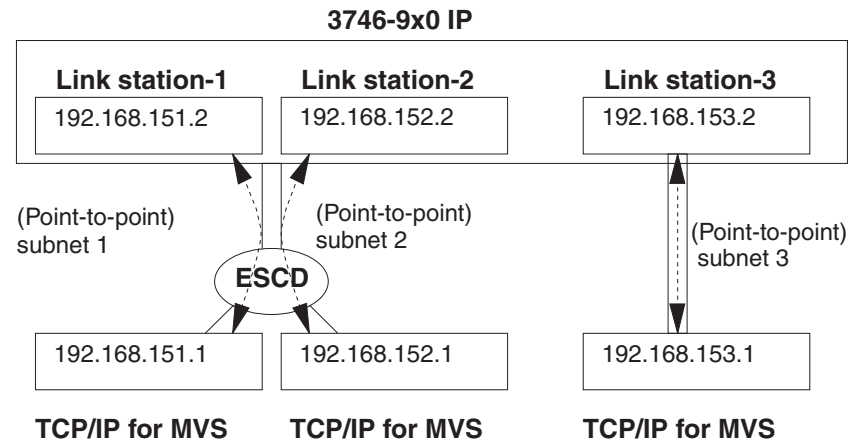


Figure 57. Multiple 3746 IP ESCON Links

Single IP addressing is possible on multiple ESCON stations to provide easier coding for multiple host access over the ESCON adapter.

On each ESCON adapter you can define up to 32 host links. For each host system you can define only one host link, unless the host system uses the ESCON Multiple Image Facility (EMIF), in which case you can define a separate host link (up to 32) to each of the logical partitions (LPs). On each host link you can define one or multiple link stations; however, the maximum number of link stations that can be defined per ESCON adapter is sixteen.

Note: Using the above figures, a single 3746-9x0 can provide ESCON attachment for up to 512 (16x32) host systems, and communicate with up to 512 host programs.

The ESCON configuration process consists of two phases:

1. A general definition part in which host links and link stations are defined, and the link stations are assigned to the appropriate protocol stack (3746 IP for IP routing).
2. The actual IP port configuration.

ESCON: General Configuration

To define an ESCON link, an ESCON port, a host link, and a link station definition are required. A host link, comprising a connection between a S/390 operating system and the 3746-9x0, is the logical equivalent of a port. On a single port, multiple host links can be defined, with multiple link stations on each host link. Figure 58 on page 120 depicts in configurations 1 and 2 an ESCON port on which a single host link and a single link station has been defined.

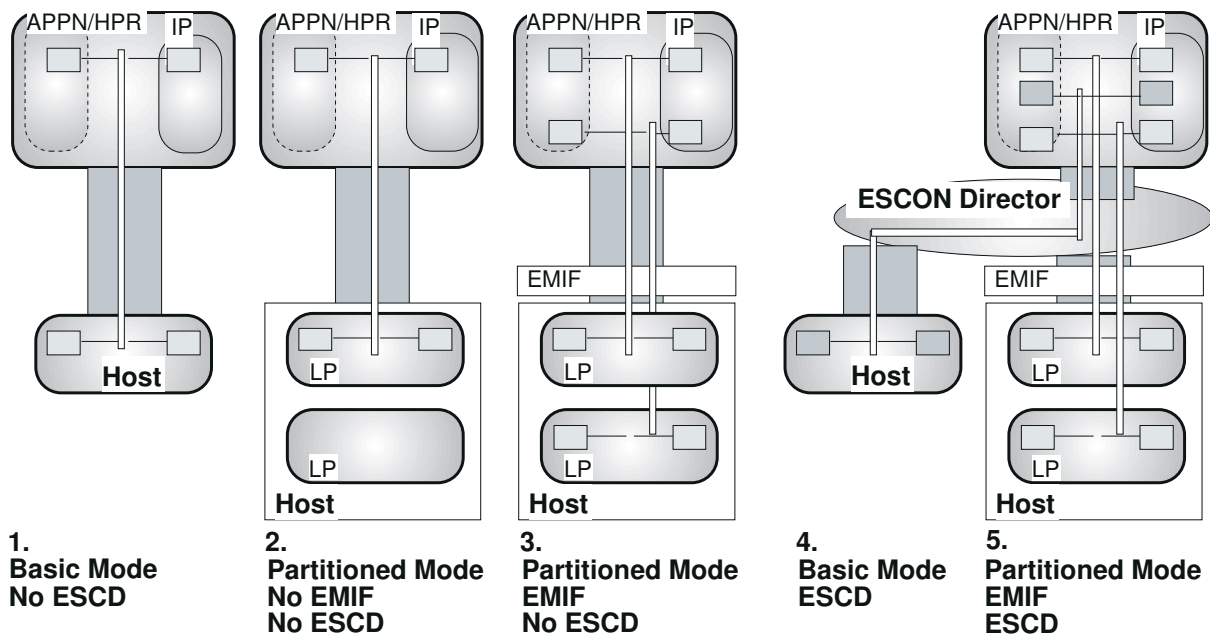


Figure 58. ESCON Host Links and Link Stations

The maximum number of host links that can be defined per ESCON adapter is 32. The actual number depends on the type of physical attachment between your host and the 3746-9x0:

Configurations 1 and 2

If your host is directly attached and running in basic (non-partitioned) mode only, or in the partitioned mode without the IBM ESCON Multiple Image Facility (EMIF), then a single host link can be defined.

Configuration 3

If your host is directly attached, is running in partitioned mode, and uses the ESCON Multiple Image Facility (EMIF), then a single host link can be defined to each LP.

Configuration 4

If your host is connected through an ESCON Director (ESCD), then a single host link is possible to each host running in basic (non-partitioned) mode.

Configuration 5

If your host is connected through an ESCON Director (ESCD), then a single host link is also possible to each LP on the host running in partitioned mode and using the ESCON Multiple Image Facility (EMIF).

Link stations map one-to-one with I/O addresses on your host. By using a specific I/O address to, for example, TCP/IP for MVS and assigning the corresponding link station to the 3746-9x0 IP functions, native IP communication between the two is possible.

ESCON: Port Sharing

As discussed in the previous section, up to 32 link stations can be defined per ESCON adapter. Each link station provides a point-to-point connection. See Figure 59 on page 121.




Figure 59. ESCON Port Sharing

ESCON ports can be defined and simultaneously activated from all protocol stacks available (IP, APPN/HPR, and for 3746-900 only; NCPs running on CCU-A and CCU-B on the attached 3745). On a 3746-900, each link station must be assigned to either 3746 IP, NCP/CCU-A, NCP/CCU-B, or the 3746 NN function; on a 3746-950, a link station must be assigned to either the 3746 IP, NCP or the 3746 NN function.

Link stations map one-to-one on I/O addresses defined on attached host systems. The I/O addresses are used by host programs to communicate with the corresponding protocol stack. The 3746-9x0 IP protocol stack can communicate only with TCP/IP for MVS (Version 3, Release 1 or 2), while the 3746 APPN/HPR and/or NCP protocol stacks can communicate with VTAM and/or TPF.

Link stations are independently assigned, and, therefore, an ESCON port and each host link can be shared between up to four (two on the 3746-950) 3746 protocol stacks. The total number of link stations per physical port cannot exceed 32.

ESCON: IP Addresses and Subnet Addresses Rules

The following section explains the rules for IP address definitions for ESCON channels.

The following general rules apply:

- One (physical) host can run multiple IP address spaces. Each IP address space is a full TCP/IP MVS stack (partition).
- Each IP address space is independent of the others.
- Each IP address space can run one or multiple channel data link control (CDLC) instances.
- Each CDLC instance of a given IP address space must run on a separate ESCON fiber.

The following rules are defined by TCP/IP for MVS:

- CDLC instances of a **single** IP address space cannot have the same next-hop IP address.
- CDLC instances of **different** IP address spaces can have the same next-hop IP address, because they are in independent IP stacks.

The following rules are defined by the 3746:

- For hosts with only one IP address space and one CDLC in the IP address space, you can define a single subnet between the hosts and the 3746 stations.




Figure 60. Four ESCP Stations with the Same IP Addresses

- For hosts with many IP address spaces and one CDLC per IP address space, you can define a single subnet between the hosts and the 3746 stations.





Figure 61. Five ESCP Stations with the Same IP Addresses

- For hosts with many IP address spaces and many CDLCs per IP address space, you **must** define multiple **subnets**, the basic rule being that two CDLCs of the same MVS IP address space cannot be in the same subnet.



Subnets:

- 10.0.0.0 used for HOST1/CDLC1, HOST2/CDLC1 and ESCP LS1 and LS3
- 22.0.0.0 used for HOST1/CDLC2, HOST2/CDLC2 and ESCP LS2 and LS4
- 33.0.0.0 used for HOST2/CDLC3 and ESCP LS5.

Figure 62. ESCP Stations on Separate Subnets

ESCON: IP Configuration

The IP configuration process for ESCON link stations consists of:

- Assigning an IP address and subnet mask to the link station. Each link station represents a separate point-to-point CDLC link, therefore, for each link station an IP address must be assigned that is part of a different subnet.
- Enabling dynamic routing protocol on this interface (or define static routes). Assigning an IP address (see “IP Addressing” on page 84) and subnet mask (see “The Subnet Mask” on page 87) is mandatory. In addition, the channel adapter slowdown limit (CASDL), the attention timer (TIMEOUT), and the delay time (DELAY) timer have to be defined.

Enabling dynamic routing protocols is optional. It is required only when the TCP/IP for MVS system has multiple IP connections, and static routing does not suffice. The only dynamic routing protocol currently supported by TCP/IP for MVS (Version 3 Release 1 and 2) is RIP.

Token-Ring Attachment

Token-ring attachment can be used to provide IP transport between the 3746 IP and IP stations or routers running TCP/IP.

For the IP communication over token-ring, the 3746-9x0 uses (connectionless) IEEE 802.2 logical link control (LLC). IEEE 802.2 LLC requires the use of a local and a remote service access point (SAP). For IP communication both local and remote SAP should be equal to X'AA', indicating sub-network access protocol (SNAP) encapsulation. See Figure 63. The (optional) **routing information field (RIF)** in the MAC header indicates the support for source route bridging, which is by default available.

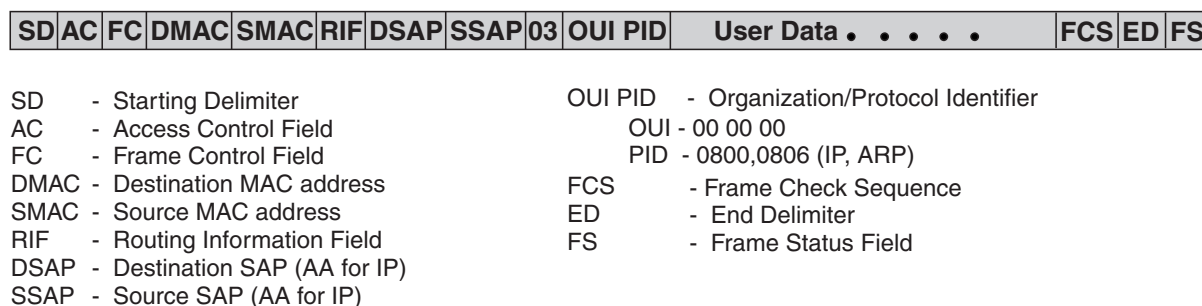


Figure 63. Token-Ring Encapsulation

On each 3746-9x0, you can install up to 16 token-ring adapters. Each token-ring adapter can control up to two token-ring interface couplers (TIC3s). Each token-ring coupler provides a single physical connection to a token-ring LAN.

Note: The 3746-9x0 supports the attachment of up to 32 token-ring LANs (31 on a 3746-900 attached to a dual CCU Model 3745).

The token-ring configuration process consists of two phases:

1. A general definition part in which token-ring port specifics are defined
2. The actual IP port configuration.

Token-Ring: General Configuration

To define a token-ring port you have to specify its IP and APPN/HPR (if relevant) names, the speed of the token-ring LAN, and the local MAC address. The APPN and HPR SAPs are relevant only if the port is also used for APPN traffic. Make sure the IP maximum transmission unit (MTU) is high enough to prevent fragmentation of IP datagrams on the 3746 IP Router. However, in a bridged environment, make sure that the MTU is smaller than the lowest value configured on any of the LAN bridges. Select IP parameters to enter the IP over token-ring configuration. For IP transport no station needs to be defined.

Token-Ring: Port Sharing

Token-ring ports can be defined on all protocol stacks. However, simultaneous activation of the port is limited to the IP protocol stack, APPN/HPR protocol, and (3746-900 only) either the NCP running on CCU-A or the NCP running on CCU-B of the attached 3745. Therefore, on the 3746-900, a token-ring port can be shared between the 3746 IP, a single NCP and the 3746 NN function; on a 3746-950, the token-ring port can be shared between the 3746 IP and the 3746 NN function. The 3746-9x0 data link control layer differentiates traffic from and to any of these

protocol stacks using the service access points (SAPs) within the token-ring MAC header.




Figure 64. Token-Ring Port Sharing

When receiving data the 3746-9x0 will accept the following source SAPs:

- X'AA' Data to or from 3746 IP stack
- X'04' SNA data to or from the NCP protocol stack (3746-900 only)
- X'C8' HPR non-ERP data to or from the NCP protocol stack (3746-900 only)
- X'xx' APPN data to or from the 3746 NN protocol stack, where xx is configured in the APPN local SAP field during token-ring port configuration
- X'yy' HPR non-ERP data to or from the 3746 NN protocol stack, where yy is configured in the APPN/HPR local SAP field during token-ring port configuration

Token-Ring: IP Configuration

The IP configuration for token-ring ports consists of:

- Assigning an IP address and subnet mask to the token-ring port.

Note that the 3746 IP Router allows you to specify multiple IP addresses and subnet masks to the same physical interface. When defining multiple IP addresses, the 3746 IP Router connects to multiple subnets. The 3746-9x0 is capable of routing between subnets on the same physical interface, although in general this is not recommended.

- Enabling dynamic routing protocol for each IP address (or define static routes).

The 3746 IP dynamic routing protocols operate independently for each IP address defined on a token-ring interface; therefore, each IP address is defined in the configuration.

Assigning an IP address and subnet mask is mandatory. At least one but optionally multiple (up to 16) IP addresses can be defined. If multiple IP addresses are defined, make sure that the addresses are part of different subnets.

The RIF timer indicates how many seconds will elapse between the last reference to an ARP entry and when it will be deleted. Increasing the timer may result in more storage being required to cache ARP entries. A zero value will disable the ARP caching function.

Enabling dynamic routing protocols is optional. It is only required when any of the token-ring attached devices have multiple IP interfaces, and static routing does not suffice. Dynamic routing protocols supported by the 3746-9x0 are RIP, OSPF, and BGP.

External IP Connection Between 3745 and 3746-900

With the introduction of NCP V7R6 it is now possible to link the NCP IP router and the 3746 IP router together via the CBC between the 3745 and 3746. Previously, an external link between the 3745 and 3746 was need to connect the two routers. To support this function from the 3746 side, the optional FC 5800 is needed. This is supported by microcode level D46130I (ECA 170) and later versions.

The internal connection or connections between the two routers are seen as token-ring point-to-point (PtP) IP connections, these connections must be defined in the NCP and 3746 CCM. Each connection uses a separate TIC3 on the 3746, this TIC3 may also be used to drive a token-ring, but beaconing on that token-ring or other problems may interfere with the IP routing function. For that reason it is recommended, but not mandatory, that each TIC3 used for an internal IP connection is used only for that function.

Figure 65 shows a representation of the internal point-to-point IP connection.

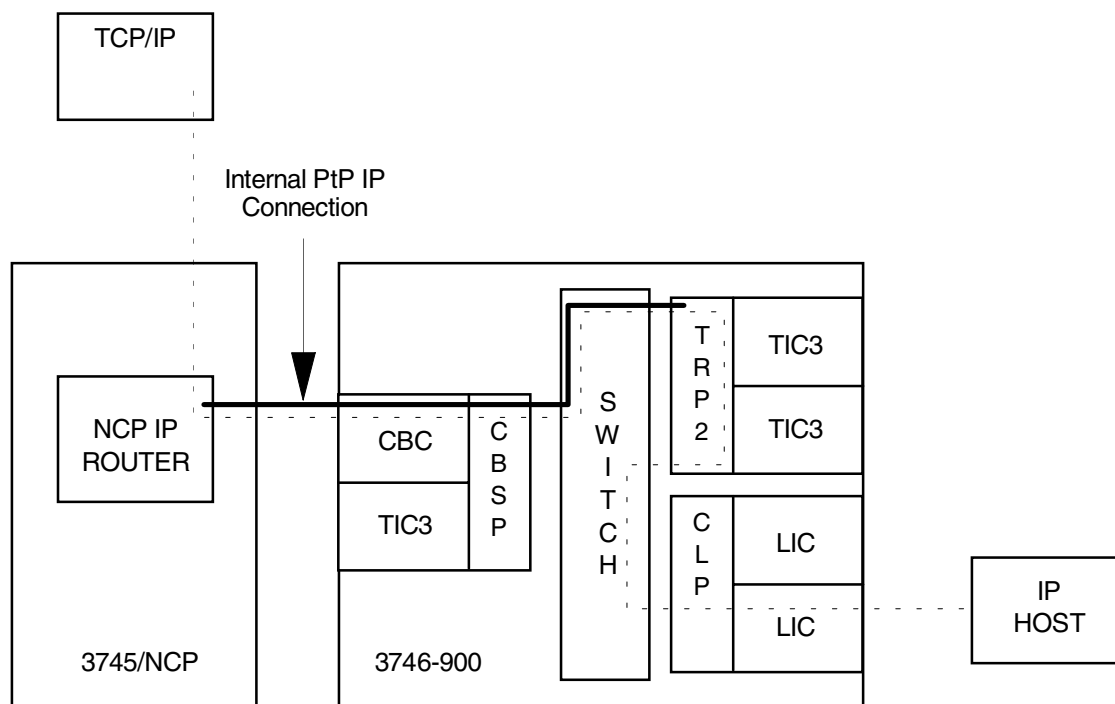


Figure 65. Internal IP PtP Connection

Internal IP Connection Definitions

The following sections explain how to define an internal point-to-point IP connection. Figure 66 shows the configuration that is being defined. The definitions shown here include all the information needed for this part of the NCP generation process: the actual PtP link is defined in the LN2880 LINE (**1**), and IP2880 PU (**2**) statements, and the IPLOCAL statement for LADDR=10.04.00.99 (**5**). The NETWORK must be IP (**3**), and PUTYPE must be 1 (**4**).

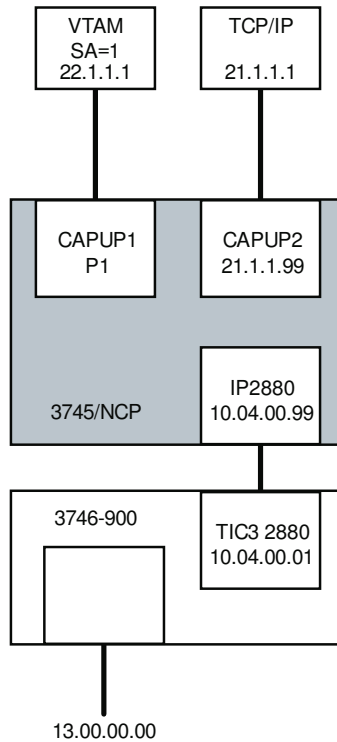


Figure 66. Example Configuration

NCP Definitions: The following examples show the NCP definitions needed to configure the internal IP PtP link from the 3745 side.


```

*****
* VTAM HOST DEFINITIONS
*****
*
HOSTA01  HOST  BFRPAD=0,DELAY=0.0,INBFRS=6,MAXBFRU=32,
              SUBAREA=1,
              STATMOD=YES,UNITSZ=1024
*
*****
* TOKEN-RING LINE 2880
*****
*
ODLCGRP1 GROUP ECLTYPE=(PHY,ANY),ADAPTER=TIC3,
              DIAL=NO,LNCTL=SDLC,TYPE=NCP
LN2880   LINE  ADDRESS=(2880,FULL) 1,
              LOCADD=400000002880,
              MAXPU=2,
              MAXTSL=16732,
              NPACOLL=YES,
              PORTADD=2,
              SPEED=9600,
              TRSPEED=16
IP2880   PU    ADDR=02,
              INTFACE=(TR2880,2048) 2,
              NPACOLL=YES,
              NETWORK=IP 3,
              PUTYPE=1 4,
              ARPTAB=(1)

*****
* CHANNEL ADAPTERS
*****
*
CAGRP1   GROUP LNCTL=CA,CA=TYPE7,NCPCA=ACTIVE,MONLINK=CONT,
              TIMEOUT=240.0,ISTATUS=ACTIVE,CASDL=420.0,DELAY=0.0
*
*****
* VTAM HOST CHANNEL ATTACHMENT
*****
*
CALNP1   LINE  ADDRESS=P1
CAPUP1   PU    PUTYPE=5,TGN=1
*
*****
* CHANNELS FOR IP CHANNEL ATTACHED ROUTERS
*****
*
CALNP2   LINE  ADDRESS=P2,CASDL=420,DELAY=0.0,MONLINK=NO
CAPUP2   PU    PUTYPE=1,INTFACE=CPU2,ARPTAB=(10,,NOTCANON)

```

```

*****
*      IP ROUTING DEFINITIONS
*****
*
      IPOWNER  HOSTADDR=21.1.1.1,NUMROUTE=(25,25,25),UDPPORT=580,
                INTERFACE=(CPU2),MAXHELLO=6,NUMDRIF=10

      IPLOCAL  LADDR=21.1.1.99,INTERFACE=CPU2,METRIC=1,
                P2PDEST=21.1.1.1,PROTOCOL=RIP
*
      IPLOCAL  LADDR=10.04.00.99 5, INTERFACE=TR2880, METRIC=1,
                P2PDEST=10.04.00.01 6, PROTOCOL=RIP, SNETMASK=255.255.0.0
*
      IPRUTE   DESTADDR=13.0.0.0 7,
                NEXTADDR=10.4.00.01 8,
                INTERFACE=TR2880,
                METRIC=1,DISP=PERM,HOSTRT=NO
*
      GENEND   GENEND
                END

```

For reference, the following list provides a short description of some of the parameters used:

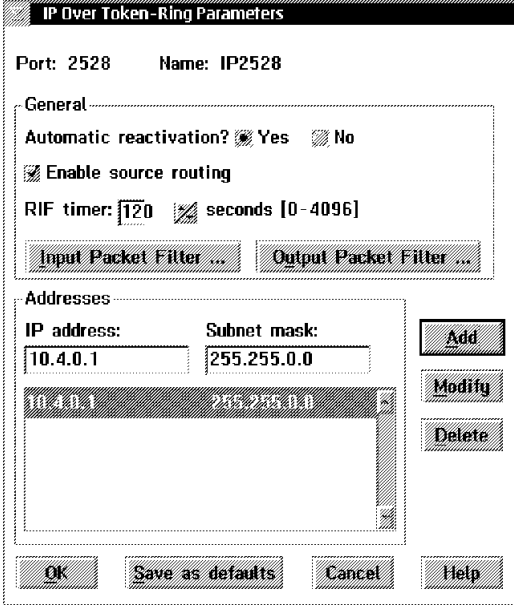
- IPOWNER** Identifies the TCP/IP MVS/VM host that manages the routing tables (NCPROUTE)
- IPLOCAL** Defines an interface to the NCP IP router.
- IPROUTE** Defines an entry in the NCP IP routing table.
- HOSTADDR** Defines the IP address of the owning IP host.
- INTERFACE** Defines the name of the IP interface to NCPROUTE.
- LADDR** Defines the IP address of the interface.
- P2PDEST** Defines IP address of the destination IP host.
- SNETMASK** Defines subnet mask for the interface.

3746 CCM Definitions:

1. The token-ring port 2880 must be configured first, and IP must be activated on that port.

Figure 67. Port Configuration

- On port 2880, an IP address must be defined. This is the IP address of the 3746 end of the PtP IP connection, 10.4.0.1. This must match the IP address specified on the P2PDEST operand on the IPLOCAL statement in NCP.



IP Over Token-Ring Parameters

Port: 2528 Name: IP2528

General

Automatic reactivation? ☒ Yes ☐ No

☒ Enable source routing

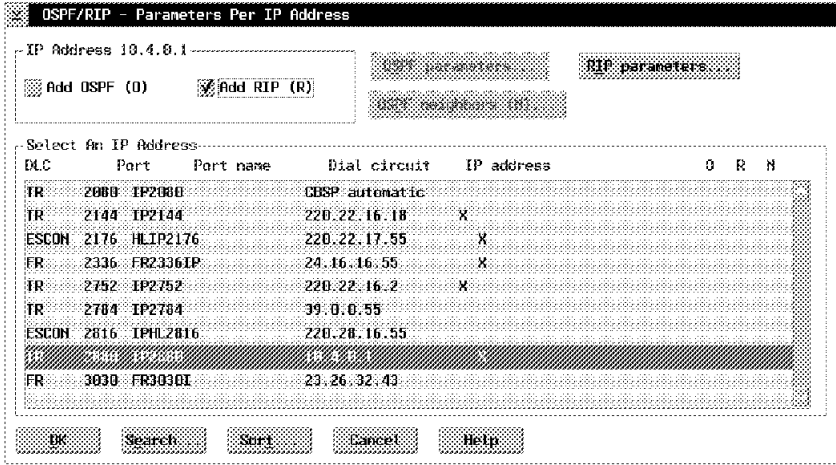
RIF timer: 120 ☒ seconds [0-4096]

Addresses

IP address: 10.4.0.1 Subnet mask: 255.255.0.0

Figure 68. IP over Token-Ring Parameters

- RIP should be activated as shown on port 2880.



OSPF/RIP - Parameters Per IP Address

IP Address 10.4.0.1

☒ Add OSPF (O) ☒ Add RIP (R)

Select An IP Address

DLC	Port	Port name	Dial circuit	IP address	O	R	N
TR	2080	IP2080	CESP automatic				
TR	2144	IP2144	220.22.16.18	X			
ESCON	2176	HLIP2176	220.22.17.55	X			
FR	2336	FR2336IP	24.16.16.55	X			
TR	2752	IP2752	220.22.16.2	X			
TR	2784	IP2784	39.0.0.55				
ESCON	2816	IPHL2816	220.28.16.55				
FR	2816	IPHL2816	112.8.38				
FR	3030	FR3030I	23.26.32.43				

Figure 69. OSPF/RIP Parameters per IP Address

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

RIP - Parameters Per IP Address

IP address: 10.4.0.1

Broadcast address style: ☒ Network ☐ Local-wire

Address fill pattern: ☒ Zeroes ☐ Ones

Interface tag (AS number): 1 ☒ numerical [1-65535]

☒ Send RIP routes ☒ Receive RIP routes

☒ Send net routes ☒ Receive net routes

☒ Send subnet routes ☒ Receive subnet routes

☒ Send host routes ☒ Receive host routes

☒ Send static routes ☐ Override static routes

☒ Send default routes ☐ Override default routes

Inbound metric: 1 ☒ numerical [1-15]

Outbound metric: 0 ☒ numerical [0-15]

☒ Enable RIP V2 ☐ Enable sending RIP V1 routes only

☐ RIP V2 authentication Key:

OK Save as defaults Cancel Help

Figure 70. RIP Parameters per IP Address

To use the internal PtP link for IP routing to reach the destination network 13.0.0.0 (DESTADDR=13.0.0.0 **7**) from NCP, the next hop (NEXTADDR **8**) points to the IP address of the 3746 end of the PtP connection 10.4.0.1.

Internal IP Connection Between 3745 and 3746-900

The previous sections explained that on a 3746-900, NCP and 3746 IP functions can coexist. Because NCP is also providing IP routing functions, it is important that to make a distinction between the 3746 IP (stand-alone) and the NCP-controlled IP (NCP-IP) routing functions.

3746 IP provides IP routing for ESCON channel, token-ring (TIC3), Ethernet, and serial line (frame relay, X.25, and PPP) attached equipment. Connectivity is provided through 3746-900 attachments only. NCP-IP provides IP routing for ESCON channel, parallel channel, token-ring (TIC1/TIC2), Ethernet, and serial line (frame relay) attached equipment. However, with the exception of ESCON, NCP-IP connectivity is provided through base attachments only, that is, for equipment not attached to the 3746-900.

3746 IP and NCP-IP perform their IP routing functions in a fully independent fashion. IP connectivity between both IP routers requires an external connection. Internal IP connections are not supported through the CBC. Figure 71 on page 133 depicts how direct IP connectivity between NCP-IP and 3746 IP can be realized through token-ring, Ethernet, and frame relay. In all cases, a port (token-ring, Ethernet, or frame relay, respectively) is required on both the 3746-900 and the 3745 base frame (or a 3746 Model L1x).




Figure 71. Direct NCP-IP 3746 Routing

Figure 72 shows an alternative for the direct 3746 IP to NCP-IP attachment. A single ESCON channel can be used to connect TCP/IP for MVS (Version 3, Release 1 or 2) to both 3746 IP and NCP-IP. The result, however, of IP connectivity between 3746 IP and NCP-IP is that all IP datagrams that are sent between 3746 IP and NCP-IP will be routed through TCP/IP for MVS.





Figure 72. NCP-IP 3746 Routing Over ESCON

Telnet Operations through the 3746 Network Node

The Telnet protocol provides a standardized interface through which a program on one host (the Telnet client) can access the resources of another host (the Telnet server) through a 3746 Network Node as though the client were a local terminal connected to the server (see Figure 73 on page 134). For example, a user on a workstation attached to a LAN can connect to a host as though the workstation were a terminal attached directly to the host.

The 3746 Network Node provides the through connection.



How Telnet appears to the host and client.



Figure 73. Telnet Connection via a 3746 Network Node

Telnet allows the LAN-attached user to log in the same way as the local terminal user.

Note: The 3746 Network Node provides the connection but does not provide the actual host/client function. Refer to the documentation for the host and client systems for the functions available on those hosts and clients. This chapter gives an overview of Telnet possibilities.

Basic Operation

The Telnet protocol is based on three ideas:

- Network Virtual Terminal (NVT) concept

An NVT is an imaginary device having a basic structure common to a wide range of real terminals. Each host maps its own terminal characteristics to those of an NVT, and assumes that every other host will do the same.

- A symmetric view of terminals and processes
- Negotiation of terminal options

The principle of negotiated options is used by the Telnet protocol, because many hosts can provide additional services, beyond those available with the NVT. Various options can be negotiated. After this minimum understanding is achieved, they can negotiate additional options to extend the capabilities of the NVT to reflect more accurately the capabilities of the real hardware in use. Because of the symmetric model used by Telnet, both the host and the client can propose additional options to be used.

Telnet operations, enabling and disabling, can be protected by passwords.

Network Virtual Terminal

The NVT has a printer (or display) and a keyboard. The keyboard produces outgoing data, which is sent over the Telnet connection. The printer receives the incoming data. The basic characteristics of an NVT, unless they are modified by mutually agreed options, are:

- The data representation is 7-bit ASCII transmitted in 8-bit bytes.
- The NVT is a half-duplex device operating in a line-buffered mode.
- The NVT provides a local echo function.

All of these can be negotiated by the two hosts.

NVT Printer

An NVT Printer has an unspecified carriage width and page length. It can handle printable ASCII characters (ASCII code 32 to 126) and understands some ASCII control characters.

Full-Screen Capability

Full-screen Telnet is possible provided that the client and server have compatible full-panel capabilities. For example, VM and MVS provide a TN3270-capable server. To use this facility, a Telnet client must support TN3270.

Command Structure

The communication between client and server is handled with internal commands, which are not accessible by users. All internal Telnet commands consist of 2-or 3-byte sequences, depending on the command type. An Interpret As Command (IAC) character is followed by a command code. This command deals with option negotiation, and has a third byte to show the code for the referenced option.

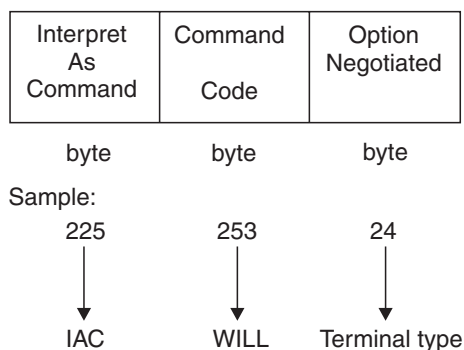


Figure 74. Telnet Command Structure

Host/Client Implementations

You can only use Telnet if it is implemented on **both** client **and** host machines, refer to your documentation for both. The documentation should also tell you the full range of Telnet functions available for those systems.

Mandatory Access Control Entry

The following entry in the access control list is mandatory and should ***not*** be removed:

```
PERMIT 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0
```

If this entry is removed, NetView/6000 cannot gain access to the SNMP agent that resides on the NNP and no SNMP management of the 3746 will be possible.

List of Abbreviations

AB	area border	CLIST	command list
ACF	advanced communications function	CLA	communication line adapter
ACF/VTAM	advanced communications function for the virtual telecommunications access method	CLP	communication line processor
ANR	automatic network routing	CM	communications manager
APPN	advanced peer-to-peer networking	CNN	composite network node
ARB	adaptive rate-based flow/congestion control	CNM	communication network management
ARC	active remote connector	COS	cost of service
ARP	address resolution protocol	CP	control point
AS	autonomous system	CR	communications rate
ASB	autonomous system border	CSU	customer service unit
ASE	autonomous system external	DCAF	distributed console access facility
ASCII	american national standard code for information interchange	DCE	data circuit-terminating equipment
AUTO	automatic	DDS	digital data service
BAN	boundary access node	DE	discard eligibility
BECN	backward explicit congestion notification	DLC	data link control
BER	box event record	DLCI	data link connection identifier
BGP	border gateway protocol	DLSw	data link switching
BOOTP	bootstrap protocol	DLUR	dependent LU requester
bps	bits per second	DLUS	dependent LU server
BRS	bandwidth reservation system	DMUX	double multiplex circuit
BSC	binary synchronous communication	DSU	data service unit
C&SM	communications and system management	DTE	data terminal equipment
CBSP	control bus and service processor	DX	duplex
CCITT	Comité Consultative International Télégraphique et Téléphonique The international telegraph and telephone consultative committee	EBCDIC	extended binary-coded decimal interchange code
CCU	central control unit	EBN	extended border node
CD	carrier detector	EC	engineering change
CDF-E	configuration data file - extended	EMIF	ESCON multiple image facility
CE	customer engineer	EN	end node
CF3745	3745 and 3746 configurator and performance model	EP	emulation program
CHPID	channel path id	EPO	emergency power OFF
CIDR	classless inter-domain routing	ESCA	ESCON channel adapter
CIR	committed information rate	ESCC	ESCON channel coupler
		ESCD	ESCON Director
		ESCON	Enterprise Systems Connection
		ESCP	ESCON processor
		FC	feature code
		FDX	full duplex
		FECN	forward explicit congestion notification

FRAD	frame-relay access device	LQ	line quality
FRFH	frame-relay frame handler	LU	logical unit
FRSE	frame-relay switching equipment	LSS	low-speed scanner
FRTE	frame-relay terminating equipment	MAC	medium access control
HCD	hardware configuration definition	MAU	medium attachment unit
HDX	half duplex	MB	megabyte (processor storage) 1MB = 2 ²⁰ bytes (1 048 576 bytes)
HI	high	Mbps	megabits per second (speed or communication volume per second) 1 Mbps = 1 000 000 (one million) bits per second
HLA	host link address	MCL	microcode change level
HONE	hands-on network environment	MES	miscellaneous equipment specification
HPR	high performance routing	MIB	management information base
HSS	high-speed scanner	MIH	missing interrupt handler
ICMP	internet control message protocol	MLC	machine level control
IML	initial microcode load	MLTG	multi-link transmission group
INN	intermediate network node or IBM information network	MOSS-E	maintenance and operator subsystem - extended
IOCP	Input/Output Configuration Program	MTP	multipoint
IP	internet, or internetwork, protocol	MUX	multiplex circuit
IPL	initial program load	MVS	multiple virtual storage
IPR	installation planning representative	NAU	network addressable unit
ITU-T	international telecommunications union - telecommunications (ex-CCITT)	NMBA	nonbroadcast multiaccess
KB	kilobyte (processor storage) 1KB = 2 ¹⁰ bytes (1 024 bytes)	NCP	Network Control Program
kbps	kilobits per second (speed or communication volume per second) 1 kbps = 1 000 (one thousand) bits per second	NDRS	non-disruptive route switching
LAA	locally administered address	NGMF	netView graphic monitor facility
LAN	local area network	NN	network node
LCB	line connection box	NNP	network node processor
LCBB	line connection box base	NPM	netView performance monitor
LCBE	line connection box expansion	NRZI	non-return-to-zero inverted
LCP	link control protocol	NVT	network virtual terminal
LDM	limited distance modem	ODLC	outboard data link control
LED	light emitting diode	OSPF	open shortest path first
LIB n	line interface board type n	PBN	peripheral border node
LIC n	line interface coupler type n	PCI	Peripheral component interconnect
LSA	link state advertisement	PEP	partitioned emulation program
LIU n	line interface coupler unit type n	PING	packet internet groper
LIV	link integrity verification	PN	peripheral node
LMI	local management interface	PPP	point-to-point protocol
LP	logical partition	PPPNCP	point-to-point network control protocol
LPDA®	link problem determination aid	PTP	point-to-point

PTT	post, telegraph, and telephone	SRC	service reference code
PU	physical unit	S/S	start-stop
PVC	permanent virtual circuit	SVC	switched virtual circuit
QUAL	quality	TC	test control
RCV	receive clock	TCM	trellis code modulation
RETAIN®	remote technical assistance information network	TCP	transmission control protocol
RFS	ready for sending	TG	transmission group
RIP	routing information protocol	THRES	threshold
RNR	receive not ready	TICn	Token-ring interface coupler type n
ROS	read-only storage	TIM	time services
RR	receive ready	TOS	type of service
RSF	remote support facility	TPF	transaction processing facility
RTP	rapid transport protocol	TRA	Token-ring adapter
RTS	request to send	TRP	Token-ring processor
SDLC	synchronous data link control	TSS	transmission subsystem
SMUX	single multiplex circuit	UDP	user datagram protocol
SNBU	switched network backup	UTP	unshielded twisted pair
SNI	SNA network interconnection	VTAM	virtual telecommunications access method
SNMP	simple network management protocol	XID	exchange station identification
SPAU	service processor access unit	XMIT	transmit

Glossary

This glossary defines new terms used in this manual.

adaptive rate-based flow and congestion control (ARB). A function of High Performance Routing (HPR) that regulates the flow of data over an RTP connection by adaptively changing the sender's rate based on feedback on the receiver's rate. It allows high link utilization and prevents congestion before it occurs, rather than recovering after congestion has occurred.

advanced communication function (ACF). A group of IBM licensed programs, principally VTAM programs, TCAM, NCP, and SSP, that use the concepts of Systems Network Architecture (SNA), including distribution of function and resource sharing.

advanced communications function for the virtual telecommunications access method (ACF/VTAM). An IBM licensed program that controls communication and the flow of data in an SNA network. It provides single-domain, multiple-domain, and interconnected network capability.

advanced peer-to-peer networking (APPN). Data communications support that routes data in a network between two or more advanced program-to-program communications (APPC) systems that do not need to be adjacent.

automatic network routing. A function of High Performance Routing (HPR) that provides a low-level routing mechanism that requires no intermediate storage.

channel adapter (CA). A communication controller hardware unit used to attach the controller to a host processor.

communication controller. A device that directs the transmission of data over the data links of a network; its operation may be controlled by a program executed in a processor to which the controller is connected or it may be controlled by a program executed within the device. For example, the IBM 3745 and 3746 Network Nodes.

communications manager. A function of the OS/2 Extended Edition program that lets a workstation connect to a host computer and use the host resources as well as the resources of the other personal computers to which the workstation is attached, either directly or through a host system. The communications manager provides application programming interfaces (APIs) so that users can develop their own applications.

configuration data file - extended (CDF-E). A 3746 Network Node MOSS-E file that contains a description

of all the hardware features (presence, type, address, and characteristics).

communications management configuration host node. The type 5 host processor in a communications management configuration that does all network-control functions in the network except for the control of devices channel-attached to a data host nodes. Synonymous with communications management host. See also data host node.

control panel. A panel that contains switches and indicators for the customer's operator and service personnel.

control program. A computer program designed to schedule and to supervise the execution of programs of the controller.

control subsystem. The part of the controller that stores and executes the control program, and monitors the data transfers over the channel and transmission interfaces.

customer engineer. See IBM service representative

data circuit-terminating equipment (DCE). The equipment installed at the user's premises that provides all the functions required to establish, maintain, and terminate a connection, and the signal conversion between the data terminal equipment (DTE) and the line. For example, a modem is a DCE.

Note: The DCE may be a stand-alone equipment or integrated in the 3745.

data terminal equipment (DTE). That part of a data station that serves as a data source, data link, or both, and provides for the data communication control function according to protocols. For example, the 3174 and PS/2s are DTEs.

data host node. In a communication management configuration, a type 5 host node that is dedicated to processing applications and does not control network resources, except for its channel adapter-attached or communication adapter-attached devices. Synonymous with data host. See also communications management configuration host node.

direct attachment. The attachment of a DTE to another DTE without a DCE.

ESCON channel. A channel having an Enterprise System Connection* channel-to-control-unit I/O interface that uses optical cables as a transmission medium.

ESCON channel adapter (ESCA). A communication controller hardware unit used to attach the controller to a host via ESCON fiber optics. An ESCA consists of an ESCON channel processor (ESCP) and an ESCON channel coupler (ESCC).

ESCON channel coupler (ESCC). A communication controller hardware unit which is the interface between the ESCON channel processor and the ESCON fiber optic cable.

ESCON channel processor (ESCP). A communication controller hardware unit which provides the channel data link control for the ESCON channel adapter.

distributed console access facility. (1) This program product provides a remote console function that allows a user at one programmable workstation (PS/2) to remotely control the keyboard input and monitor the display of output of another programmable workstation. The DCAF program does not affect the application programs that are running on the workstation that is being controlled. (2) An icon that represents the Distributed Console Access Facility.

enterprise systems connection (ESCON). A set of IBM products and services that provides a dynamically connected environment within an enterprise.

Host. See host processor

host processor. (1) A processor that controls all or part of a user application network. (2) In a network, the processing unit where the access method for the network resides. (3) In an SNA network, the processing unit that contains a system services control point (SSCP). (4) A processing unit that executes the access method for attached communication controllers.

High performance routing (HPR). An extension of APPN that provides faster traffic throughput, lower delays, and lower storage overheads.

IBM service representative. An individual in IBM who does maintenance services for IBM products or systems. Also called the IBM *Customer Engineer*.

initial microcode load (IML). The process of loading the microcode into an adapter, the MOSS, or the service processor.

internet. (1) A wide area network connecting disparate networks using the internetwork protocol (IP) (2) A public domain wide area network connecting thousands of disparate networks in industry, education, government and research. The Internet uses TCP/IP as the standard for transmitting information.

internet address. The numbering system used in IP internetwork communications to specify a particular

network, or a particular host on that network with which to communicate.

internet control message protocol (ICMP). A protocol used by a gateway to communicate with a source host, for example, to report an error in a datagram. It is an integral part of the Internetwork Protocol (IP).

internetwork protocol. A protocol that routes data from its source to its destination in an internet environment. It is also called the *Internet Protocol*.

internetwork. Any wide area network connecting more than one network.

initial program load (IPL). The initialization procedure that causes the 3745 control program (NCP) to begin operation.

LAN-attached console. A PS/2 attached to the token-ring LAN that has the service processor attached. It is used to operate remotely the MOSS and MOSS-E functions.

IP router. A device that enables an Internetwork Protocol (IP) host to act as a gateway for routing data between separate networks.

line interface coupler (LIC). A circuit that attaches up to four transmission cables to the controller (from DTEs, DCEs or telecommunication lines).

locally administered address. In a local area network, an adapter address that the user can assign to override the universally administered address.

maintenance and operator subsystem - extended (MOSS-E). The licensed internal code loaded on the service processor hard disk to provide maintenance and operator facilities to the user and IBM service representative.

microcode. A program that is loaded in a processor (for example, the MOSS processor) to replace a hardware function. The microcode is not accessible to the customer.

modem (modulator-demodulator). See DCE.

multiple virtual storage (MVS). Multiple Virtual Storage, consisting of MVS/System Product Version 1 and the MVS/370 Data Facility Product operating on a System/370™ processor.

NetView. An IBM licensed program used to monitor a network, manage it, and diagnose its problems.

nonswitched line. A connection between systems or devices that does not have to be made by dialing. The

connection can be point-to-point or multipoint. The line can be leased or private. Contrast with *switched line*.

ping. A simple IP application that sends one or more messages to a specified destination host requesting a reply. Usually used to verify that the target host exists, or that its IP address is a valid address.

remote console. A PS/2 attached to the 3746 Network Node either by a switched line (with modems) or by one of the communication lines of the user network.

remote technical assistance information network (RETAIN).

service processor. The processor attached to a 3745, 3746-900, and 3746-950 via a token-ring LAN.

remote support facility (RSF). RSF provides IBM maintenance assistance when requested via the public switched network. It is connected to the IBM RETAIN database system.

service representative. See IBM service representative

services. A set of functions designed to simplify the maintenance of a device or system.

switched line. A transmission line with which the connections are established by dialing, only when data transmission is needed. The connection is point-to-point and uses a different transmission line each time it is established. Contrast with *nonswitched line*.

synchronous data link control (SDLC). A discipline for managing synchronous, code-transparent, serial-by-bit information transfer over a link connection. Transmission exchanges may be duplex or half-duplex over switched or nonswitched links. The configuration of the link connection may be point-to-point, multipoint,

or loop. SDLC conforms to subsets of the Advanced Data Communication Control Procedures of the American National Standards Institute and High-Level Data Link Control (HDLC) of the International Standards Organization.

synchronous transmission. Data transmission in which the sending and receiving instruments are operating continuously at substantially the same frequency and are maintained, through correction, in a desired phase relationship.

Token-ring adapter (TRA) type 3. 3746-900 and 3746-950 line adapter for IBM Token-Ring Network, composed of one token-ring processor card (TRP2), and two Token-Ring interface couplers type 3 (TIC 3s).

Token-ring interface coupler type 2 (TIC2). A circuit that attaches an IBM Token-Ring network to the 3745.

Token-Ring Interface Coupler type 3 (TIC3). A circuit that attaches an IBM Token-Ring network to the 3746-900 or 3746-950.

user access area. A specific area in the controller where the customer can install, remove, change, or swap couplers and cables without IBM assistance.

universally administered address. In a local area network, the address permanently encoded in an adapter at the time of manufacture. All universally administered addresses are unique.

user application network. A configuration of data processing products, such as processors, controllers, and terminals, for data processing and information exchange. This configuration may use circuit-switched, packet-switched, and leased-circuit services provided by carriers or PTT. Also called a *user network*.


V.24, V.35, and X.21. ITU-T (ex-CCITT) recommendations on transmission interfaces.

Bibliography

Customer Documentation for the 3745 (All Models) and 3746 (Model 900)

Table 20 (Page 1 of 6). Customer Documentation for the 3745 Models X10 and X1A, and 3746 Model 900

This customer documentation has the following formats:



Finding Information

3745 Models A and 3746 Books

All of the books in the 3745 Models A and 3746 library are available on the CD-ROM that contains the Licensed Internal Code (LIC) for the machine.

Evaluating and Configuring



GA33-0092

IBM 3745 Communication Controller Models 210, 310, 410, and 610

Introduction

Gives an introduction of the IBM Models 210 to 610 capabilities.

For Models A, refer to the *Overview*, GA33-0180.



GA33-0180

IBM 3745 Communication Controller Models A and 170² IBM 3746 Nways Multiprotocol Controller Models 900 and 950

Overview

Gives an overview of connectivity capabilities within SNA, APPN, and IP networking.



GA27-4234

IBM 3745 Communication Controller Models A² IBM 3746 Nways Multiprotocol Controller Models 900 and 950


Planning Series:

Overview, Installation, and Integration

Provides information for:

- Overall 3746 planning
- Installation and upgrade scenarios
- Controller and service processor network integration
- Related MOSS-E and CCM worksheets for these tasks.

Table 20 (Page 2 of 6). Customer Documentation for the 3745 Models X10 and X1A, and 3746 Model 900




GA27-4235

IBM 3745 Communication Controller Models A²
IBM 3746 Nways Multiprotocol Controller
Models 900 and 950

Planning Series:
Serial Line Adapters

Provides information for:

- Serial line adapter descriptions
- Serial line adapter line weights and connectivity
- Types of SDLC support
- Configuring X.25 lines
- Performance tuning for frame-relay, PPP, X.25, and NCP lines.
- ISDN adapter description and configuration.



GA27-4236

IBM 3745 Communication Controller Models A²
IBM 3746 Nways Multiprotocol Controller
Models 900 and 950

Planning Series:
Token Ring and Ethernet

Provides information for:

- Token-ring adapter description and configuration
- Ethernet adapter description and configuration.




GA27-4237

IBM 3745 Communication Controller Models A²
IBM 3746 Nways Multiprotocol Controller
Models 900 and 950

Planning Series:
ESCON Channels

Provides information for:

- ESCON adapter descriptions
- ESCON configuration and tuning information
- ESCON configuration examples.



GA27-4238

IBM 3745 Communication Controller Models A²
IBM 3746 Nways Multiprotocol Controller
Models 900 and 950

Planning Series:
Physical Planning

Provides information for:

- 3746 and MAE physical planning details
- 3746 and MAE cable information
- Explanation of installation sheets
- 3746 plugging sheets.

Table 20 (Page 3 of 6). Customer Documentation for the 3745 Models X10 and X1A, and 3746 Model 900

	GA27-4239	<p>IBM 3745 Communication Controller Models A² IBM 3746 Nways Multiprotocol Controller Models 900 and 950</p> <p>Planning Series: Management Planning</p>
		<p>Provides information for:</p> <ul style="list-style-type: none"> • Overview for 3746 • 3746 APPN/HPR, IP router, and X.25 • NetView Performance Monitor (NPM), remote consoles, and RSF • MAE APPN/HPR management.
	GA27-4240	<p>IBM 3745 Communication Controller Models A² IBM 3746 Nways Multiprotocol Controller Models 900 and 950</p> <p>Planning Series: Multiaccess Enclosure Planning</p>
		<p>Provides information for:</p> <ul style="list-style-type: none"> • MAE adapters details • MAE ESCON planning and configuration • ATM and ISDN support.
	GA27-4241	<p>IBM 3745 Communication Controller Models A² IBM 3746 Nways Multiprotocol Controller Models 900 and 950</p> <p>Planning Series: Protocol Descriptions</p>
		<p>Provides information for:</p> <ul style="list-style-type: none"> • Overview and details about APPN/HPR and IP.
	On-line information	<p>IBM 3745 Communication Controller Models A² IBM 3746 Nways Multiprotocol Controller Models 900 and 950</p> <p>Planning Series: Controller Configuration and Management Worksheets</p>
		<p>Provides planning worksheets for ESCON, Multiaccess Enclosure, serial line, and token-ring definitions.</p>
Preparing Your Site		
	GC22-7064	<p>IBM System/360™, System/370™, 4300 Processor</p> <p>Input/Output Equipment Installation Manual-Physical Planning (Including Technical News Letter GN22-5490)</p>
		<p>Provides information for physical installation for the 3745 Models 130 to 610. For 3745 Models A and 3746 Model 900, refer to the <i>Planning Guide</i>, GA33-0457.</p>
	GA33-0127	<p>IBM 3745 Communication Controller Models 210, 310, 410, and 610</p> <p>Preparing for Connection</p>
		<p>Helps for preparing the 3745 Models 210 to 610 cable installation. For 3745 Models A refer to the <i>Connection and Integration Guide</i>, SA33-0129.</p>

Table 20 (Page 4 of 6). Customer Documentation for the 3745 Models X10 and X1A, and 3746 Model 900

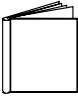
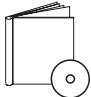
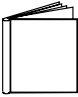
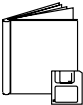
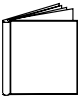
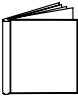
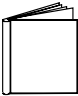
Preparing for Operation		
	GA33-0400	<p>IBM 3745 Communication Controller All Models³ IBM 3746 Nways Multiprotocol Controller Models 900 and 950</p> <p>Safety Information¹</p> <p>Provides general safety guidelines.</p>
	SA33-0129	<p>IBM 3745 Communication Controller All Models³ IBM 3746 Nways Multiprotocol Controller Model 900</p> <p>Connection and Integration Guide¹</p> <p>Contains information for connecting hardware and integrating network of the 3745 and 3746-900 after installation.</p>
	SA33-0416	<p>Line Interface Coupler Type 5 and Type 6 Portable Keypad Display</p> <p>Migration and Integration Guide</p> <p>Contains information for moving and testing LIC types 5 and 6.</p>
	SA33-0158	<p>IBM 3745 Communication Controller All Models³ IBM 3746 Nways Multiprotocol Controller Model 900</p> <p>Console Setup Guide¹</p> <p>Provides information for:</p> <ul style="list-style-type: none"> Installing local, alternate, or remote consoles for 3745 Models 130 to 610 Configuring user workstations to remotely control the service processor for 3745 Models A and 3746 Model 900 using: <ul style="list-style-type: none"> DCAF program Telnet Client program Java Console support.
Customizing Your Control Program		
	SA33-0178	<p>Guide to Timed IPL and Rename Load Module</p> <p>Provides VTAM procedures for:</p> <ul style="list-style-type: none"> Scheduling an automatic reload of the 3745 Getting 3745 load module changes transparent to the operations staff.
Operating and Testing		
	SA33-0098	<p>IBM 3745 Communication Controller All Models⁴</p> <p>Basic Operations Guide¹</p> <p>Provides instructions for daily routine operations on the 3745 Models 130 to 610.</p>
	SA33-0177	<p>IBM 3745 Communication Controller Models A² IBM 3746 Nways Multiprotocol Controller Model 900</p> <p>Basic Operations Guide¹</p> <p>Provides instructions for daily routine operations on the 3745 Models 17A to 61A, and 3746 Model 900 operating as an SNA node (using NCP), APPN/HPR Network Node, and IP Router.</p>

Table 20 (Page 5 of 6). Customer Documentation for the 3745 Models X10 and X1A, and 3746 Model 900

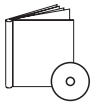

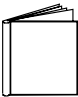
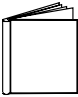
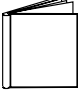

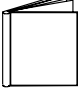
	SA33-0097	<p>IBM 3745 Communication Controller All Models³</p> <p>Advanced Operations Guide¹</p> <p>Provides instructions for advanced operations and testing, using the 3745 MOSS console.</p>
	On-line Information	<p>Controller Configuration and Management Application</p> <p>Provides a graphical user interface for configuring and managing a 3746 APPN/HPR Network Node and IP Router, and its resources. It is also available as a stand-alone application, using an OS/2 workstation. Defines and explains all the 3746 Network Node and IP Router configuration parameters through its online help.</p>
	SH11-3081	<p>IBM 3746 Nways Multiprotocol Controller Models 900 and 950</p> <p>Controller Configuration and Management: User's Guide⁵</p> <p>Explains how to use CCM and gives examples of the configuration process.</p>
	GA33-0479	<p>IBM 3745 Communication Controller Models A IBM 3746 Nways Multiprotocol Controller Models 900 and 950</p> <p>NetView Console APPN Command Reference Guide</p> <p>Explains how to use the RUN COMMAND from the NetView S/390 Program and gives examples.</p>
Managing Problems		
	SA33-0096	<p>IBM 3745 Communication Controller All Models³</p> <p>Problem Determination Guide¹</p> <p>A guide to perform problem determination on the 3745 Models 130 to 61A.</p>
	On-line Information	<p>Problem Analysis Guide</p> <p>An online guide to analyze alarms, events, and control panel codes on:</p> <ul style="list-style-type: none"> • IBM 3745 Communication Controller Models A² • IBM 3746 Nways Multiprotocol Controller Models 900 and 950.
	SA33-0175	<p>IBM 3745 Communication Controller Models A² IBM 3746 Expansion Unit Model 900 IBM 3746 Nways Multiprotocol Controller Model 950</p> <p>Alert Reference Guide</p> <p>Provides information about events or errors reported by alerts for:</p> <ul style="list-style-type: none"> • IBM 3745 Communication Controller Models A² • IBM 3746 Nways Multiprotocol Controller Models 900 and 950.

Table 20 (Page 6 of 6). Customer Documentation for the 3745 Models X10 and X1A, and 3746 Model 900

- ¹ Documentation shipped with the 3745.
- ² 3745 Models 17A to 61A.
- ³ 3745 Models 130 to 61A.
- ⁴ Except 3745 Models A.
- ⁵ Documentation shipped with the 3746-900.

Additional Customer Documentation for the 3745 Models 130, 150, 160, and 170

Table 21. Additional Customer Documentation for the 3745 Models 130 to 170

This customer documentation has the following format:



Finding Information

3745 Models A and 3746 Books

All of the books in the 3745 Models A and 3746 library are available on the CD-ROM that contains the Licensed Internal Code (LIC) for the machine.

Evaluating and Configuring



GA33-0138

IBM 3745 Communication Controller Models 130, 150, 160, and 170

Introduction

Gives an introduction about the IBM Models 130 to 170 capabilities, including Model 160.

For Model 17A refer to the *Overview*, GA33-0180.

Preparing Your Site



GA33-0140

IBM 3745 Communication Controller Models 130, 150, 160, and 170

Preparing for Connection

Helps for preparing the 3745 Models 130 to 170 cable installation.

For 3745 Model 17A refer to the *Connection and Integration Guide*, SA33-0129.

¹ Documentation shipped with the 3745.

Customer Documentation for the 3746 Model 950

Table 22 (Page 1 of 4). Customer Documentation for the 3746 Model 950

This customer documentation has the following formats:



Finding Information

3745 Models A and 3746 Books

All of the books in the 3745 Models A and 3746 library are available on the CD-ROM that contains the Licensed Internal Code (LIC) for the machine.

Preparing for Operation



GA33-0400

IBM 3745 Communication Controller All Models¹
IBM 3746 Expansion Unit Model 900
IBM 3746 Nways Multiprotocol Controller Model 950

Safety Information²

Provides general safety guidelines.

Evaluating and Configuring



GA33-0180

IBM 3745 Communication Controller Models A and 170³
IBM 3746 Nways Multiprotocol Controller
Models 900 and 950

Overview

Gives an overview of connectivity capabilities within SNA, APPN, and IP networking.



GA27-4234


IBM 3745 Communication Controller Models A²
IBM 3746 Nways Multiprotocol Controller
Models 900 and 950

Planning Series: Overview, Installation, and Integration

Provides information for:

- Overall 3746 planning
- Installation and upgrade scenarios
- Controller and service processor network integration
- Related MOSS-E and CCM worksheets for these tasks.

Table 22 (Page 2 of 4). Customer Documentation for the 3746 Model 950



GA27-4235

IBM 3745 Communication Controller Models A²
IBM 3746 Nways Multiprotocol Controller
Models 900 and 950

Planning Series:
Serial Line Adapters

Provides information for:

- Serial line adapter descriptions
- Serial line adapter line weights and connectivity
- Types of SDLC support
- Configuring X.25 lines
- Performance tuning for frame-relay, PPP, X.25, and NCP lines.
- ISDN adapter description and configuration.




GA27-4236

IBM 3745 Communication Controller Models A²
IBM 3746 Nways Multiprotocol Controller
Models 900 and 950

Planning Series:
Token Ring and Ethernet

Provides information for:

- Token-ring adapter description and configuration
- Ethernet adapter description and configuration.



GA27-4237

IBM 3745 Communication Controller Models A²
IBM 3746 Nways Multiprotocol Controller
Models 900 and 950

Planning Series:
ESCON Channels

Provides information for:

- ESCON adapter descriptions
- ESCON configuration and tuning information
- ESCON configuration examples.



GA27-4238

IBM 3745 Communication Controller Models A²
IBM 3746 Nways Multiprotocol Controller
Models 900 and 950

Planning Series:
Physical Planning

Provides information for:

- 3746 and MAE physical planning details
- 3746 and MAE cable information
- Explanation of installation sheets
- 3746 plugging sheets.

Table 22 (Page 3 of 4). Customer Documentation for the 3746 Model 950



GA27-4239

IBM 3745 Communication Controller Models A²
IBM 3746 Nways Multiprotocol Controller
Models 900 and 950

Planning Series:
Management Planning

Provides information for:

- Overview for 3746
- 3746 APPN/HPR, IP router, and X.25
- NetView Performance Monitor (NPM), remote consoles, and RSF
- MAE APPN/HPR management.



GA27-4240

IBM 3745 Communication Controller Models A²
IBM 3746 Nways Multiprotocol Controller
Models 900 and 950

Planning Series:
Multiaccess Enclosure Planning

Provides information for:

- MAE adapters details
- MAE ESCON planning and configuration
- ATM and ISDN support.



GA27-4241

IBM 3745 Communication Controller Models A²
IBM 3746 Nways Multiprotocol Controller
Models 900 and 950

Planning Series:
Protocol Descriptions

Provides information for:

- Overview and details about APPN/HPR and IP.



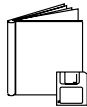

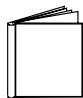
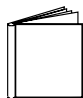

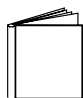
On-line information

IBM 3745 Communication Controller Models A²
IBM 3746 Nways Multiprotocol Controller
Models 900 and 950

Planning Series:
Controller Configuration and Management Worksheets

Provides planning worksheets for ESCON, Multiaccess Enclosure, serial line, and token-ring definitions.

Table 22 (Page 4 of 4). Customer Documentation for the 3746 Model 950

Operating and Testing		
	SA33-0356	<p>IBM 3746 Nways Multiprotocol Controller Model 950</p> <p>User's Guide²</p> <p>Explains how to:</p> <ul style="list-style-type: none"> • Carry out daily routine operations on Nways controller • Install, test, and customize the Nways controller after installation • Configure user's workstations to remotely control the service processor using: <ul style="list-style-type: none"> – DCAF program – Telnet client program – Java Console support.
	On-line information	<p>Controller Configuration and Management Application</p> <p>Provides a graphical user interface for configuring and managing a 3746 APPN/HPR network node and IP Router, and its resources. It is also available as a stand-alone application, using an OS/2 workstation. Defines and explains all the 3746 Network Node and IP Router configuration parameters through its on-line help.</p>
	SH11-3081	<p>IBM 3746 Nways Multiprotocol Controller Models 900 and 950</p> <p>Controller Configuration and Management: User's Guide²</p> <p>Explains how to use CCM and gives examples of the configuration process.</p>
	GA33-0479	<p>IBM 3745 Communication Controller Models A IBM 3746 Nways Multiprotocol Controller Models 900 and 950</p> <p>NetView Console APPN Command Reference Guide</p> <p>Explains how to use the RUN COMMAND from the NetView S/390 Program and gives examples.</p>
Managing Problems		
	On-line information	<p>Problem Analysis Guide</p> <p>An on-line guide to analyze alarms, events, and control panel codes on:</p> <ul style="list-style-type: none"> • IBM 3745 Communication Controller Models A³ • IBM 3746 Nways Multiprotocol Controller Models 900 and 950.
	SA33-0175	<p>IBM 3745 Communication Controller Models A³ IBM 3746 Expansion Unit Model 900 IBM 3746 Nways Multiprotocol Controller Model 950</p> <p>Alert Reference Guide</p> <p>Provides information about events or errors reported by alerts for:</p> <ul style="list-style-type: none"> • IBM 3745 Communication Controller Models A³ • IBM 3746 Nways Multiprotocol Controller Models 900 and 950.
<p>¹ Models 130 to 61A. ² Documentation shipped with the 3746-950 ³ 3745 Models 17A to 61A.</p>		

Required Documentation

The following documents are indispensable for planning for your 3745/3746 controllers:

- *3745 Communication Controller Models A and 170, 3746 Nways Multiprotocol Controller Models 900 and 950: Overview*, GA33-0180
- *3745 Communication Controller All Models, 3746 Nways Multiprotocol Controller Model 900: Console Setup Guide*, SA33-0158.

Be sure to use the latest editions of the above documents.

Related Documentation

The following documents are also helpful for **planning** for your 3745/3746 controllers:

- *Planning for Integrated Networks*, SC31-8062
- *Planning and Reference for NetView, NCP, and VTAM*, SC31-7122.
- *Virtual Telecommunications Access Method V3 R4: Resource Definition Reference*, SC31-6438

The following Enterprise Systems Connection (**ESCON**) documents may be helpful:

- *Introducing the Enterprise Systems Connection*, GA23-0383
- *Enterprise Systems Connection Migration*, GA23-0383
- *Planning for Enterprise Systems Connection Links*, GA23-0367
- *Introducing Enterprise Systems Connection Directors*, GA23-0363.

The following IBM International Technical Support Centers “redbooks” are generally very helpful:

- *Frame Relay Guide*, GG24-4463
- *3746-900 and NCP Version 7 Release 2*, GG24-4464.

The following Network Control Program (**NCP**) documents may be helpful:

- For NCP V6 R2:
 - *Network Control Program V6 R2: Migration Guide*, SC31-6216
 - *Network Control Program V6 R2, ACF/SSP V3 R8, EP R11: Resource Definition Guide*, SC31-6209-01
 - *Network Control Program V6 R2, ACF/SSP V3 R8, EP R11: Resource Definition Reference*, SC31-6210-01
 - *Network Control Program V6 R2: Planning and Implementation Guide*, GG24-4012
 - *Network Control Program V6 R2, ACF/SSP V3 R8, EP R11: Library Directory*, SC31-6215.
- For NCP V6 R3:
 - *Network Control Program V6 R3: Migration Guide*, SC31-6217
 - *Network Control Program V6 R3, ACF/SSP V3 R9, EP R11: Resource Definition Guide*, SC31-6209-02
 - *Network Control Program V6 R3, ACF/SSP V3 R9, EP R11: Resource Definition Reference*, SC31-6210-02 Guide,
 - *Network Control Program V6 R3, ACF/SSP V3 R9, EP R11: Library Directory*, SC31-6218.
- For NCP V7 R1:
 - *Network Control Program V7 R1: Migration Guide*, SC31-6219
 - *Network Control Program V7 R1, ACF/SSP V4 R1, EP R12: Resource Definition Guide*, SC31-6223-00
 - *Network Control Program V7 R1, ACF/SSP V4 R1, EP R12: Resource Definition Reference*, SC31-6224-00
 - *Network Control Program V7 R1, ACF/SSP V4 R1, EP R12: Library Directory*, SC31-6220.

- For NCP V7 R2:
 - *Network Control Program V7 R2, ACF/SSP V4 R2, EP R12: Generation and Loading Guide*, SC31-6221.
 - *Network Control Program V7 R2: Migration Guide*, SC31-6258-00
 - *Network Control Program V7 R2, ACF/SSP V4 R2, EP R12: Resource Definition Guide*, SC31-6223-01
 - *Network Control Program V7 R2, ACF/SSP V4 R2, EP R12: Resource Definition Reference*, SC31-6224-01
 - *Network Control Program V7 R2, ACF/SSP V4 R2, EP R12: Library Directory*, SC31-6259.
- For NCP V7 R3:
 - *Network Control Program V7 R3: Migration Guide*, SC31-6258-01
 - *Network Control Program V7 R3, ACF/SSP V4 R3, EP R12: Resource Definition Guide*, SC31-6223-02
 - *Network Control Program V7 R3, ACF/SSP V4 R3, EP R12: Resource Definition Reference*, SC31-6224-02
 - *Network Control Program V7 R3, ACF/SSP V4 R3, EP R12: Library Directory*, SC31-6262.
- For NCP V7 R4:
 - *Network Control Program V7 R4: Migration Guide*, SC30-3786
 - *Network Control Program V7 R4, ACF/SSP V4 R4, EP R12: Resource Definition Guide*, SC31-6223-03
 - *Network Control Program V7 R4, ACF/SSP V4 R4, EP R12: Resource Definition Reference*, SC31-6224-03
 - *Network Control Program V7 R4, ACF/SSP V4 R4, EP R12: Library Directory*, SC30-3785.
- For NCP V7 R5:
 - *Network Control Program V7 R5: Migration Guide*, SC30-3833
 - *Network Control Program V7 R5, ACF/SSP V4 R4, EP R12: Resource Definition Guide*, SC31-6223-04
 - *Network Control Program V7 R5, ACF/SSP V4 R4, EP R12: Resource Definition Reference*, SC31-6224-04
 - *Network Control Program V7 R5, ACF/SSP V4 R4, EP R12: Library Directory*, SC30-3832.
- For NCP V7 R6:
 - *Network Control Program V7 R6: Migration Guide*, SC30-3833-01
 - *Network Control Program V7 R6, ACF/SSP V4 R4, EP R14: Resource Definition Guide*, SC31-6223-06
 - *Network Control Program V7 R6, ACF/SSP V4 R4, EP R14: Resource Definition Reference*, SC31-6224-06
 - *Network Control Program V7 R6, ACF/SSP V4 R4, EP R14: Library Directory*, SC30-3785.
- For NCP V7 R7:
 - *Network Control Program V7 R7: Migration Guide*, SC30-3889
 - *Network Control Program V7 R7, ACF/SSP V4 R4, EP R14: Resource Definition Guide*, SC31-6223-07
 - *Network Control Program V7 R7, ACF/SSP V4 R4, EP R14: Resource Definition Reference*, SC31-6224-07
 - *Network Control Program V7 R7, ACF/SSP V4 R4, EP R14: Library Directory*, SC30-3971.
- For NCP V7 R8:
 - *Network Control Program V7 R8: Migration Guide*, SC30-4024
 - *Network Control Program V7 R8, ACF/SSP V4 R8, EP R14: Resource Definition Guide*, SC31-6223-09

- *Network Control Program V7 R8, ACF/SSP V4 R8, EP R14: Resource Definition Reference*, SC31-6224-09
- *Network Control Program V7 R8, ACF/SSP V4 R8, EP R14: Library Directory*, SC30-4025.

The following **OS/2** document may be of some help:

IBM Extended Services® for OS/2 Programming Services and Advanced Problem Determination for Communications, SO4G-1007.

For the Distributed Console Access Facility (**DCAF**) Version 1.3 the following documents are needed:

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

To learn more about the **APPN** architecture, including high-performance routing (HPR), adaptive rate based flow and congestion control (ARB), dependent LU requesters/servers (DLURs/DLUSs), and other subjects, refer to:

- *Inside APPN - The Essential Guide to the Next-Generation SNA*, SG24-3669.
- *APPN Architecture and Protocol Implementations Tutorial* SG24-3669.

The following Virtual Telecommunications Access Method (**VTAM**), may be helpful:

- *Virtual Telecommunications Access Method V4R3: Resource Definition Reference*, SC31-6438.

For help with **TCP/IP**, refer to:

- *TCP/IP for MVS: Performance Tuning Guide*, SC31-7188.

To learn about token-ring configurations and the **IEEE 802.2** standard, refer to:

- *Token-Ring Network Architecture Reference*, SC30-3374.

These latest NetView documents may be helpful:

- *TME 10 NetView for OS/390 Version 1: Planning Guide*, GC31-8226
- *TME 10 NetView for OS/390 Version 1: Tuning Guide*, SC31-8240.

The following NetView Performance Monitor (**NPM**) documents are available:

- *NetView Performance Monitor: Concepts and Planning V2R2*, GH19-6961-01
- *NetView Performance Monitor: Concepts and Planning V2R3*, GH19-6961-02
- *NetView Performance Monitor: Concepts and Planning V2R4*, GH19-6961-03
- *NetView Performance Monitor: Concepts and Planning V3R1*, GH19-4221-00.

Index

Numerics

3746 IP routing 132
3746 router (IP) 91
3746-900 in an SNA/APPN network 66
3746-950 in an SNA/APPN network 66
802.2 LLC Type 2 71

A

access control, IP 114
adaptive rate-based congestion control
 See ARB
address (IP)
 aggregation 94
 class 85
 general principles 84
 number 86
 structure 85
 summarization 94
 support 92
Address Resolution Protocol (ARP), IP 100
address space manager 12
addresses, multiple IP 95
advanced program-to-program communication 1
alive timer 26
alive timer, HPR 31
ANR 20
ANR labels 29
applications, IP 109
APPN
 external connection between a 3745-CNN and a
 3746-900 61
 internal Connection between a 3745 and a
 3746-900 55
 Internal link 55
 interoperation with HPR 32
 Migration to HPR 32
 network node 66
 shared topology with HPR 32
APPN border node 45
APPN or LEN node 10
ARB 22
area border (AB) router support, IP 103
ARP cache table 100
ARP cache table, IP 100
ASCII format VTAM keywords 42
ASM 35
authorized end node 3

B

base and towers, HPR 18
BCD format VTAM keywords 42
BIND 36, 37
Bisynch 3270 sessions 34
BOOTP, bootstrap protocol, IP 112
Border Gateway Protocol Version 4 (BGP), IP 107
border node 5, 8
 extended border node 8
boundary function support, APPN/HPR 19
boundary node 5, 6
broadcast, IP
 general principles 88
 limited 93
 local-wire 93
 network-directed 94
 OSPF 103
 subnet-directed 93

C

CBSP2 34
CD-ROM Online documentation xxiii
central resource registration 2
changes since last edition xviii
Classless Inter-Domain Routing (CIDR), IP 94
CLP 34
CNN 6, 7
commands (Telnet), structure 135
common prefix, IP address 94
communicating across incompatible networks 83
composite network node
 See CNN
composite network node (CNN) 43
composite node 5, 6
configuration control and management software
 (CCM) 38
configuration services 11, 35
configuration, 3746-900 requirements 55
connection network 7, 29, 74
connection-oriented 71
connectivity, network node 39
consolidation of IP routers 95
contention
 loser 33
 winner 33
control flows over RTP tower, HPR 19
control point
 See CP
controlling access (IP) 114

CoS 2, 21
CP 2, 11
CP-CP session 6, 8
CPSVRMGR session 74
CS 35
customer tasks xxv

D

default route (IP)
 BGP 109
 OSPF protocol 105
definitions, internal APPN link 56
Dependent Logical Unit Requester (DLUR) 67
dependent LU requester/server 33
dependent LU server
 See DLUS
dependent sessions 45
direct DLC 71
directory services 11
 See also DS
DLC 13, 35
DLCADDR, VTAM keyword PATH statement 42
DLCs support, HPR 39
DLSw 75
DLSw MAC address 72
DLUR 33, 35, 67, 74
DLUS 33
DLUS/DLUR
 dependent LU requester 40
 VTAM DLCADDR keyword in PATH statement 42
domain 17
domains (IP) 89, 91
domains (IP), generic 91
dotted notation, IP address 85
drop-in migration to HPR 32
DS 35
DSPU 74
dynamic routing protocol support 92
dynamic switched definition facility, VTAM 41

E

EBCDIC format VTAM keywords 42
end node 3
enhanced session address 18
equal-cost multipath routing, IP OSPF 104
error recovery 28, 39
 error recovery, HPR 40
 ESCON, HPR 40
 frame relay, HPR 39
 HPR 40
 IP, HPR 40
 priority, HPR 30
 SDLC, HPR 40
 token-ring, HPR 39

ESCC 34
ESCON 34
 addresses, IP 121
 configuration, IP 124
 error recovery, HPR 40
 IP 118
 port sharing 120
 subnet addresses 121
ESCP2 34
Ethernet 72
Ethernet LAN 72
examples
 ANR node, HPR 58
 APPN/HPR networks 58
 APPN/RTP node, HPR 59
 APPN/RTP node, SNI 60
 DLUR/RTP node, HPR 59
 internal APPN link 55
extended border node 5, 8
extended border node (PBN) 43
extended recovery facility 74

F

FID2 28
FID2 PIU 18
filters, IP 113
frame relay 72
 error recovery, HPR 39
 links 74
Full-Screen Capability, Telnet 135
fully qualified domain names (IP) 90

G

generic domains (IP) 91

H

hexadecimal format VTAM keywords 42
hierarchical name space (IP) 90
High-Performance Routing (HPR) 68
host number (IP) 86
HPR 17, 63, 68
 ANR Node example 58
 APPN control point protocols and algorithms 32
 APPN/RTP node example 59, 60
 base and towers 18
 base functions 18
 boundary function support, APPN/HPR 19
 control flows over RTP tower 19
 DLCs support 39
 DLUR/RTP node example 59
 drop-in migration from APPN 32
 error recovery 39
 ESCON 40
 frame relay 39

HPR (*continued*)

error recovery (*continued*)

- IP 40

- SDLC 40

- token-ring 39

- HPR-only route calculation 31

- interoperation with APPN 32

- migration 32

- migration from APPN 32

- network examples 58

- Nondisruptive path switch 19

- Priority 30

- Rapid Transport Protocol (RTP) 19

- route calculation 30

- RTP tower 19

- shared topology with APPN 32

- timers 31

- Alive timer 31

- Path switch timer 31

- Short_req timer 31

- HPR MLTG 63

I

- IEEE 802.2 LLC Type 2 72

- importing routes, IP OSPF 104

- inactivity timeout 72

- incompatible networks, communicating across 83

- independent sessions 45

- inter-operability of routing protocols, IP 107

- interchange node 5, 6

- interfaces, IP 117

- intermediate session routing 2, 12

- internal APPN connection between a 3745 and a 3746-900 55

- internal APPN link 55

- internal APPN link definitions 56

- Internet Control Message Protocol (ICMP), IP 98

- interoperability of routing protocols, IP 107

IP

- 3746 IP routing 132

- 3746 Router 91

- access control 114

- address

- aggregation 94

- summarization 94

- support 92

- address class 85

- Address Resolution Protocol (ARP) 100

- address structure 85

- addressing 84

- area border (AB) router support 103

- ARP cache table 100

- BOOTP, bootstrap protocol 112

- Border Gateway Protocol Version 4 (BGP) 107

- broadcast, OSPF 103

IP (*continued*)

- broadcasting 88, 93

- Classless Inter-Domain Routing (CIDR) 94

- common prefix, address 94

- configuration

- ESCON 121, 124

- token-ring 126

- consolidation of routers 95

- default route

- BGP 109

- OSPF 105

- defined 83

- dividing networks 87

- domain name space 89

- domains 89

- dotted notation 85

- dynamic routing protocol support 92

- equal-cost multipath routing, OSPF 104

- error recovery, HPR 40

- ESCON 118

- filters 113

- fully qualified domain names 90

- general description 83

- generic domains 91

- hierarchical name space 90

- host number 86

- importing routes, OSPF 104

- interfaces 117

- internal applications 109

- Internet Control Message Protocol (ICMP) 98

- interoperability of routing protocols 107

- limited broadcast 94

- maximum connectivity 92

- multiaccess

- broadcast 93

- nonbroadcast 93

- multicasting 89

- multiple addressing 95

- NCP-IP 132

- network number 86

- network-directed broadcast 94

- nonbroadcast multiaccess, OSPF 103

- numbered point-to-point 103

- open shortest path first protocol (OSPF) 103

- OSPF interface support 103

- overlaying networks 84

- password authentication, OSPF 104

- PING 110

- point-to-point (PtP) 93

- port sharing

- ESCON 120

- token-ring 125

- protocols 101

- Proxy-ARP 101

- RIP 107

- route

- acceptance policy 102

IP (*continued*)

route (*continued*)

- advertisement policy 102
- advertisement, OSPF 104
- aggregation 94, 97
- table explosion problem 94

router 84

router addresses 86

routing

- information protocol (RIP) 101
- policies, OSPF 104
- protocols 101

security 112

spreading the message 88

static routes 109

stub area support 103

subnet mask 87

subnet-directed broadcast 93

subnets 86

supernetting 94

Telnet Operations 133

- Command structure 135
- Full-Screen Capability 135
- Negotiation of options 134
- Network Virtual Terminal (NVT) 135
- NVT printer 135

token-ring 125

traceroute 111

Transmission Control Protocol (TCP) 98

transparent communication across networks 83

unicasting 88

unnumbered point-to-point 103

User Datagram Protocol (UDP) 98

ISDN 74

L

LEN end node 3

level 2 protocol identifier 72, 73

level 3 protocol identifier APPN HPR). 73

level 3 protocol identifier SNA-APPN/FID2). 72

LFSID 12, 36

LFSID swapping 20, 37

LIC 34

limited broadcast, IP 94

limited resource 29, 74

Local APPN SAP address parameter 72

Local HPR SAP address parameter 72

local-form session identifier

See LFSID

locally administered MAC address 75

logical unit 12

long-lived RTP connection 19

M

management services 12

mask, subnet IP address 87

maximum connectivity, IP 92

microcode levels required for year 2000 readiness xvii

migration 32

drop-in, to HPR 32

HPR from APPN

SNI connections to 3746 Models 900 and 950 43

MLTG 27, 69

MLTG ANR labels 29

MOSS-E 34

multiaccess

broadcast networks, IP 93

nonbroadcast networks, IP 93

multicasting (IP) 89

Multilink Transmission Group (MLTG) 69

multilink transmission groups

See MLTG

multiple DLUR/DLUS 41

multiprotocol encapsulation 72

N

NAU 15

NCE 19

NCP-IP 132

net ID 15

Network Layer Packets (NLP) 39

network node connectivity 39

network node processor

See NNP

network node server 3

network node, APPN 66

Network Routing Facility

See NRF

Network Terminal Option

See NTO

Network Virtual Terminal (NVT), Telnet 135

network-directed broadcast, IP 94

network-qualified name 15

network, APPN/HPR

examples 59, 60

network, APPN/HPR examples 58

NLP 18

NNP 34, 35, 37, 38

NOF 10, 35

Non-activation XID 29

nonbroadcast multiaccess, IP OSPF 103

nondisruptive path switch, HPR 27

nondisruptive path switching, HPR 19

Not so stubby area 105

NPP 38

NRF 74

NTO 74
numbered point-to-point, IP 103
NVT printer, Telnet 135

O

open shortest path first protocol (OSPF), IP 103
OSPF interface support, IP 103

P

parallel TGs 27, 74
password authentication, IP OSPF 104
path switching, HPR timer 31
PC 13, 35
peripheral border node (PBN) 43
peripheral node 5, 6
PING, IP 110
point-to-point (PtP) 93
port sharing
 ESCON 120
 token-ring 125
PPP 72, 73, 74
primary LU dependent sessions 45
priority, HPR 30
procedure correlation identification (PCID) 41
Protocols, IP
 Address Resolution Protocol (ARP) 100
 Border Gateway Protocol Version 4 (BGP) 107
 inter-operability of routing protocols 107
 Internet Control Message Protocol (ICMP) 98
 interoperability of routing protocols 107
 open shortest path first protocol (OSPF) 103
 RIP V1 101
 routing protocols 101
 Transmission Control Protocol (TCP) 98
 User Datagram Protocol (UDP) 98
Proxy-ARP, IP 101

Q

Q.933 encoding 72, 73
QLLC 72

R

Rapid Transport Protocol (RTP) 19
re-sequencing 28
remote DLSw 71
RFC 1490 bridged frame format 72
RFC 1490 routed frame format 72
RIP
 metrics 102
route calculation, HPR 30
route setup protocol 18
route test 73

route, IP
 acceptance policy 102
 advertisement
 OSPF 104
 policy 102
 aggregation (IP) 97
 policies, IP OSPF 104
 table explosion problem 94
router, IP
 addresses 86
 aggregation (IP) 94
 defined 84
Routing Information Protocol (RIP), IP 101
RTP 18, 19, 21, 28
 end points 28
 tower, control flows 19
 tower, HPR 19

S

SAP address 72
SATF 7
SCM 35
SDLC
 error recovery, HPR 40
 protocol 72
secondary LU 45
 dependent sessions 45
security, IP 112
service processor 34, 38
session
 connector 36, 37
 services 35
session connector 12
session identifier 16
session services 12
Session Services Extension (SSE) 45
session-level security 74
shared topology, APPN/HPR 32
Short_req timer, HPR 31
single-link TGs 27
SNA
 connection problems 43
 Network Interconnection (SNI) 43
SNI connection migration to 3746 Models 900 and 950 43
SNMP 73, 98, 117
SP 34
spreading the message 88
SS 35
SSCP 6, 17
SSCP-SSCP session 6
static routes, IP 109
store-and-forward 28
stub area support, IP 103

- subarea connection problems 43
- subnet (IP)
 - general principles 86
 - mask 87
 - mask example 87
- subnet-directed broadcast, IP 93
- subnetwork 8
- supernetting, IP 94

T

- Telnet
 - Command structure 135
 - Full-Screen Capability 135
 - Negotiation of options 134
 - Network Virtual Terminal (NVT) 135
 - NVT printer 135
 - Operations 133
- Telnet Operations 133
- TIC3 34
- token-ring 72
 - configuration requirements 55
 - error recovery, HPR 39
 - IP 125
 - IP configuration 126
 - port sharing 125
- topology and routing services 11
 - See also* TRS
- topology subnetwork 8
- towers and base, HPR 18
- traceroute 111
- transaction program 11
- Transmission Control Protocol (TCP), IP 98
- transmission header 16
- TRP2 34
- TRS 29, 35, 73

U

- unauthorized end node 3
- unicasting (IP) 88
- unnumbered point-to-point, IP 103
- user API 74
- User Datagram Protocol (UDP), IP 98
- User-Defined Parameters 64

V

- virtual routing node 5, 7
- VR-TG 28
- VRN 75
- VTAM
 - ASCII format keywords 42
 - BCD format keywords 42
 - DLCADDR keyword in PATH statement 42
 - EBCDIC format keywords 42

- VTAM (*continued*)
 - hexadecimal format keywords 42

X

- X.25 SVC 74
- XID 72
- XRF 74
- XRF/CRYPTO 74

Y

- Year 2000
 - microcode levels required xvii
 - readiness xvii

Tell Us What You Think!

3745 Communication Controller Model A
3746 Nways Multiprotocol Controller
Models 900 and 950
Planning Series:

Protocol Descriptions

Publication No. GA27-4241-03

We hope you find this publication useful, readable, and technically accurate, but only you can tell us! Your comments and suggestions will help us improve our technical publications. Please take a few minutes to let us know what you think by completing this form. If you are in the USA, you can mail this form postage free or fax it to us at 1-800-253-3520. Elsewhere, your local IBM branch office or representative will forward your comments or you may mail them directly to us.

Overall, how satisfied are you with the information in this book?	Satisfied	Dissatisfied
	<input type="checkbox"/>	<input type="checkbox"/>

How satisfied are you that the information in this book is:	Satisfied	Dissatisfied
Accurate	<input type="checkbox"/>	<input type="checkbox"/>
Complete	<input type="checkbox"/>	<input type="checkbox"/>
Easy to find	<input type="checkbox"/>	<input type="checkbox"/>
Easy to understand	<input type="checkbox"/>	<input type="checkbox"/>
Well organized	<input type="checkbox"/>	<input type="checkbox"/>
Applicable to your task	<input type="checkbox"/>	<input type="checkbox"/>

Specific comments or problems:

Please tell us how we can improve this book:

Thank you for your comments. If you would like a reply, provide the necessary information below.


Name	Address
Company or Organization	
Phone No.	




Fold and Tape

Please do not staple

Fold and Tape



NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES



BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

Design & Information Development
IBM Corporation
Software Reengineering
Department G71A/ Bldg 503
P.O. Box 12195
Research Triangle Park, NC 27709-9990



Fold and Tape

Please do not staple

Fold and Tape



Printed in the United States of America
on recycled paper containing 10%
recovered post-consumer fiber.

GA27-4241-03

