

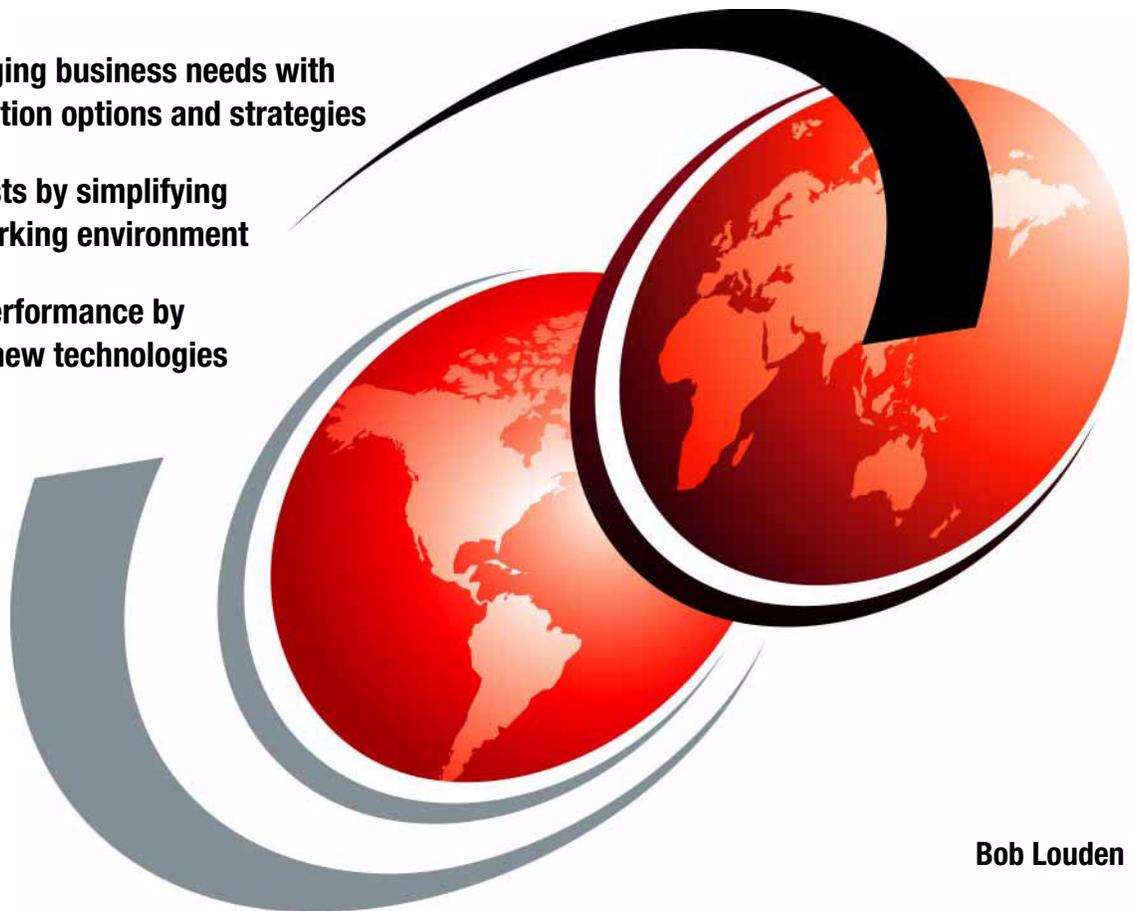
IBM Communication Controller Migration Guide:

Updated for Communication Controller for Linux on
System z9 and zSeries (V1.2)

Meet changing business needs with
clear migration options and strategies

Reduce costs by simplifying
your networking environment

Enhance performance by
exploiting new technologies



Bob Loudon



International Technical Support Organization

**IBM Communication Controller Migration Guide:
Updated for Communication Controller for Linux on
System z9 and zSeries (V1.2)**

January 2006

Note: Before using this information and the product it supports, read the information in “Notices” on page xi.

Third Edition (January 2006)

This edition applies to the functions available across IBM communication controller products including:

- Communication Controller for Linux on zSeries (abbreviated as CCL V1.1), 5724-J38
- Communication Controller for Linux on System z9 and zSeries (CCL V1.2), 5724-J38
- 3745 Communication Controller
- 3746 Nways Multiprotocol Controller
- Network Control Program (NCP), 5648-063
- Emulation Program (EP), 5735-XXB
- Network Terminal Option (NTO), 5735-XX7
- Non-SNA Interconnection (NSI), 5668-951
- Network Routing Facility (NRF), 5668-963
- X.25 NCP Packet Switching Interface (NPSI), 5688-035
- MERVA Extended Connectivity, 5655-110
- X.25 SNA Interconnection (XI) and Network Supervisory Function (NSF), 5685-035
- NtuneMON and NtuneNCP, 5648-D72
- System Support Programs (SSP), 5655-041
- Teleprocessing Network Simulator (TPNS), 5688-121

© Copyright International Business Machines Corporation 2002, 2005, 2006

Note to U.S Government Users – Documentation related to restricted rights – Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Contents

Noticesxi
Trademarks	xii
Preface	xiii
Third edition, 2006, SNA update	xiv
The structure of this book.....	xvi
The team that wrote this redbook.....	xvii
Special notice.....	xviii
Become a published author	xix
Comments welcome.....	xix
Summary of changes	xxi
January 2006, Third Edition	xxi
Part 1. Getting started	1
Chapter 1. What to do	3
1.1 Process overview	4
1.2 Task descriptions	5
1.2.1 Physical inventory (Task 1).....	5
1.2.2 Logical and functional inventory (Task 2)	5
1.2.3 Reconcile and optimize (Task 3).....	7
1.2.4 Controller strategic planning (Task 4).....	8
1.2.5 Functional alternative migrations (Tasks 5-7).....	9
Chapter 2. Controller optimization examples	11
2.1 Client example 1: ZYX Company	12
2.2 Task 1 - ZYX Company physical inventory	12
2.3 Task 2 - ZYX Co. logical and functional inventory.....	21
2.4 Task 3 - ZYX Company reconcile and optimize	25
2.4.1 Technology results	27
2.4.2 Business results	43
2.4.3 The next step	43
2.5 Task 4 - ZYX Company controller strategic planning	44
2.6 Task 5 - ZYX Co. functional alternative migration.....	44
2.7 Client example 2: M Manufacturing Company	45
2.8 Task 1 - M Company and A Company physical inventory.....	47
2.9 Task 2 - M Company and A Company logical and functional inventory ..	60
2.10 Task 3 - M Manufacturing Company reconcile and optimize	72

2.10.1 Technology results	73
2.10.2 Business results	87
2.10.3 The next step	87
2.11 Task 4 - M Manufacturing Company controller strategic planning	88
2.12 Task 5 - M Manufacturing Company functional alternative migration	88
Chapter 3. Hot topics - SNI, EP, BSC, and X.25	89
3.1 SNA Network Interconnection (SNI)	90
3.2 Emulation Program (EP)	90
3.3 Binary synchronous communication (BSC)	90
3.4 X.25 communications	91
Part 2. Functional alternatives reference	93
Chapter 4. Communication controller hardware	101
4.1 Communication controller hardware overview	102
4.2 Physical interface functions and alternatives	104
4.2.1 Communication lines	106
4.2.2 Token ring (and duplicate addressing)	108
4.2.3 Ethernet, FDDI, and ATM	124
4.2.4 Channel attachments	128
4.3 Network Node Processor functions and alternatives	131
4.3.1 APPN network node (including HPR and DLUR)	131
4.3.2 IP routing	137
4.4 Multi-Access Enclosure functions and alternatives	138
4.4.1 TN3270 server	138
4.4.2 Network Dispatcher	138
Chapter 5. Network Control Program (NCP)	141
5.1 NCP product overview	142
5.2 Link-level protocol functions and alternatives	143
5.2.1 Airlines Line Control (ALC)	144
5.2.2 BSC 3270 terminal connection to SNA applications	146
5.2.3 Frame relay	147
5.2.4 Integrated Services Digital Network (ISDN)	151
5.2.5 LAN 802.2, 802.3, and 802.5	151
5.2.6 Start-stop terminal connection to TCAM applications	152
5.2.7 Synchronous Data Link Control (SDLC)	152
5.2.8 X.25 (with licensed support feature on 3746)	152
5.3 Advanced capability functions and alternatives	158
5.3.1 SNA class of service (COS)	159
5.3.2 Multi-link transmission group (MLTG)	161
5.3.3 SNA subarea addressing, routing, and boundary function (BF)	162
5.3.4 SNA Network Interconnection (SNI)	167

5.3.5 APPN (and LEN) Composite Network Node (CNN)	176
5.3.6 IP routing	180
5.3.7 Extended Recovery Facility (XRF)	180
Chapter 6. X.25 NCP Packet Switching Interface (NPSI)	187
6.1 NPSI product overview	188
6.2 Functions and alternatives	189
6.2.1 Supporting SNA (PSH and QLLC) communication over X.25	190
6.2.2 Supporting non-SNA communication over X.25	196
Chapter 7. MERVA Extended Connectivity	203
7.1 MERVA product overview	204
7.2 Functions and alternatives	205
Chapter 8. X.25 SNA Interconnection (XI) and Network Supervisory Function (NSF)	207
8.1 XI and NSF product overview	208
8.2 Functions and alternatives	209
8.2.1 Transporting X.25 traffic	209
8.2.2 NSF-based chargeback for X.25 transport service	210
Chapter 9. Emulation Program (EP)	211
9.1 EP product overview	212
9.2 Functions and alternatives	213
9.2.1 BSC 3270 terminal connection to non-SNA applications	214
9.2.2 BSC RJE connection to non-SNA applications	217
9.2.3 Start-stop terminal connection to non-SNA applications	223
Chapter 10. Network Terminal Option (NTO)	225
10.1 NTO product overview	226
10.2 Functions and alternatives	227
10.2.1 Start-stop terminal connection to SNA host applications	227
10.2.2 BSC RJE connection to SNA host applications	229
10.2.3 Peer-to-peer connection of non-SNA devices	235
Chapter 11. Network Routing Facility (NRF)	237
11.1 NRF product overview	238
11.2 Functions and alternatives	239
11.2.1 Peer-to-peer connection of SNA devices (before PU type 2.1)	239
11.2.2 Peer-to-peer connections involving non-SNA devices	245
Chapter 12. Non-SNA Interconnection (NSI)	251
12.1 NSI product overview	252
12.2 Functions and alternatives	253
12.2.1 Non-SNA NJE to NJE connections between hosts	253

12.2.2	Transporting BSC traffic across your network.	253
Chapter 13.	Teleprocessing Network Simulator (TPNS)	255
13.1	Teleprocessing Network Simulator (TPNS).	256
13.2	Functions and alternatives	257
13.2.1	Simulating subarea traffic	257
Part 3.	Strategic solution technologies.	259
Chapter 14.	Net390 architecture.	261
14.1	What is Net390?	262
14.2	Why is Net390 strategic?	262
14.3	What you need to understand about Net390.	262
14.4	Strategies for exploiting Net390	268
Chapter 15.	OSA-Express.	269
15.1	What is OSA-Express?	270
15.1.1	An historical perspective	270
15.1.2	An architectural perspective	271
15.1.3	Summary of features.	272
15.2	Why is OSA-Express strategic?	273
15.3	What you need to understand about OSA-Express	274
15.3.1	TCP/IP support	274
15.3.2	SNA support	275
15.3.3	Layer 2 support.	276
15.3.4	Management.	277
15.3.5	TPF support	279
15.4	Strategies for exploiting OSA-Express	279
Chapter 16.	Communication Controller for Linux on System z9 and zSeries (CCL)	283
16.1	What is CCL?	284
16.2	Why is CCL important?	285
16.3	What you need to understand about CCL.	286
16.3.1	Layer 2 support.	290
16.3.2	CCL V1.2 OSA for NCP (OSN) CDLC channel support	291
16.3.3	CCL DLSw support	292
16.3.4	CCL V1.2 IP TG support.	294
16.4	Strategies for exploiting CCL.	295
Chapter 17.	TCP/IP	299
17.1	What is TCP/IP?	300
17.2	Why is TCP/IP strategic?	301
17.3	What you need to understand about TCP/IP.	301

17.3.1	Routing infrastructure	301
17.3.2	TCP/IP security	302
17.3.3	Quality of Service	303
17.3.4	High availability and load distribution	304
17.3.5	Web-based applications	304
17.3.6	Print solutions	305
17.4	Strategies for exploiting TCP/IP	305
17.4.1	TCP/IP convergence	305
17.4.2	Secure your network	308
17.4.3	Quality of Service	309
17.4.4	High availability and load distribution	310
Chapter 18.	TN3270	311
18.1	What is TN3270?	312
18.2	Why is TN3270 strategic?	312
18.3	What you need to understand about TN3270	313
18.3.1	Secure TN3270 sessions	317
18.3.2	Telnet printer support	319
18.4	Strategies for exploiting TN3270	323
18.4.1	TN3270 server placement	323
18.4.2	High availability and load balancing	326
Chapter 19.	Enterprise Extender (EE)	331
19.1	What is Enterprise Extender (EE)?	332
19.2	Why is EE strategic?	332
19.2.1	Cost-effectiveness and resource convergence	332
19.2.2	Flow and congestion control	332
19.2.3	Class of service	333
19.2.4	Exploits ubiquitous Internet connectivity	333
19.2.5	Session availability	333
19.2.6	Why not Data Link Switching (DLSw)?	334
19.3	What you need to understand about EE	335
19.3.1	Transporting SNA over IP	335
19.3.2	Enterprise Extender architecture	336
19.3.3	Flow and congestion control	337
19.3.4	Enterprise Extender implementation	338
19.4	Strategies for exploiting EE	338
19.4.1	The IP backbone as an APPN Connection Network	338
19.4.2	EE within the sysplex	339
19.4.3	Connecting branches to hosts	340
19.4.4	Securing Enterprise Extender links with IPSec	343
Chapter 20.	Message queuing (MQ)	345
20.1	What is MQ?	346

20.2 Why is MQ strategic?	346
20.3 What you need to understand about MQ	347
20.3.1 Asynchronous messaging	348
20.3.2 Adapters	348
20.4 Strategies for exploiting MQ	349
Part 4. Appendixes	351
Appendix A. Physical inventory worksheets	353
Instructions	354
3745 and attached frames physical inventory	356
3746-950 physical inventory.	384
Appendix B. Logical and functional inventory worksheets	391
Instructions	392
Appendix C. Avoiding VTAM network addressing problems	403
Understanding VTAM network addressing	405
SNA sessions	405
Applications (APPLs)	406
Dependent PUs and LUs	406
Independent LUs (CDRSCs).	409
VTAM enhanced addressing	410
VTAM V4R2	414
CS for OS/390 V2R5 DLUR-attached dependent LUs	416
CS for OS/390 V2R7.	417
z/OS CS V1R4 EE logical LINEs and PUs	420
z/OS CS V1R6	420
Maximizing VTAM use of enhanced addressing	422
Enable VTAM enhanced addressing for sessions	422
Use VTAM application cloning	423
Migrate your network from subarea to APPN	423
Migrate dependent LUs to DLUR-attachment.	424
Migrate your APPN network to EE	425
Glossary	427
Abbreviations and acronyms	449
Related publications	453
IBM Redbooks	453
Other resources	453
Referenced Web sites	454
How to get IBM Redbooks	456

Index	457
--------------------	-----

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. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described 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.

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.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

Advanced Peer-to-Peer Networking®	MVS™ MVS/ESA™	System/360™ System/370™
AIX®	NetView®	System z9™
AS/400®	Nways®	Tivoli®
CICS®	OS/390®	Virtualization Engine™
CICS/ESA®	Parallel Sysplex®	VM/ESA®
DB2®	POWER™	VSE/ESA™
ESCON®	RACF®	VTAM®
FICON®	Redbooks™	WebSphere®
HiperSockets™	Redbooks (logo)  ™	z/OS®
HACMP™	Resource Link™	z/VM®
IBM®	RMF™	z/VSE™
IMS™	RS/6000®	z9™
MQSeries®	S/390®	zSeries®

The following terms are trademarks of other companies:

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

Microsoft, Windows NT, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States, other countries, or both.

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

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

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

Preface

IBM® communication controllers have reliably carried the bulk of the world's business traffic for more than 30 years. Beginning with the introduction of the IBM 3705 Communication Controller in 1973, IBM controllers have been enhanced over the years to the point that the functional capabilities of the current products, the 3745 Communication Controller and the 3746 Nways® Multiprotocol Controller—along with their numerous software product adjuncts—surpass the capabilities of any other data networking equipment ever developed. Beyond the SNA architecture PU Type 4, beyond APPN, and even beyond IP routing, these controllers support an extraordinary set of functions and protocols. Because of their long history and their functional richness, IBM controllers continue today to play a critical role in the networks of most of the largest companies in the world.

Over the past decade, however, focus has shifted from SNA networks and applications to TCP/IP and Internet technologies. In some cases, SNA application traffic now runs over IP-based networks using technologies such as TN3270 and Data Link Switching (DLSw). In other cases, applications have been changed, and processes reengineered, using TCP/IP rather than SNA. Consequently, for some organizations, the network traffic that traverses IBM communication controllers has declined to the point where it is in their business interest to find functional alternatives for the remaining uses of their controllers so they can consolidate and possibly eliminate them from their networking environments.

This IBM Redbook provides you with a starting point to help in your efforts to optimize your communication controller environment, whether simply consolidating them or migrating from them altogether. It is a compendium of the best ideas from the most knowledgeable people. We discuss alternative means to provide the communication controller functions that you use or ways to eliminate the need for those functions outright. Where multiple options exist, we discuss the relative advantages and disadvantages of each.

Important: Consider the broader context of your organization's IT infrastructure strategy before making the decision to consolidate or eliminate communication controllers. For example, if your organization is engaged in a long-term project to reengineer business processes and Web-enable key business applications, it may be better to wait until most users have moved to the new Web-based applications before investing the time and resources required to migrate traffic from your communication controller. Likewise, your IT infrastructure strategy must be formulated in the context of the needs and initiatives of your organization.

This book focuses on the newer communication controller hardware platforms, the 3745 Communication Controller and the 3746 Nways Multiprotocol Controller, but most of the discussions apply likewise to the older controller hardware (3705, 3725, and 3720) environments.

Often, environments involving IBM communication controllers provide functions vital to the successful operation of the organization. Deciding whether to do a migration and, if so, deciding which existing components are worth preserving is the starting point for the planning process. If you do plan a migration involving IBM communication controllers, note that the list of hardware and software customization parameters involved is long. Each parameter value, whether explicitly coded or defaulted, may be important. So review your configuration carefully to determine whether and how your existing parameters relate to any proposed migration alternative. Then proceed with careful testing of the best alternative to determine whether its cost, function, performance, manageability, serviceability, security, and availability meet your organization's needs.

Important: Except where explicitly stated, the alternatives described in this book have not been tested by IBM.

Third edition, 2006, SNA update

While the strategic importance of TCP/IP and related Internet technologies is readily apparent to IBM and all of information technology, IBM continues to show unprecedented commitment to helping SNA clients preserve their investment while migrating to strategic technologies. For example, IBM has provided the following information in formal statements of direction:

- ▶ VTAM® in z/OS® Communications Server: IBM intends to support VTAM in z/OS Communications Server for the foreseeable future. Clients have a substantial investment in 3270 and SNA applications. We continue to support and enhance VTAM's capabilities while integrating it with new technologies. IBM has no plans at this time to discontinue SNA support in z/OS

Communications Server. (IBM United States Software Announcement 204-179, August 10, 2004)

- ▶ 3745/3746 support: IBM continues to provide service for the 3745 and 3746 controllers consistent with normal IBM Technical Services service plans and guidelines, which are typically five years after the end of marketing date. While its plans are subject to change at its sole discretion, IBM has no current plans to end service for these products before 2010. (IBM United States Software Announcement 205-030, February 15, 2005)
- ▶ ACF/NCP and associated software products: IBM intends to support and maintain ACF/NCP and associated software products (EP, NPSI, NTuneMon) beyond the 3745 and 3746 product support dates. (IBM United States Software Announcement 205-030, February 15, 2005)

Important: All statements regarding the future direction or intent of IBM are subject to change or withdrawal without notice, and represent goals and objectives only.

In addition, on February 15, 2005, IBM announced the IBM Communication Controller for Linux® on zSeries® V1.1 (abbreviated as CCL V1.1) which enables you to run the NCP and NRF program products originally designed to be run on IBM 3745 Communication Controller hardware on Linux for zSeries. *In many cases, the CCL option will provide the most transparent migration from your IBM 3745 Communication Controllers.* (IBM United States Software Announcement 205-030, February 15, 2005)

Also, on October 25, 2005, IBM announced the IBM Communication Controller for Linux on System z9™ and zSeries V1.2 (abbreviated as CCL V1.2) which includes performance enhancements and adds support for:

- ▶ The X.25 NCP Packet Switching Interface (NPSI) program product
- ▶ Open Systems Adapter for NCP (OSN) CDLC channel connectivity to Communications Server (VTAM) and TPF
- ▶ Data Link Switching (DLSw)

(IBM United States Software Announcement 205-267, October 25, 2005)

With that said, however, it is still the case that SNA is now more than 30 years old and is gradually on its way out throughout the industry. Additionally, the ongoing complexity inherent in supporting both strategic TCP/IP-based technologies and SNA results in higher costs, often with lower performance and availability. Finally, over the next few years support will become increasingly problematic for other elements of your SNA environment including non-IBM SNA software and hardware products and declining overall SNA skills. Consequently, it is now more important

than ever to revisit your strategy with respect to SNA rather than wait and react as specific SNA elements become difficult and ultimately impossible to support.

The structure of this book

Figure 1 illustrates the organization of this book.

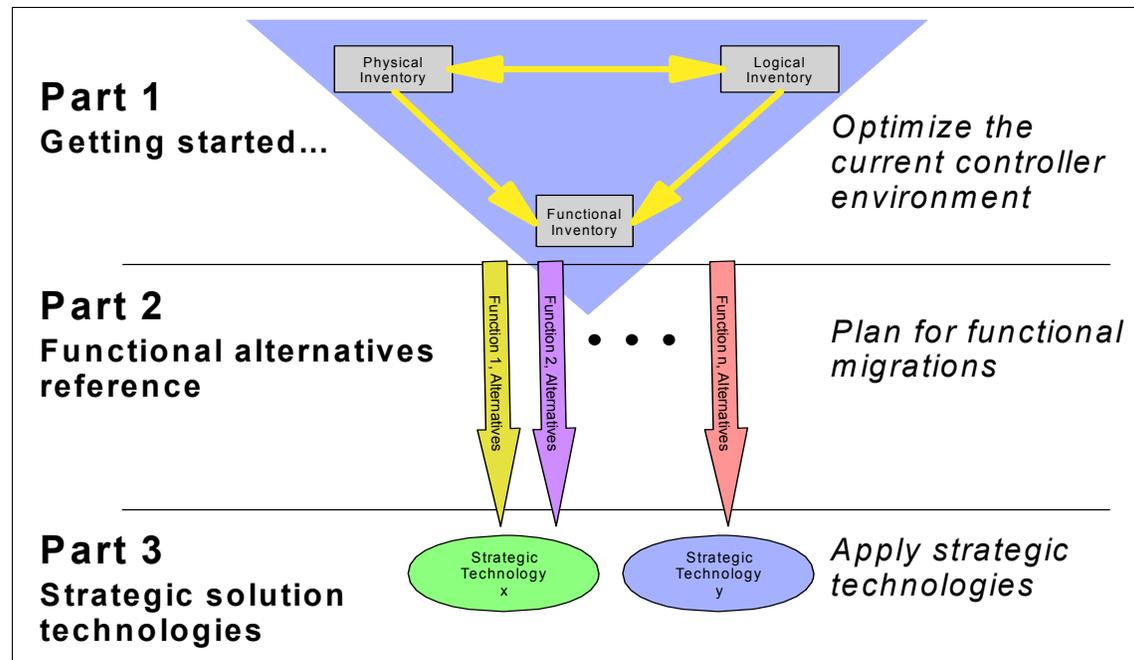


Figure 1 The structure of the book

The book is divided into three parts:

- ▶ **Part 1. Getting started** covers the steps that you will need to take, and the skills that you will require, to optimize and migrate your communication controller environment. A key part of the “getting started” process will be to identify the communication controller functions used by your organization so that you can go directly to the appropriate sections in Part 2 of this book.
- ▶ **Part 2. Functional alternatives reference** reviews the capabilities of each of the IBM communication controller products, explores the alternatives for replacing or eliminating the need for each, and discusses the relative advantages and disadvantages of the various alternatives. This part of the book is organized into groups of chapters that cover:
 - Communication controller hardware

- Network Control Program (NCP)
 - X.25 related products
 - Other controller software products, including products such as Emulation Program (EP), Network Terminal Option (NTO), and Network Routing Facility (NRF)
- **Part 3. Strategic solutions technologies** provides in-depth discussions of the most strategically important alternative solution technologies.

Icons in the margins of this redbook are used as follows:



Important: Denotes a potential problem area where, in our experience, organizations have encountered significant problems in the past.



Product limitation: Denotes a potential problem area where a product limitation could affect your use of a specific product capability.



CCL key capability: Identifies a functional area where the CCL product may play an important role.

The team that wrote this redbook

The author of this redbook is:

Bob Louden, a consultant in the IGS, IT Services, Network Services Delivery practice. His 23 years as a networking professional with IBM have enabled him to develop strong professional and consulting skills, including leadership, project management, problem solving, and decision analysis. Bob's technical expertise includes wide-area and local-area networking and SNA and TCP/IP protocols. More important, however, is his ability to use his technology understanding to develop business solutions. For more than 15 years, Bob has applied professional, technical, and business expertise to help clients develop effective strategies and optimal networking solutions.

For his contributions as co-author of the first edition of this book, special thanks are due to **Robert Silverman**. Robert retired from IBM at the end of 2004 and is currently working at JPMorgan Chase on network performance engineering, management, and automation. While at IBM, Robert provided sales support for IBM network hardware and software. Robert was an original developer of the

Apparent Network Speed Analysis tool and has several IBM patents issued and pending in the field of network management.

Thanks to **Haechul Shin**, Business Development Manager - Enterprise Networking Solutions, for his sponsorship and drive to make the second and third edition updates of this redbook possible.

I am especially grateful for the significant expertise and contributions of content to the second and third editions of the book from:

- ▶ **Alfred Christensen**, Programming Consultant, Enterprise Networking Solutions.
- ▶ **Johnathan Harter**, a senior software engineer who joined VTAM Development in 1987, is now recognized as a leading authority on APPN, and continues to be an advocate for SNA networking as the Chief Programmer for the Communication Controller for Linux on zSeries (V1.1) product.
- ▶ **Jerry Stevens**, Enterprise Networking Solutions, Lead Developer for the Communication Controller for Linux on System z9 and zSeries (V1.2) product.
- ▶ **Joe Welsh**, an Application Integration Middleware IT Consultant who joined VTAM Development in 1989, is responsible for performing consulting engagements focused on SNA, APPN/HPR, Enterprise Extender, and TCP/IP for Communications Server for z/OS at large organizations around the world. This includes providing: education and training, network designs and migration plans, product strategy and direction, problem determination, and installation and migration assistance.

I would also like to acknowledge contributions to the first edition of the book from: Robert Brinkman, Nancy W. Gates, James Goethals, Shelly Howrigan, Adolfo Rodriguez, Sebastian Altemir, Steve Culbreth, Ted Gary, Vincent Herbay, Richard H. Le Sesne, Jean-Pierre Marce, Jim Robinson, Wolfgang Singer, Christian Thiery, Geert Van de Putte, and Bill White.

Finally, I would like to thank the ITSO staff for their outstanding support and guidance; particularly, Cecilia Bardy and Leslie Parham, Editors.

Special notice

This publication is intended to provide you with a starting point to help you in your efforts to migrate traffic from your communication controllers. The information in this publication is not intended as the specification of any programming interfaces that are provided by the IBM communication controllers. See the PUBLICATIONS section of the IBM Programming Announcement for those

controllers for more information about what publications are considered to be product documentation.

Become a published author

Join us for a two- to seven-week residency program! Help write an IBM Redbook dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You'll team with IBM technical professionals, Business Partners and/or customers.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you will develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our Redbooks™ to be as helpful as possible. Send us your comments about this or other Redbooks in one of the following ways:

- ▶ Use the online **Contact us** review redbook form found at:

ibm.com/redbooks

- ▶ Send your comments in an e-mail to:

redbook@us.ibm.com

- ▶ Mail your comments to:

IBM Corporation, International Technical Support Organization
Dept. HZ8 Building 662
P.O. Box 12195
Research Triangle Park, NC 27709-2195

Summary of changes

This section describes the technical changes made in this edition of the book and in previous editions. This edition may also include minor corrections and editorial changes that are not identified.

Summary of Changes
for SG24-6298-02
for IBM Communication Controller Migration Guide
as created or updated on May 9, 2006.

January 2006, Third Edition



The primary driver for the second and third edition updates to the *IBM Communication Controller Migration Guide* is the introduction of an important new product family: the Communication Controller for Linux on zSeries V1.1 and its follow-on product, the Communication Controller for Linux on System z9 and zSeries V1.2. (Both share IBM product number 5724-J38 and are abbreviated throughout this book as CCL). The CCL product enables you to run the Network Control Program (NCP), the Network Routing Facility (NRF), and the X.25 NCP Packet Switching Interface (NPSI) program products in your mainframe (migrating them from your IBM 3745 Communication Controllers); therefore, the CCL opens up a broad range of communication controller migration options that were not previously available. These options are easily identified by the distinctive CCL margin icon (as shown to the left of this paragraph) and are discussed in greater detail, where applicable, in the chapters of this book.

The CCL product is introduced and discussed in Chapter 16, “Communication Controller for Linux on System z9 and zSeries (CCL)” on page 283.

Most significant CCL capabilities

As “beauty is in the eye of the beholder,” what is significant (or insignificant) to your organization depends on the specifics of your business, IT environment, plans, and strategy. The CCL functional capabilities listed below, however, provide alternatives where either no viable solution existed before (XRF, for example) or where prior alternatives posed significant implementation challenges (for example, SNI).

- ▶ SNA Network Interconnection (SNI): CCL provides an attractive option because it eliminates the need for coordinated inter-organizational changes

required by every other alternative (discussed in 5.3.4, “SNA Network Interconnection (SNI)” on page 167).

- ▶ Token ring to Ethernet migration: The CCL can handle SNA traffic on Ethernet LAN interfaces whereas the 3745 only supports SNA on token-ring LANs (discussed in “Migrating from token ring to Ethernet” on page 117).
- ▶ Extended Recovery Facility (XRF): Prior to the introduction of CCL, no functional alternative existed other than to make significant application changes (discussed in 5.3.7, “Extended Recovery Facility (XRF)” on page 180).
- ▶ Network Routing Facility (NRF): Prior to the introduction of CCL, no functional alternative existed other than to make significant application changes (discussed in 11.2.1, “Peer-to-peer connection of SNA devices (before PU type 2.1)” on page 239).
- ▶ X.25 NCP Packet Switching Interface (NPSI): Using X.25 over TCP/IP (XOT) technology and CCL V1.2 NPSI program product support, your mainframe environment can support *all of the X.25 capabilities* currently supported by NPSI in your IBM 3745 communication controllers. Routers are used to support the actual X.25 physical interfaces and XOT technology transports the X.25 traffic over your TCP/IP network into your mainframes. Advantages of the CCL NPSI approach over the previously available host XOT software alternatives include comprehensive support of NPSI functional capabilities and easier migration (discussed in “Supporting NPSI in your mainframe using CCL” on page 200).

Other CCL capabilities

- ▶ “Replacing token-ring connectivity” on page 110: Token-ring OSA can now support SNA *either* to VTAM or to a CCL NCP.
- ▶ “Implementing CCL Ethernet high availability” on page 118: Provides a high-level design for CCL high availability.
- ▶ 4.3, “Network Node Processor functions and alternatives” on page 131: You can now migrate SNA boundary function requirements from NNP DLUR to CCL NCP support.
- ▶ 5.3.2, “Multi-link transmission group (MLTG)” on page 161: CCL supports token-ring and Ethernet LAN MLTGs.
- ▶ 5.3.3, “SNA subarea addressing, routing, and boundary function (BF)” on page 162: CCL gives you a way to migrate devices from your IBM 3745 Communication Controller hardware without having to move device addresses from your NCP subarea into a VTAM subarea and while preserving your NCP boundary function.
- ▶ 5.3.7, “Extended Recovery Facility (XRF)” on page 180: CCL allows you to continue to run NCP, supporting your XRF environment, on your mainframe.

- ▶ 16.3.1, “Layer 2 support” on page 290: Enables CCLs to use fiber OSA connectivity and can support multiple MAC addresses over a single OSA port.
- ▶ 16.3.2, “CCL V1.2 OSA for NCP (OSN) CDLC channel support” on page 291: Provides QDIO-based emulated CDLC channel connectivity to a CCL NCP. This provides more efficient connectivity between CCLs and VTAMs located in the same mainframe system, enables “local” (channel-attached) NCP configurations for those VTAMs, and provides the only supported connectivity between CCLs and TPF.
- ▶ 16.3.3, “CCL DLSw support” on page 292: Implements a standards-compliant DLSw peer in the CCL Linux environment that can help reduce (or eliminate) the need for data center DLSw routers, thereby improving availability.
- ▶ 16.3.4, “CCL V1.2 IP TG support” on page 294: Provides an efficient means of connecting two CCL NCPs across an IP network.

Other changes

- ▶ Updated 5.3.7, “Extended Recovery Facility (XRF)” on page 180 to include the alternative of migrating from XRF to the use of **Parallel Sysplex®** technologies, specifically Generic Resources.
- ▶ Updated the **OSA-Express** chapter to reflect ongoing enhancements in OSA technology (see Chapter 15, “OSA-Express” on page 269)
- ▶ Added discussion to the **TCP/IP** chapter (in 17.4.1, “TCP/IP convergence” on page 305) regarding options for migrating from the use of SNA altogether
- ▶ Significantly revised the **TN3270** chapter (see Chapter 18, “TN3270” on page 311) to reflect a new TN3270 server option: Communications Server for Linux on zSeries
- ▶ Rewrote the **VTAM enhanced addressing** appendix to reflect the significant ongoing efforts to improve VTAM’s management of SNA network element addresses (see Appendix C, “Avoiding VTAM network addressing problems” on page 403)



Part 1

Getting started

This part of the book discusses our thoughts on the best way to get started optimizing and potentially migrating functions from your IBM communication controllers.

In deciding *when* to migrate, *what* to migrate, or even *if* you should migrate functions from your communication controllers, you face the initial challenge of knowing what functions your communication controllers are currently providing. Communication controllers have evolved over a long period and, as such, are very robust. The very nature of the environments controllers typically run in require them to be masters of many functions and can make it difficult to determine all the vital roles they play.

This part of the book provides a process for optimizing your communication controller environment. This process has been used by IBM sales and consulting professionals on numerous engagements. It offers a step-by-step approach to determine the:

- ▶ *Physical* configuration of the controller (what is installed)
- ▶ *Logical* usage of the controller (what is being used)
- ▶ *Functional* roles of the controller (how it is being used)

With this information, you can proceed on two fronts:

- ▶ Cleaning up and optimizing what you currently have
- ▶ Formulating and executing your strategy with respect to the future role of the controller in your environment.

Based on past experience, there is a common approach for success in controller optimization projects:

- ▶ Optimize your current communication controller environment. This usually means a consolidation of controller resources on less hardware. This optimization concentrates on improving performance, function, and cost. A key benefit of optimization is the immediate saving of operational costs such as NCP software license and maintenance charges, which can sometimes be very substantial.
- ▶ Initiate the migration of the highest-priority functions. The functions initially migrated to other solutions can be determined by ease of migration and the benefits to be derived (such as cost savings, performance enhancements, and availability improvements). For example, many organizations have migrated from large numbers of low-speed lines to LAN infrastructures. Such a migration is fairly low risk and has the potential to save significant costs in the controller environment as well as line costs.
- ▶ Develop and execute a strategy for the long-term strategic migrations.



What to do

This chapter presents an effective methodology for optimizing and migrating functions from your communication controller environment. An overall plan of action is discussed. Then each individual task is discussed in further detail identifying the skills required to accomplish it.

1.1 Process overview

The high-level project plan illustrated in Figure 1-1 below presents an approach for optimizing your current communication controller environment and preparing for the functional migrations that are appropriate for your organization.

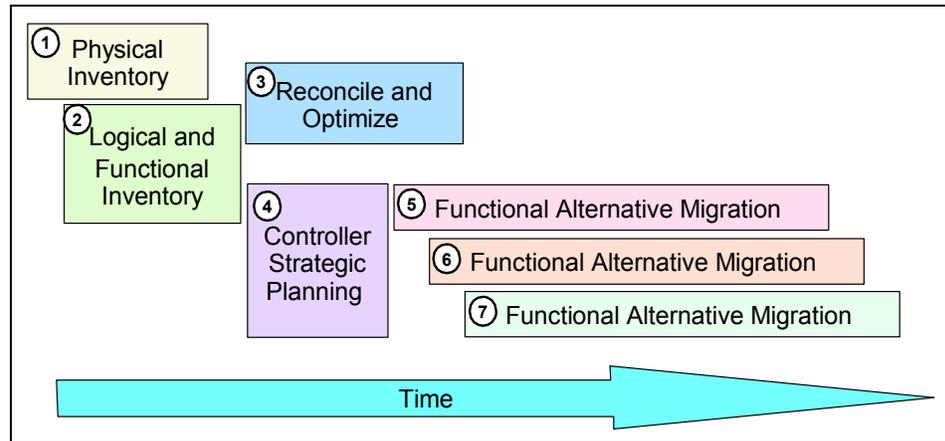


Figure 1-1 High-level controller migration project plan

There are a few key points that should be noted about this proposed project plan:

- ▶ Before you begin making changes to your controller environment, you should have a clear understanding of what you have and how your organization is using it. **Task 1: Physical Inventory**, and **Task 2: Logical and Functional Inventory** provide a structured approach for assessing your current environment.
- ▶ Many organizations have reaped substantial financial and network performance benefits simply by cleaning up their current communication controller environment. This includes activities such as reducing operating costs by consolidating controllers, removing unused lines and software, and reducing NCP software license tier levels. **Task 3: Reconcile and Optimize** is focused on achieving such near-term benefits.
- ▶ While you're optimizing your current controller environment, you should also be considering your organization's overall strategy with respect to the communication controller functions that your organization still uses. How will your communication controllers be affected by your organization's business and IT strategies? Are there benefits to migrating certain functions from your communication controllers? Should you plan on migrating all functions from your controllers? Which functions should be moved first, and which last? Questions such as these should be the focus of **Task 4: Controller Strategic Planning**.

- ▶ Finally, based on your communication controller strategic plans, you should begin your functional migrations (represented by **Tasks 5, 6, and 7**).

1.2 Task descriptions

The following sections describe specifically what is involved in performing the tasks shown in Figure 1-1 on page 4.

1.2.1 Physical inventory (Task 1)

The overall goal of the physical inventory is to identify and verify what you have installed on your communication controller equipment through on-site visual inspections. Organizations are rarely fully aware of the interfaces installed on their controllers (for which they may be paying maintenance fees).

Note: Even for older model controllers (such as IBM 3705, 3720, and 3725 controllers), it is important to carefully explore and clearly understand the equipment that you have installed as well as how it is being used. However, a physical inventory is not necessary for these devices because they are no longer supported by IBM, and you will almost always be better off migrating from them. For older model controllers, then, a logical and functional inventory will be sufficient.

Task

Through on-site visual inspections, identify all communication controllers and their installed interfaces. The worksheets located in Appendix A, “Physical inventory worksheets” on page 353 may be used to guide the inventory process. When the controller physical inventory is complete, verify it against the IBM records to ensure accurate maintenance billing.

Skills and resources

To develop an accurate physical inventory of your installed equipment, you will need people who know about communication controller hardware. These people could be members of your technical staff and/or IBM Customer Engineers (CEs). Note that IBM may charge you a fee for this service.

1.2.2 Logical and functional inventory (Task 2)

While the overall goal of the physical inventory is to identify and verify what you *have installed* on your communication controller equipment, the goal of the logical and functional inventory is to understand the resources and functions that

are *actually being used* in your controller environment and *how they are being used*. In Task 3 (Reconcile and optimize), the logical and functional inventory is used along with the physical inventory to identify controller hardware that is no longer needed. The logical and functional inventory also provides important information on how communication controllers are currently being used for Task 4 (Controller strategic planning).

Task

Your logical and functional inventory can start with a review of your NCP generation statements. For those controller resources that are still in use, understand how they serve the needs of your organization. For example:

- ▶ Are they lines to a remote office or to another data center?
- ▶ Or, are they SNI connections to other organizations?

What software and functions are configured in the communication controllers (such as NCP, EP, NRF, or NTO)? In particular, how many lines, LAN interfaces, and channel connections are actually in use. The worksheets located in Appendix B, “Logical and functional inventory worksheets” on page 391 may be used to guide the inventory process. Once the controller logical and functional inventory is complete, verify your software usage against the IBM records to ensure accurate software charges.

Important: Include *only* those resources that you have verified are still in use. In the past, we have experienced situations where resources still existed in communication controller configurations but were actually redundant or no longer necessary.

Tools such as NTuneMON, NetView®, and NPM can be helpful in determining whether resources are still in use. Also, for users of the Network Node Processor for APPN and/or IP resources, on the 3746-900 and 3746-950, the Controller Configuration and Management (CCM) tool may also be useful.

Skills and resources

To develop your controller logical and functional inventory, you will need people who know about communication controller software. These people could be members of your technical staff, such as your communications systems programmers, and/or IBM Global Services communication specialists.

1.2.3 Reconcile and optimize (Task 3)

The overall goal is to optimize the current environment both in terms of performance and cost. Benefits of this task include reducing:

- ▶ **Maintenance charges:** by consolidating machines, cleaning up unused hardware, and using more current features
- ▶ **Software license charges:** by consolidating machines, cleaning up NCP generation statements, getting lower NCP tier charges by cleaning up unused hardware, and using more current software and hardware
- ▶ **Other operating costs:** by identifying communication lines that are no longer in use and streamlining processes

Many organizations have been quite surprised by what is installed on their controllers but no longer being used.

Task

Review the physical inventory information in light of what you learned from the logical and functional inventory process.

- ▶ Identify physical interfaces that are installed but no longer in use.
- ▶ Identify any installed software components (such as NRF or NTO) that are no longer being used.
- ▶ With an understanding of hardware costs, maintenance costs, NCP tier license charges, expected growth in the environment, and availability requirements, you can develop and execute a plan to clean out, consolidate, and optimize your controller hardware and software environments.

Note: An optimal controller consolidation plan requires significant expertise in IBM controller environments. The communication controller configurator tool can be used to find the optimal configuration of controllers to take advantage of what you currently own, to identify what you should invest in upgrading, and to determine what equipment should be eliminated.

For example, you may be able to enhance performance while also reducing your NCP tier license charges by minimizing the interfaces on the 3745 base frame and putting as many LANs, lines, and channels as possible onto a 3746 Model 900. Also, using newer physical interfaces might help you get the maximum amount of connectivity and performance for your money.

- ▶ Clean up NCP generation statements.

- ▶ Update IBM records (both hardware and software) to ensure accurate maintenance and software license charges.
- ▶ Verify your telecommunications charges against actual usage and circuits. Cancel unnecessary telecommunications circuits.

Skills and resources

To reconcile and optimize your controller environment, you will need people with expert knowledge of communication controller hardware and software (including hardware costs, maintenance costs, NCP tier license charges, and the communication controller configuration tool). These people could include members of your technical staff but they will almost certainly need to be assisted by communications specialists from IBM or IBM-authorized business partners.

1.2.4 Controller strategic planning (Task 4)

The physical inventory as well as the logical and functional inventory provide a solid foundation for your controller strategic planning. Essentially, they define the “as is” of your controller environment while your strategic plan should establish the “to be” or target environment. Only with such a long-range vision of your controller environment and a clear understanding of how that vision supports the overall IT and organizational strategies, can you make optimal strategic decisions regarding your controller environment.

Task

Your controller strategic planning should include the following activities:

- ▶ Review your current physical and logical controller environment.
- ▶ Review the functional roles that your communication controllers play.
- ▶ Understand your broader business and IT environment, plans, strategies, and expectations for the future.
- ▶ Formulate your strategy with respect to the future role of communication controllers in your environment.
- ▶ Develop cost-effective strategic plans for controller consolidations and functional migrations without compromising the IT organization’s ability to meet the ongoing IT and networking needs of the business. This plan should include:
 - Consideration of expected near-term increases in lines or traffic in the controller environment
 - Identification of which functions should be migrated and the appropriate timing of each migration

- Specification of near-term activities which should be initiated in order to prepare for long-term migrations

The communication controller functional alternatives part of this book is intended to help you to understand the alternatives available for the controller functions that you use.

Skills and resources

Developing strategic plans for your controller environment requires a coordinated effort involving people with knowledge ranging from executive-level understanding of your organization's business and IT strategies to detailed understanding of your IT infrastructure and communication controller environment. This effort must involve a broad team of managers and technical staff from your organization and may, additionally, include IBM consultants.

1.2.5 Functional alternative migrations (Tasks 5-7)

The specific tasks and resources required will depend on the particular migration being pursued and may be determined by reviewing the functional alternatives discussed in the next part of this book. For each migration, however, you will need to perform high-level design, detailed design, planning, and implementation activities.



Controller optimization examples

A small effort in optimizing your communication controller environment could reap substantial benefits both in terms of reducing costs and enhancing network performance. This chapter discusses two case studies involving controller optimization projects.

This chapter provides client case studies to illustrate:

- ▶ Physical inventory
- ▶ Logical and functional inventory
- ▶ Results, both technology and business benefits, of a controller optimization project

These case studies are actual client scenarios seen by IBM Sales and Consulting professionals; however, we altered the names of the companies to protect client confidentiality. It is our belief that these studies will provide a guide for getting started on your own optimization project.

2.1 Client example 1: ZYX Company

ZYX Company represents a commercial financial institution. They have deployed communication controllers in their infrastructure for many years. Their normal course of action was to add features to their existing controllers or additional communication controllers to meet their growth requirements. Over time many technologies changed in their infrastructure environment. ZYX Company is now investigating updates to their communication controller environment. At the beginning of their optimization project they had four individual 3745-210 communication controllers, each with an attached 3746-A11 and 3746-L13 expansion unit, all four configured similarly.

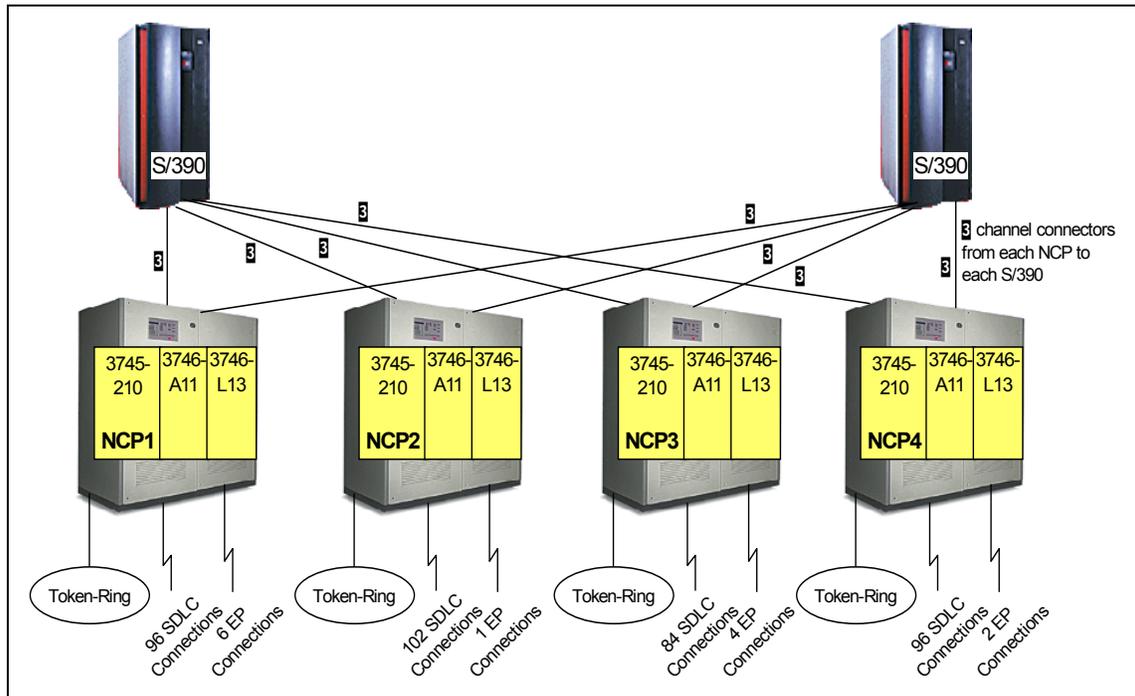


Figure 2-1 ZYX Company before optimization

2.2 Task 1 - ZYX Company physical inventory

One way that ZYX Company can determine the physical inventory of their communication controllers is to enlist the support of an IBM Customer Engineer (CE). The IBM CE can inventory the features installed on the communication controller via a combination of physically identifying installed features along with using the Maintenance and Operations SubSystem (MOSS). Actual physical

inventory and machine configurations (installed features) tend to change over time. Since the physical inventory will determine the current actual configuration, the IBM CE is an excellent resource to use. Note that IBM may charge you a fee for this service.

The following figures represent samples of ZYX Company's physical inventory sheets for one of their 3745-210 complexes:

3745 Physical Inventory	
Part 1: Machine Overview	
3745 Model	<u> 210 </u> Serial Number <u> 00001 </u>
Customer Designation (Ex. NCP Name)	<u> NCP1 </u>
Indicate all attached 3746 expansion frames.	
3746-A11	SN <u> 01111 </u>
3746-A12	SN <u> </u>
3746-L13	SN <u> 02222 </u>
3746-L14	SN <u> </u>
3746-L15	SN <u> </u>
3746-900	SN <u> </u>
Controller Expansion (Rack) Qty	<u> 0 </u> (0,1 or 2)
(Please inventory each attached frame separately on 3746 Model Sheets.)	
How much memory (per CCU if Model 41x or 61x)?	
4M	<u> </u>
8M	<u> x </u>
16M	<u> </u>

Figure 2-2 ZYX Company Model 210 machine overview before optimization

3745 Physical Inventory

Part 1: Machine Overview

3745 Model **210** Serial Number **00001**

Console Information

3151 or Equivalent **x**

Service Processor _____

Type:

9577 _____

9585 _____

3172 P/N 41H7520 _____

3172 P/N 55H7630 _____

7585 _____

6275 _____

6563 _____

Indicate serial numbers of other 37XXs that share this console.

Figure 2-3 ZYX Company Model 210 console information before optimization

3745 Physical Inventory

Part 2: Base 3745 Configuration

Section A: 3745 Models x1x

3745 Model 210 Serial Number 00001

Channel Board

Bus Group 1	Pos. 1 <u>CADS</u>	Pos. 2 <u>CADS</u>	Pos. 3 <u>CADS</u>	Pos. 4 <u>CADS</u>
	Pos. 5 <u>CADS</u>	Pos. 6 <u>CADS</u>	Pos. 7 _____	Pos. 8 _____

Legend:

CADS = Channel Adapter Data Streaming
BCCA = Buffer Chaining Channel Adapter
TPS = Two Processor Switch

Figure 2-4 ZYX Company Model 210 channel board before optimization

3745 Physical Inventory

Part 2: Base 3745 Configuration

Section A: 3745 Models x1x

3745 Model **210** Serial Number **00001**

TSS Board

Bus Group 1	Pos. 1 Adap: <u> TRA2 </u>		Pos. 2 Adap: <u> TRA2 </u>		Pos. 3 Adap: <u> HSS </u>		Pos. 4 Adap: <u> LSS </u>	
	Port <u> TIC </u>	Port <u> TIC </u>	Port <u> TIC </u>	Port <u> TIC </u>	Port <u> Active </u>	Port <u> </u>	Port <u> #1 210 </u>	Port <u> </u>
Bus Group 2	Pos. 5 Adap: <u> TRA2 </u>		Pos. 6 Adap: <u> </u>		Pos. 7 Adap: <u> HSS </u>		Pos. 8 Adap: <u> LSS </u>	
	Port <u> TIC </u>	Port <u> TIC </u>	Port <u> </u>	Port <u> </u>	Port <u> Active </u>	Port <u> </u>	Port <u> #2 210 </u>	Port <u> </u>

Legend:

Adapters

- HSS = High Speed Scanner
- LSS = Low Speed Scanner
- ELA = Ethernet/IEEE 802.3 Adapter
- TRA1 = Token-Ring Adapter Type 1
- TRA2 = Token-Ring Adapter Type 2

Port Cards

- TIC = Token-Ring Interface Coupler
- LAN = Ethernet LAN Attached
- #x fff = Where 'x' is the Line Unit Area driven by the installed LSS and 'fff' is the model number of frame where the Line Unit Area resides
- (Example: #2 L13 - would indicate Line Unit Area 2 in the 3746-L13)
- (Example: #3 210 - would indicate Line Unit Area 3 in the 3745-210)

Figure 2-5 ZYX Company Model 210 TSS board before optimization

3745 Physical Inventory

Part 2: Base 3745 Configuration

Section A: 3745 Models x1x

3745 Model 210 Serial Number 00001

Line Unit

Area 1				Area 2			
LIC Type <u>1</u>							

Additional Line Units

Area 3				Area 4			
LIC Type <u>1</u>							

Figure 2-6 ZYX Company Model 210 line unit before optimization

3745 Physical Inventory
Part 2: Base 3745 Configuration

Section A: 3745 Models x1x

3745 Model **210** Serial Number **00001**

Line Unit

Area 5				Area 6			
LIC Type							
<u> 1 </u>	<u> 1 </u>	<u> 1 </u>	<u> 1 </u>	<u> 3 </u>	<u> 3 </u>	<u> 3 </u>	<u> 3 </u>

Additional Line Units

Area 7				Area 8			
LIC Type							
<u> 3 </u>							

Figure 2-7 ZYX Company Model 210 line unit before optimization

Part 3: 3746 Adapter Expansion Configuration

3746 Model **A11** Serial Number **01111**

Channel Board

3746 Models A11 Only

Bus Group 1	Pos. 9 _____	Pos. 10 _____	Pos. 11 _____	Pos. 12 _____
	Pos. 13 _____	Pos. 14 _____	Pos. 15 _____	Pos. 16 _____

Legend:

CADS = Channel Adapter Data Streaming

BCCA = Buffer Chaining Channel Adapter

TPS = Two Processor Switch

Figure 2-8 ZYX Company Model A11 channel board before optimization

Part 3: 3746 Adapter Expansion Configuration

3746 Model **A11** Serial Number **01111**

TSS Board

3746 Models A11 Only

Bus Group 1	Pos. 9 Adap: <u> LSS </u>	Pos. 10 Adap: <u> LSS </u>	Pos. 11 Adap: <u> LSS </u>	Pos. 12 Adap: <u> LSS </u>
	LIC area <u> #3 210 </u>	LIC area <u> #5 210 </u>	LIC area <u> #7 210 </u>	LIC area <u> #1 L13 </u>
Bus Group 2	Pos. 13 Adap: <u> LSS </u>	Pos. 14 Adap: <u> LSS </u>	Pos. 15 Adap: <u> LSS </u>	Pos. 16 Adap: <u> LSS </u>
	LIC area <u> #4 210 </u>	LIC area <u> #6 210 </u>	LIC area <u> #8 210 </u>	LIC area <u> #2 L13 </u>
Bus Group 1	Pos. 17 Adap: _____	Pos. 18 Adap: _____	Pos. 19 Adap: _____	Pos. 20 Adap: _____
	LIC area # _____			
Bus Group 2	Pos. 21 Adap: _____	Pos. 22 Adap: _____	Pos. 23 Adap: _____	Pos. 24 Adap: _____
	LIC area # _____			

Legend:

Adapters

LSS = Low Speed Scanner
ALC = Airline Line Control Scanner

Port Cards

#x fff = Where 'x' is the Line Unit Area driven by the installed LSS and 'fff' is the model number of frame where the Line Unit Area resides
(Example: #2 L13 - would indicate Line Unit Area 2 in the 3746-L13)
(Example: #3 210 - would indicate Line Unit Area 3 in the 3745-210)

Figure 2-9 ZYX Company Model A11 TSS board before optimization

Part 4: 3746 Line Expansion Configuration

3746 Model **L13** Serial Number **02222**

Line Unit

Area 1				Area 2			
LIC Type							
<u> 3 </u>							

Additional Line Units

Area 3				Area 4			
LIC Type							
<u> </u>							

Area 5				Area 6			
LIC Type							
<u> </u>							

Area 7				Area 8			
LIC Type							
<u> </u>							

Figure 2-10 ZYX Company 3746 Model L13 line unit before optimization

2.3 Task 2 - ZYX Co. logical and functional inventory

ZYX Company uses their Network Control Program (NCP) generation statements to determine what active logical resources are defined in their NCP which relate to active physical resources in their communication controllers. In addition, NTuneMON can be used to determine the actual logical count of resources (for example, lines, PUs and LUs) that are active within the NCP program and relate back to live physical lines and devices. Additionally, this logical and functional inventory provides a clearer picture of the functions that

Part 2: NCP Owned Resource Section Lines

Serial Lines: Indicate Serial Line Groups.

Protocol	Speed	Line Count	SNI Count	Autocall Count
Example: SDLC	56K	5	3	0
Example: EP	9.6	12	0	6
SDLC	19,200	74	10	0
SDLC	56,000	20	8	0
SDLC	1,544,000	2	2	0
EP	14,400	6	0	0

Legends:	
Protocol	The line group protocol (SDLC, BSC3270, EP, Frame Relay, X.25, etc.).
Speed	The speed of the line group.
Line Count	The number of lines of a certain speed and protocol.
SNI Count	The count of any SNI lines within the group.
Autocall Count	The count of any Autocall lines within the group.

Figure 2-12 ZYX Company NCP owned resource section lines for NCP1 before optimization

Part 2: NCP Owned Resource Section

Token-Ring

Downstream PU Count (DSPU)	Logical Unit Count (LUDRPOOL)	TICs In Use
<i>Example: 25,000</i>	<i>50,000</i>	<i>4</i>
8,000	32,000	6
Note: Counts in this table are for all Token-Ring resources within this NCP.		

Are you currently using duplicate TIC addresses to balance and back up your Token-Ring SNA traffic?

Yes _____

No _____

Ethernet LAN (NCP Owned Resources)

What type of traffic is on your Ethernet LAN connection?

SNA _____

IP _____

If you are running IP, how many routes are you supporting? _____

Figure 2-13 ZYX Company token-ring and Ethernet LAN for NCP1 before optimization

Part 2: NCP Owned Resource Section Channels

Channel Type	LPAR Count
Bus and Tag	6
ESCON	0
<p>Note: If you are currently not using ESCON, is your mainframe ESCON capable?</p> <p>Yes <u> x </u></p> <p>No <u> </u></p>	

Figure 2-14 ZYX Company channels for NCP1 before optimization

2.4 Task 3 - ZYX Company reconcile and optimize

As ZYX Company examined their physical and logical inventories, they soon began to realize that there were more hardware features physically installed on the communication controllers than required to support their business environment. As ZYX Company explored further, they determined that the amount of physical resources (lines) required to support their environment did not require the four individual 3745-210 communication controllers currently installed in their infrastructure. In support of this further exploration, NTuneMON also indicated that many of the resources defined in the NCP generation statements were not actually being used and had not been activated. Taking this a step further, these resources could be removed from the NCP source, as well as physically removed from the communication controller. Additionally ZYX Company decided that the use of the older communication controller technology (3745-210) may not be the best use of their IT budget dollars. ZYX Company began to examine the “best-fit” alternatives to fit their current infrastructure requirements.

To ZYX Company, the following points were clear:

- ▶ Reduced requirement for low-speed lines. The communication controllers were originally installed to support low-speed line concentration. Over time, many of these low-speed lines have been migrated to token-ring resources. In lieu of physical line support via Line Interface Couplers (LICs) and Low Speed Scanners (LSS) these physical and logical resources are now being handled by the Token-Ring Interface Coupler (TICs) on the communication controllers.

ZYX Company analysis showed that while the resources have migrated from low-speed lines to token-ring access, the low-speed line hardware was still physically installed on the communication controllers. These physical line resources (LICs and scanners) were the main reason that additional 3746 expansion frames were required. These additional expansion frames resulted in wasted floor space, power requirements, and higher NCP charges and maintenance charges.

- ▶ ESCON® is the preferred channel technology. ZYX Company's installed 3745-210 communication controllers did not have the capability to support ESCON channel technology. In their current environment, ZYX Company deployed standard parallel channel technology to connect the communication controllers to the S/390® server environment. These parallel channel connections are also known as Bus and Tag technology. In addition, ZYX Company realized they could simplify their channel connectivity scenario as well as significantly improve channel throughput by using ESCON technology. It also became apparent as ZYX Company evaluated their mainframe environment, including future IBM zSeries technology, that future channel connectivity requirements would be based on ESCON channel technologies. If ZYX Company were to continue to support their current communication controller environment, they would be required to add legacy parallel channels to their mainframe technology acquisitions.
- ▶ Performance characteristics were not sufficient. As ZYX Company continued their evaluations, the performance of the communication controllers were evaluated. The installed 3745-210 models had multiple performance limitations. These performance limitations were token-ring performance and resource limitations, channel (parallel vs. ESCON), line speeds and concentration, along with memory limitations. Some of these performance limitations were relative to inactive resources within the NCP, and their impact on NCP storage along with the number of active resources supported.
- ▶ EP is still needed. As with many financial institutions, ZYX Company realized the need for continued support for Emulation Program (EP) within their environment. Some of these needs included branch banking devices, wire transfer, institution-to-institution connections, and other applications that use legacy device types for communication.
- ▶ Large SNI need. ZYX Company has a large amount of SNA Network Interconnect (SNI) connections to other financial institutions, businesses, and vendors. SNI network connections are very common and are the “workhorse” for connections between SNA networks. The two key attributes for SNI connections are network resource independence (no naming convention issues along with resource resolution) and access security (can be open or predefined to limit access, and/or a combination of both). In most cases, SNI is not easily replaced because it requires coordination with SNI partners. In

ZYX Company's evaluation, it was determined that they would need to continue with SNI connections for the foreseeable future.

- ▶ Current failover scenarios are not adequate. Over the years ZYX Company added more and more SNA resources without adequately adding failover hardware. This is one area that ZYX Company needed to evaluate: "At what cost can we provide backup resource, what resources are critical to our business, and what features and configurations are available on the communication controller technology to enable failover and backup in our environment?"

2.4.1 Technology results

ZYX Company began evaluating the above factors against the costs of alternative technology. ZYX Company chose to optimize their environment in the following manner:

- ▶ Reengineering the 3745 Communication Controller environment to utilize the newer 3746-900 technology.
- ▶ Migrating a portion of their resources to TN3270 technology supported on the mainframe environment.
- ▶ Consolidating from four current communication controllers to two running one active NCP on each with a backup (twin-backup).
- ▶ Making the primary focus of the new environment to be token-ring concentration, SNI, and EP.

Their new communication controller environment consists of two similarly configured 3745-61As with attached 3746-900 frames.

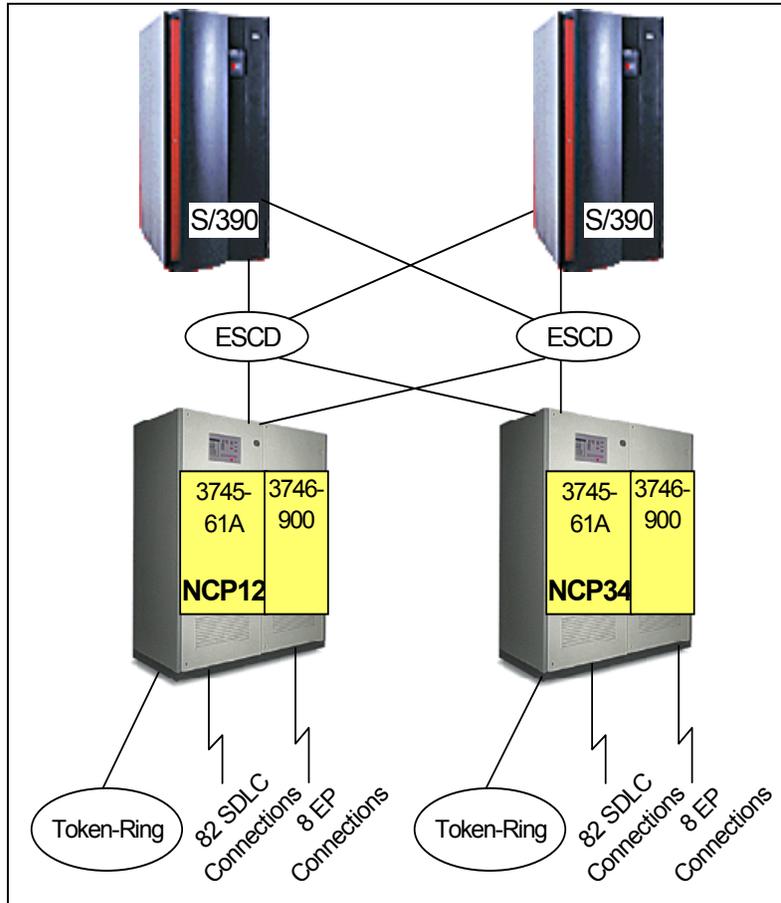


Figure 2-15 ZYX Company after optimization

The 3746-900 allows for the migration to ESCON channel technology for all their NCP traffic and has improved performance characteristics over the old 3745-210 environment with parallel channels. The 3746-900 also allowed for the consolidation of physical token-ring and line connections, doing more, faster, in a smaller footprint. The two 3745-61As with two 3746-900s meant a significant floor space reduction over the original installation.

All resources were migrated to the advanced hardware of the 3746-900 with the exception of EP. EP requires older generation hardware including Bus and Tag channels. These resources were physically placed on the base 3745-61A.

Each of these two complexes will have an active NCP running on Central Control Unit-A (CCU-A), and a backup NCP loaded (CCU-B). This will help satisfy ZYX Company in their failover/backup scenario. To migrate to this environment, ZYX

Company had to consolidate their current four NCPs into two NCPs. In NCP consolidation, one must consider the additional storage requirements in consolidating multiple NCPs and associated resources into a single NCP. In ZYX Company's example, the choice of a 3745-61A communication controller with its maximum storage configuration allowed for the ease of consolidation.

The new physical and logical and functional inventory sheets for one 3745-61A and 3746-900 complex are illustrated below:

<h2>3745 Physical Inventory</h2> <h3>Part 1: Machine Overview</h3>	
3745 Model	<u> 61A </u> Serial Number <u> 00055 </u>
Customer Designation (Ex. NCP Name)	<u> NCP12 </u>
Indicate all attached 3746 expansion frames.	
3746-A11	SN <u> </u>
3746-A12	SN <u> </u>
3746-L13	SN <u> </u>
3746-L14	SN <u> </u>
3746-L15	SN <u> </u>
3746-900	SN <u> 90900 </u>
Controller Expansion (Rack) Qty	<u> 1 </u> (0,1 or 2)
(Please inventory each attached frame separately on 3746 Model Sheets.)	
How much memory (per CCU if Model 41x or 61x)?	
4M	<u> </u>
8M	<u> </u>
16M	<u> x </u>

Figure 2-16 ZYX Company 3745 Model 61A machine overview after optimization

3745 Physical Inventory

Part 1: Machine Overview

3745 Model 61A Serial Number 00055

Console Information

3151 or Equivalent

Service Processor **x**

Type:

9577

9585

3172 P/N 41H7520

3172 P/N 55H7630

7585

6275

6563 **x**

Indicate serial numbers of other 37XXs that share this console.

Figure 2-17 ZYX Company 3745 Model 61A console information after optimization

3745 Physical Inventory

Part 2: Base 3745 Configuration

Section A: 3745 Models x1x

3745 Model 61A Serial Number 00055

Channel Board

Bus Group 1	Pos. 1	Pos. 2	Pos. 3	Pos. 4
	<u>CADS</u>	_____	_____	_____
Bus Group 2	Pos.	Pos.	Pos. 7	Pos. 8
	_____	_____	_____	_____

Legend:

CADS = Channel Adapter Data Streaming
BCCA = Buffer Chaining Channel Adapter
TPS = Two Processor Switch

Figure 2-18 ZYX Company 3745 Model 61A channel board after optimization

3745 Physical Inventory

Part 2: Base 3745 Configuration

Section A: 3745 Models x1x

3745 Model **61A** Serial Number **00055**

TSS Board

Bus Group 1	Pos. 1		Pos. 2		Pos. 3		Pos. 4	
	Adap: <u> LSS </u>		Adap: _____		Adap: _____		Adap: _____	
	Port #1 61A	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____
Bus Group 2	Pos. 5		Pos. 6		Pos. 7		Pos. 8	
	Adap: _____		Adap: _____		Adap: _____		Adap: _____	
	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____

Legend:

Adapters

- HSS = High Speed Scanner
- LSS = Low Speed Scanner
- ELA = Ethernet/IEEE 802.3 Adapter
- TRA1 = Token-Ring Adapter Type 1
- TRA2 = Token-Ring Adapter Type 2

Port Cards

- TIC = Token-Ring Interface Coupler
- LAN = Ethernet LAN Attached
- #x = Where 'x' is the Line Unit Area driven by the installed LSS
(Example: #2 L13 - would indicate Line Unit Area 2 in the 3746-L13)
- (Example: #3 210 - would indicate Line Unit Area 3 in the 3745-210)

Figure 2-19 ZYX Company 3745 Model 61A TSS board after optimization

3745 Physical Inventory
Part 2: Base 3745 Configuration

Section A: 3745 Models x1x

3745 Model **61A** Serial Number **00055**

Line Unit

Area 1				Area 2			
LIC Type	LIC Type	LIC Type	LIC Type	LIC Type	LIC Type	LIC Type	LIC Type
<u> 1 </u>	<u> 1 </u>	<u> </u>					

Additional Line Units

Area 3				Area 4			
LIC Type							
<u> </u>							

Figure 2-20 ZYX Company 3745 Model 61A line unit after optimization

3745 Physical Inventory

Part 5: 3746-900 Configuration

3746 Model 900 Serial Number 90900

Overview

Extended Microcode Options:

Extended Functions 1 (5800) _____
Extended Functions 2 (5802) _____
Extended Functions 3 (5801) _____
Extended Functions 4 (5810) x _____
Extended Functions 5 (5812) x _____
Extended Functions 6 (5813) x _____
X.25 (5030) _____
IP (5033) _____

Multiaccess Enclosure (MAE) present?

Yes _____ No x

MAE Microcode Options:

Extended Functions 1 (5804) _____
Extended Functions 2 (5805) _____
Extended Functions 3 (5807) _____
TN3270 Server (5806) _____

Network Node Processor (NNP)

Qty 1 (0,1 or 2)

NNP Type:

Type 1 (3172) _____
Type 2 (7585) _____
Type 3 (6275) _____
Type 4 (6563) x _____
Type 5 (6578) _____

Figure 2-21 ZYX Company 3746 Model 900 machine overview after optimization

3745 Physical Inventory

Part 5: 3746-900 Configuration

3746 Model **900** Serial Number **90900**

Base Enclosure Top View

P	N	M	L	K	J	H	G	F	E	D	C	Rear Side
Port	Port	Port	Port	Port	Port	Port	Port_	Port	Port	Port	Port	
_____	<u>LIC12</u>	<u>LIC113</u>	<u>LIC112</u>	<u>TIC3</u>	<u>TIC3</u>	_____	<u>CBC</u>	TIC3	CBC	N/A	N/A	
Processor		Processor		Processor		Processor		Processor CBSP Type <u> 3 </u>		Processor Power Supply		
_____		<u>CLP3</u>		<u>TRP3</u>		<u>TRP3</u>		_____		_____		
Slot 6 (P)		Slot 5 (M)		Slot 4 (K)		Slot 3 (H)		Slot 2 (F)		Slot 1 (D)		Front Side

- Notes: 1. Indicate CBSP Type in Slot 2 (F)
 2. For 3745 Models 41A and 61A, Slot 3 (H) must be a TRPx and Port G must by a CBC

Processors:

TRP = Token-Ring Processor Type 1
 TRP2 = Token-Ring Processor Type 2
 TRP3 = Token-Ring Processor Type 3
 ESCP = ESCON Processor Type 1
 ESCP2 = ESCON Processor Type 2
 ESCP3 = ESCON Processor Type 3
 CLP = Communication Line Processor
 CLP3 = Communication Line Processor Type 3
 SIE = Switch Interface Extension (MAE connection)

Ports:

TIC3 = Token-Ring Coupler Type 3
 ETH = Ethernet Port/Ethernet-TR Bridge
 ESCC = ESCON Coupler Type 1
 ESCC2 = ESCON Coupler Type 2
 LIC11x = Line Interface Coupler Type 11 where 'x' is
 the Line Connection Box (LCB) ID where the
 ARCS are installed that correspond to this LIC
 LIC12 = Line Interface Coupler Type 12

Figure 2-22 ZYX Company 3746 Model 900 base enclosure top view after optimization

3745 Physical Inventory

Part 5: 3746-900 Configuration

3746 Model 900 Serial Number 90900

Expansion Enclosure 1 Top View

P	N	M	L	K	J	H	G	F	E	D	C	Rear Side
Port	Port	Port	Port	Port	Port	Port	Port	Port	Port	Port	Port	
_____	_____	<u>LIC12</u>	<u>LIC111</u>	_____	_____	_____	<u>ESCC2</u>	_____	<u>ESCC2</u>	<u>TIC3</u>	<u>TIC3</u>	
Processor		Processor		Processor		Processor		Processor		Processor		
_____		<u>CLP3</u>		_____		<u>ESCP3</u>		<u>ESCP3</u>		<u>TRP3</u>		
Slot 12 (P)		Slot 11 (M)		Slot 10 (K)		Slot 9 (H)		Slot 8 (F)		Slot 7 (D)		Front Side

Processors:

TRP = Token-Ring Processor Type 1
 TRP2 = Token-Ring Processor Type 2
 TRP3 = Token-Ring Processor Type 3
 ESCP = ESCON Processor Type 1
 ESCP2 = ESCON Processor Type 2
 ESCP3 = ESCON Processor Type 3
 CLP = Communication Line Processor
 CLP3 = Communication Line Processor Type 3
 SIE = Switch Interface Extension (MAE connection)

Ports:

TIC3 = Token-Ring Coupler Type 3
 ETH = Ethernet Port/Ethernet-TR Bridge
 ESCC = ESCON Coupler Type 1
 ESCC2 = ESCON Coupler Type 2
 LIC11x = Line Interface Coupler Type 11 where 'x' is
 the Line Connection Box (LCB) ID where the
 ARCS are installed that correspond to this LIC
 LIC12 = Line Interface Coupler Type 12

Figure 2-23 ZYX Company 3746 Model 900 expansion enclosure 1 top view after optimization

3745 Physical Inventory

Part 5: 3746-900 Configuration

3746 Model 900 Serial Number 90900

LCB ID 1 Line Connection Box Layout

	0	+1	+2	+3	+4	+5	+6	+7	+8	+9	+10	+11	+12	+13	+14	
ARC Type	V.35	V.35	V.35	V.35	V.24	To LCBE										
Attach	DCE															
Length	15	15	15	15	12	12	12	12	12	12	12	12	12	12	12	

Line Connection Box Expansion Layout

	+16	+17	+18	+19	+20	+21	+22	+23	+24	+25	+26	+27	+28	+29	+30	
ARC Type	V.35	V.35	V.35	V.35	V.24	To LCB										
Attach	DCE															
Length	15	15	15	15	12	12	12	12	12	12	12	12	12	12	12	

	Length:	
ARC Types:	.6 = .6 Meters	
V.35	1.2 = 1.2 Meters	
V.24	2.4 = 2.4 Meters	
X.21	5 = 5 Meters	
	10 = 10 Meters	
	12 = 12 Meters	
	15 = 15 Meters	
	ST = Stub cable to 3745 cable	
		Notes:
		Fill out one sheet for each LCB/LCBE combo installed

Note: A 3746-900 may have as many as 32 LCB/LCBE pairs. Fill out one of these sheets for each one.

Figure 2-24 ZYX Company 3746 Model 900 line connection box ID 1 after optimization

3745 Physical Inventory

Part 5: 3746-900 Configuration

3746 Model 900 Serial Number 90900

LCB ID 2 Line Connection Box Layout

	0	+1	+2	+3	+4	+5	+6	+7	+8	+9	+10	+11	+12	+13	+14	
ARC Type	V.35	V.35	V.35	V.24			To LCBE									
Attach	DCE															
Length	15	15	15	12	12	12	12	12	12	12	12	12	12			

Line Connection Box Expansion Layout

	+16	+17	+18	+19	+20	+21	+22	+23	+24	+25	+26	+27	+28	+29	+30	
ARC Type	V.35	V.35	V.35	V.24			To LCB									
Attach	DCE															
Length	15	15	15	12	12	12	12	12	12	12	12	12	12			

Length:

- .6 = .6 Meters
- 1.2 = 1.2 Meters
- 2.4 = 2.4 Meters
- 5 = 5 Meters
- 10 = 10 Meters
- 12 = 12 Meters
- 15 = 15 Meters
- ST = Stub cable to 3745 cable

ARC Types:
 V.35
 V.24
 X.21

Attach:
 DCE
 DTE

Notes:

Fill out one sheet for each LCB/LCBE combo installed

Note: A 3746-900 may have as many as 32 LCB/LCBE pairs. Fill out one of these sheets for each one.

Figure 2-25 ZYX Company 3746 Model 900 line connection box ID 2 after optimization

3745 Physical Inventory

Part 5: 3746-900 Configuration

3746 Model 900 Serial Number 90900

LCB ID 3 Line Connection Box Layout

	0	+1	+2	+3	+4	+5	+6	+7	+8	+9	+10	+11	+12	+13	+14	
ARC Type	V.35	V.35	V.35	V.24			To LCBE									
Attach	DCE															
Length	15	15	15	12	12	12	12	12	12	12	12	12	12			

Line Connection Box Expansion Layout

	+16	+17	+18	+19	+20	+21	+22	+23	+24	+25	+26	+27	+28	+29	+30	
ARC Type	V.35	V.35	V.35	V.24			To LCB									
Attach	DCE															
Length	15	15	15	12	12	12	12	12	12	12	12	12	12			

	Length:	
ARC Types: V.35 V.24 X.21	Attach: DCE DTE	.6 = .6 Meters 1.2 = 1.2 Meters 2.4 = 2.4 Meters 5 = 5 Meters 10 = 10 Meters 12 = 12 Meters 15 = 15 Meters ST = Stub cable to 3745 cable
		Notes: Fill out one sheet for each LCB/LCBE combo installed

Note: A 3746-900 may have as many as 32 LCB/LCBE pairs. Fill out one of these sheets for each one.

Figure 2-26 ZYX Company 3746 Model 900 line connection box ID 3 after optimization

Part 2: NCP Owned Resource Section Lines

Serial Lines: Indicate Serial Line Groups.

Protocol	Speed	Line Count	SNI Count	Autocall Count
Example: SDLC	56K	5	3	0
Example: EP	9.6	12	0	6
SDLC	19,200	60	30	0
SDLC	56,000	20	16	0
SDLC	1,544,000	2	2	0
EP	14,400	8	0	0

Legends:	
Protocol	The line group protocol (SDLC, BSC3270, EP, Frame Relay, X.25, etc.).
Speed	The speed of the line group.
Line Count	The number of lines of a certain speed and protocol.
SNI Count	The count of any SNI lines within the group.
Autocall Count	The count of any Autocall lines within the group.

Figure 2-28 ZYX Company NCP owned resource section for NCP12 after optimization

Part 2: NCP Owned Resource Section

Token-Ring

Downstream PU Count (DSPU)	Logical Unit Count (LUDRPOOL)	TICs In Use
<i>Example: 25,000</i>	<i>50,000</i>	<i>4</i>
8,000	32,000	4
Note: Counts in this table are for all Token-Ring resources within this NCP.		

Are you currently using duplicate TIC addresses to balance and backup your Token-Ring SNA traffic?

Yes _____

No _____

Ethernet LAN (NCP Owned Resources)

What type of traffic is on your Ethernet LAN connection?

SNA _____

IP _____

If you are running IP, how many routes are you supporting? _____

Figure 2-29 ZYX Company token-ring and Ethernet LAN for NCP12 after optimization

Part 2: NCP Owned Resource Section Channels

Channel Type	LPAR Count
Bus and Tag	1
ESCON	6
<p>Note: If you are currently not using ESCON, is your mainframe ESCON capable?</p> <p>Yes _____</p> <p>No _____</p>	

Figure 2-30 ZYX Company channels for NCP12 after optimization

2.4.2 Business results

ZYX Company has realized a significant reduction in their operating expense from this project. The following represent realized savings.

- ▶ Floor space reduction of 50%
- ▶ Maintenance savings of 10% per year
- ▶ NCP Software savings of 85% (\$70,272 per year)
- ▶ Reduced line costs by 50%
- ▶ Reduced environmental costs (HVAC)

The most notable savings for ZYX Company were the software charges for NCP. Because of the usage tier structure of NCP, it is far less expensive to utilize the 3746-900 communication controller hardware than the hardware of the base 3745. To further reduce NCP charges, ZYX Company decided to utilize the NCP usage tier 'C' option. This option allowed some of the NCP charges to be waived. To realize this usage tier 'C' option, ZYX Company invested in the Network Node Processor (NNP) feature on the 3746-900 communication controller, new from IBM. For more information about usage tier 'C,' contact your IBM representative or authorized business partner.

2.4.3 The next step

Now that some significant results were achieved in the optimization exercise, ZYX Company has begun to examine their strategic plan.

2.5 Task 4 - ZYX Company controller strategic planning

ZYX Financial would like to further reduce their dependence on controllers. They have identified four major functional areas that the controllers continue to provide:

- ▶ Token-ring gateway to the mainframe
- ▶ Low-speed line concentration
- ▶ SNI Gateway functions
- ▶ EP Support

Of these functional areas, ZYX Company decided to make an effort to migrate their low-speed line concentration and token-ring gateway functions away from the controllers within three years. They chose to migrate to Ethernet technology and utilize the Open System Adapters (OSA) currently available as a hardware feature on their IBM zSeries Server and the z/OS operating system. Refer to Chapter 15, “OSA-Express” on page 269 for more information on this type of migration.

The other two functions, SNI and EP, are harder for ZYX Company to migrate. They have chosen to make a long-term study into how to migrate these functions. These two areas in particular are usually difficult migration issues for most organizations. See Chapter 3, “Hot topics - SNI, EP, BSC, and X.25” on page 89 for more information on SNI and EP.

2.6 Task 5 - ZYX Co. functional alternative migration

ZYX Company has begun the token-ring and low-speed line migration. Their first step was to move to newer mainframe technology with the IBM zSeries. This provided ZYX Company with OSA gigabit Ethernet technology to a mainframe with sufficient CPU processing power to support a large SNA to IP migration effort.

2.7 Client example 2: M Manufacturing Company

M Manufacturing Company is a medium-size manufacturing company. Their name has been altered. They have had 3745 Communication Controllers installed for many years, and have stayed current with newer 37xx communication controller technology and currently utilize 3746-900 communication controllers in their network infrastructure.

M Manufacturing Company recently acquired another smaller company, A Company. In planning the acquisition, M Manufacturing Company was faced with merging a second network infrastructure that had older 3725 and 3745 Communication Controllers at its core.

M Manufacturing Company decided to undertake an optimization approach to refresh their own environment as well as incorporating A Company resources. At the beginning of their optimization project, M Manufacturing Company had two individual 3745-31A communication controllers, each with 3746-900 expansion frames.

The acquired A Company currently had two 3725s and one 3745-170 communication controller. All of these communication controllers were local controllers, located in their respective data centers.

M Manufacturing Company determined that the best course of action was to consolidate the two data centers into a single data center and utilize mainframe technology already in place at the M Manufacturing Company data center location. As such, it was determined that refreshing their communication controller environment to include both their existing resources and those of A Company was the best approach. Their plan was to remove all resources from A Company data center.

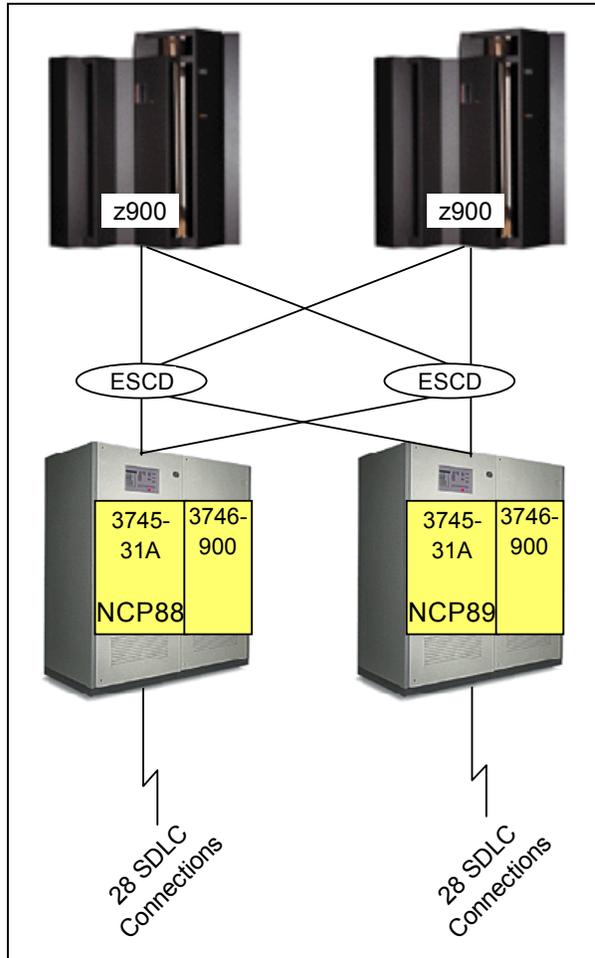


Figure 2-31 M Manufacturing Company before optimization

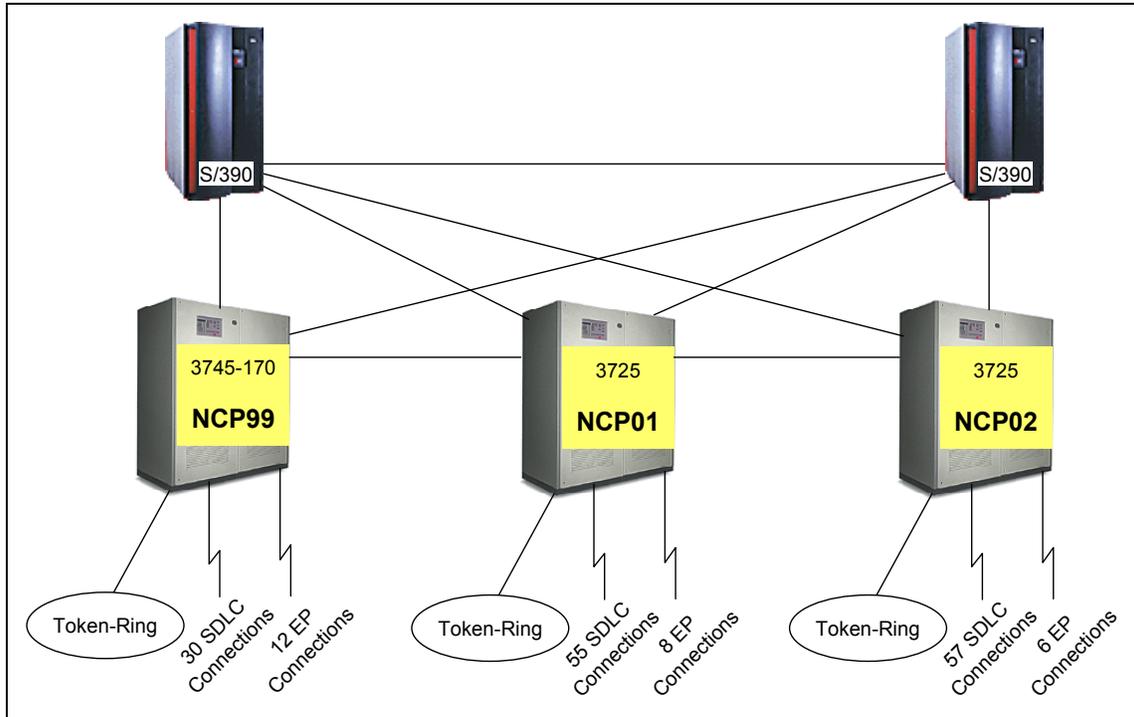


Figure 2-32 A Company before merger with M Manufacturing Company and optimization

2.8 Task 1 - M Company and A Company physical inventory

M Manufacturing Company engaged their IBM Customer Engineer (CE) to determine what was physically installed on their 3745-31A and 3746-900 communication controllers. They also contracted with the local CE of the acquired company, A Company, to inventory their existing 3745-170. A sample of their physical inventory sheets is illustrated below:

3745 Physical Inventory

Part 1: Machine Overview

3745 Model 31A Serial Number 00110

Customer Designation (Ex. NCP Name) NCP88

Indicate all attached 3746 expansion frames.

3746-A11 SN _____

3746-A12 SN _____

3746-L13 SN _____

3746-L14 SN _____

3746-L15 SN _____

3746-900 SN 91234

Controller Expansion (Rack) Qty 1 (0,1 or 2)

(Please inventory each attached frame separately on 3746 Model Sheets.)

How much memory (per CCU if Model 41x or 61x)?

4M _____

8M _____

16M x

Figure 2-33 M Manufacturing Company 3745 Model 31A machine overview before optimization

3745 Physical Inventory

Part 1: Machine Overview

3745 Model 31A Serial Number 00110

Console Information

3151 or Equivalent _____

Service Processor x

Type:

9577 _____

9585 _____

3172 P/N 41H7520 _____

3172 P/N 55H7630 _____

7585 _____

6275 x

6563 _____

Indicate serial numbers of other 37XXs that share this console.

Figure 2-34 M Manufacturing Company 3745 Model 31A console information before optimization

3745 Physical Inventory

Part 5: 3746-900 Configuration

3746 Model 900 Serial Number 91234

Overview

Extended Microcode Options:

Extended Functions 1 (5800) _____
Extended Functions 2 (5802) _____
Extended Functions 3 (5801) _____
Extended Functions 4 (5810) x
Extended Functions 5 (5812) x
Extended Functions 6 (5813) _____
X.25 (5030) _____
IP (5033) _____

Multiaccess Enclosure (MAE) present?

Yes _____ No x

MAE Microcode Options:

Extended Functions 1 (5804) _____
Extended Functions 2 (5805) _____
Extended Functions 3 (5807) _____
TN3270 Server (5806) _____

Network Node Processor (NNP)

Qty 0 (0,1 or 2)

NNP Type:

Type 1 (3172) _____
Type 2 (7585) _____
Type 3 (6275) _____
Type 4 (6563) _____
Type 5 (6578) _____

Figure 2-35 M Manufacturing Company 3746 Model 900 machine overview before optimization

3745 Physical Inventory

Part 5: 3746-900 Configuration

3746 Model **900** Serial Number **91234**

Base Enclosure Top View

P	N	M	L	K	J	H	G	F	E	D	C	Rear Side
Port	Port	Port	Port	Port	Port	Port	Port_	Port	Port	Port	Port	
_____	<u>LIC12</u>	<u>LIC112</u>	<u>LIC111</u>	_____	<u>ESCC2</u>	_____	<u>CBC</u>	TIC3	CBC	N/A	N/A	
Processor		Processor		Processor		Processor		Processor CBSP Type <u> 3 </u>		Processor Power Supply		
_____		<u>CLP3</u>		<u>ESCP3</u>		<u>TRP3</u>						
Slot 6 (P)		Slot 5 (M)		Slot 4 (K)		Slot 3 (H)		Slot 2 (F)		Slot 1 (D)		Front Side

- Notes: 1. Indicate CBSP Type in Slot 2 (F)
 2. For 3745 Models 41A and 61A, Slot 3 (H) must be a TRPx and Port G must by a CBC

Processors:

TRP = Token-Ring Processor Type 1
 TRP2 = Token-Ring Processor Type 2
 TRP3 = Token-Ring Processor Type 3
 ESCP = ESCON Processor Type 1
 ESCP2 = ESCON Processor Type 2
 ESCP3 = ESCON Processor Type 3
 CLP = Communication Line Processor
 CLP3 = Communication Line Processor Type 3
 SIE = Switch Interface Extension (MAE connection)

Ports:

TIC3 = Token-Ring Coupler Type 3
 ETH = Ethernet Port/Ethernet-TR Bridge
 ESCC = ESCON Coupler Type 1
 ESCC2 = ESCON Coupler Type 2
 LIC11x = Line Interface Coupler Type 11 where 'x' is
 the Line Connection Box (LCB) ID where the
 ARCS are installed that correspond to this LIC
 LIC12 = Line Interface Coupler Type 12

Figure 2-36 M Manufacturing Company 3746 Model 900 base enclosure top view before optimization

3745 Physical Inventory

Part 5: 3746-900 Configuration

3746 Model 900 Serial Number 91234

LCB ID 1 Line Connection Box Layout

	0	+1	+2	+3	+4	+5	+6	+7	+8	+9	+10	+11	+12	+13	+14	
ARC Type	<u>V.35</u>				<u>V.35</u>	<u>V.24</u>	<u>V.24</u>	<u>V.24</u>	<u>V.24</u>	<u>V.24</u>	<u>V.24</u>					To LCBE
Attach	<u>DCE</u>				<u>DCE</u>											
Length	<u>15</u>				<u>15</u>	<u>12</u>	<u>12</u>	<u>12</u>	<u>12</u>	<u>12</u>	<u>12</u>					

Line Connection Box Expansion Layout

	+16	+17	+18	+19	+20	+21	+22	+23	+24	+25	+26	+27	+28	+29	+30	
ARC Type	<u>V.35</u>				<u>V.24</u>					To LCB						
Attach	<u>DCE</u>				<u>DCE</u>											
Length	<u>15</u>				<u>12</u>											

Length:

- .6 = .6 Meters
- 1.2 = 1.2 Meters
- 2.4 = 2.4 Meters
- 5 = 5 Meters
- 10 = 10 Meters
- 12 = 12 Meters
- 15 = 15 Meters
- ST = Stub cable to 3745 cable

ARC Types:

- V.35
- V.24
- X.21

Attach:

- DCE
- DTE

Notes:

Fill out one sheet for each LCB/LCBE combo installed

Note: A 3746-900 may have as many as 32 LCB/LCBE pairs. Fill out one of these sheets for each one.

Figure 2-37 M Manufacturing Company 3746 Model 900 line connection box ID 1 before optimization

3745 Physical Inventory

Part 5: 3746-900 Configuration

3746 Model 900 Serial Number 91234

LCB ID 2 Line Connection Box Layout

	0	+1	+2	+3	+4	+5	+6	+7	+8	+9	+10	+11	+12	+13	+14	
ARC Type	V.35	V.35	V.35	V.35	V.35	V.24	V.24	V.24	V.24	V.24	V.24					To LCBE
Attach	DCE															
Length	15	15	15	15	15	12	12	12	12	12	12					

Line Connection Box Expansion Layout

	+16	+17	+18	+19	+20	+21	+22	+23	+24	+25	+26	+27	+28	+29	+30	
ARC Type																To LCB
Attach																
Length																

Length:

ARC Types:	Attach:	.6 = .6 Meters 1.2 = 1.2 Meters 2.4 = 2.4 Meters 5 = 5 Meters 10 = 10 Meters 12 = 12 Meters 15 = 15 Meters ST = Stub cable to 3745 cable
V.35	DCE	
V.24	DTE	
X.21		

Notes:

Fill out one sheet for each LCB/LCBE combo installed

Note: A 3746-900 may have as many as 32 LCB/LCBE pairs. Fill out one of these sheets for each one.

Figure 2-38 M Manufacturing Company 3746 Model 900 line connection box ID 2 before optimization

3745 Physical Inventory

Part 1: Machine Overview

3745 Model 170 Serial Number 54321

Customer Designation (Ex. NCP Name) NCP99

Indicate all attached 3746 expansion frames.

3746-A11 SN _____
3746-A12 SN _____
3746-L13 SN _____
3746-L14 SN _____
3746-L15 SN _____
3746-900 SN _____

Controller Expansion (Rack) Qty 0 (0,1 or 2)

(Please inventory each attached frame separately on 3746 Model Sheets.)

How much memory (per CCU if model 41x or 61x)?

4M _____

8M x

16M _____

Figure 2-39 A Company 3745 Model 170 machine overview before optimization

3745 Physical Inventory

Part 1: Machine Overview

3745 Model **170** Serial Number **54321**

Console Information

3151 or Equivalent **x**

Service Processor

Type:

9577

9585

3172 P/N 41H7520

3172 P/N 55H7630

7585

6275

6563

Indicate serial numbers of other 37XXs that share this console.

Figure 2-40 A Company 3745 Model 170 console information before optimization

3745 Physical Inventory

Part 2: Base 3745 Configuration

Section E: 3745 Models 17x

3745 Model 170 Serial Number 54321

Channel Board

Pos. 8	Pos. 7	Pos. 6	Pos. 5
<u>CADS</u>	<u>CADS</u>	<u>CADS</u>	<u>CADS</u>

Legend:

CADS = Channel Adapter Data Streaming
BCCA = Buffer Chaining Channel Adapter
TPS = Two Processor Switch

Figure 2-41 A Company 3745 Model 170 channel board before optimization

3745 Physical Inventory

Part 2: Base 3745 Configuration

Section E: 3745 Models 17x

3745 Model 170 Serial Number 54321

TR Board

Pos. 1 <u>TRA2</u>	
Port <u>TIC</u>	Port <u>TIC</u>

Legend:

Adapters

TRA1 = Token-Ring Adapter Type 1

TRA2 = Token-Ring Adapter Type 2

Port Cards

TIC = Token-Ring Interface Coupler

Figure 2-42 A Company 3745 Model 170 TR board before optimization

3745 Physical Inventory

Part 2: Base 3745 Configuration

Section E: 3745 Models 17x

3745 Model 170 Serial Number 54321

TSS Board

Pos. 12 Adap: <u>LSS</u>		Pos. 11 Adap: <u>LSS</u>		Pos. 10 Adap: <u>LSS</u>		Pos. 9 Adap: <u>LSS</u>		Pos. 4 Adap: <u>LSS</u>		Pos. 3 Adap: _____	
Port <u>#4</u>	Port _____	Port <u>#3</u>	Port _____	Port <u>#2</u>	Port _____	Port <u>#1</u>	Port _____	Port <u>#5</u>	Port _____	Port _____	Port _____

Legend:

Adapters

HSS = High Speed Scanner
 LSS = Low Speed Scanner
 ELA = Ethernet/IEEE 802.3 Adapter

Port Cards

LAN = Ethernet LAN Attached
 #x = Where 'x' is the Line Unit Area
 driven by the installed LSS
 (Example: #2 L13 - would indicate
 Line Unit Area 2 in the 3746-L13)

Figure 2-43 A Company 3745 Model 170 TSS board before optimization

3745 Physical Inventory

Part 2: Base 3745 Configuration

Section E: 3745 Models 17x

3745 Model 170 Serial Number 54321

Line Unit

Area 1				Area 2			
LIC Type							
<u>1</u>							

Additional Line Units

Area 3				Area 4			
LIC Type							
<u>3</u>							

Figure 2-44 A Company 3745 Model 170 line unit before optimization

3745 Physical Inventory							
Part 2: Base 3745 Configuration							
Section E: 3745 Models 17x							
3745 Model <u> 170 </u>				Serial Number <u> 54321 </u>			
Additional Line Units							
Area 5				Area 6			
LIC Type	LIC Type	LIC Type	LIC Type	LIC Type	LIC Type	LIC Type	LIC Type
<u> 3 </u>	<u> 3 </u>	<u> </u>	<u> </u>	<u> </u>	<u> </u>	<u> </u>	<u> </u>

Figure 2-45 A Company 3745 Model 170 line unit areas 5 and 6 before optimization

Only a logical inventory was required for A Company’s installed 3725 communication controllers because they did not plan to retain or optimize the 3725 controllers. The 3725 communication controllers have been withdrawn from service and retaining those machines would incur an unacceptable level of risk and, consequently, the 3725s were not part of M Manufacturing Company’s long-term strategy.

2.9 Task 2 - M Company and A Company logical and functional inventory

M Manufacturing Company used their Network Control Program (NCP) generation statements to determine what logical resources were defined in their NCP, and corresponded with actively used physical resources in their communication controllers. In addition, NTuneMON can be used to determine the actual logical count of resources (for example, lines, PUs, and LUs) that are active within the NCP program and relate back to live physical lines and devices. Additionally, this logical and functional inventory provides a clearer picture of the functions that NCP was providing in their business environment. Following are the logical and functional inventory sheets for M Manufacturing Company and A Company:

Part 2: NCP Owned Resource Section Lines

Serial Lines: Indicate Serial Line Groups.

Protocol	Speed	Line Count	SNI Count	Autocall Count
<i>Example: SDLC</i>	56K	5	3	0
<i>Example: EP</i>	9.6	12	0	6
SDLC	19,200	20	0	0
SDLC	56,000	5	4	0
SDLC	256,000	2	1	0
SDLC	1,544,000	1	1	0

Legends:	
Protocol	The line group protocol (SDLC, BSC3270, EP, Frame Relay, X.25, etc.).
Speed	The speed of the line group.
Line Count	The number of lines of a certain speed and protocol.
SNI Count	The count of any SNI lines within the group.
Autocall Count	The count of any Autocall lines within the group.

Figure 2-47 M Manufacturing Company NCP owned resource section for NCP88 before optimization

Part 2: NCP Owned Resource Section

Token-Ring

Downstream PU Count (DSPU)	Logical Unit Count (LUDRPOOL)	TICs In Use
<i>Example: 25,000</i>	<i>50,000</i>	<i>4</i>
0	0	0
Note: Counts in this table are for all Token-Ring resources within this NCP.		

Are you currently using duplicate TIC addresses to balance and backup your Token-Ring SNA traffic?

Yes _____

No _____

Ethernet LAN (NCP Owned Resources)

What type of traffic is on your Ethernet LAN connection?

SNA _____

IP _____

If you are running IP, how many routes are you supporting? _____

Figure 2-48 M Manufacturing Company token-ring and Ethernet LAN for NCP88 before optimization

Part 2: NCP Owned Resource Section Channels

Channel Type	LPAR Count
Bus and Tag	0
ESCON	4
Note: If you are currently not using ESCON, is your mainframe ESCON capable? Yes _____ No _____	

Figure 2-49 M Manufacturing Company channels for NCP88 before optimization

Part 2: NCP Owned Resource Section Lines

Serial Lines: Indicate Serial Line Groups.

Protocol	Speed	Line Count	SNI Count	Autocall Count
<i>Example: SDLC</i>	56K	5	3	0
<i>Example: EP</i>	9.6	12	0	6
SDLC	19,200	20	2	0
SDLC	56,000	10	5	0
EP	14,400	12	0	0

Legends:	
Protocol	The line group protocol (SDLC, BSC3270, EP, Frame Relay, X.25, etc.).
Speed	The speed of the line group.
Line Count	The number of lines of a certain speed and protocol.
SNI Count	The count of any SNI lines within the group.
Autocall Count	The count of any Autocall lines within the group.

Figure 2-51 A Company NCP owned resource section for NCP99 before optimization

Part 2: NCP Owned Resource Section

Token-Ring

Downstream PU Count (DSPU)	Logical Unit Count (LUDRPOOL)	TICs In Use
<i>Example: 25,000</i>	<i>50,000</i>	<i>4</i>
1,000	3,600	2
Note: Counts in this table are for all Token-Ring resources within this NCP.		

Are you currently using duplicate TIC addresses to balance and back up your Token-Ring SNA traffic?

Yes _____

No _____

Ethernet LAN (NCP Owned Resources)

What type of traffic is on your Ethernet LAN connection?

SNA _____

IP _____

If you are running IP, how many routes are you supporting? _____

Figure 2-52 A Company token-ring and Ethernet LAN for NCP99 before optimization

Part 2: NCP Owned Resource Sectionyes, Channels

Channel Type	LPAR Count
Bus and Tag	4
ESCON	0
Note: If you are currently not using ESCON, is your mainframe ESCON capable?	
Yes <input checked="" type="checkbox"/>	
No <input type="checkbox"/>	

Figure 2-53 A Company channels for NCP99 before optimization

Part 2: NCP Owned Resource Section Lines

Serial Lines: Indicate Serial Line Groups.

Protocol	Speed	Line Count	SNI Count	Autocall Count
<i>Example: SDLC</i>	<i>56K</i>	<i>5</i>	<i>3</i>	<i>0</i>
<i>Example: EP</i>	<i>9.6</i>	<i>12</i>	<i>0</i>	<i>6</i>
SDLC	9,600	30	0	0
SDLC	19,200	25	4	0
EP	9,600	8	0	0

Legends:	
Protocol	The line group protocol (SDLC, BSC3270, EP, Frame Relay, X.25, etc.).
Speed	The speed of the line group.
Line Count	The number of lines of a certain speed and protocol.
SNI Count	The count of any SNI lines within the group.
Autocall Count	The count of any Autocall lines within the group.

Figure 2-55 A Company NCP owned resource section for NCP01 before optimization

Part 2: NCP Owned Resource Section

Token-Ring

Downstream PU Count (DSPU)	Logical Unit Count (LUDRPOOL)	TICs In Use
<i>Example: 25,000</i>	<i>50,000</i>	<i>4</i>
500	763	2
Note: Counts in this table are for all Token-Ring resources within this NCP.		

Are you currently using duplicate TIC addresses to balance and backup your Token-Ring SNA traffic?

Yes _____

No x

Ethernet LAN (NCP Owned Resources)

What type of traffic is on your Ethernet LAN connection?

SNA _____

IP _____

If you are running IP, how many routes are you supporting? _____

Figure 2-56 A Company token-ring and Ethernet LAN for NCP01 before optimization

Part 2: NCP Owned Resource Section Channels

Channel Type	LPAR Count
Bus and Tag	2
ESCON	0
<p>Note: If you are currently not using ESCON, is your mainframe ESCON capable?</p> <p>Yes <input checked="" type="checkbox"/> _____</p> <p>No <input type="checkbox"/> _____</p>	

Figure 2-57 A Company channels for NCP01 before optimization

2.10 Task 3 - M Manufacturing Company reconcile and optimize

Upon examination of the completed inventory sheets for both companies, some basic strategies began to emerge. First, M Manufacturing Company had already undertaken a large project to implement Enterprise Extender (EE). With the success of this project, the workload on the 3746-900s had diminished. It was determined that the A Company workload could also take advantage of the EE technology and thus reduce the dependence on the communication controllers as well.

The following points emerged:

- ▶ Continue Enterprise Extender project. By putting a portion of the A Company workload into the EE project and continuing their own migration to EE, M Manufacturing Company could actually reengineer their existing 3745 and 3746 communication controller infrastructure and migrate the additional A Company workload. Their EE project utilizes the Open Systems Adapter (OSA) as the main access to the mainframe environment.
- ▶ Continued SNI presence. The combined SNI requirements had been decreasing over the years of both companies. However, there were still approximately 20 SNI connections to be maintained. Their decision was not to attempt a technology shift away from SNI technology at this time. While they

saw the benefits of the eventual migration from SNI to EE, many of their SNI partners were not yet capable of doing any type of APPN connection to M Manufacturing Company. The communication controllers would still be maintained to support these requirements.

- ▶ ESCON still the preferred channel technology for channel-attached devices. As it became obvious that the communication controller environment could be reduced, there was discussion of retaining only a base 3745 for the SNI and EP traffic. However, parallel channels were not available for anything other than EP. M Manufacturing Company needed to keep their ESCON presence by retaining their 3746-900 technology.
- ▶ EP needed in the short term. M Manufacturing Company had successfully migrated from EP over the past 18 months; however, A Company had not been quite as successful. Based on the consolidation workload it was decided that EP would remain as a short-term requirement. M Manufacturing Company could migrate the EP resources to PC-based alternatives, but chose to wait for one year until the personnel could be made available for the change.

2.10.1 Technology results

M Manufacturing Company began evaluating the above factors against the costs of alternative technology. They chose to optimize their environment and add the A Company workload in the following manner:

- ▶ Consolidating from two communication controllers to one running one active NCP with a standby (twin-standby).
- ▶ Focusing on their in-progress Enterprise Extender project.
- ▶ Reengineering the new environment to handle the SNI and EP resources only.
- ▶ Inactivating and de-installing A Company controllers (two 3725s and one 3745-170).

Their new environment includes one 3745-61A communication controller with an attached 3746-900 frame.

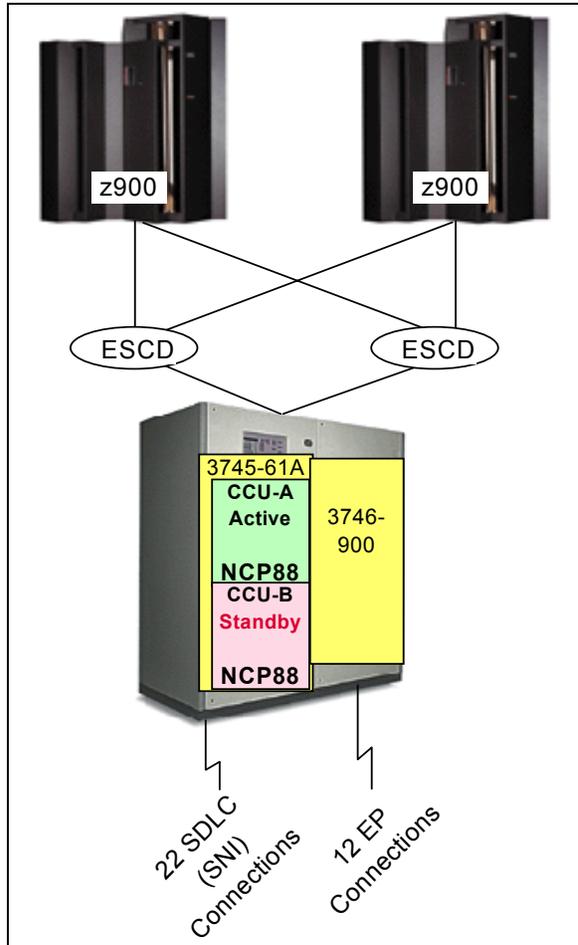


Figure 2-58 Merged M Manufacturing Company and A Company after optimization

The 3746-900 allowed M Manufacturing Company to maintain ESCON channel connectivity for all their SNI traffic. By migrating all but the SNI and EP workload away from the controllers they also could consolidate into one NCP, allowing critical personnel to support the work effort on their EE project.

All resources were to remain on the 3746-900 frame with the exception of EP. EP requires older generation hardware including Bus and Tag channels. These resources were physically placed on the base 3745-61A communication controller.

For backup, M Manufacturing Company chose to implement a 3745-61A communication controller. This allowed them to run their active NCP on one

logical side of the controller (CCU-A) and have a standby NCP on the other logical side of the controller (CCU-B).

M Manufacturing Company already owned two 3745-31A communication controllers, each with a corresponding 3746-900 frame. They chose to keep one of these communication controllers and perform the following physical upgrades:

- ▶ Upgrade the existing 3745-31A to a 3745-61A.
- ▶ Add a Bus and Tag channel and LIC type 1s to the 3745 base for EP.
- ▶ De-install the extra token-ring and line interfaces from the 3746-900 that were now no longer needed.

The physical and logical and functional inventory sheets for the new environment are illustrated below:

3745 Physical Inventory

Part 1: Machine Overview

3745 Model 61A Serial Number 00110

Customer Designation (Ex. NCP Name) NCP88

Indicate all attached 3746 expansion frames.

3746-A11 SN _____

3746-A12 SN _____

3746-L13 SN _____

3746-L14 SN _____

3746-L15 SN _____

3746-900 SN 91234

Controller Expansion (Rack) Qty 1 (0,1 or 2)

(Please inventory each attached frame separately on 3746 Model Sheets.)

How much memory (per CCU if Model 41x or 61x)?

4M _____

8M _____

16M x

Figure 2-59 M Manufacturing Company 3745 Model 61A machine overview after optimization

3745 Physical Inventory

Part 1: Machine Overview

3745 Model 61A Serial Number 00110

Console Information

3151 or Equivalent

Service Processor **x**

Type:

9577

9585

3172 P/N 41H7520

3172 P/N 55H7630

7585

6275 **x**

6563

Indicate serial numbers of other 37XXs that share this console.

Figure 2-60 M Manufacturing Company 3745 Model 61A console information after optimization

3745 Physical Inventory

Part 2: Base 3745 Configuration

Section A: 3745 Models x1x

3745 Model **61A** Serial Number **00110**

Channel Board

Bus Group 1	Pos. 1 <u> CADS </u>	Pos. 2 <u> </u>	Pos. 3 <u> </u>	Pos. 4 <u> </u>
Bus Group 2	Pos. 5 <u> </u>	Pos. 6 <u> </u>	Pos. 7 <u> </u>	Pos. 8 <u> </u>

Legend:

CADS = Channel Adapter Data Streaming
 BCCA = Buffer Chaining Channel Adapter
 TPS = Two Processor Switch

Figure 2-61 M Manufacturing Company 3745 Model 61A channel board after optimization

3745 Physical Inventory

Part 2: Base 3745 Configuration

Section A: 3745 Models x1x

3745 Model 61A Serial Number 00110

TSS Board

Pos. 12 Adap: <u>LSS</u>		Pos. 11 Adap: _____		Pos. 10 Adap: _____		Pos. 9 Adap: _____		Pos. 4 Adap: _____		Pos. 3 Adap: _____	
Port #1 61A	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____

Legend:

Adapters

HSS = High Speed Scanner
 LSS = Low Speed Scanner
 ELA = Ethernet/IEEE 802.3 Adapter

Port Cards

LAN = Ethernet LAN Attached
 #x fff = Where 'x' is the Line Unit Area
 driven by the installed LSS and 'fff' is
 the model number of frame where the
 Line Unit Area resides
 (Example: #2 L13 - would indicate
 Line Unit Area 2 in the 3746-L13)
 (Example: #3 210 - would indicate Line
 Unit Area 3 in the 3745-210)

Figure 2-62 M Manufacturing Company 3745 Model 61A TSS board after optimization

3745 Physical Inventory

Part 2: Base 3745 Configuration

Section A: 3745 Models x1x

3745 Model 61A Serial Number 00110

Line Unit

Area 1				Area 2			
LIC Type							
<u>1</u>	<u>1</u>	<u>1</u>	_____	_____	_____	_____	_____

Additional Line Units

Area 3				Area 4			
LIC Type							
_____	_____	_____	_____	_____	_____	_____	_____

Figure 2-63 M Manufacturing Company 3745 Model 61A line unit after optimization

3745 Physical Inventory

Part 5: 3746-900 Configuration

3746 Model 900 Serial Number 91234

Overview

Extended Microcode Options:

Extended Functions 1 (5800) _____
Extended Functions 2 (5802) _____
Extended Functions 3 (5801) _____
Extended Functions 4 (5810) x
Extended Functions 5 (5812) x
Extended Functions 6 (5813) _____
X.25 (5030) _____
IP (5033) _____

Multiaccess Enclosure (MAE) present?

Yes _____ No x

MAE Microcode Options:

Extended Functions 1 (5804) _____
Extended Functions 2 (5805) _____
Extended Functions 3 (5807) _____
TN3270 Server (5806) _____

Network Node Processor (NNP)

Qty 0 (0,1 or 2)

NNP Type:

Type 1 (3172) _____
Type 2 (7585) _____
Type 3 (6275) _____
Type 4 (6563) _____
Type 5 (6578) _____

Figure 2-64 M Manufacturing Company 3746 Model 900 machine overview after optimization

3745 Physical Inventory

Part 5: 3746-900 Configuration

3746 Model 900 Serial Number 91234

Base Enclosure Top View

P	N	M	L	K	J	H	G	F	E	D	C	Rear Side
Port	Port	Port	Port	Port	Port	Port	Port_	Port	Port	Port	Port	
_____	_____	_____	<u>LIC111</u>	_____	<u>ESCC2</u>	_____	<u>CBC</u>	TIC3	CBC	N/A	N/A	
Processor		Processor		Processor		Processor		Processor CBSP		Processor Power Supply		
_____		<u>CLP3</u>		<u>ESCP3</u>		<u>TRP3</u>		Type <u>3</u>				
Slot 6 (P)		Slot 5 (M)		Slot 4 (K)		Slot 3 (H)		Slot 2 (F)		Slot 1 (D)		Front Side

- Notes: 1. Indicate CBSP Type in Slot 2 (F)
 2. For 3745 Models 41A and 61A, Slot 3 (H) must be a TRPx and Port G must be a CBC

Processors:

TRP = Token-Ring Processor Type 1
 TRP2 = Token-Ring Processor Type 2
 TRP3 = Token-Ring Processor Type 3
 ESCP = ESCON Processor Type 1
 ESCP2 = ESCON Processor Type 2
 ESCP3 = ESCON Processor Type 3
 CLP = Communication Line Processor
 CLP3 = Communication Line Processor Type 3
 SIE = Switch Interface Extension (MAE connection)

Ports:

TIC3 = Token-Ring Coupler Type 3
 ETH = Ethernet Port/Ethernet-TR Bridge
 ESCC = ESCON Coupler Type 1
 ESCC2 = ESCON Coupler Type 2
 LIC11x = Line Interface Coupler Type 11 where 'x' is the Line Connection Box (LCB) ID where the ARCS are installed that correspond to this LIC
 LIC12 = Line Interface Coupler Type 12

Figure 2-65 M Manufacturing Company 3746 Model 900 base enclosure top view after optimization

3745 Physical Inventory

Part 5: 3746-900 Configuration

3746 Model 900 Serial Number 91234

LCB ID 1 Line Connection Box Layout

	0	+1	+2	+3	+4	+5	+6	+7	+8	+9	+10	+11	+12	+13	+14	
ARC Type	V.35				V.35	V.35	V.35	V.35	V.24	V.24	V.24	V.24	V.24	V.24		To LCBE
Attach	DCE				DCE											
Length	15				15	15	15	15	12	12	12	12	12	12		

Line Connection Box Expansion Layout

	+16	+17	+18	+19	+20	+21	+22	+23	+24	+25	+26	+27	+28	+29	+30	
ARC Type	V.35				V.35	V.35	V.35	V.35	V.24	V.24	V.24	V.24	V.24	V.24		To LCB
Attach	DCE				DCE											
Length	15				15	15	15	15	12	12	12	12	12	12		

Length:

- .6 = .6 Meters
- 1.2 = 1.2 Meters
- 2.4 = 2.4 Meters
- 5 = 5 Meters
- 10 = 10 Meters
- 12 = 12 Meters
- 15 = 15 Meters
- ST = Stub cable to 3745 cable

ARC Types:

- V.35
- V.24
- X.21

Attach:

- DCE
- DTE

Notes:

Fill out one sheet for each LCB/LCBE combo installed

Note: A 3746-900 may have as many as 32 LCB/LCBE pairs. Fill out one of these sheets for each one.

Figure 2-66 M Manufacturing Company 3746 Model 900 line connection box ID1 after optimization

Part 2: NCP Owned Resource Section Lines

Serial Lines: Indicate Serial Line Groups.

Protocol	Speed	Line Count	SNI Count	Autocall Count
Example: SDLC	56K	5	3	0
Example: EP	9.6	12	0	6
SDLC	19,200	12	12	0
SDLC	56,000	8	8	0
SDLC	256,000	2	2	0
EP	14,400	12	0	0

Legends:	
Protocol	The line group protocol (SDLC, BSC3270, EP, Frame Relay, X.25, etc.).
Speed	The speed of the line group.
Line Count	The number of lines of a certain speed and protocol.
SNI Count	The count of any SNI lines within the group.
Autocall Count	The count of any Autocall lines within the group.

Figure 2-68 M Manufacturing Company NCP owned resource section for NCP88 after optimization

Part 2: NCP Owned Resource Section

Token-Ring

Downstream PU Count (DSPU)	Logical Unit Count (LUDRPOOL)	TICs In Use
<i>Example: 25,000</i>	<i>50,000</i>	<i>4</i>
0	0	0
Note: Counts in this table are for all Token-Ring resources within this NCP.		

Are you currently using duplicate TIC addresses to balance and back up your Token-Ring SNA traffic?

Yes _____

No _____

Ethernet LAN (NCP Owned Resources)

What type of traffic is on your Ethernet LAN connection?

SNA _____

IP _____

If you are running IP, how many routes are you supporting? _____

Figure 2-69 M Manufacturing Company token-ring and Ethernet LAN for NCP88 after optimization

Part 2: NCP Owned Resource Section Channels

Channel Type	LPAR Count
Bus and Tag	1
ESCON	8
<p>Note: If you are currently not using ESCON, is your mainframe ESCON capable?</p> <p>Yes _____</p> <p>No _____</p>	

Figure 2-70 M Manufacturing Company channels for NCP88 after optimization

2.10.2 Business results

M Manufacturing Company realized the following reduction in their operating expense from this project:

- ▶ Floor space reduction of 50%.
- ▶ Maintenance savings of 65% (\$26,500 per year).
- ▶ NCP Software savings of 80% (\$54,888 per year).
- ▶ Proceeds from the sale of the de-installed equipment completely paid for the upgrades to the remaining 3745 and 3746-900.

2.10.3 The next step

Since M Manufacturing Company has seen significant results from their optimization exercise, they now focused on their strategic plan.

2.11 Task 4 - M Manufacturing Company controller strategic planning

At the end of the optimization project, M Manufacturing Company had two major functional areas that the controller supported:

- ▶ SNI Gateway
- ▶ EP

Of these functional areas, M Manufacturing Company planned to migrate the EP resources to their existing PC solution within 18 months. This migration would leave only the SNI functions on the controller.

2.12 Task 5 - M Manufacturing Company functional alternative migration

As with many clients, M Manufacturing Company could not foresee a plan at this time to migrate completely away from SNI. Some of their SNI partners were, as they were, moving to EE and these would be migrated to the Border Node function. However, M Manufacturing Company determined they would need to maintain their communication controller environment for the few partners who insisted on maintaining their SNA Interconnection. As a long-term strategy, M Manufacturing Company decided to reevaluate the SNI situation after three years. See Chapter 3, “Hot topics - SNI, EP, BSC, and X.25” on page 89 for more information on SNI and EP.



Hot topics - SNI, EP, BSC, and X.25

Although most organizations use just a small subset of the numerous functions supported by IBM communication controllers, each functional capability is probably being used by someone; consequently, this book is intended to be comprehensive. Based on our experience, however, four functional areas will be the most important to the majority of readers. They are:

- ▶ SNA Network Interconnection (SNI)
- ▶ Emulation Program (EP)
- ▶ Binary synchronous communication (BSC)
- ▶ X.25 communications

This chapter is intended to help you to find your way to these important discussions in the book.

3.1 SNA Network Interconnection (SNI)

The SNI function is one of many important capabilities supported by the Network Control Program (NCP) program product. SNI enables controlled interconnection of autonomous SNA networks for many important business-to-business communication activities. Functional alternatives for SNI are discussed in 5.3.4, “SNA Network Interconnection (SNI)” on page 167.

3.2 Emulation Program (EP)

Although EP was created long ago in order to support migration from controller environments that existed even before IBM introduced SNA communication controllers, many organizations still depend on EP support. Functional alternatives for EP are discussed in Chapter 9, “Emulation Program (EP)” on page 211.

3.3 Binary synchronous communication (BSC)

The IBM communication controllers support many different kinds of BSC communication:

- ▶ BSC 3270 connections to SNA applications are supported by NCP and are discussed in 5.2.2, “BSC 3270 terminal connection to SNA applications” on page 146.
- ▶ BSC 3270 connections to non-SNA applications (such as CICS® using BTAM) are supported by EP and are discussed in 9.2.1, “BSC 3270 terminal connection to non-SNA applications” on page 214.
- ▶ BSC RJE connections to SNA applications are supported by the Network Terminal Option (NTO) program product and are discussed in 10.2.2, “BSC RJE connection to SNA host applications” on page 229.
- ▶ BSC RJE connections to non-SNA applications (such as JES using RTAM) are supported by EP and are discussed in 9.2.2, “BSC RJE connection to non-SNA applications” on page 217.
- ▶ Interconnections between peer non-host-based BSC RJE devices are supported by NTO working in conjunction with the Network Routing Facility (NRF) program product and are discussed in 11.2.2, “Peer-to-peer connections involving non-SNA devices” on page 245.

3.4 X.25 communications

The IBM communication controllers support many different kinds of X.25 communication:

- ▶ X.25 LLC type 3, QLLC, connections can be supported either by NCP with the 3746-900 (with special X.25 licensed support feature microcode) or via the X.25 NCP Packet Switching Interface (NPSI) program product. The alternatives for these functional capabilities are discussed in 5.2.8, “X.25 (with licensed support feature on 3746)” on page 152 and in 6.2.1, “Supporting SNA (PSH and QLLC) communication over X.25” on page 190 respectively.
- ▶ A very wide variety of non-SNA X.25-related capabilities including PCNE, DATE, GATE, and PAD support are provided by NPSI and are discussed in 6.2.2, “Supporting non-SNA communication over X.25” on page 196.
- ▶ Alternatives for X.25 SNA Interconnection (XI) ability to transport X.25 traffic across your SNA network are discussed in Chapter 8, “X.25 SNA Interconnection (XI) and Network Supervisory Function (NSF)” on page 207.
- ▶ X.25 connectivity to the S.W.I.F.T. financial messaging network is discussed in Chapter 7, “MERVA Extended Connectivity” on page 203.



Part 2

Functional alternatives reference

This part of the book provides a review of the capabilities of each of the IBM communication controller products, exploration of the alternatives for replacing or eliminating the need for each, and consideration of the relative advantages and disadvantages of the various alternatives.

There is a large, integrated set of IBM products which together form the IBM communication controllers. Each of those products is covered in the chapters of this part of the book. They are organized as follows:

- ▶ Communication controller hardware:
 - 3745 Communication Controller, 3746 Expansion Frame, and 3746 Nways Multiprotocol Controller Model 900 Network Node and Model 950
- ▶ Network Control Program (NCP):
 - Network Control Program (NCP), 5648-063
- ▶ X.25-related software products:
 - X.25 NCP Packet Switching Interface (NPSI), 5688-035
 - X.25 SNA Interconnection (XI) and Network Supervisory Function (NSF), 5685-035
- ▶ Other controller software products:
 - Emulation Program (EP), 5735-XXB
 - Network Terminal Option (NTO), 5735-XX7
 - Network Routing Facility (NRF), 5668-963
 - Non-SNA Interconnection (NSI), 5668-951
 - MERVA Extended Connectivity, 5655-110
 - Teleprocessing Network Simulator (TPNS), 5688-121.

While this book focuses on the newer communication controller hardware platforms, the 3745 Communication Controller and the 3746 Nways Multiprotocol Controller, most of the software function discussions apply likewise to the older controller hardware (3705, 3725, and 3720) environments.



Note: Three of the above program products, NCP, NPSI, and NRF, can now be run on your mainframe using the Communication Controller for Linux on System z9 and zSeries (CCL) product, 5724-J38. In many cases, the CCL option will provide the most transparent migration from your IBM 3745 communication controllers. The CCL option is discussed in greater detail, where applicable, in the chapters that follow.

See Chapter 16, “Communication Controller for Linux on System z9 and zSeries (CCL)” on page 283 for additional information about the CCL.

Important migration-strategy alternatives

Each of the chapters in this part of the book presents our best ideas for technical solutions for the product functions discussed. However, there are some important migration strategy alternatives that you should also consider. They include:

- ▶ Leaving things as they are (at least for now)

Continuing to run an IT service without change can save your organization the cost and potential disruptions implicit in any migration. One situation in which such a strategy may make sense is when the use of the IT service in question is expected to decline substantially over time (perhaps due to business process reengineering efforts). This strategy, however, has the risk that at some point in time the hardware and software foundation of your configuration will become unsupported. Running on an unsupported configuration means that recovery from an outage may be difficult or impossible. You should make sure that your use of the configuration will end before support for its underlying products is terminated, or that loss of the service will not severely affect your business operation.

- ▶ Eliminating an IT service

Through the years, many IT organizations have evolved to have multiple, somewhat redundant, means of delivering a specific service (for example, file transfer). Eliminating an IT service through consolidation of redundant services not only reduces IT costs but may also permit investment that will enhance the quality of strategic services. In some cases, usage of a particular IT service has declined to the point where continuance of that service cannot be cost justified. Such re-evaluation of the IT services “portfolio” should be an ongoing part of your organization’s overall IT strategy.

- ▶ Replacing an IT service with a completely new solution

While you may be using this book in an effort to optimize your organization’s use of IBM communication controllers, there are probably other, strategically focused, efforts underway in your organization that will affect the same IT services. While the option of leaving things as they are, above, suggests *waiting* for the more strategic projects to complete, it might also be appropriate to try and *expedite* such projects by adding to their business case the potential savings from the optimization and consolidation of your communication controller environment.

Functional alternatives summary

Table 3-1 summarizes the communication controller functions and alternatives discussed in this book.

Table 3-1 Functional alternatives summary

Function	Alternatives
Controller physical interfaces	
4.2.1 Communication lines	<ul style="list-style-type: none"> ▶ Migrating devices to LAN attachment ▶ Migrating cluster controllers to LAN attachment ▶ Moving lines to routers
4.2.2 Token-ring (and duplicate addressing)	<ul style="list-style-type: none"> ▶ Replacing controller token-ring gateway <ul style="list-style-type: none"> – OSA (to VTAM or to CCL) – Channel-attached server – Channel-attached router ▶ Migrating everything to Ethernet
4.2.3 Ethernet, FDDI, and ATM	<ul style="list-style-type: none"> ▶ OSA ▶ Channel-attached server ▶ Channel-attached router
4.2.4 Channel attachments (host-to-host)	<ul style="list-style-type: none"> ▶ Parallel channel ▶ Enterprise Systems Connectivity (ESCON) ▶ Fibre Channel (FICON®) ▶ Channel-to-Channel (CTC) ▶ Open Systems Adapter (OSA) ▶ Cross Coupling Facility (XCF)
Network Node Processor functions	
4.3.1 APPN network node (including HPR and DLUR)	<ul style="list-style-type: none"> ▶ Replacing the APPN functions either with routers or with servers ▶ Migrating DLUR clients to CCL NCP ▶ Eliminating the APPN functions <ul style="list-style-type: none"> – Using LAN interfaces on the hosts and supporting SNA boundary function in VTAM – Migrating from the devices that depend upon the SNA support using technologies such as TN3270
4.3.2 IP routing	OSA for host access and routers or L3 switches otherwise

Function	Alternatives
Multi-Access Enclosure functions	
4.4.1 TN3270 server	<ul style="list-style-type: none"> ▶ Migrating TN3270 server to your mainframe host ▶ Replacing with another outboard solution
4.4.2 Network Dispatcher	<ul style="list-style-type: none"> ▶ Replacing the MAE platform with a server ▶ Leveraging Sysplex Distributor ▶ Implementing another outboard load-balancing solution, preferably one that supports Server/Application State Protocol (SASP)
NCP Link-level protocol functions	
5.2.1 Airlines Line Control (ALC)	<ul style="list-style-type: none"> ▶ Mapping of Airline Traffic Over TCP/IP (MATIP) ▶ Message Queuing (MQ) with MQSeries® for TPF
5.2.2 BSC 3270 terminal connection to SNA applications	<ul style="list-style-type: none"> ▶ Replace the BSC 3270 devices ▶ Use protocol conversion devices
5.2.3 Frame relay	<ul style="list-style-type: none"> ▶ Frame switches and routers ▶ IP-based infrastructure
5.2.4 Integrated Services Digital Network (ISDN)	Routers with ISDN interfaces
5.2.5 LAN 802.2, 802.3, and 802.5	Covered in “Controller physical interfaces” on page 96
5.2.6 Start-Stop terminal connection to TCAM applications	<ul style="list-style-type: none"> ▶ Replace the Start-Stop devices ▶ Use protocol conversion devices
5.2.7 Synchronous Data Link Control (SDLC)	Covered in “Controller physical interfaces” on page 96
5.2.8 X.25 SNA QLLC (with licensed support feature on 3746)	<ul style="list-style-type: none"> ▶ Migrating to router X.25 support ▶ Using routers to migrate from X.25 ▶ Integrated Communication Adapter (ICA)

Function	Alternatives
NCP Advanced capabilities	
5.3.1 SNA class of service (COS)	<ul style="list-style-type: none"> ▶ Implement support for SNA COS-based prioritization (using APPN) ▶ Increase available bandwidth on network links so that the total traffic rarely exceeds the available bandwidth
5.3.2 Multi-link transmission group (MLTG)	<ul style="list-style-type: none"> ▶ Continue to run NCP using CCL ▶ Migrating to high-speed, high-availability, IP-based infrastructure
5.3.3 SNA subarea addressing, routing, and boundary function (BF)	<ul style="list-style-type: none"> ▶ Continue to run NCP using CCL ▶ Migrating users to TN3270 ▶ Implement DLUR ▶ Implement Enterprise Extender (for SNA routing)
5.3.4 SNA Network Interconnection (SNI)	<ul style="list-style-type: none"> ▶ Continue to run NCP using CCL ▶ Converting to SNI adjacent network connectivity ▶ Migrating to APPN Border Node (BN)
5.3.5 APPN (and LEN) Composite Network Node (CNN)	<ul style="list-style-type: none"> ▶ Using APPN or DLSw function in your routers to move the application traffic between APPN nodes in your network ▶ Using EE to allow your APPN host application traffic to use your IP infrastructure
5.3.6 IP routing	OSA for host access and routers or L3 switches otherwise
5.3.7 Extended Recovery Facility (XRF)	<ul style="list-style-type: none"> ▶ Continue to run NCP using CCL ▶ Migrate from XRF using generic resources
NPSI	
6.2.1 Supporting SNA (PSH and QLLC) communication over X.25	<ul style="list-style-type: none"> ▶ Migrating to router X.25 support ▶ Using routers to migrate from X.25 ▶ Integrated Communication Adapter (ICA)
6.2.2 Supporting non-SNA communication over X.25	<ul style="list-style-type: none"> ▶ Running NCP and NPSI using CCL V1.2 ▶ Host-based X.25 over TCP/IP (XOT)

Function	Alternatives
XI and NSF	
7.2.1 Transporting X.25 traffic	<ul style="list-style-type: none"> ▶ Migrating end devices from X.25 ▶ Migrating to an X.25 service ▶ Migrating to router transport of X.25
7.2.2 NSF-based charge back for X.25 transport service	No alternative
EP	
8.2.1 BSC 3270 terminal connection to non-SNA applications	<ul style="list-style-type: none"> ▶ Preserving BSC 3270 using an ICA or Hydra 3000 ▶ Preserving BTAM applications using VM console ▶ Preserving BSC 3270 devices using protocol conversion
8.2.2 BSC RJE connection to non-SNA applications	<ul style="list-style-type: none"> ▶ BSC RJE to SNA RJE protocol conversion ▶ Preserving BSC RJE connections using an ICA or Hydra 3000 ▶ Migrating from RJE to message queuing (MQ)
8.2.3 Start-stop terminal connection to non-SNA applications	No alternative short of changing the application, terminals, or both
NTO	
9.2.1 Start-stop terminal connection to SNA host applications	<ul style="list-style-type: none"> ▶ Protocol conversion ▶ Using X.25-specific protocol conversion functions
9.2.2 BSC RJE connection to SNA host applications	<ul style="list-style-type: none"> ▶ BSC RJE to SNA RJE protocol conversion ▶ Migrating to SNA RJE ▶ Preserving BSC RJE connections using an ICA or Hydra 3000 ▶ Migrating from RJE to message queuing (MQ)
9.2.3 Peer-to-peer connection of non-SNA devices	Discussed under “NRF” on page 100

Function	Alternatives
NRF	
10.2.1 Peer-to-peer connection of SNA devices (before PU type 2.1)	<ul style="list-style-type: none"> ▶ Running NCP and NRF using CCL ▶ Host-based SNA message routing ▶ Using a server for message routing ▶ Message Queuing (MQ)
10.2.2 Peer-to-peer connections involving non-SNA devices	<ul style="list-style-type: none"> ▶ Router tunneling ▶ NRF replacement with protocol conversion <ul style="list-style-type: none"> – Using an SNA-based message routing program – Using an MQ-compatible client format (to use an MQ-based solution) ▶ Using a protocol conversion platform for message routing
NSI	
11.2.1 Non-SNA NJE to NJE connections between hosts	<ul style="list-style-type: none"> ▶ Host-based SNA message routing ▶ Using a server for message routing ▶ Message Queuing (MQ)
11.2.2 Transporting BSC traffic across your network	Router tunneling in conjunction with ICA or Hydra 3000
MERVA	
12.2.1 Connecting to the S.W.I.F.T. network	<ul style="list-style-type: none"> ▶ MECO/PC or other PC-based alternatives ▶ Migrating to SWIFTNet
Teleprocessing Network Simulator (TPNS)	
13.2.1 Simulating subarea traffic	No alternative



Communication controller hardware

This chapter discusses alternatives to supporting the devices (such as controllers, terminals, and printers) that connect to physical interfaces on your communication controllers. In general, the alternatives include finding different ways to support those interfaces (such as replacing 3745 token-ring with OSA token-ring), or migrating devices to new forms of attachment (such as migrating from SDLC to LAN attachment).

This chapter focuses on the following categories of communication controller physical interfaces:

- ▶ Communication lines
- ▶ Token-ring (particularly its unique duplicate addressing support)
- ▶ Ethernet, FDDI, and ATM
- ▶ Channels

Note: In addition to the technical solutions discussed in this chapter, you should consider the alternatives of leaving things as they are, eliminating the IT service in question, or replacing that service with a completely new solution. See “Important migration-strategy alternatives” on page 95 for a discussion of each of those alternatives.

4.1 Communication controller hardware overview

IBM developed communication controllers to serve two important networking needs:

- ▶ As gateways for connecting mainframe host servers to networks offloading network functions such as line control and device support, thereby saving host CPU cycles
- ▶ As line concentration devices that could be placed in remote locations to save on communications line charges by consolidating the traffic from many remote devices into a few, usually higher speed, communication lines



Figure 4-1 IBM communication controllers

The 3745 Communication Controllers such as the one shown on the left in Figure 4-1 support:

- ▶ Low and high-speed line attachments
- ▶ Token-ring LAN attachments at 4 Mbps or 16 Mbps
- ▶ Ethernet Version 2 or IEEE 802.3 LAN interfaces
- ▶ Parallel channel host connections

The 3745 Models 210/A, 310/A, 410/A, and 610/A are compatible with the 3746 Models Axx/Lxx Expansion Units, offering modular growth for up to 16 parallel host connections and 896 line attachments.

The 3746 Model 950, such as the one shown on the right in Figure 4-1, along with its Multi-Access Enclosure add support for:

- ▶ Additional low and high-speed lines
- ▶ Additional token-ring interfaces

- ▶ ATM
- ▶ ISDN PRI
- ▶ HSSI (T3/E3 speeds)
- ▶ FDDI
- ▶ 10/100-Mbps Ethernet
- ▶ Parallel channel
- ▶ ESCON channel

Figure 4-2 shows how a communication controller, built out with all of its expansion frames, would look.

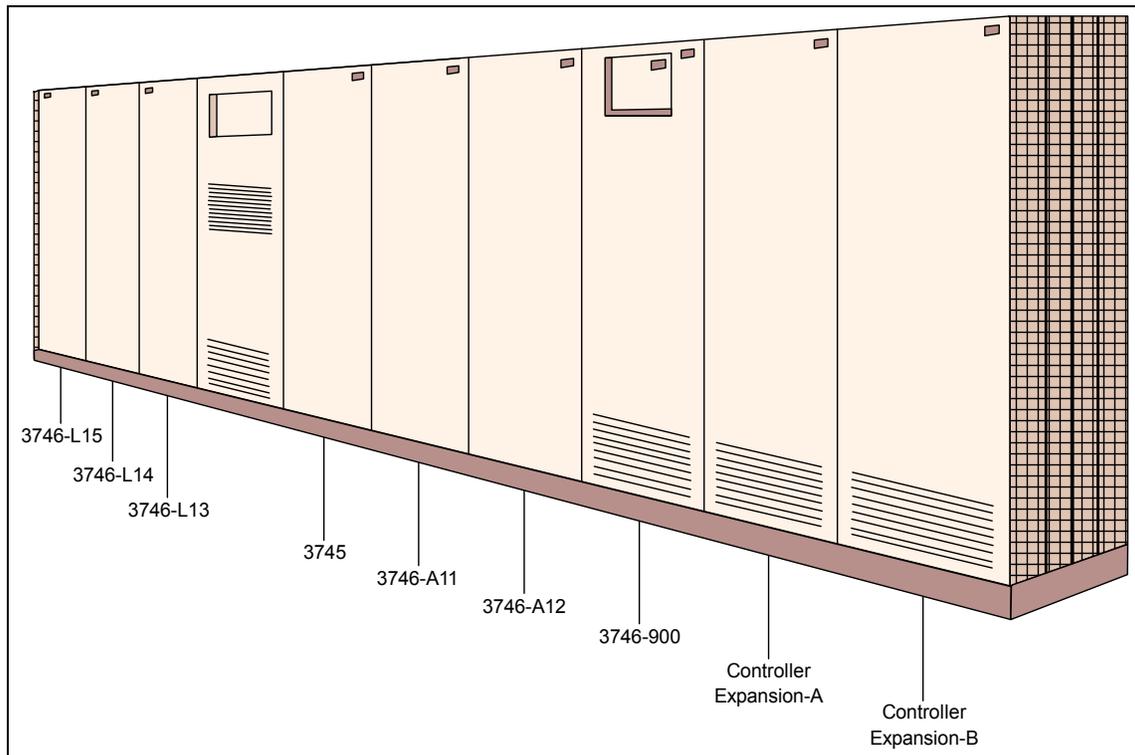


Figure 4-2 3745 with all expansion frames

Other than physical interfaces, most of the functional capabilities of the communication controller are provided by the software that runs on it—for example, the Network Control Program—and are discussed in subsequent chapters. The two exceptions to this are:

- ▶ The 3746 Nways Multiprotocol Controller Model 900 and 950 network node processor, discussed in 4.3, “Network Node Processor functions and alternatives” on page 131

- ▶ The Multi-Access Enclosure (MAE), discussed in 4.4, “Multi-Access Enclosure functions and alternatives” on page 138

4.2 Physical interface functions and alternatives

This section covers the migration alternatives for the physical interfaces available on the communication controllers. In terms of functional alternatives, it does not matter which hardware component, the 3745, the 3746, or the MAE, supports a given physical interface; consequently, they are all addressed together in this chapter. The various link-level protocols supported on those interfaces are covered in the chapters for the software products that support them. For example, the Airlines Line Control (ALC) protocol that is supported via NCP software is covered in Chapter 5, “Network Control Program (NCP)” on page 141.

The simple network diagram shown in Figure 4-3 on page 105 illustrates the use of many of the different kinds of physical attachments that are provided by communication controllers. In the environment illustrated, terminals located in the remote offices, the regional office, and the data center all use the 3745 Communication Controller for access to host-based application programs. A communication controller in the regional site provides communication line cost savings by consolidating the traffic from terminals located in the regional office—as well as those in nearby remote offices (Office C and Office D)—into a single line to the data center.

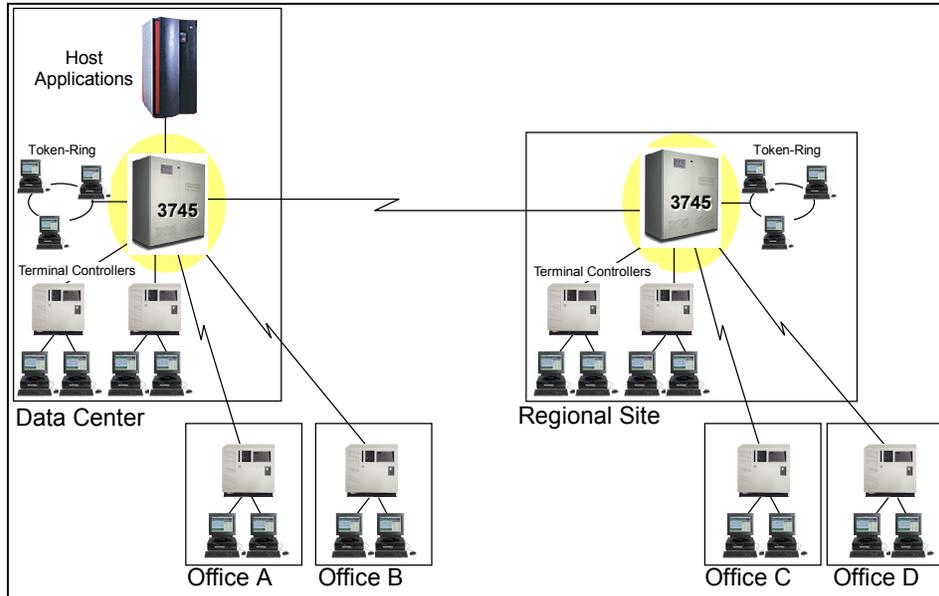


Figure 4-3 A simple network diagram

From looking at the network in Figure 4-3, you might conclude that the communication controllers could simply be replaced by routers. Indeed, routers support all of the different physical interfaces shown—as well as many others—and they provide certain multi-protocol capabilities well beyond those of communication controllers. However, there are a number of issues to consider, regardless of the physical interface support you require, before simply attempting to replace one box with another.

- ▶ *Not all link-level protocols supported by communication controllers are supported by routers.* For example, it may be difficult to find a router that supports the breadth of X.25 protocols (particularly the non-SNA logical link control types) that are supported in communication controllers. Also, support for unique link-level protocols such as start-stop, binary synchronous communication (BSC), or Airline Line Control (ALC) may be limited.
- ▶ *Similarly, there are advanced capabilities supported by communication controllers that may be difficult to find in routers.* For example, SNA routing and boundary function (BF) are communication controller capabilities that can only be replaced by routers that support Advanced Peer-to-Peer Networking® (APPN) and dependent LU requester (DLUR) functions. SNA network interconnection (SNI) is another important communication controller capability that is not available in any router.
- ▶ *Finally, such an infrastructure change should only be made after careful consideration of the strategic infrastructure needs and plans of your*

organization. For example, perhaps there are (or should be) plans in place for migrating the terminal-controller-attached fixed-function terminals in Figure 4-3 on page 105 to LAN-attached intelligent workstations. If so, completing the desktop migration *first* will save the costs of buying router line interfaces—and bigger routers to support them—not needed in the long term. Similarly, the purchase of expensive host channel interfaces for routers may be avoided by migrating to a more strategic direct host attachment to the network (see 4.2.4, “Channel attachments” on page 128).

4.2.1 Communication lines

In a network like the one diagrammed in Figure 4-3 on page 105, low-speed lines are typically used for the connections to end devices such as terminal controllers, while low or high-speed lines are used (depending upon the volume of traffic) for the connections between communication controllers.

Alternatives

The following alternatives offer ways to migrate the communication line connections between communication controllers and end devices (such as terminal controllers):

- ▶ Fixed-function terminals can be migrated to LAN-attached intelligent workstations. Such “desktop migrations” can be difficult and expensive because they often involve hundreds or thousands of devices, but they have the advantage of enhancing the organization’s ability to support strategic new applications such as applications based upon Internet technologies.
- ▶ Devices such as terminal controllers that are connected by communication lines and provide connectivity for large numbers of fixed-function terminals can be migrated to LAN attachment. This type of migration might be easier and less expensive than a desktop migration, although it may not provide the same advantage of positioning the organization to support strategic new applications as would migrating desktops to LAN attachment.
- ▶ The communication lines can be moved from the communication controller onto a router where, using technologies such as SDLC-to-LLC2 conversion, the router can make the devices on those lines appear to the rest of the network to be LAN-attached. As with a terminal controller migration, this alternative will not provide the same advantage of positioning the organization to support strategic new applications as would migrating the desktops.

Each of these alternatives require that an interconnected LAN infrastructure (also called an “*intranet*”) be in place to enable communication with the host systems.



Important: These alternatives address the migration of physical interfaces *only*. Each terminal controller or other end device in an SNA network entails a certain amount of SNA boundary function workload that must be provided somewhere in the network. As devices are migrated from 3745 link attachment to LAN or router attachment, you must carefully consider where the boundary function will be provided for those devices (such as in a host-attached 3745 or in the host itself), as well as the impact of that increased boundary function workload. Alternatives for handling SNA boundary function workload are discussed in 5.3.3, “SNA subarea addressing, routing, and boundary function (BF)” on page 162.



Note: The IBM Communication Controller for Linux on System z9 and zSeries (CCL) product enables you to run NCP on your mainframe, which gives you a way to migrate devices from your IBM 3745 Communication Controller hardware while preserving your NCP boundary function. See Chapter 16, “Communication Controller for Linux on System z9 and zSeries (CCL)” on page 283 for additional information about CCL.

Figure 4-4 shows an example migration scenario of the network shown in Figure 4-3 on page 105.

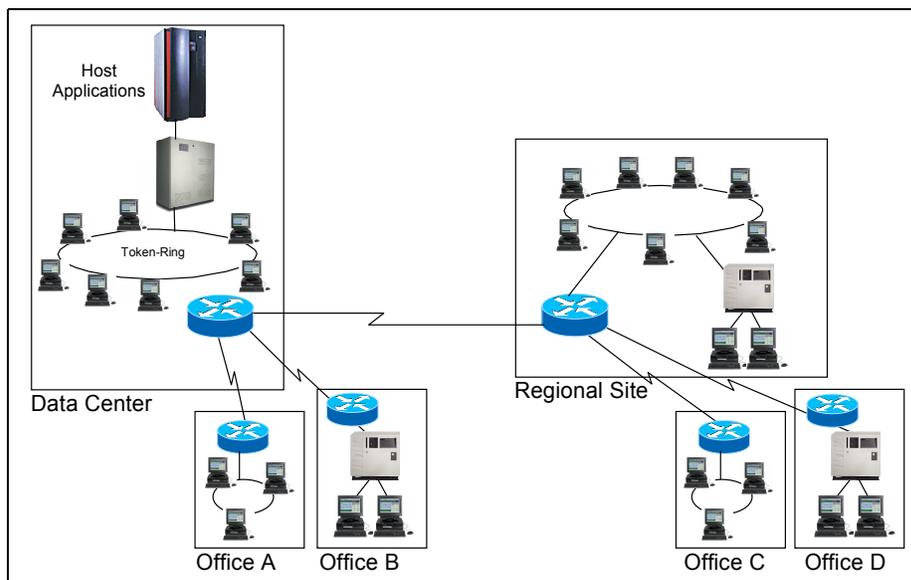


Figure 4-4 An example network migration scenario

In this example, fixed-function terminals in the data center, regional site, and in offices A and C were migrated to LAN-attached intelligent workstations. The regional site also shows an example of LAN-attaching a terminal controller. In offices B and D, the communications lines from terminal controllers were connected into routers, and the terminal controllers were left unchanged. Note also that, because the regional-site communication controller has been replaced with a router, the SNA boundary function formerly provided by that controller will now become additional workload in the controller in the data center.

Note: As you plan to migrate communication lines from your controllers, carefully explore *all* of your wide-area networking connectivity options. While routers may be the logical WAN connectivity devices for your locations, you may find newer service offerings such as frame relay services or virtual private network (VPN) offerings that are more cost-effective than your current network structure. Such service offerings can even include the provision and management of the routers you use to access the network.

Migration considerations

The interconnected LAN infrastructure must be in place first. Most organizations have such networks today in support of their intranet and TCP/IP application traffic. Peripheral devices can be migrated onto the router network via either the desktop migration, the terminal controller migration, or by moving the line interfaces as discussed above. Once all peripheral devices have been moved from a remote 3745, it can be redeployed or removed as appropriate, thereby eliminating the need for the low or high-speed lines used between the communication controllers.

Recommendation

Pursue the migration of end devices (such as the migration of fixed-function terminals to intelligent workstations) first because such migrations will usually offer greater strategic value to your organization. Use routers to connect those devices for which such a migration is impractical.

4.2.2 Token ring (and duplicate addressing)

For more than a decade, token-ring interfaces on communication controllers have provided optimal high-bandwidth connectivity between hosts and LAN-attached devices. Network migrations like the one in the sample network migration illustrated in Figure 4-4 on page 107 have increased the importance of the token-ring interface simply because, in many networks, most of the SNA traffic to and from mainframe hosts now goes through communication controller token-ring interfaces. In order to achieve very high availability for these important

interfaces, many organizations have implemented a duplicate addressing design like the one shown in Figure 4-5.

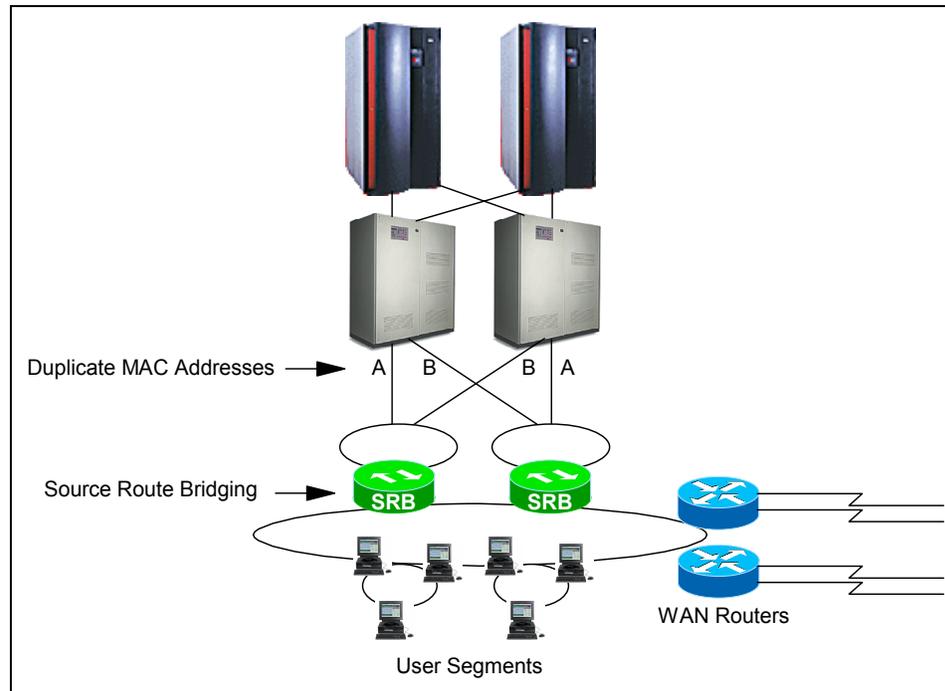


Figure 4-5 Duplicate addressing

Duplicate addressing takes advantage of the fact that token-ring source route bridging, in conjunction with LAN logical link control type 2 (LLC2), can allow duplicate LAN Media Access Control (MAC) addresses to be active on the network at the same time. In Figure 4-5, each of the 3745s has MAC addresses A and B on each of two different LAN segments. This configuration allows the immediate recovery of sessions impacted by the failure of any single bridge, LAN segment, token-ring adapter, 3745, or host channel. Additionally, some degree of load balancing is achieved across the 3745s during LLC2 connection setup. For a more detailed discussion of the duplicate token-ring gateway design, see *New and Improved! IBM Multisegment LAN Design Guidelines*, GG24-3398-02, available on the Web at:

<http://www.redbooks.ibm.com/pubs/pdfs/redbooks/gg243398.pdf>



Important: Communication controllers save host CPU cycles by offloading certain network functions, including SNA device support (also known as the SNA boundary function). Removing communication controllers will consequently result in increased host CPU cycles and storage requirements. The amount of increase will depend on:

- ▶ The number of devices in the network and their transaction rates, and
- ▶ Whether a dependent LU requester (DLUR) is implemented in the network to support SNA devices.

Alternatives for handling SNA boundary function workload are discussed in 5.3.3, “SNA subarea addressing, routing, and boundary function (BF)” on page 162.



Important: Large networks have encountered problems such as SNA session setup failures or the inability to activate additional resources as the result of running out of SNA element addresses in a VTAM subarea. Migrating devices from your communication controllers can add to this problem, because the device addressing can end up moving from the communication controller subarea into a VTAM subarea.

Alternatives for handling the SNA subarea function are discussed in 5.3.3, “SNA subarea addressing, routing, and boundary function (BF)” on page 162.

Replacing token-ring connectivity

Migration from your communication controller token-ring interfaces requires either:

- ▶ Replacing the token-ring connectivity that is currently provided by your communication controller
- ▶ Migrating all of the devices that communicate with your host through your communication controller via the token-ring infrastructure to another form of connectivity — Ethernet being the best choice.



Important: When the first edition of this book was published in early 2002, token-ring technology was nearing end of life yet still viable. Now, however, few (if any) token-ring networking products can be purchased and token-ring environments are becoming increasingly difficult to support. Consequently, rather than replacing token-ring connectivity, consider migrating your remaining LAN connectivity to Ethernet wherever possible.

Ways to migrate host connectivity to Ethernet are discussed in “Migrating from token ring to Ethernet” on page 117.

The alternatives for replacing the token-ring connectivity that is currently provided by your communication controller include:

- ▶ Open Systems Adapter (OSA) token-ring interfaces
- ▶ Channel-attached servers with token-ring interfaces
- ▶ Channel-attached routers with token-ring interfaces

Open Systems Adapter for token ring

Figure 4-6 shows the use of Open Systems Adapter (OSA) token-ring connectivity as an alternative to using the communication controller token-ring interfaces. In this solution, the token-ring connections are moved from the communication controllers onto OSA cards in the hosts. Note, however, that OSA interfaces are supported only on certain IBM mainframe hosts. OSA-2 is supported in G2 and higher, OSA-Express is supported in G5 and higher. For token ring, only OSA-Express adapters can be ordered for zSeries processors.

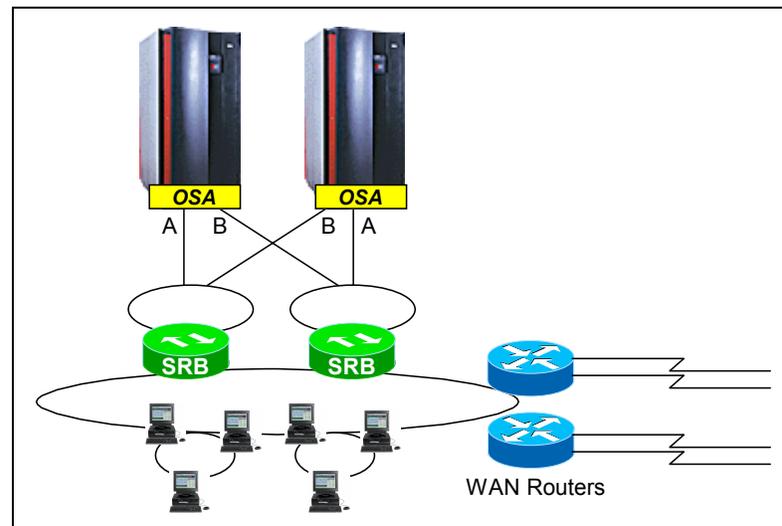


Figure 4-6 OSA token ring

Because the duplicate addressing scheme does not depend upon any specific support in communication controllers, it will work in an OSA environment as well.



Important: If you use OSA token-ring adapters with duplicate addressing for high availability for SNA, it is recommended that you not share the same ports for TCP/IP traffic due to problems that can occur because of the duplicate addresses in IP environments. There is no problem, however, with using one port of an OSA adapter for SNA, with duplicate addressing, and the other port of the OSA for TCP/IP.



Note: The introduction of the Communication Controller for Linux on System z9 and zSeries (CCL) product adds an important new option for the OSA token-ring alternative. Using OSA with CCL for SNA traffic preserves the original NCP subarea, thereby eliminating the need to use additional element addresses out of the VTAM subarea and reducing concerns about running out of VTAM element addresses (mentioned in 5.3.3, “SNA subarea addressing, routing, and boundary function (BF)” on page 162).

See Chapter 16, “Communication Controller for Linux on System z9 and zSeries (CCL)” on page 283 for additional information on CCL.

Migration to the OSA token ring (to either VTAM or to CCL) may be facilitated by configuring the OSA with the same token-ring locally administered address (LAA) that is being used by the communication controller token-ring interface; however, you must attach the OSA interface to a separate token-ring segment from the communication controller. While both the OSA and the communication controller token-ring interfaces are active, some connections will find the OSA while others will use the communication controller. After you are comfortable with the operation of the OSA, you can disable the communication controller token-ring interface and let all connections come through the OSA.

Important: If you stay with token-ring technology, VTAM and CCL can support only one locally administered MAC address (LAA) per OSA port and, consequently, require a separate OSA token-ring port for each communication controller token-ring port. However, if you migrate from token ring to Ethernet OSA, CCL can support multiple MAC addresses on a single OSA port using OSA Layer 2 function (discussed in “Migrating from token ring to Ethernet” on page 117).

Advantages of migrating to OSA token ring include:

- ▶ **Availability:** By attaching mainframe hosts directly to the network using an OSA adapter, you can eliminate intermediate boxes such as

channel-attached routers or servers, which should yield better availability for host access.

- ▶ **Cost:** Depending upon your specific environment, an OSA solution will probably be a less expensive solution than channel-attached routers or servers.
- ▶ **Performance:** The OSA-Express token-ring interface can support a 100 Mbps token ring with certain token-ring switches.
- ▶ **Simplicity:** An OSA solution, when used in conjunction with CCL, can preserve subarea network characteristics.

Considerations of migrating to OSA token ring include:

- ▶ **Scalability:** In addition to the concerns regarding SNA boundary function workload and VTAM subarea element addresses (5.3.3, “SNA subarea addressing, routing, and boundary function (BF)” on page 162), carefully consider the OSA configuration that would be required to meet the needs of your environment. The OSA-Express token ring supports up to 4096 PUs per port.
- ▶ **OSA interfaces are supported only on certain IBM mainframe hosts:** OSA-2 is supported in G2 and higher, OSA Express is supported in G5 and higher. For token ring, only OSA-Express adapters can be ordered for zSeries processors. Token ring is not supported on IBM System z9 servers.
- ▶ **Availability:** Migrating to OSA token ring continues the use and dependence upon end-of-life token-ring technology.

Channel-attached servers

Figure 4-7 on page 114 shows the use of channel-attached servers as token-ring gateways for the host environment. In this solution, the token-ring connections are moved from the communication controllers onto token-ring adapters in servers.

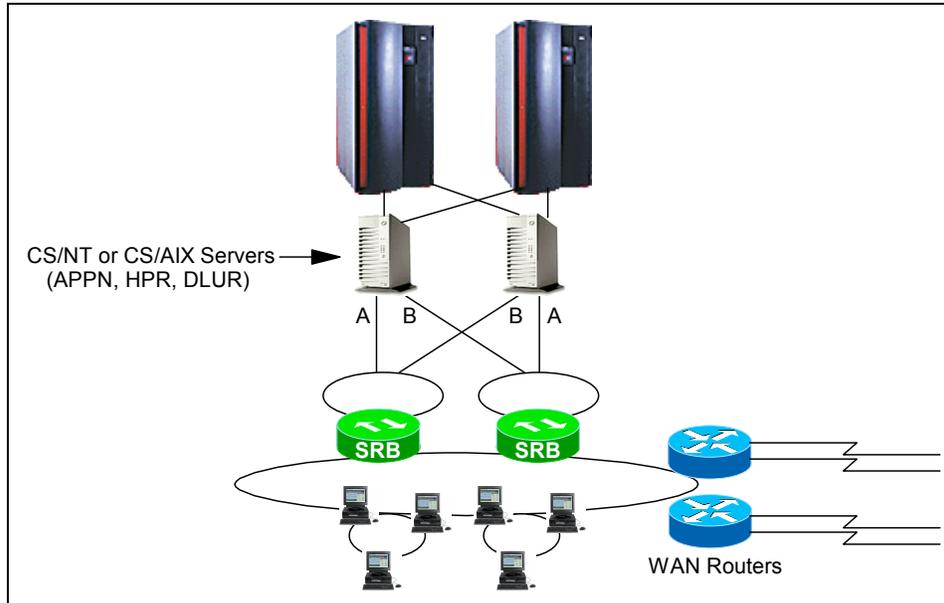


Figure 4-7 Channel-attached servers as token-ring gateways

Because the duplicate addressing scheme does not depend upon any specific support in communication controllers, it will work in a channel-attached server environment as well.



Important: Depending upon the functions running in the servers, server CPU capacity could become a performance problem. Large numbers of users or high transaction rates could drive up the number of required channel-attached servers and, consequently, the cost and complexity of the solution.

Migration to the channel-attached servers as token-ring gateways may be facilitated by configuring the server with the same token-ring locally administered address (LAA) that is being used by the communication controller token-ring interface; however, you must attach the server token-ring interface to a separate token-ring segment from the communication controller. While both the server and the communication controller token-ring interfaces are active, some connections will find the server while others will use the communication controller. After you are comfortable with the operation of the server gateway, you can disable the communication controller token-ring interface and let all connections come through the server.

Advantages of migrating to channel-attached servers as token-ring gateways include:

- ▶ Viable alternative for those environments where the host processor does not support OSA.
- ▶ Well suited to environments where most traffic is to and from intermediate servers that provide key functions such as:
 - APPN, HPR, and DLUR capabilities
 - Web front-end to host applications
 - The application tier of a three-tier client-server solution

Considerations of migrating to channel-attached servers as token-ring gateways include:

- ▶ Scalability: In addition to the concerns regarding SNA boundary function workload and VTAM subarea element addresses (5.3.3, “SNA subarea addressing, routing, and boundary function (BF)” on page 162), server CPU utilization issues could drive up the number of required servers and, therefore, cost and complexity.
- ▶ Availability: Migrating to channel-attached servers as token-ring gateways continues the use and dependence upon end-of-life technologies, including token-ring and ESCON.

Channel-attached routers

Figure 4-8 shows the use of channel-attached routers as token-ring gateways for the host environment. In this solution, the token-ring connections are moved from the communication controllers onto token-ring adapters in routers.

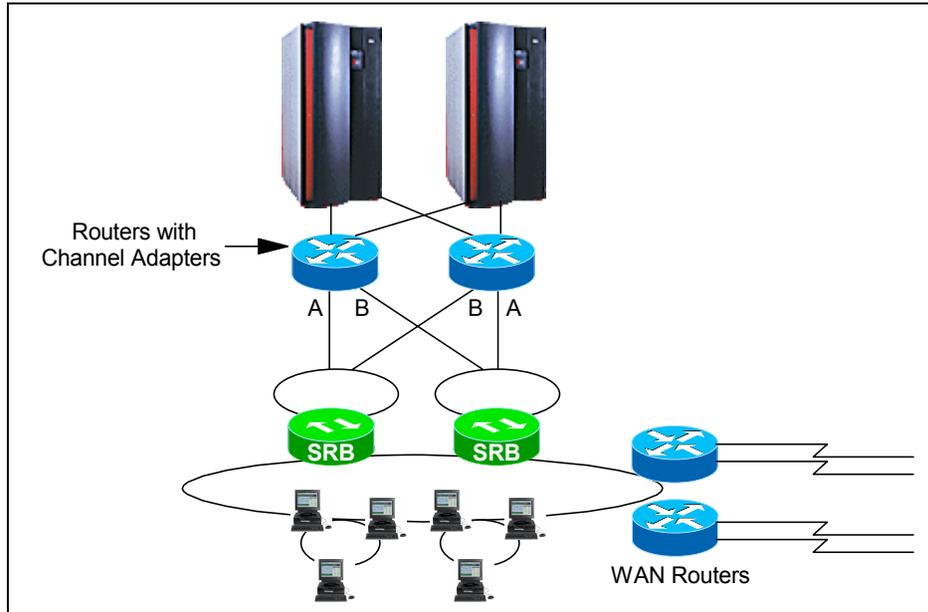


Figure 4-8 Channel-attached routers as token-ring gateways

Because the duplicate addressing scheme does not depend upon any specific support in communication controllers, it will work in a channel-attached router environment as well.

Migration to the channel-attached router as token-ring gateways may be facilitated by configuring the router with the same token-ring locally administered address (LAA) that is being used by the communication controller token-ring interface; however, you must attach the router token-ring interface to a separate token-ring segment from the communication controller. While both the router and the communication controller token-ring interfaces are active, some connections will find the router while others will use the communication controller. After you are comfortable with the operation of the router gateway, you can disable the communication controller token-ring interface and let all connections come through the router.

Advantages of migrating to channel-attached routers as token-ring gateways include:

- ▶ Offer a viable alternative for those environments where the host processor does not support OSA
- ▶ Likely to provide higher throughput between the host and the network than channel-attached servers
- ▶ Can connect many token-ring segments through a single channel interface (if you need to connect from the host to many separate token-ring segments—either for scalability or due to a fundamental aspect of your campus token-ring design)

Considerations of migrating to channel-attached routers as token-ring gateways include:

- ▶ **Scalability:** In addition to the concerns regarding SNA boundary function workload and VTAM subarea element addresses (5.3.3, “SNA subarea addressing, routing, and boundary function (BF)” on page 162), carefully consider the channel-attached router configuration that would be required to meet the needs of your environment.
- ▶ **Cost:** Router channel interfaces can be expensive.
- ▶ **Availability:** Migrating to channel-attached routers as token-ring gateways continues the use and dependence upon end-of-life technologies, including token ring and ESCON.

Migrating from token ring to Ethernet

Due to the typically lower costs, the high bandwidths available, and the IT industry focus on Ethernet, many organizations today have plans to migrate their token-ring environments to Ethernet. In an Ethernet environment, you may be able to exploit the very powerful and unique capabilities of the OSA-Express host interfaces. OSA-Express offers the fastest and most efficient means of connecting host systems to the network, and is discussed in Chapter 15, “OSA-Express” on page 269.

The alternatives for providing host connectivity to an Ethernet environment include:

- ▶ Open Systems Adapter (OSA) Ethernet interfaces
- ▶ Channel-attached servers with Ethernet interfaces
- ▶ Channel-attached routers with Ethernet interfaces

These alternatives are discussed in 4.2.3, “Ethernet, FDDI, and ATM” on page 124.

Note: Replacing 3745 token-ring connectivity with any of the Ethernet alternatives requires the migration of *all* devices which communicate directly through the token ring to the host. See “Migration strategy” on page 121 for a discussion of how such a migration can be accomplished—both incrementally and at a pace that is appropriate for your organization.



Note: While the OSA-Express directly supports IP connectivity on the Gigabit Ethernet interface, for SNA application traffic to use the OSA-Express for Gigabit Ethernet you must either:

- ▶ Transform the SNA traffic into IP using a technology such as TN3270 (discussed in Chapter 17, “TCP/IP” on page 299)
- ▶ Carry the SNA traffic over IP using a technology such as Enterprise Extender (discussed in Chapter 19, “Enterprise Extender (EE)” on page 331) or
- ▶ Use OSA Layer 2 support to a CCL NCP (discussed in 16.3.1, “Layer 2 support” on page 290).



Important: Duplicate addressing exploits the unique nature of token-ring source route bridging and LLC-2 to support duplicate MAC addresses. Therefore, it cannot simply be copied for an Ethernet network. To provide high availability for SNA application traffic similar to that of duplicate token-ring MAC addressing when using Ethernet host access, you must either:

- ▶ Implement duplicate MAC addresses to isolated Ethernet segments and use routers to direct traffic to them (discussed in “Implementing CCL Ethernet high availability” on page 118).
- ▶ Transform the SNA traffic into IP using a technology such as TN3270 (discussed in Chapter 17, “TCP/IP” on page 299).
- ▶ Carry the SNA traffic over IP using a technology such as Enterprise Extender (discussed in Chapter 19, “Enterprise Extender (EE)” on page 331) or
- ▶ Use Advanced Peer-to-Peer Networking (APPN) High Performance Routing (HPR) over the Ethernet infrastructure.



Implementing CCL Ethernet high availability

IBM communication controllers only support token ring for LAN-based SNA connectivity. However, token-ring environments are nearing end of life and becoming increasingly difficult to support. Consequently, organizations are trying to migrate their remaining token-ring connectivity to Ethernet. The CCL product

can help because it supports direct Ethernet connectivity into your NCP. If availability is important for your SNA applications, you may need to implement a highly available Ethernet structure similar to the token-ring duplicate addressing structures of the past (as described in 4.2.2, “Token ring (and duplicate addressing)” on page 108).

Figure 4-9 shows one possible Ethernet high availability configuration for a CCL implementation.

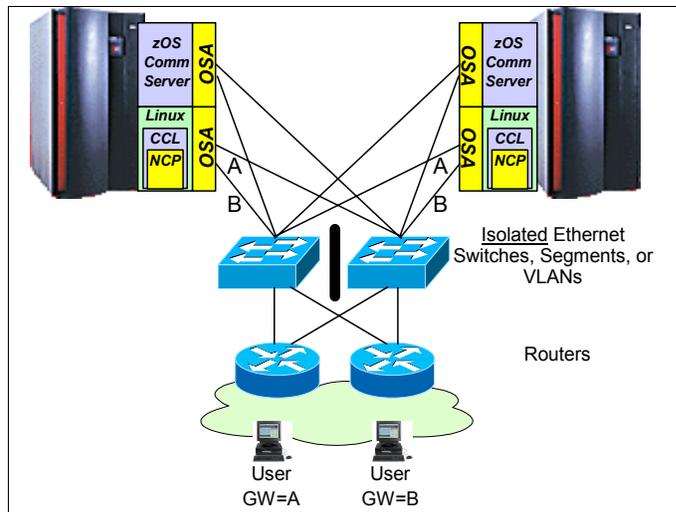


Figure 4-9 CCL Ethernet high availability

In the figure, the OSA ports of the two CCL NCPs are assigned identical MAC addresses, A and B, and connected to each of two isolated Ethernet segments such that there is only one instance of each address on a segment. Each CCL will define a Logical Link Control (LLC) Service Access Point (SAP) of X'04' for connections to users over their redundant MAC address A and B. Additionally, for connections to z/OS Communications Server (VTAM), each CCL NCP will define either:

- ▶ A unique SAP address (for example, X'08' or X'0C')
- or
- ▶ A unique MAC address (using OSA Layer 2 multiple MAC address support)

Half of the user workstations are configured for destination MAC address A for their host access and half use MAC address B. If the users are connected via remote Data Link Switching (DLSw)+ peer routers, DLSw+ load balancing (round robin), or enhanced load balancing (circuit count or circuit weight) can be used to distribute the connections across the available peer routers (and paths to the redundant CCL NCP MAC addresses).

Each VTAM has two transmission groups (TGs) defined to each of the NCPs (through each of the two Ethernet switches). The PATH statements should be coded to allow VTAM to use whichever TG is active at the time.

This configuration permits the following recovery scenarios:

- ▶ If the two z/OS CS systems are configured as CMC and backup CMC, both CCLs would actually be activated and owned by the CMC. In the event of a CMC failure, the backup CMC would take over the CCL NCPs without disruption of sessions other than those going to the failed CMC host.
- ▶ Failure of either CCL NCP will disrupt only those sessions passing through it. Even if it did not recover for a while, users could immediately reconnect finding the remaining CCL NCP.
- ▶ Failure of either Ethernet switch will disrupt only those sessions passing through it. Users can immediately reconnect finding the same MAC address (on the other CCL NCP) across the remaining Ethernet switch.
- ▶ Failure of either central-site router will disrupt only those sessions passing through it. Users can immediately reconnect through the other central-site router.

This CCL Ethernet high availability structure should therefore provide comparable availability characteristics to the previously popular token-ring duplicate addressing.

Migration from 3745 token ring connectivity to CCL Ethernet connectivity can be facilitated by configuring the CCL NCP with the same MAC address as the communication controller token-ring interface.

Note: Though NCP only supports SNA over token ring, CCL transparently converts Ethernet frames to token ring for the NCP. So the CCL NCP is configured with a token-ring MAC address even though it is actually using an Ethernet OSA port. The Ethernet OSA port, however, will be configured (through the OSA/Support Facility) with the address bits flipped to accommodate the differences between noncanonical token-ring MAC addresses and canonical Ethernet MAC addresses.

While both the CCL Ethernet interface and the communication controller token-ring interfaces are active, router load balancing can be used to distribute connections between the CCLs and the 3745 communication controllers. After you are comfortable with the operation of the CCL infrastructure, you can disable the communication controller token-ring interfaces and let all connections come through the CCLs.

Note: Multiple communication controller token-ring interfaces potentially can be replaced by fewer CCL OSA ports because the CCL can support multiple MAC addresses on a single Ethernet port using OSA Layer 2 function. A CCL NCP can support up to eight MAC addresses (corresponding to the NCP support limitation of up to eight token-ring interfaces). Those eight MAC addresses could all be assigned to a single OSA port, distributed across eight different OSA ports, or any combination in between. For example, an optimal availability configuration for a single CCL NCP with eight MAC address might be to have four MAC address, on each of two OSA ports, on separate OSA adapters.

If, for migration reasons, you need to support even more than eight MAC addresses, you can run additional CCL NCPs (either in the same Linux image or in a separate ones depending on your performance and scalability requirements), distributing those MAC addresses across however many OSA ports you require for availability (probably just a couple).

Duplicate MAC addresses, however, must be distributed across different Linux images. (If you set up duplicate MAC addresses within one Linux image, all traffic will be routed to the same CCL NCP TIC, defeating the point of duplicate MAC addressing.)

Migration strategy

Figure 4-10 on page 122 shows the network connectivity part of an architectural model—known as Net390—developed by IBM Global Services through numerous engagements to help organizations migrate toward strategic host connectivity. The Net390 architecture is discussed in Chapter 14, “Net390 architecture” on page 261).

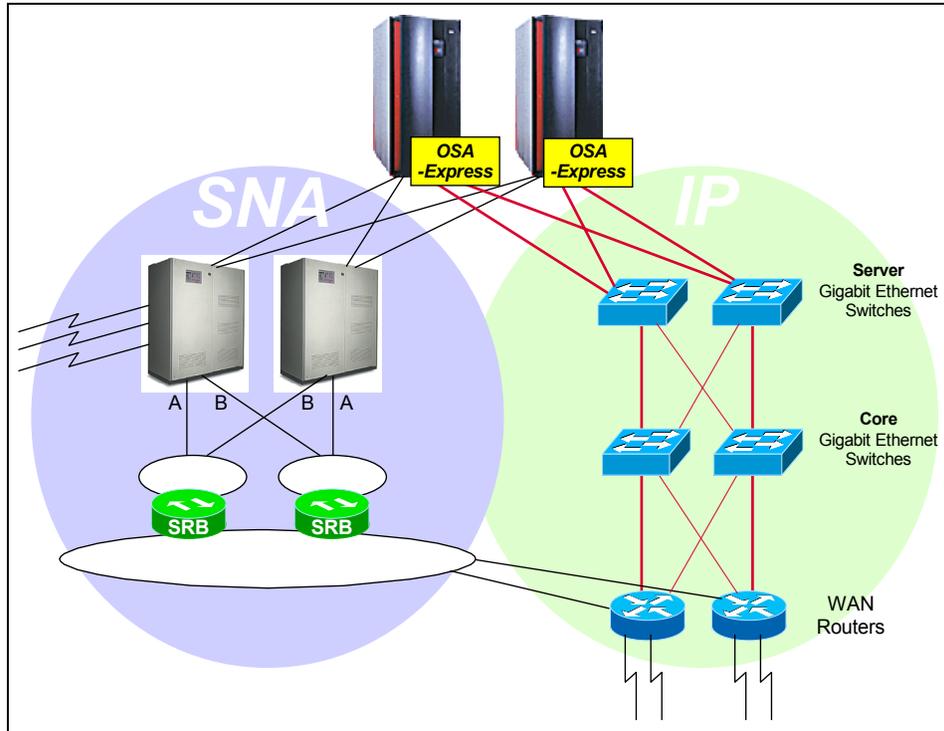


Figure 4-10 Net390 network connectivity

The Net390 network connectivity model provides high-speed, highly available, IP-centric, Gigabit Ethernet-based access to your host environment through direct OSA-Express attachment to your network—all while also continuing to support the current 3745 SNA token-ring infrastructure. By supporting both for a time, your users can be migrated toward the strategic host access incrementally and at a pace that is appropriate for your organization.



Note: The introduction of the Communication Controller for Linux on System z9 and zSeries (CCL) product adds important new options to the Net390 model. The function of the 3745 Communication Controllers shown in Figure 4-10 can be migrated to CCL running in the mainframe environment and the token-ring infrastructure shown can then be replaced with Ethernet (as described in “Migrating from token ring to Ethernet” on page 117.)

See Chapter 16, “Communication Controller for Linux on System z9 and zSeries (CCL)” on page 283 for additional information about CCL.

Technologies that are essential to the Net390 network connectivity shown in Figure 4-10 on page 122 and to the migration of SNA traffic to strategic Ethernet connectivity are discussed in the “Strategic solution technologies” part of this book. They include:

- ▶ OSA-Express (discussed in Chapter 15, “OSA-Express” on page 269)
- ▶ CCL (discussed in Chapter 16, “Communication Controller for Linux on System z9 and zSeries (CCL)” on page 283)
- ▶ TN3270 (discussed in Chapter 18, “TN3270” on page 311)
- ▶ Enterprise Extender (discussed in Chapter 19, “Enterprise Extender (EE)” on page 331)

Recommendations

Direct host-to-network attachment through OSA-Express Ethernet interfaces (Illustrated in Figure 4-10 on page 122) provides the optimal means of connecting your host environment to the network for TCP/IP. For SNA traffic to be optimally supported over the OSA-Express, however, it should be carried either via TN3270 or Enterprise Extender. With that infrastructure in place, the highest performance and least expensive connectivity between hosts and intermediate servers (such as the Web servers discussed in “Channel-attached servers” on page 113) is through the Gigabit Ethernet switches rather than host channel interfaces.

Not all host environments can use OSA-Express adapters. OSA-Express is supported only on IBM S/390 G5 and G6 and zSeries processors. In other environments, channel-attached servers may be appropriate if most traffic transits such servers anyway, and if the total workload required to support your environment can be accommodated by a few servers. Otherwise, channel-attached routers may be your best alternative for host network connectivity.

Likewise, not all host environments support Enterprise Extender, either due to product limitations (such as VTAM for VM/ESA® or VSE/ESA™ environments) or due to the organization’s choice not to implement such APPN technologies. In those cases, the CCL product can provide a means to migrate from communication controller hardware and token-ring connectivity while continuing to support mission-critical SNA requirements.

If the strategic infrastructure needs or plans of your organization include an eventual migration from token ring to Ethernet, coordinate your communication controller migration plans with those for your LAN (as described in “Migration strategy” on page 121) to avoid the costs and disruptions of two separate migration projects.

As you migrate from your 3745 token-ring connectivity, carefully plan for and monitor the changes in your host CPU utilization.

4.2.3 Ethernet, FDDI, and ATM

As shown in Figure 4-11, the Ethernet interfaces on the 3745 support TCP/IP communication either between Ethernet LANs over an SNA network, or between a 3745-attached Ethernet and a host.

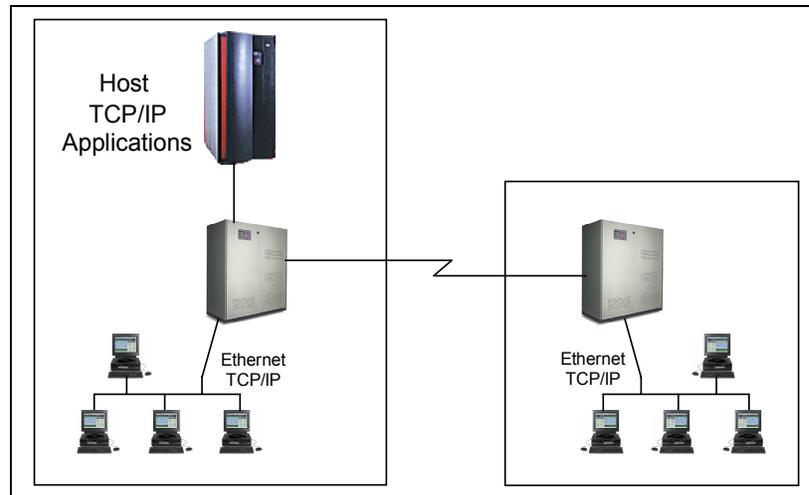


Figure 4-11 3745 Ethernet support

You can accommodate the traffic between Ethernet LANs by simply attaching those segments into routers. Routers were originally developed specifically to carry TCP/IP protocols over networks such as interconnected Ethernet LANs. Many organizations already have such an infrastructure in place to support intranet and TCP/IP applications.

In addition to Ethernet, communication controller MAE interfaces support FDDI and ATM access over channel interfaces into hosts. The alternatives for supporting communication between communication-controller-attached LANs and a host are the same as the alternatives discussed above for token ring:

- ▶ Open Systems Adapter (OSA) Ethernet, FDDI, or ATM interfaces
- ▶ Channel-attached servers with Ethernet, FDDI, or ATM interfaces
- ▶ Channel-attached routers with Ethernet, FDDI, or ATM interfaces

Open Systems Adapter for Ethernet, FDDI, or ATM

Figure 4-12 shows the use of Open Systems Adapter (OSA) LAN connectivity as an alternative to using the communication controller LAN interfaces. In this solution, LAN connections are moved from the communication controllers onto OSA cards in the hosts. Note, however, that OSA interfaces are supported only on certain IBM mainframe hosts. OSA-2 is supported in G2 and higher, OSA Express is supported in G5 and higher.

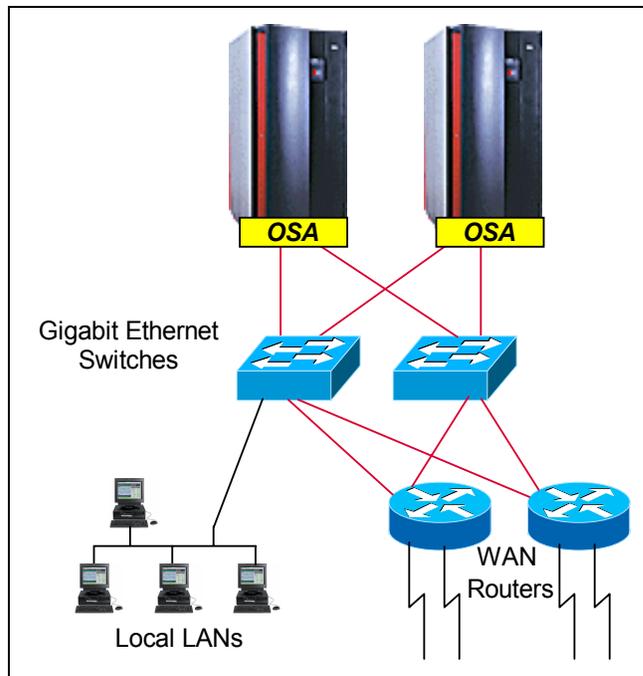


Figure 4-12 OSA Ethernet

Advantages of migrating to an OSA LAN interface include:

- ▶ **Strategic:** OSA interfaces, particularly the OSA-Express for gigabit Ethernet, are the IBM strategic means for connecting mainframe hosts to networks. They are discussed in Chapter 15, “OSA-Express” on page 269.
- ▶ **Performance:** For TCP/IP traffic, OSA-Express for Gigabit Ethernet offers the fastest and most efficient means of connecting host systems to the network, and is discussed in 15.3.1, “TCP/IP support” on page 274.
- ▶ **Availability:** By attaching mainframe hosts directly to the network using an OSA adapter, you can eliminate intermediate boxes such as channel-attached routers or servers, which should yield better availability for host access.

- ▶ Cost: Depending upon your specific environment, an OSA solution will probably be a less expensive solution than channel-attached routers or servers.

A consideration of migrating to an OSA LAN interface is:

- ▶ OSA interfaces are supported only on certain IBM mainframe hosts. OSA-2 is supported in G2 and higher, OSA-Express is supported in G5 and higher.

Channel-attached servers

Figure 4-13 shows the use of channel-attached servers as LAN gateways for the host environment. In this solution, the LAN connections are moved from the communication controllers onto LAN adapters in servers.

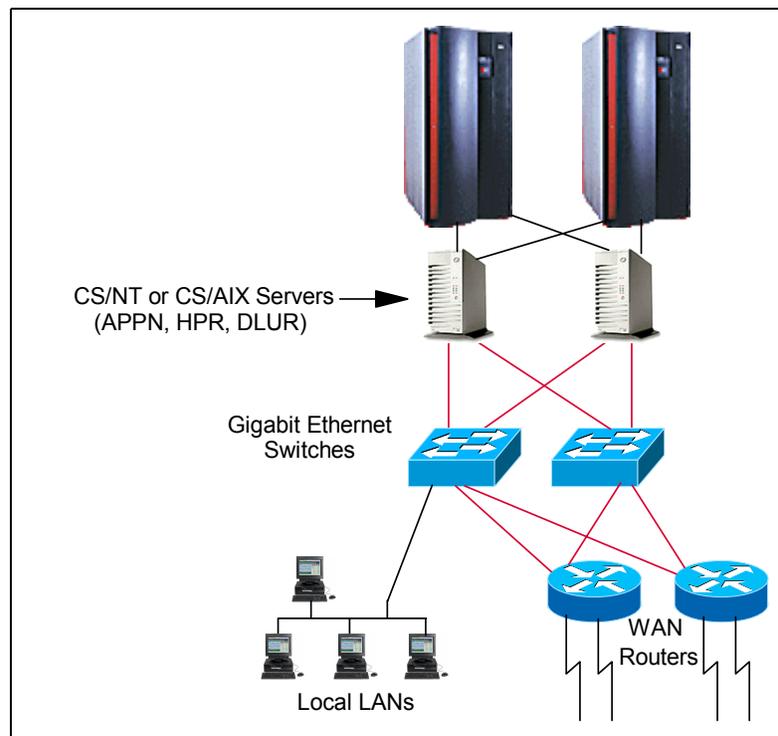


Figure 4-13 Channel-attached servers as Ethernet gateways



Important: Depending on the functions running in the servers, server CPU capacity could become a performance problem. Large numbers of users or high transaction rates could drive up the number of required channel-attached servers and, consequently, the cost and complexity of the solution.

Advantages of migrating to channel-attached servers as LAN gateways include:

- ▶ Viable alternative for those environments where the host processor does not support OSA.
- ▶ Well-suited to environments where most traffic is to and from intermediate servers that provide key functions such as:
 - Web front-end to host applications
 - The application tier of a three-tier client-server solution

A consideration of migrating to channel-attached servers as LAN gateways is:

- ▶ Scalability: Server CPU utilization issues could drive up the number of required servers and, therefore, cost and complexity.

Channel-attached routers

Figure 4-14 shows the use of channel-attached routers as LAN gateways for the host environment. In this solution, the LAN connections are moved from the communication controllers onto LAN adapters in routers.

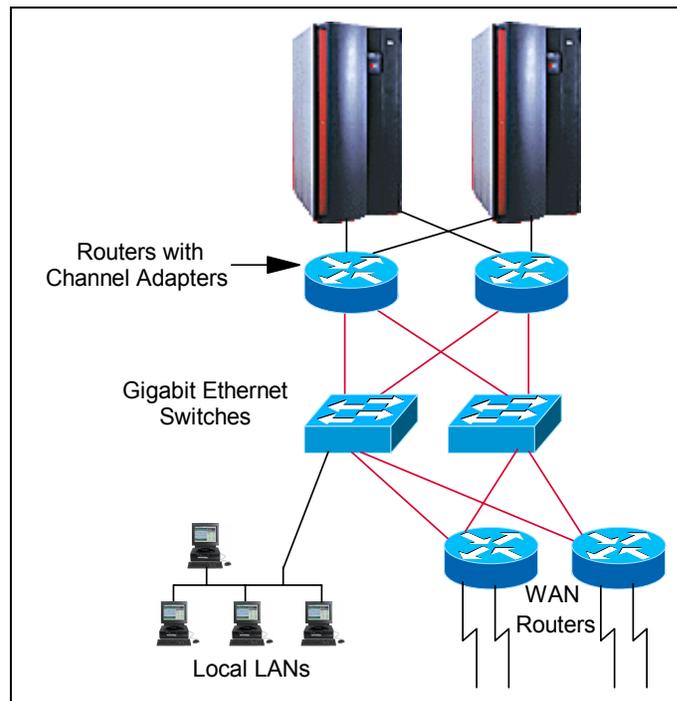


Figure 4-14 Channel-attached routers as Ethernet gateways

Advantages of migrating to channel-attached routers as LAN gateways include:

- ▶ Offers a viable alternative for those environments where the host processor does not support OSA
- ▶ Likely to provide higher throughput between the host and the network than channel-attached servers
- ▶ Connects many LANs through a single channel interface (if you need to support many different LANs—either for scalability or due to a fundamental aspect of your campus design)

A consideration of migrating to channel-attached routers as LAN gateways is:

- ▶ Cost: Router channel interfaces can be expensive.

Recommendations

Direct attachment through OSA-Express particularly for gigabit Ethernet interfaces (as shown in Figure 4-12 on page 125) provides the optimal means of connecting your host environment to the network. With that infrastructure in place, the highest performance and least expensive connectivity between hosts and intermediate servers (such as the Web servers discussed in “Channel-attached servers” on page 126) is through the Gigabit Ethernet switches rather than host channels.

However, not all host environments can use OSA-Express adapters. OSA-Express is supported only on S/390 G5 and G6 and zSeries processors. In other environments, channel-attached servers may be appropriate if most traffic transits such servers anyway, and if the total workload required to support your environment can be accommodated by a few servers. Otherwise, channel-attached routers may be your best alternative for host network connectivity.

4.2.4 Channel attachments

The previous sections of this chapter focused on the role of the communication controller in providing client access from communication lines or LANs to host applications. Such client-access support depends, in turn, on channel connectivity from the communication controller to the host. Consequently, while those sections focused on the type of client connectivity (communication lines or LANs) the discussions necessarily addressed the use of the channel as well.

This section of the book discusses the use of communication controller channel connectivity for the support of host-to-host communications—either between different channel-attached hosts on the same communication controller, or over a network link to another communication controller, as illustrated in Figure 4-15 on page 129.

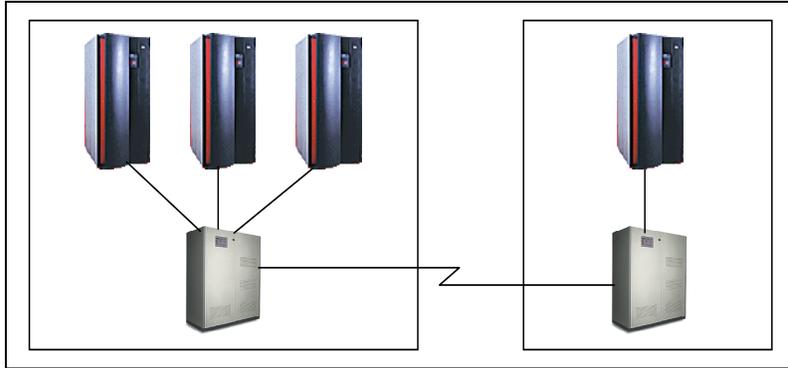


Figure 4-15 Host-to-host communications

Some common reasons for using communication controllers for host-to-host communication include:

- ▶ Interconnecting different SNA networks using SNI
- ▶ Interconnecting multiple data centers
- ▶ Providing alternate routing in case a channel-to-channel (CTC) connection fails

The role of a communication controller as an SNI gateway is discussed in 5.3.4, “SNA Network Interconnection (SNI)” on page 167.

When it comes to interconnecting hosts, there are numerous alternatives including:

- ▶ Parallel channel
- ▶ Enterprise Systems Connectivity (ESCON)
- ▶ Fibre Channel (FICON)
- ▶ Channel-to-Channel (CTC)
- ▶ Open Systems Adapter (OSA)
- ▶ Cross-System Coupling Facility (XCF)

Each of these host connectivity alternatives is discussed in the *zSeries Connectivity Handbook*, SG24-5444-01, available on the Web at:

<http://www.redbooks.ibm.com/pubs/pdfs/redbooks/sg245444.pdf>

These are not covered here. Because of its importance, however, we briefly discuss the use of the OSA-Express gigabit Ethernet for host-to-host communication below.

OSA-Express gigabit Ethernet for host-to-host communication

Figure 4-16 shows hosts interconnected via OSA-Express gigabit Ethernet interfaces through gigabit Ethernet switches. Those “server switches” are connected to “core switches” which, in turn, are interconnected via a high-speed wide area network: your IP-based “Intranet backbone”. The backbone network could be based upon any of a number of high-speed WAN technologies including routers or ATM switches.

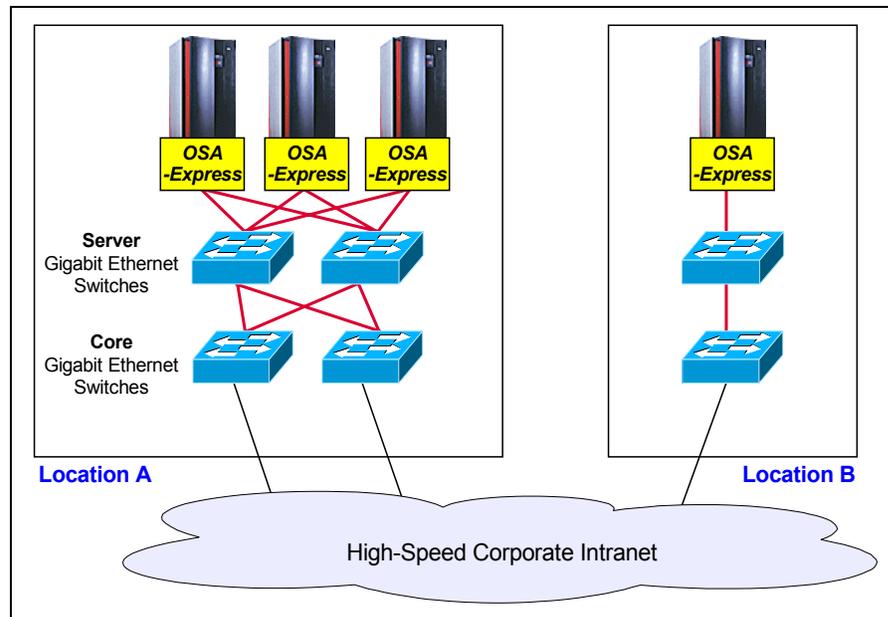


Figure 4-16 OSA-Express gigabit Ethernet for host-to-host communication

In the data center campus environment, this solution provides 1 Gbps connectivity between hosts. Communication between data centers can leverage your investment in a high-speed IP intranet backbone. By using Enterprise Extender between the hosts, this very high-speed infrastructure can be used by SNA traffic as well as TCP/IP traffic.

Recommendation

While you will need to consider all of the host-to-host connectivity options available to you, give special consideration to OSA-Express gigabit Ethernet for host-to-host communication because of its simplicity, versatility, performance, low-cost, and potential strategic importance.

4.3 Network Node Processor functions and alternatives

The network node processor (NNP) is a key component of the 3746 Nways Multiprotocol Controller Models 900 and 950. The NNP was introduced to give communication controllers stand-alone APPN network node (NN) and high-performance IP routing capabilities.

4.3.1 APPN network node (including HPR and DLUR)

Recall that, previously, a communication controller could only support APPN as part of a “composite network node” (CNN), but VTAM and the NCP together gave the appearance of a single network node (NN). With the NNP, as shown in Figure 4-17, the 3746 can be used as a stand-alone NN providing:

- ▶ APPN routing for SNA traffic, including high-performance routing (HPR) support
- ▶ SNA boundary function for non-APPN SNA devices, called *dependent LUs*, with an APPN extension called Dependent LU Requester (DLUR)

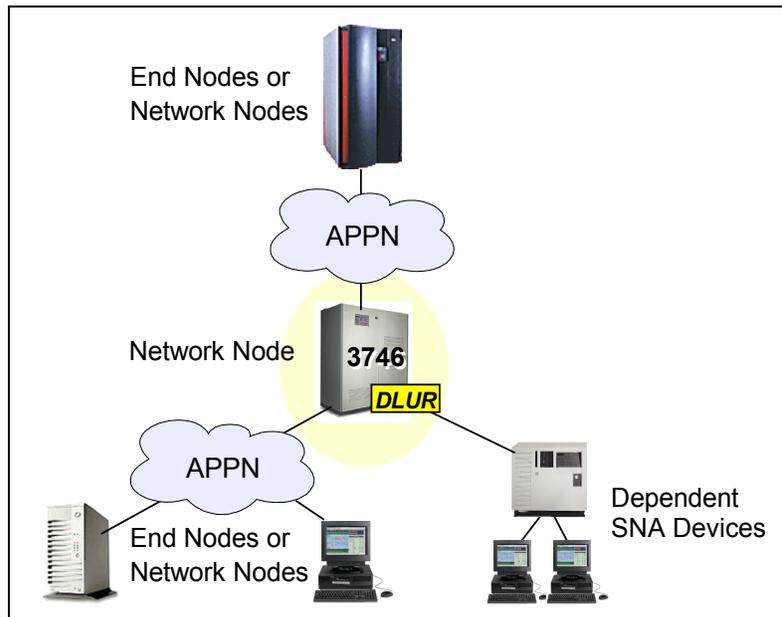


Figure 4-17 3746 NNP APPN capabilities

There are a couple of things worth noting about the 3746 NNP APPN environment illustrated above:

- ▶ While the figure shows only one host, there could be many hosts. Using APPN routing, the communication controller can route SNA traffic directly between the endpoints of the SNA sessions. As discussed in “Eliminating the APPN functions” on page 135, without the APPN routing, SNA session traffic may have to be routed between the session endpoints by one of your host systems.
- ▶ The communication controllers may be channel attached to the host systems in the data center, in a remote location, or both. It should be noted, however, that the original, architectural, intent for the DLUR function was that it be distributed and implemented as close as possible to the dependent LUs requiring DLUR services. Because of the processing and memory required to support the DLUR function, centralizing it can present significant scalability concerns. DLUR scalability is discussed further in “Replacing the APPN functions” on page 132.

Your migration alternatives for the NNP APPN capabilities include:

- ▶ Replacing the APPN functions either with routers or with servers running APPN-capable communication software
- ▶ Eliminating the APPN functions altogether by either:
 - Using LAN interfaces on the hosts (such as OSA interfaces) and a bridged-LAN infrastructure (including, if needed, bridging or DLSw over your WAN) and supporting SNA boundary function in VTAM or in a CCL NCP
 - Migrating from the devices that depend upon the SNA support using technologies such as TN3270

Each of these alternatives is discussed below.

Replacing the APPN functions

In addition to the 3746 NNP, APPN functions have been implemented in:

- ▶ Routers: Such as the SNA Switching Services (SNASw) in Cisco routers. For more information, search on SNASw at:
<http://www.cisco.com/>
- ▶ Servers: Running software such as IBM Communications Server software. Communications Server runs on AIX®, Windows®, z/OS, and Linux. For more information, visit:
<http://www.ibm.com/software/network/commsrver/>

Consequently, as shown in Figure 4-18, you could replace the APPN functions of your 3746 NNP with either routers or servers.

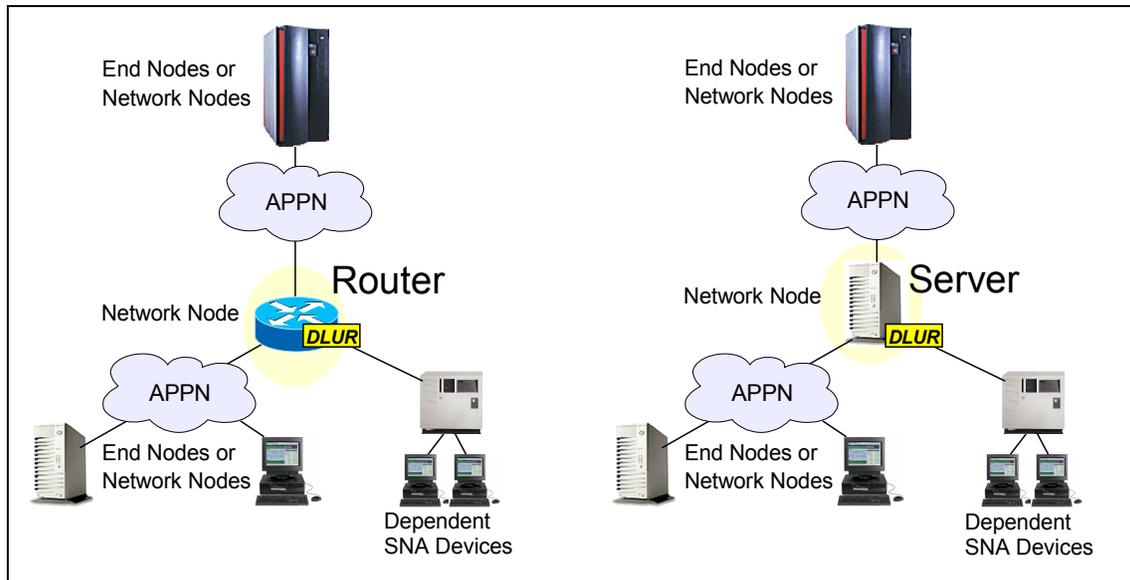


Figure 4-18 Replacing NNP APPN functions with routers or servers

In some respects, because APPN is a networking function, it makes more sense to run it in the networking boxes (specifically, routers). However, while routers provide very inexpensive solutions for the base capabilities of IP routing, router software and hardware upgrade costs can be very expensive for APPN functions. Consequently, if you already have servers in the environment that could run the APPN-capable communication software, it may be a more cost-effective solution for you to use the routers strictly for IP routing and use the servers for the APPN functions.

Note: Enterprise Extender technology (discussed in Chapter 19, “Enterprise Extender (EE)” on page 331) available in IBM Communications Server software can be particularly useful for server-based APPN environments because it enables SNA traffic to use your IP-based infrastructure without any special network support (such as DLSw or bridging).

Regardless of whether you use servers or routers for the replacement of your NNP APPN functions, a key design issue is where in your environment you place the APPN function. The simplest approach, recommended in past situations by router vendors, is to locate the APPN function in the data center and use a router

function such as bridging or DLSw to carry the SNA traffic to the data center. This approach has the following considerations:

- ▶ **Cost:** While DLSw support in a router usually does not cost as much as APPN, you spend more on routers in the data center: first to handle the DLSw workload, then to provide the APPN support. If you have multiple data centers, and especially if you use them to back each other up, you duplicate those more expensive routers for each data center. In comparison, using a distributed-APPN solution, you need more expensive routers (or servers) for each remote location (and there may be many remote locations); however, you spend less on routers in the data centers.
- ▶ **Availability:** With the hybrid DLSw and APPN approach, your single points of failure include each DLSw endpoint (branch and data center), the APPN router that provides DLUR function, and the host. By placing the APPN function in the remote location and using EE technology, your only single points of failure will be the remote router and the host and there are ways to design for high availability both for the remote routers and for the host.
- ▶ **Scalability:** DLSw usually requires only a small amount of workload in a remote router; however, in the data center where sessions from many remote locations converge, the DLSw workload can be very substantial often requiring multiple “DLSw peering routers”. Likewise, APPN workload and especially DLUR while manageable when distributed, can be very significant if consolidated into a data center. As the number of remote SNA devices increases, it can become difficult to design a centralized APPN DLUR environment to handle the workload especially when your design also needs to consider high availability.
- ▶ **Performance:** A bridged or DLSw-based solution cannot distinguish between different kinds of SNA traffic, such as batch print jobs and interactive transaction system users. EE technology enables your network to differentiate between the SNA batch and interactive traffic and can use distinct UDP ports and set the IP differentiated services bits so that the IP network can give appropriate priority throughout. Also, if you have multiple data centers, using bridging or DLSw technology can cause traffic going to one data center to take a suboptimal route through the other data center.

Consequently, while it may appear more complex and may even be more expensive, consider distributing your APPN functions to your remote locations, ideally, using EE technology to transport the SNA traffic over your IP infrastructure.



Restriction: VTAM for VM/ESA and VSE/ESA does not support Enterprise Extender.

Eliminating the APPN functions

You can use a combination of mainframe TN3270 services, a CCL NCP, and z/OS Communication Server (VTAM) for the elimination of the APPN functions provided by your 3746 NNP:

- ▶ You can use LAN interfaces on your mainframe hosts (such as OSA interfaces) and a bridged-LAN infrastructure (including, if needed, bridging or DLSw over your WAN) and support SNA boundary function in VTAM or in a CCL NCP.



Note: The IBM Communication Controller for Linux on System z9 and zSeries (CCL) product enables you to run NCP on your mainframe, which gives you a way to migrate devices from your NNP hardware to CCL NCP boundary function. See Chapter 16, “Communication Controller for Linux on System z9 and zSeries (CCL)” on page 283 for additional information on CCL.

- ▶ You can migrate from the devices that depend upon the SNA support, for example, using TN3270 rather than SNA 3270 support.

These options are illustrated in Chapter 4-19, “Accommodating the elimination of NNP APPN functions” on page 135.

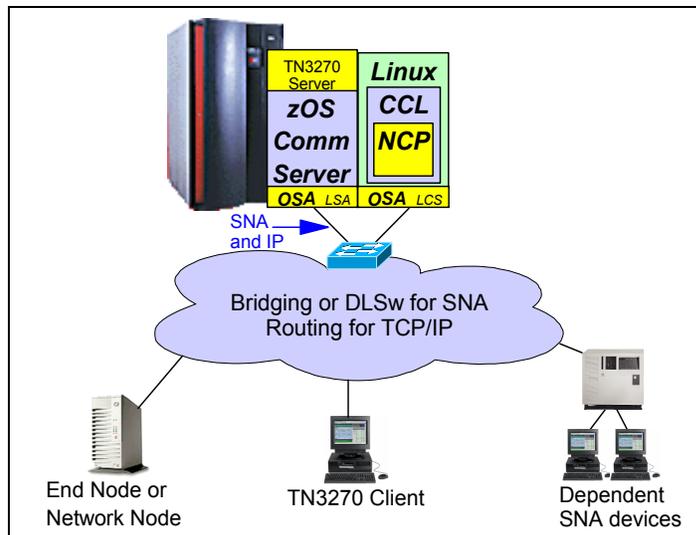


Figure 4-19 Accommodating the elimination of NNP APPN functions

Chapter 4-19, “Accommodating the elimination of NNP APPN functions” on page 135 shows the combination of three important functional capabilities: a mainframe TN3270 server, a CCL NCP, and z/OS Communication Server

(VTAM) for the elimination of the APPN functions provided by your 3746 NNP. By implementing a TN3270 server, you may be able to migrate some of your SNA users from SNA 3270 to TCP/IP-based TN3270 client function. Strategically, this could help you to move toward an all-TCP/IP environment. While the TN3270 server can be implemented in routers or servers, in most environments today, the mainframe-host-based TN3270 server has emerged as the best option. For more information on TN3270, see Chapter 18, “TN3270” on page 311.

Any APPN hosts in your environment can use the bridged infrastructure (as a “connection network”) to provide efficient, logically meshed connectivity to each other APPN host in your environment.

Supporting your dependent SNA devices may be your most difficult challenge. If your SNA environment is small, you may be able to support it simply by directing SNA users into VTAM boundary function. For larger environments, a CCL NCP may be required as well. See the important notes below.



Important: Communication controllers save host CPU cycles by offloading certain network functions, including SNA device support (also known as the SNA boundary function). Removing communication controllers will consequently result in increased host CPU cycles and storage requirements. The amount of increase will depend on:

- ▶ The number of devices in the network and their transaction rates, and
- ▶ Whether a dependent LU requester (DLUR) is implemented in the network to support SNA devices.

Alternatives for handling SNA boundary function workload are discussed in 5.3.3, “SNA subarea addressing, routing, and boundary function (BF)” on page 162.



Important: Large networks have encountered problems such as SNA session setup failures or the inability to activate additional resources as the result of running out of SNA element addresses in a VTAM subarea. Migrating devices from your communication controllers can add to this problem, because the device support (boundary function) ends up moving from the communication controller subarea into the VTAM subarea. Alternatives for handling the SNA subarea function are discussed in 5.3.3, “SNA subarea addressing, routing, and boundary function (BF)” on page 162. For information on VTAM addressing enhancements, see Appendix C, “Avoiding VTAM network addressing problems” on page 403.



Note: The CCL product gives you a way to migrate devices from your NNP hardware to CCL NCP boundary function and reduces the risk of encountering problems with running out of SNA element addresses in a VTAM subarea. Though using CCL might increase the host CPU cycles overall, CCL CPU cycles may be less expensive than other mainframe CPU cycles because Linux partitions can use Integrated Facility for Linux (IFL) processors. See Chapter 16, “Communication Controller for Linux on System z9 and zSeries (CCL)” on page 283 for additional information about CCL.

Recommendation

Pursue strategic migrations first. Strategic migrations such as migrating from SNA 3270 to TN3270 or Web-enabling applications and deploying Web browsers can simultaneously work to position your organization to support new applications while also reducing your dependency on SNA. Through such migrations, you may be able to gradually reduce the SNA environment to the point that removing your NNP APPN support will have minimal impact. If however, you will be required to support a substantial and mission-critical SNA environment for the foreseeable future, then pursue one of the following two approaches:

- ▶ If the remaining SNA environment will be too large to be supported by a single VTAM or is very important (in terms of performance or availability), then replace the NNP APPN with a CCL NCP or a router or server APPN implementation.
- ▶ Otherwise, let VTAM support the remaining SNA users.

4.3.2 IP routing

The 3746 Nways Multiprotocol Controller also supports IP routing (independent of that provided by the NCP). If you are using your 3746 to route IP traffic into your host environment, migrate to the OSA-Express adapter (discussed in Chapter 15, “OSA-Express” on page 269). If you are using your 3746 for other IP routing, consider migrating that function to routers or to layer-3 (L3) switches. L3 switches are switches that have very high-speed IP routing capabilities. In general, today’s routers and L3 switches can provide more robust IP routing (including such capabilities as IP differentiated services based prioritization), at a lower cost, and with better performance.

4.4 Multi-Access Enclosure functions and alternatives

The Multi-Access Enclosure (MAE) was really just a nice marketing term for an IBM 2216 Nways Multiprotocol Router installed inside of an IBM 3746 cabinet. The router did, however, provide two potentially valuable capabilities: TN3270 server and Network Dispatcher.

4.4.1 TN3270 server

If you are using your Multi-Access Enclosure (MAE) as a TN3270 server, your alternatives include migrating the TN3270 server function to your mainframe host or replacing the MAE TN3270 server with another outboard solution such as a:

- ▶ Server

Typically, such outboard TN3270 servers are RISC or Intel®-based machines running software such as the IBM Communications Server or the Microsoft® SNA server.

- ▶ Router

For example, Cisco Systems, Inc. supports TN3270 server functionality on their Channel Interface Processor (CIP) and Channel Port Adapter (CPA).

In most environments today, the mainframe-host-based TN3270 server has emerged as the most straightforward, scalable, reliable, and cost-effective option. For more information on TN3270, see Chapter 18, “TN3270” on page 311.

4.4.2 Network Dispatcher

If you are using your Multi-Access Enclosure (MAE) to support Network Dispatcher function for high-availability IP services, your alternatives include:

- ▶ Replacing the MAE as your platform for running the Network Dispatcher component of WebSphere® Edge Server

The WebSphere Edge Server software from IBM can be run on any platform that supports the IBM AIX, Red Hat Linux, Sun™ Solaris™, or Microsoft Windows NT® operating systems. Therefore, replacing the MAE with different hardware is probably your simplest alternative.

- ▶ Migrating to Sysplex Distributor function in your mainframe host

If all of the IP services that you are currently serving with the Network Dispatcher are being provided from your mainframe host Parallel Sysplex environment, the Sysplex Distributor will provide you the optimal, and most strategic, alternative to Network Dispatcher.

- ▶ Replacing the MAE and Network Dispatcher with another outboard load-balancing solution, preferably one that leverages Server/Application State Protocol (SASP)

The Server/Application State Protocol (SASP) defines the interface that is used between the external load balancers and the z/OS Load Balancing Advisor. SASP was introduced as part of the IBM Enterprise Workload Manager (EWLM) solution that is part of the IBM Virtualization Engine™ and is available on several platforms. Cisco Systems has announced support for SASP in their Content Switching Module (CSM).

Your simplest alternative is to replace the MAE with a different hardware platform; however, it may be worthwhile for you to explore other, possibly more strategic, alternatives such as the Sysplex Distributor or a SASP-based distributed approach. For more information on designing for IP high availability, see Chapter 17, “TCP/IP” on page 299.



Network Control Program (NCP)

IBM Product Number: 5648-063

Advanced Communications Function for Network Control Program (ACF/NCP, or simply NCP) is licensed software originally developed to run on the 3745 family of communication controllers. This chapter first discusses functions and alternatives of NCP link-level protocol support such as SDLC, frame relay, and ALCI. Then we discuss the functions and alternatives for advanced capabilities such as SNA class of service (COS), SNA subarea addressing, routing, and boundary function (BF), and SNA Network Interconnection (SNI) support.



Note: NCP can now be run on your mainframe using the Communication Controller for Linux on System z9 and zSeries (CCL) product, 5724-J38. In many cases, the CCL option will provide the most transparent migration from your IBM 3745 Communication Controllers.

See Chapter 16, “Communication Controller for Linux on System z9 and zSeries (CCL)” on page 283 for additional information on CCL.

Note: In addition to the technical solutions discussed in this chapter, you should also consider the alternatives of leaving things as they are, eliminating the IT service in question, or replacing that service with a completely new solution. See “Important migration-strategy alternatives” on page 95 for a discussion of each of those alternatives.

5.1 NCP product overview

For the most part, the communication controller hardware provides little more than physical interfaces—the bottom layer of the seven-layer OSI reference model. As illustrated in Figure 5-1, NCP is a software product that IBM initially developed more than 30 years ago to run on communication controller hardware to provide function that is roughly equivalent to the next three layers of the OSI model: the data link, network, and transport layers.

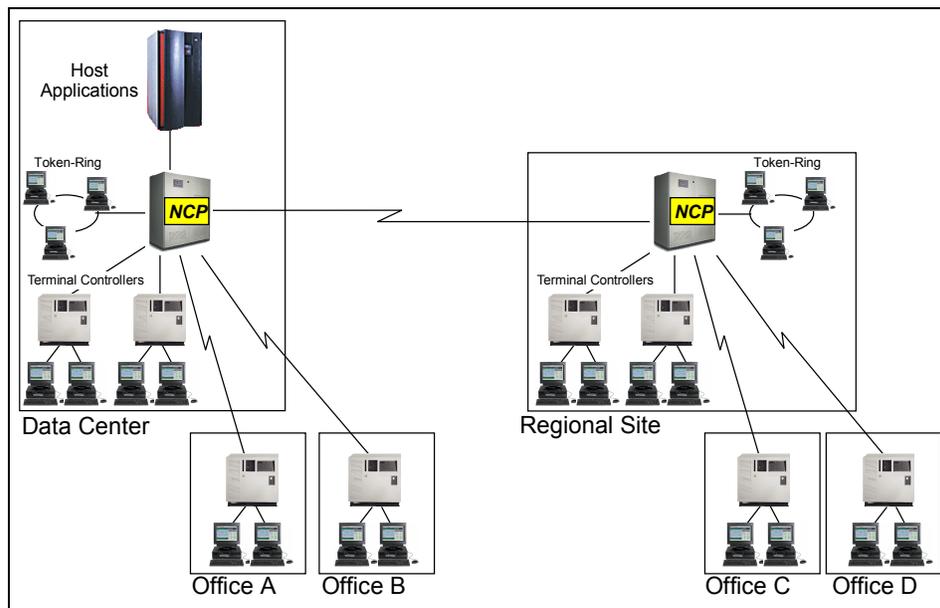


Figure 5-1 The role of NCP in a network

The NCP software implements an SNA architecture Physical Unit type 4 (PU 4) and provides the efficient routing and transmission of data across its:

- ▶ Peripheral links: Connections to peripheral devices such as terminal controllers or intelligent workstations

- ▶ Intermediate network node (INN) links: Connections to other communication controllers and to hosts

For the purposes of this book, NCP functions have been divided into:

- ▶ Link-level protocol functions and alternatives, such as SDLC, frame relay, and ALC
- ▶ Advanced capability functions and alternatives, such as SNA class of service (COS), SNA subarea addressing, routing, and boundary function (BF), and SNA Network Interconnection (SNI) support

Each of these functional areas are discussed in the sections that follow.

For more information about NCP, see:

- ▶ *NCP, SSP and EP Resource Definition Guide*, SC31-6223
- ▶ *NCP, SSP and EP Resource Definition Reference*, SC31-6224

5.2 Link-level protocol functions and alternatives

Link-level protocol support includes the functions necessary to deliver data over a physical link. In some cases, the communication controller must split data into a number of smaller frames for sending it on the physical link. The communication controller must also keep track of acknowledgment frames, perform error checking, and retransmit frames not received correctly. Essentially, the link-level support must provide an error-free virtual channel for the layers above.

With NCP, functions beyond simple data link control (such as protocol conversion) are frequently coupled with NCP link-level protocol support. For example, NCP support of BSC3270 includes SNA terminal emulation. Such NCP capabilities will be covered in this section as well.

NCP supports the following link-level protocols:

- ▶ Airlines Line Control (ALC)
- ▶ BSC 3270 terminal connection to SNA applications
- ▶ Frame relay
- ▶ Integrated Services Digital Network (ISDN)
- ▶ LAN 802.2, 802.3, and 802.5
- ▶ Start/Stop (S/S) terminal connection to TCAM applications
- ▶ Synchronous Data Link Control (SDLC)
- ▶ X.25 (with licensed support feature on 3746)

Each of these link-level protocols is discussed below.

5.2.1 Airlines Line Control (ALC)

Airlines Line Control (ALC) is a 6-bit synchronous data link protocol that operates over full-duplex lines and is widely used in the travel industry. ALC was developed as a very efficient link-level protocol for use with the IBM Transaction Processing Facility (TPF), an operating system used for hosting airline and other reservation systems. As illustrated in Figure 5-2, the NCP Airlines Line Control Interconnection (ALCI) feature provides communication controller support for ALC connections.

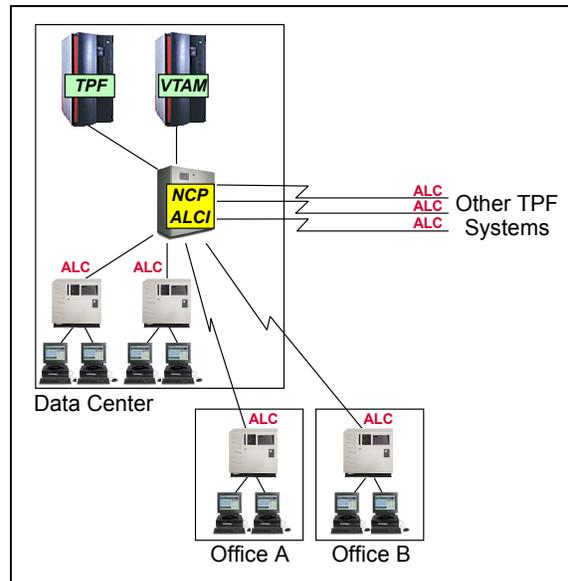


Figure 5-2 Communication controller ALC support

There are two complementary approaches to addressing your migration from ALC support in your communication controllers:

- ▶ Mapping of Airline Traffic Over TCP/IP (MATIP)

The MATIP standard was developed by a group of airlines and major airline system vendors. The standard is published as an informational RFC by the Internet Engineering Task Force (IETF) as RFC 2351. MATIP provides a means of transporting ALC traffic over a TCP/IP infrastructure. Products that support MATIP include:

- The IBM Transaction Processing Facility (TPF) at:
<http://www.ibm.com/software/ts/tpf/index.html>
- Cisco Systems, Inc. routers at:
<http://www.cisco.com>

- Datalex Alliance MATIP Gateway software at:
<http://www.datalex.com>
- JBM Electronics, Inc. protocol conversion devices at:
<http://www.jbmelectronics.com>

► Message queuing (MQ)

The MQSeries for TPF product supports both MQ Client and Queue Manager function and enables reliable, high-performance, messaging-based communication between TPF systems and applications residing on more than 35 different platforms. MQ provides independence from the networking infrastructure and, therefore, can enable your migration to a higher-performance and more strategic IP-based network. Consequently, MQ is the strategic communications API for TPF.

A hybrid MATIP and MQ solution is illustrated in Figure 5-3.

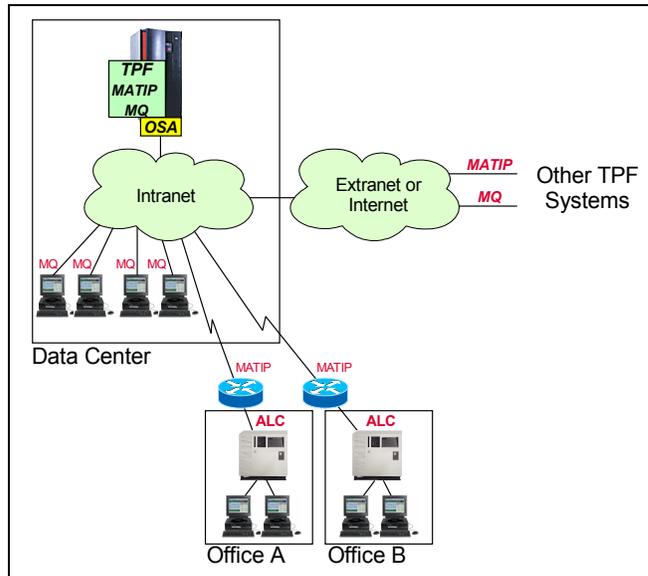


Figure 5-3 MATIP and MQ alternative to communication controller ALCI support

In this example, the devices in the data center were migrated using MQ while the devices in Offices A and B were preserved using routers with MATIP support. Other TPF systems can connect over private IP connections (extranet) or over the Internet (perhaps using VPN technology for security) and, while MQ would be the most strategic alternative, they can connect using either MQ or MATIP.

Where practical, migrate current users of ALC to MQ because it enables very high-performance, IP-based access to your TPF environment. In situations where you must continue to support existing ALC devices or connections, use MATIP.

5.2.2 BSC 3270 terminal connection to SNA applications

It is easy to be confused about BSC 3270 support in the communication controller. This is because both NCP and EP can support BSC 3270 devices, which means that they may or may not be part of the SNA network. BSC 3270 devices connected to communication controller ports that are under the control of NCP (and VTAM) are part of the SNA network, and it is this type of connection that is covered in this section. For EP-supported BSC 3270 devices, see 9.2.1, “BSC 3270 terminal connection to non-SNA applications” on page 214.

NCP support of BSC 3270 involves quite a bit more than just link-level protocol support. NCP and VTAM actually provide protocol conversion on behalf of BSC 3270 devices including emulating SNA PUs and LUs for those devices. Interestingly, that is why it was never possible to support the SSCP takeover function for BSC 3270 devices. To migrate from NCP support of BSC 3270, you have two options:

- ▶ Replace the BSC 3270 devices with devices, such as intelligent workstations, that support SNA 3270 or TN3270

Such desktop migrations can be difficult and expensive because they often involve many devices, but they have the advantage of enhancing the organization’s ability to support strategic new applications such as applications based upon Internet technologies. Unfortunately, however, if it were easy to migrate from your BSC 3270 devices, your organization probably would have already done so.

- ▶ Use protocol conversion devices to convert the BSC 3270 traffic into SNA 3270 or TN3270

If you have BSC 3270 devices that you need to continue to support, you may be able to use a BSC 3270 protocol conversion solution. According to JBM Electronics Co. (<http://www.jbmelectronics.com>), their gateway products can provide BSC 3270 to TN3270 conversion.

Alternatively, according to INETCO Systems (<http://www.inetco.com>), their INETCO Connect product can provide BSC 3270 to SNA 3270 conversion.

Your first choice should be to replace the BSC3270 devices with newer devices because such device upgrades can help to enable the support of strategic new applications.

5.2.3 Frame relay

Frame relay is a fast packet-switched networking technology. End stations dynamically share the network media and the available bandwidth sending variable-length packets through the network. These packets are switched between the various network segments until the destination is reached. As with most other communication controller functions, there is more to NCP frame relay support than just supporting attachment to carrier frame relay service networks. Figure 5-4 illustrates some of the different ways communication controllers support frame relay.

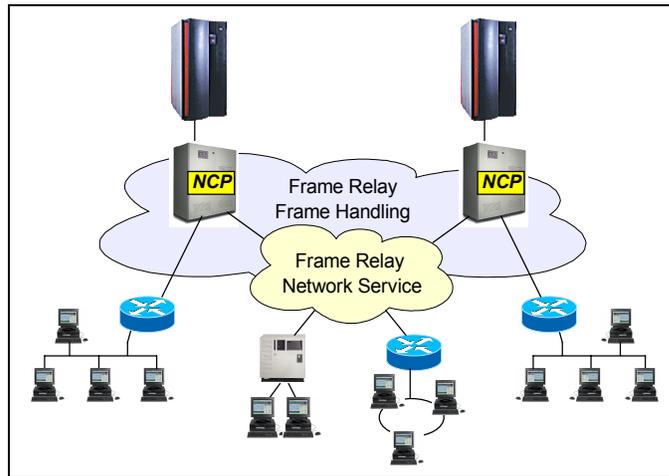


Figure 5-4 Communication controller frame relay support

The illustration shows frame relay connectivity over both a “public” frame relay service network and over a “private” frame relay network using the communication controller frame relay frame handling function. The public frame relay service network includes connectivity:

- ▶ Between communication controllers: Such links are called frame relay intermediate network node (INN) links. These links could be between data centers or between separate organizations and may also be using SNI (discussed in 5.3.4, “SNA Network Interconnection (SNI)” on page 167).
- ▶ From SNA devices to communication controllers: Such connections typically use frame relay boundary network node (BNN) links. Using frame relay connectivity for SNA devices allows many SNA devices to be consolidated into a few communication controller physical interfaces.
- ▶ From routers to communication controllers: For SNA, routers typically use boundary access node (BAN) links. With BAN, routers use frame relay RFC 1490 bridged frame format to forward LAN-based SNA traffic to the

communication controller. When the communication controller receives the frame, it treats it as if it were received on a token-ring interface.

The “private” frame relay, using the communication controller frame relay frame handling function, provides connectivity between routers. The routers may either be connected to communication controllers or to public frame relay services and can route any protocol supported by the routers across the public and private frame relay networks. Essentially, frame relay enables communication controllers to provide a multi-protocol network backbone.

Your alternatives for communication controller frame relay support include:

- ▶ Replacing your communication controller frame relay support with a combination of frame switches and routers
- ▶ Migrating your private frame relay to an IP-based infrastructure (which still uses public frame relay services as appropriate)

Each alternative is discussed below.

Frame switches and routers

Figure 5-5 on page 149 shows an example of replacing the communication controller frame relay support with frame relay switches and routers. The frame relay switches are required to replace the frame relay frame handler function of the communication controllers and the routers are required to take the frames from the frame relay network and move them into the host access infrastructure. The host access infrastructure can, in turn, be any of the host LAN access alternatives discussed in 4.2.2, “Token ring (and duplicate addressing)” on page 108 or in 4.2.3, “Ethernet, FDDI, and ATM” on page 124, whichever is most appropriate for your environment.

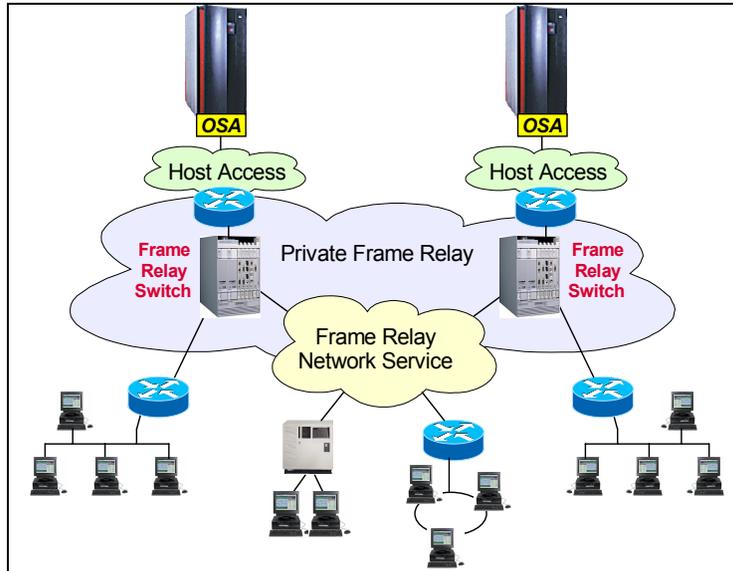


Figure 5-5 Replacing communication controller frame relay with frame switches and routers

Using frame relay switches could enable you to preserve most of your current logical frame relay network design and can help considerably in simplifying the router network design, especially for very large networks. However, introducing new frame relay switching equipment requires new technical staff skills and new management tools.



Important: Communication controllers save host CPU cycles by offloading certain network functions, including SNA device support (also known as the SNA boundary function). Removing communication controllers will consequently result in increased host CPU cycles and storage requirements. The amount of increase will depend on:

- ▶ The number of devices in the network and their transaction rates
- ▶ Whether a dependent LU requester (DLUR) is implemented in the network to support SNA devices.

Alternatives for handling SNA boundary function workload are discussed in 5.3.3, “SNA subarea addressing, routing, and boundary function (BF)” on page 162.

IP-based infrastructure

Figure 5-6 shows an example of migrating your private frame relay to an IP-based infrastructure. The frame handler function of the communication controllers could be replaced using IP routing and those routers can also deliver the appropriate traffic into the host access infrastructure. The host access infrastructure can, in turn, be any of the host LAN access alternatives discussed in 4.2.2, “Token ring (and duplicate addressing)” on page 108 or in 4.2.3, “Ethernet, FDDI, and ATM” on page 124, whichever is most appropriate for your environment.

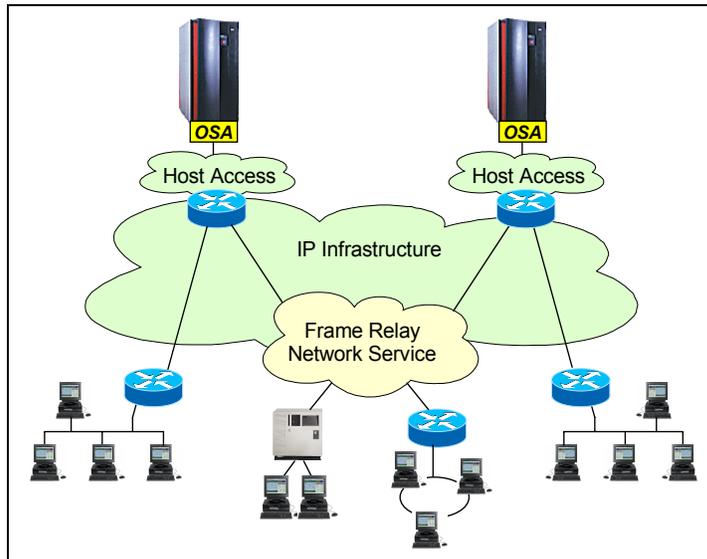


Figure 5-6 IP Infrastructure alternative to frame relay

This alternative could have a substantial impact on your router network design and configuration because, with the communication controller frame relay frame handler function, the communication controllers handled much of the connectivity requirements of the network and provided the routers with a simpler, logical view, of the network.



Important: Communication controllers save host CPU cycles by offloading certain network functions, including SNA device support (also known as the SNA boundary function). Removing communication controllers will consequently result in increased host CPU cycles and storage requirements. The amount of increase will depend on:

- ▶ The number of devices in the network and their transaction rates
- ▶ Whether a dependent LU requester (DLUR) is implemented in the network to support SNA devices.

Alternatives for handling SNA boundary function workload are discussed in 5.3.3, “SNA subarea addressing, routing, and boundary function (BF)” on page 162.

Recommendation

A homogeneous IP-router-based network infrastructure is attractive because you may be able to reduce the number of different networking technologies that your organization must support. However, for very large networks, a switched frame relay network may provide better scalability.

5.2.4 Integrated Services Digital Network (ISDN)

Integrated Services Digital Network (ISDN) is a digital end-to-end telecommunication network that supports multiple services including both voice and data. ISDN services are used in public and private network architectures. Starting with NCP V7R5, the 3746 Model 900 connectivity subsystem (CSS) supports an interface into ISDN environments for SNA sessions routed through the CSS from other communication controllers and some routers. The NCP ISDN function supports only the Euro-ISDN standard.

Your best alternative for migrating from communication controller ISDN support is to use routers with ISDN interfaces.

5.2.5 LAN 802.2, 802.3, and 802.5

The NCP software provides protocol support for the communication controller LAN interfaces. Functional alternatives to communication controller LAN interfaces are discussed in 4.2.2, “Token ring (and duplicate addressing)” on page 108 or in 4.2.3, “Ethernet, FDDI, and ATM” on page 124.

5.2.6 Start-stop terminal connection to TCAM applications

NCP supports the connection of start-stop terminals specifically to TCAM applications.

Note: For VTAM, the start-stop protocol must be converted or enveloped into an SNA data stream using protocol conversion such as that provided by:

- ▶ An IBM 3708 or other protocol converter
- ▶ The NTO program product
- ▶ The NCP Packet Switching Interface (NPSI) program product (through its PCNE 3767 support)

Since TCAM is SNA-based, one can reach TCAM from any SNA device. Consequently, if you still have start-stop devices connecting into TCAM today, there are probably reasons that your organization is trying to preserve those devices. For example, they may be some sort of specialized hardware device. If so, your best alternative for preserving start-stop devices is protocol conversion.

Start-stop terminals can be supported by a protocol converter that converts start-stop to SNA 3767 protocol. One example of a product that appears to provide such support is the INETCO Connect product from INETCO Systems, at:

<http://www.inetco.com>

5.2.7 Synchronous Data Link Control (SDLC)

The NCP software provides SDLC protocol support for communication controller line interfaces. Functional alternatives to communication controller line interfaces are discussed in 4.2.1, “Communication lines” on page 106.

5.2.8 X.25 (with licensed support feature on 3746)

Starting with NCP V7R4, NCP running in a communication controller with a 3746 Model 900 communication subsystem (with the 3746-900 X.25 licensed support feature microcode) can connect to X.25 SNA devices through a public switched data network (PSDN) *without* the need for the X.25 NCP Packet Switching Interface (NPSI) Program Product. Only SNA Qualified Logical Link Control (QLLC) Data Terminal Equipment (DTE) connections are supported by this function. Non-SNA devices still require the NPSI Program Product.

Note: X.25 is also supported by the X.25 NCP Packet Switching Interface (NPSI) product. If you have an LLC type 3, QLLC, connection attached to a port on a 3746-900, you can choose whether that port is supported by NCP and NPSI or by NCP and the X.25 feature microcode. However, LLC types 0, 2, 4, and 5 are only supported by NPSI.

For further information on NPSI support of X.25, see Chapter 6, “X.25 NCP Packet Switching Interface (NPSI)” on page 187.

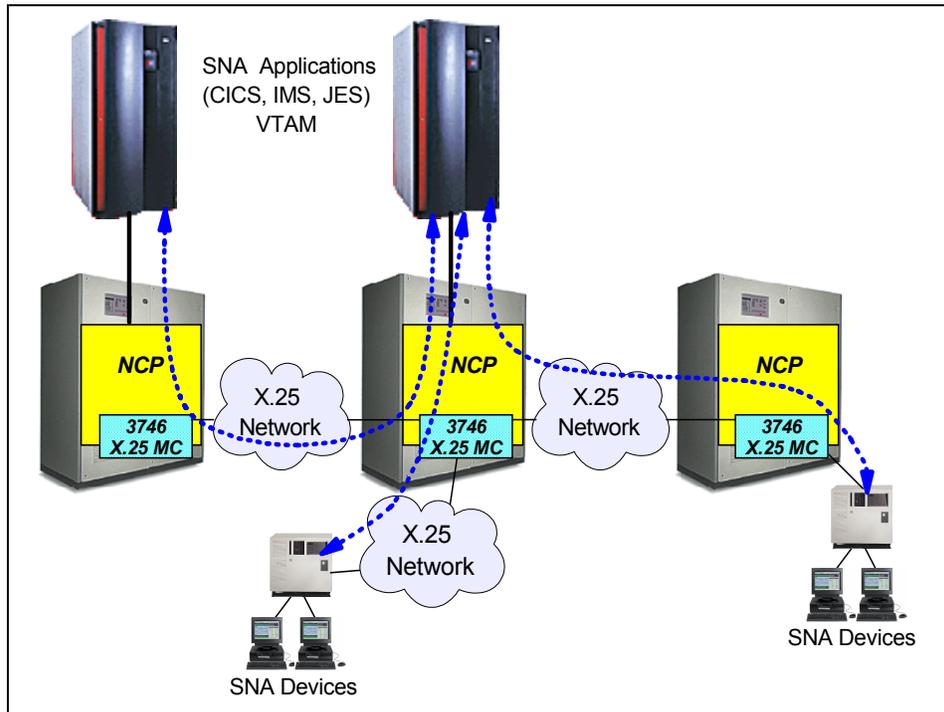


Figure 5-7 The role of NCP and the 3746 X.25 microcode

As illustrated in Figure 5-7, the NCP software (with the 3746-900 or 3746-950 X.25 licensed support feature microcode) can support:

- ▶ SNA devices (such as terminals and printers) connecting by X.25 network facilities to a communication controller
- ▶ Remote communication controllers connecting by X.25 to data center controllers
- ▶ Data center to data center (including between companies) SNA communication over X.25 networks

Important: In some cases, X.25 connections are used for interconnection between *different organizations*. Consequently, any attempt to change such an X.25 environment will require coordinating the change with those other organizations. The interorganizational coordination of such changes often poses much more difficult challenges than do technical issues. Some approaches that may help include:

- ▶ Strive to work within international and/or industry standards such as message queuing (MQ), electronic data interchange (EDI), and extensible markup language (XML).
- ▶ Work closely with industry consortia and extranets. For example:
 - The automobile manufacturing industry has the Automotive Network Exchange.
 - The IVANS network provides similar interconnectivity for the insurance industry.
- ▶ Take advantage of value-added network service providers by letting them deal with the individual connectivity needs of your trading partners.
- ▶ Explore secure virtual private network (VPN) connectivity across an IP connection or across the Internet.

QLLC is an SNA end-to-end implementation. The SNA logical units are completely unaware of the X.25 component in the network. This means that CICS, TSO, IMS™, printers, terminals, and intelligent devices all behave as if the connections were traditional SNA with channels, SDLC links, etc.

Migrating to router X.25 support

As illustrated in Figure 5-8 on page 155, some router implementations support QLLC. For example, Cisco's Web site (<http://www.cisco.com>) depicts sample configurations in which an organization's QLLC mainframe, mini, or control unit connects over X.25 to a router that supports QLLC. The router then connects to other SNA devices over an SNA link, LAN, or channel, thus providing full end-to-end communication for SNA over X.25 using QLLC without a communication controller.

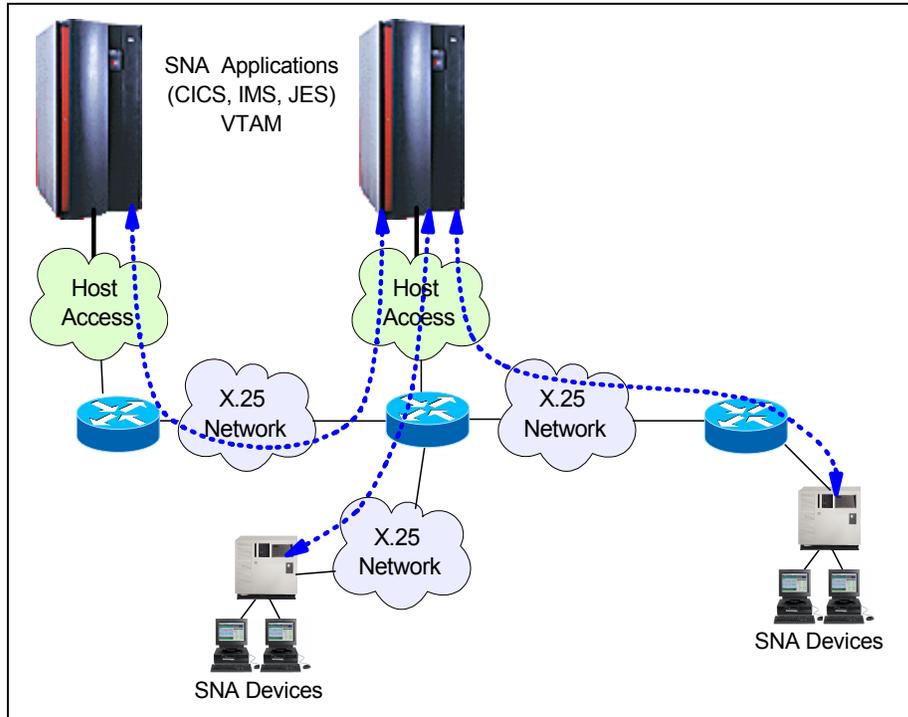


Figure 5-8 Router replacement of NCP support for X.25

Advantages of a router-based solution include:

- ▶ Your organization's X.25 end devices and network can remain in place.
- ▶ You have the flexibility of implementing native SNA or SNA over IP for the non-X.25 portion of the network without using a communication controller.

A consideration of a router-based solution is:

- ▶ Unless you implement DLUR function in the routers, router replacement of communication controllers for X.25 support will increase SNA boundary function requirements either in other communication controllers or in the host.

Note: As you plan to migrate X.25 from your controllers, carefully explore *all* of your wide-area networking connectivity options. While routers may be the logical WAN connectivity devices for your locations, you may find newer service offerings such as frame relay services or virtual private network (VPN) offerings that are more cost-effective than your current network structure. Such service offerings can even include the provision and management of the routers you use to access the network.

Using routers to migrate from X.25

The routed network structure discussed in “Migrating to router X.25 support” on page 154 actually opens up yet another alternative: that of using the routed infrastructure to migrate away from your X.25 network. Illustrated in Figure 5-9, you may be better off migrating from X.25 as you implement your routed network. Note that you may be able to use router QLLC support (such as that provided by routers from Cisco Systems) to avoid requiring changes to your remote X.25 attached devices.

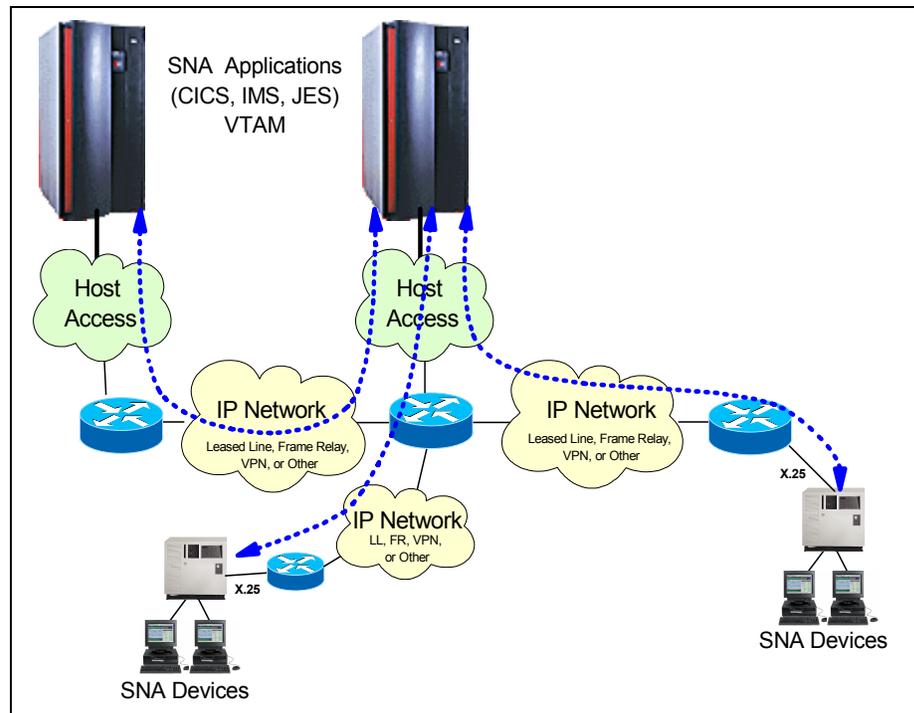


Figure 5-9 Using routers to migrate from X.25

The routers can use a technology such as DLSw or SNA Switching Services to deliver the SNA traffic to the host systems across the routed IP network infrastructure.

An advantage of a router-based migration from X.25 is:

- ▶ You can migrate toward a more strategic wide-area networking infrastructure with minimal impact to devices in the network.

A consideration of a router-based migration from X.25 is:

- ▶ Unless you implement DLUR function in the routers, router replacement of communication controllers will increase SNA boundary function requirements either in other communication controllers or in the host.

Connecting into an Integrated Communication Adapter (ICA)

Certain IBM mainframes, such as the 9221, support an integrated communications adapter (ICA) hardware feature. The ICA supports binary synchronous communication (BSC), synchronous data link control (SDLC), and X.25 link attachments directly through mainframe systems software, such as VTAM or BTAM. This means that the ICA supports X.25 links into the mainframe without NPSI or NCP (as shown in Figure 5-10). The X.25 support is provided directly by VTAM under VM and VSE only and is for SNA X.25 QLLC only. For additional information, see the *VTAM Network Implementation Guide*, SC31-6494.

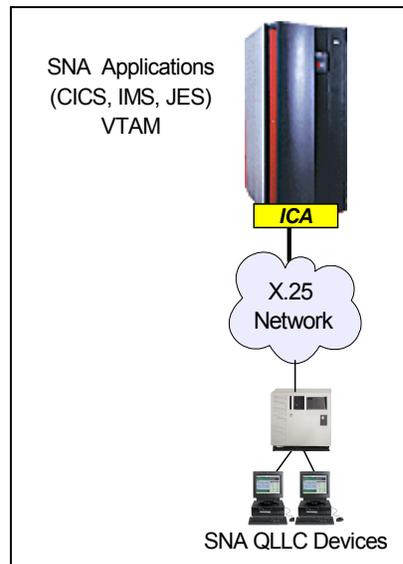


Figure 5-10 ICA X.25 support

Note: One of the key reasons for the development of communication controller products was to offload communication line handling from the host. If you are moving lines from a communication controller to an ICA, be careful to consider the possible increase in host CPU cycles it may require.

Advantages of an ICA-based solution include:

- ▶ Your organization's X.25 end devices and network can remain in place.
- ▶ It provides mainframe SNA application access with no need for a communication controller or router.

Considerations of an ICA-based solution include:

- ▶ The last S/390 system to support an ICA was the 9221, which is no longer marketed by IBM.
- ▶ The 9221 is not supported by the current VM operation system (z/VM®). It is supported by VM/ESA; however, VM/ESA is no longer marketed by IBM.
- ▶ Only VM and VSE systems support the ICA, so if you are running applications on z/OS, you will need to ensure connectivity from the system with the ICA to the application LPAR.
- ▶ Because ICAs support limited numbers and speeds of links when compared with a communication controller, scalability may be an issue with an ICA solution.
- ▶ One of the key reasons for the development of communication controller products was to offload communication line handling from the host. By moving lines from a communication controller to an ICA, you may experience a substantial host CPU cycle increase.
- ▶ ICA replacement of communication controllers for X.25 support will increase SNA boundary function requirements in the host.

Recommendation

A router-based solution will probably be your most cost-effective and strategic alternative. However, if you already have an ICA and the available ports and speeds are sufficient for your requirements, then the ICA may be an acceptable alternative.

5.3 Advanced capability functions and alternatives

Following are the key advanced NCP capabilities discussed in this section:

- ▶ SNA class of service (COS)
- ▶ Multi-link transmission group (MLTG)
- ▶ SNA subarea addressing, routing, and boundary function (BF)
- ▶ SNA Network Interconnection (SNI)
- ▶ APPN (and LEN) Composite Network Node (CNN)
- ▶ IP routing
- ▶ Extended Recovery Facility (XRF)

5.3.1 SNA class of service (COS)

SNA class of service (COS) designates the transport network characteristics of a session. It includes such characteristics as security, transmission priority, and bandwidth. For SNA sessions, the class of service must be determined before route selection can take place. During session initiation, the class of service for the session is derived from a mode name specified in the session-initiation request. The route then selected for the session depends on the class of service for that session.

In an SNA network, different classes of service can be specified, based upon the needs of the end users in the network. A range of classes of service can be provided to accommodate each users' session requirements. For example, the following classes of service can exist in an SNA network:

- ▶ A class that provides response times suitable for high-priority interactive sessions
- ▶ A class that provides response times suitable for low-priority interactive sessions
- ▶ A class that provides routes that have the best availability
- ▶ A class suitable for batch processing
- ▶ A class suitable for high-security transmissions

The original concepts of SNA COS were pioneered by IBM almost 30 years ago. Much has changed in networking since then. For one thing, there is greater variety and better availability of high-speed networking technologies, making the practice of mixing very diverse network links (such as mixing satellite links with terrestrial links) less common in today's networks. For another, SNA traffic is no longer the predominant traffic on the network; rather, SNA is just one of the protocols using the network. Consequently, COS-based route selection may not be as important as it once was; however, COS-based transmission priority can still be critical to certain SNA networking environments.



Important: SNA COS-based transmission priority may be providing critical function in your network. SNA COS ensures consistent, low, response times for interactive traffic in networks where:

- ▶ There are links that carry a combination of SNA batch (such as print or file transfer) and interactive (such as credit card, banking, or insurance transactions) traffic
- ▶ There are times during which the total traffic exceeds the available network bandwidth on such a link

Depending upon the functional capabilities of the router network, organizations that have migrated such links to routers have encountered significant variability in interactive response times during peak traffic times. There are two options for resolving such issues:

- ▶ Implement support for SNA COS-based transmission priority (for example, Cisco SNA Switching Services (SNASw) supports SNA COS)
- ▶ Increase available bandwidth on network links so that the total traffic rarely exceeds the available bandwidth

For most networking protocols such as TCP/IP, it is pretty easy to figure out the appropriate priority for a particular frame of traffic. Each frame header indicates the kind of traffic that is carried in the frame. With TCP/IP, for example, each TELNET (terminal emulation) frame has a TCP port number of 23 in its header, while FTP (file transfer) data frames use port 20, and SMTP (e-mail) uses port 25. Through configuration, you tell a router to which priority queue each kind of traffic should go. By looking into a frame header for the port number (or equivalent), the router can put each frame into the appropriate queue.

SNA protocols, however, can present a challenge for routers. Unlike other protocols, SNA is connection-oriented: That is, a connection (also called a session) is always established at the start of any data exchange. Once the session is established, SNA frames carry only a session identifier and no indication of the kind of traffic being sent. That means for a router to tell the difference between interactive terminal traffic and batch print or file transfer traffic, it must either:

- ▶ Participate in the SNA protocols used in session establishment and keep track of the appropriate transmission priorities of individual sessions, or
- ▶ Have the appropriate transmission priority determined outside of the network and indicated in each frame, which is what Enterprise Extender (EE) does.

It is important to note that, without either APPN routing code in the network or the use of EE at the edges of the network, the best that a router can do is to recognize that a particular frame is SNA traffic and treat all SNA traffic (both batch and interactive) the same. Consequently, if SNA COS does provide critical function in your network, consider implementing either APPN in the network—such as that implemented in Cisco SNASw—or EE (discussed in Chapter 19, “Enterprise Extender (EE)” on page 331).

5.3.2 Multi-link transmission group (MLTG)

An MLTG is a logical group of physical links for SNA traffic between two communication controllers. An MLTG can include various transmission media using different data link protocols (frame relay, SDLC, ISDN B-channel) and token-ring LANs. The traffic is automatically distributed over the physical links of the MLTG. If a physical link fails, MLTG provides automatic and non-disruptive data re-routing over other links of the MLTG.



Note: The IBM Communication Controller for Linux on System z9 and zSeries (CCL) product enables you to run NCP on your mainframe and supports token-ring and Ethernet LAN MLTGs. See Chapter 16, “Communication Controller for Linux on System z9 and zSeries (CCL)” on page 283 for additional information about CCL.

MLTG provides value in two ways:

- ▶ MLTG enables the aggregation of bandwidth in your network.

In the past, some organizations encountered situations where they could not get a telecommunications link between two locations with sufficient bandwidth to meet the needs of a particular application. This happened, for example, if an organization was sending very large print jobs to a printing center. MLTG provided them a way to aggregate multiple links into a single logical high-bandwidth link. Such issues are less of a problem now because of the relatively recent emergence of high-speed networking services.

- ▶ MLTG provides higher availability in case of a link outage.

The ability to use multiple links in an MLTG, and even links of different media including frame relay, SDLC, ISDN B-channel, and token-ring LANs, enables you to build very highly available networks.



Important: MLTGs will not work across Data Link Switching (DLSw) networks because the local acknowledgement of logical link control, type 2 (LLC2) protocols interferes with MLTG sequence numbering. If you are using MLTG to provide higher bandwidth, then you may be able to simply use a higher-bandwidth link from your carrier. If you are using MLTG to improve availability, then similar availability can be achieved through redundancy in the DLSw network.

Today, however, a high-speed IP-based infrastructure built on routing and switching technologies can provide orders-of-magnitude greater throughput than is possible using communication controllers. Also, your IP-based infrastructure can be designed to provide very high availability for all traffic in your network rather than just your SNA traffic. A model high-speed host-to-host networking infrastructure is discussed in “OSA-Express gigabit Ethernet for host-to-host communication” on page 130. Key technologies for building your high-speed IP-based infrastructure and leveraging it for your SNA traffic include:

- ▶ The Net390 architecture (discussed in Chapter 14, “Net390 architecture” on page 261)
- ▶ OSA-Express (discussed in Chapter 15, “OSA-Express” on page 269)
- ▶ Enterprise Extender (discussed in Chapter 19, “Enterprise Extender (EE)” on page 331)

5.3.3 SNA subarea addressing, routing, and boundary function (BF)

Perhaps the single most important function of IBM communication controllers is their support of subarea SNA networking. The controller hardware, and the NCP software that runs in it, were developed from the start to implement the SNA architecture Physical Unit type 4 (PU 4) subarea node. As shown in Figure 5-11 on page 163, PU 4 subarea nodes interact with host nodes (PU 5) and peripheral nodes (PU 2 and PU 2.1) to provide the SNA functions that route and control the flow of data in a subarea network.

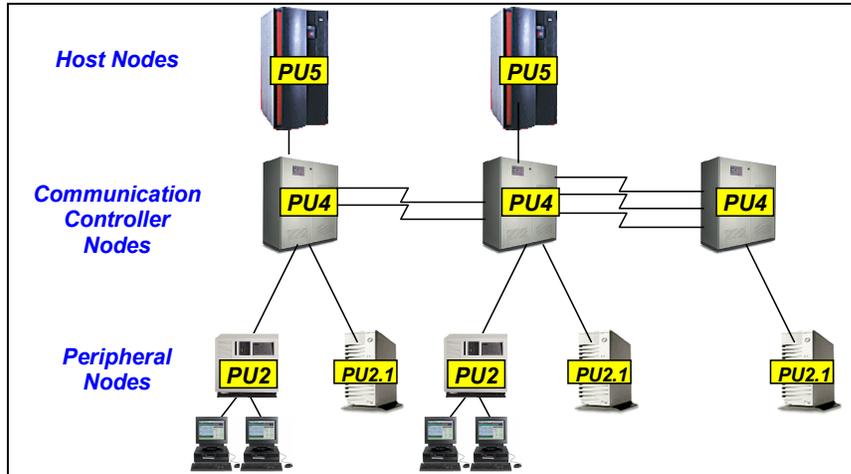


Figure 5-11 The SNA architecture of a subarea network

Communication controllers play key roles in three particular areas of subarea SNA networking:

- ▶ Addressing
- ▶ Routing
- ▶ Boundary function (BF)

Each of these is discussed further below.

Addressing

Network addresses uniquely identify the system services control points (SSCPs), logical units (LUs), physical units (PUs), links, and link stations in a subarea network. The address of each network addressable unit (NAU) consists of a subarea address and an element address. Subarea addresses are used to route message units between subareas. Element addresses identify network resources within a subarea.

Note: Whereas each subarea address is a unique number in the network, an element address is unique only within each subarea. This is the primary reason that it has always been very difficult for routers to examine SNA messages and make appropriate routing decisions. Consequently, it is often said that SNA is an “unroutable protocol” except, of course, for communication controllers, which have been routing SNA traffic for 30 years.

NCP initially assigns element addresses during system generation. The sequence in which NCP encounters the resource-definition statements for

network resources determines the order in which it assigns the element addresses for the resources.

SNA addressing uses a two-byte element address—thereby supporting up to 65,535 NAUs in a subarea. Consequently, for large organizations, the communication controller can play an important role in subdividing the network into groups of less than 65,535 NAUs.



Important: Large networks have encountered problems such as SNA session setup failures or the inability to activate additional resources as the result of running out of SNA element addresses in a VTAM subarea. Migrating devices from your communication controllers can add to this problem, because the device addresses can end up moving from the NCP subarea into a VTAM subarea. See Appendix C, “Avoiding VTAM network addressing problems” on page 403 for additional information.



Note: The IBM Communication Controller for Linux on System z9 and zSeries (CCL) product enables you to run NCP on your mainframe, which gives you a way to migrate devices from your IBM 3745 Communication Controller hardware without having to move device addresses from your NCP subarea into a VTAM subarea. See Chapter 16, “Communication Controller for Linux on System z9 and zSeries (CCL)” on page 283 for additional information about CCL.

In order to mitigate element address scalability limitations, VTAM implemented the “enhanced addressing” function to support network addresses greater than 65,535 (called *high-order network addresses*). VTAM enhanced addressing was first implemented in VTAM V4R2, and has been expanded in subsequent releases. Depending on the release, VTAM uses these addresses for various types of session-capable resources (APPLs, LUs and CDRSCs) and non-session-capable resources (LINEs and PUs). See Appendix C, “Avoiding VTAM network addressing problems” on page 403 for additional information.



Important: VTAM cannot use high-order element addresses for dependent LUs that attach via an XCA major node such as when native SNA traffic is “bridged” into VTAM through a router or an Open Systems Adapter (OSA).

The best ways to take advantage of VTAM enhanced addressing are to:

- ▶ Migrate users to TN3270
- ▶ Implement DLUR for peripheral devices

- ▶ Migrate hosts to APPN so that more sessions appear to use single-domain routes and both network addresses will be assigned by the same VTAM
- ▶ Migrate APPN connectivity to Enterprise Extender (EE) so that, in addition to using more high-order addresses for sessions, the EE Logical LINEs and PUs can also use high-order addresses (as opposed to other types of LINEs and PUs which still use low-order addresses)

See Appendix C, “Avoiding VTAM network addressing problems” on page 403 for more information.

Routing

As mentioned above, communication controllers are SNA routers. Messages are first sent to the appropriate SNA subarea, then the element address is used to send it to the appropriate device. In multi-host environments, if the communication controller has connectivity to the destination host (for example, channel attachment), NCP can route an SNA message directly to that host without requiring the message to transit to any other host.



Important: In some cases, removal of a communication controller will cause its routing responsibility to move into a host, costing extra host cycles for SNA routing and creating session dependencies on the availability of the routing host.

In the SNA architecture, all communication controller nodes and host nodes can perform subarea SNA routing. Advanced Peer-to-Peer Networking, a mid-1980s enhancement of SNA, enabled certain peripheral devices (PU 2.1) to also route SNA traffic.

In order to provide optimal routing for your SNA traffic, consider using APPN (with or without Enterprise Extender, which is discussed in Chapter 19, “Enterprise Extender (EE)” on page 331). APPN is supported by some routers such as those from Cisco Systems.

Boundary function (BF)

Subarea nodes to which peripheral nodes are attached perform a boundary function for their attached peripheral devices. Boundary function includes:

- ▶ Mapping from the subarea addressing (subarea/element address pair) used on links between subareas (INN links) to the local addressing used on peripheral links.
- ▶ Providing protocol support for attached peripheral nodes, including performing session sequence numbering, providing session-level pacing

support, and segmenting messages for links with large error rates or nodes with limited buffer space.



Important: Communication controllers save host CPU cycles by offloading certain network functions, including SNA boundary function. Removing communication controllers will consequently result in increased host CPU cycles and storage requirements. The amount of increase will depend on:

- ▶ The number of devices in the network and their transaction rates
- ▶ Whether a dependent LU requester (DLUR) is implemented in the network to support SNA devices.



Note: The IBM Communication Controller for Linux on System z9 and zSeries (CCL) product enables you to run NCP on your mainframe, which gives you a way to migrate devices from your IBM 3745 Communication Controller hardware while preserving your NCP boundary function. Although using CCL might increase the host CPU cycles overall, CCL CPU cycles may be less expensive than other mainframe CPU cycles because Linux partitions can use Integrated Facility for Linux (IFL) processors. For additional information about CCL, see Chapter 16, “Communication Controller for Linux on System z9 and zSeries (CCL)” on page 283.



Important: For devices that require SNA boundary function support, *every* message to or from them *must* pass through their respective boundary functions. Consequently, if your migration from a communication controller results in boundary function moving into a host, all SNA messages from devices that use that boundary function must transit that host even when the session is to an application in a host in a completely different data center. Such inefficient session routing can cause both performance and availability problems.

The best ways to avoid boundary function-related problems are to:

- ▶ Migrate users to TN3270 (discussed in detail in Chapter 18, “TN3270” on page 311)
- ▶ Implement DLUR for peripheral devices (discussed in “Dependent LU Requester (DLUR) placement” on page 341)
- ▶ Preserve your NCP boundary function using CCL (discussed in Chapter 16, “Communication Controller for Linux on System z9 and zSeries (CCL)” on page 283)

5.3.4 SNA Network Interconnection (SNI)

The SNA network interconnection (SNI) function was developed in the early 1980s to:

- ▶ Solve network address space scalability problems by enabling the subdivision of large networks that were running out of SNA addresses
- ▶ Facilitate controlled interconnection of separate SNA networks such as those of separate organizations.

SNI enabled the automation of intercompany exchanges of information such as interbank check clearing or the transfer of payroll data from an employer to a bank for employee direct deposit. The whole of the Electronic Data Interchange (EDI) industry (the predecessor to today's e-commerce) developed with value-added network (VAN) service providers utilizing SNI capabilities to enable electronic transfers of information such as inventory automation for manufacturing and retail between companies. Outsourcing companies (the predecessor to today's application service providers, or ASPs) emerged using SNI to connect to client companies to provide services such as payroll services, high-volume printing and mailing services, or patient admissions, discharge, and transfer management for hospitals.

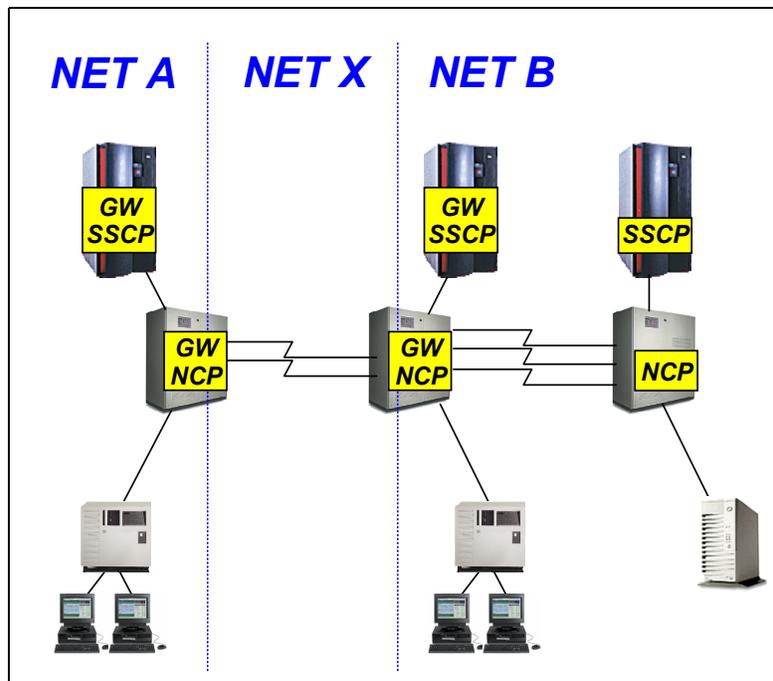


Figure 5-12 SNA Network Interconnection (SNI)

As shown in Figure 5-12 on page 167, the communication controller provides line concentration for SNI links and runs a function called the gateway NCP (GWNCP) which works with a gateway SSCP (GWSSCP) to authorize session initiations and to provide any necessary name and address translations. In the figure, NET A and NET B are the networks of two independent organizations while NET X is an empty (or “null”) network that provides insulation between the two networks. Essentially, your SNI gateway NCPs are your organization’s “firewalls” for SNA networking: providing privacy and security for your network but remaining transparent to the SNA search process and to LU-LU sessions. In order to use SNI, there must be at least one GWNCP and one GWSSCP. If SNI is being used to interconnect autonomous organizations, each organization will usually want to control its own privacy and security and, therefore, have its own GWNCPs and GWSSCPs, as in Figure 5-12 on page 167.

Important: SNI is predominantly used for interconnection between *different* organizations. Consequently, attempting to change your SNI environment may require coordinating the change with those other organizations. The inter-organizational coordination of such changes often poses much more difficult challenges than do technical issues. Some approaches that may help include:

- ▶ Strive to work within international and/or industry standards such as message queuing (MQ), electronic data interchange (EDI), and extensible markup language (XML).
- ▶ Work closely with industry consortia and extranets. For example:
 - The automobile manufacturing industry has the Automotive Network Exchange.
 - The IVANS network provides interconnectivity for the insurance industry.
- ▶ Take advantage of value-added network services providers by letting them deal with the individual connectivity needs of your trading partners.
- ▶ Explore the use of virtual private network (VPN) connectivity across a private IP connection or across the Internet.

Important: The VTAM session management exits are mechanisms for executing user-written programs during session establishment in order to facilitate security checking and the collection of charge-back information if desired. Many organizations have developed such exit code. If you are migrating from SNI in your network, you must carefully consider the impact of losing such functions.



Note: The IBM Communication Controller for Linux on System z9 and zSeries (CCL) product enables you to run NCP on your mainframe, which gives you a way to migrate SNI connectivity from your IBM 3745 Communication Controller hardware without having to fundamentally change your SNI environment. Most important, *CCL minimizes the requirement to coordinate changes with SNI partner organizations*. The CCL alternative is discussed in “Continue to run NCP using CCL” on page 169.

Alternatives for migrating from your 3745 communication controller SNI function include:

- ▶ Continue to run NCP—supporting your SNI connectivity—using CCL
- ▶ Converting to SNI adjacent network connectivity
- ▶ Migrating to APPN Border Node (BN)

Each of these is discussed below. Because SNI is such a difficult function from which to migrate, we also discuss ideas for consolidating SNI as a way to reduce costs and simplify your environment, though not totally migrating from SNI.

Important: In this part of the book, we identify *functional alternatives* to key IBM 3745 Communication Controller capabilities. For example, APPN Border Node (BN) provides a direct alternative to SNI. However, wherever practical, you should try to migrate from the use of SNA altogether. In certain cases, such as with many currently available commercial file transfer software packages, the change from SNA to TCP/IP can be accomplished just by reconfiguring the software at each end of the communication.

SNA application migration to TCP/IP is discussed further in 17.4.1, “TCP/IP convergence” on page 305.



Continue to run NCP using CCL

The IBM Communication Controller for Linux on System z9 and zSeries (CCL) product allows you to continue to use SNI connectivity by running NCP on your mainframe. CCL provides the *only* alternative that will allow you to migrate SNI connectivity from your IBM 3745 Communication Controller hardware *without having to coordinate changes with your SNI partners (other than perhaps an outage window for the migration)*. The CCL product is discussed further in Chapter 16, “Communication Controller for Linux on System z9 and zSeries (CCL)” on page 283. Figure 5-13 on page 170 illustrates how CCL can be used to support SNI connectivity.

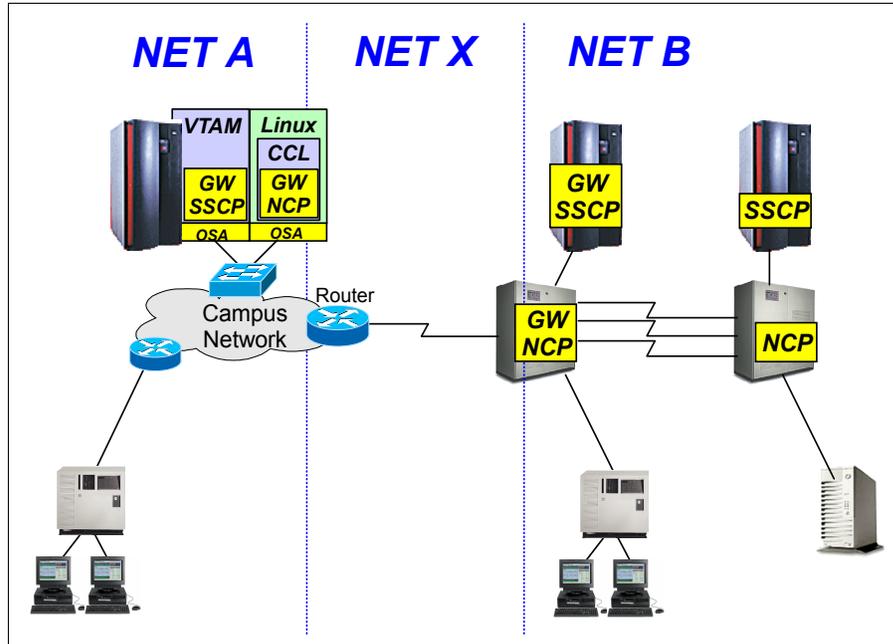


Figure 5-13 Continue to run NCP using CCL

In the figure, serial connectivity to a local device (lower left) has been migrated to a link aggregation router. Migration of 3745 Communication Controller serial connectivity to aggregation routers is discussed further in 16.3, “What you need to understand about CCL” on page 286.

Similarly, wide area network (WAN) link connectivity to the NET B GWNCP has been migrated to a router (not necessarily a different router from the one used for the local device). Adding another router (not shown in the figure) at the NET B site and using Data Link Switching (DLSw) over the WAN link presents an attractive additional option to this configuration because:

- ▶ It allows NET B to consolidate serial links and direct the SNI traffic either through a LAN interface on their 3745 or into a CCL GWNCP of their own.
- ▶ It creates network connectivity that can carry both SNA and IP traffic and, therefore, supports current SNA applications while enabling the incremental migration of application traffic (such as file transfer) from SNA to IP.



Important: MLTGs will not work across Data Link Switching (DLSw) networks because the local acknowledgement of logical link control, type 2 (LLC2) protocols interferes with MLTG sequence numbering. If you are using MLTG to improve availability, then similar availability can be achieved through redundancy in the DLSw network. If you are using MLTG to provide higher bandwidth, then you may be able to simply use a higher-bandwidth link from your carrier.



Important: DLSw cannot support SNA class of service because it has no way to distinguish between different types of SNA traffic. Therefore, if the SNI connection concurrently carries both batch and interactive traffic then, in order to prevent unacceptable delay for high-priority traffic (such as interactive traffic), you must ensure that sufficient network bandwidth is available so that the total traffic rarely exceeds the available bandwidth. (For a more detailed discussion of SNA class of service support, see 5.3.1, “SNA class of service (COS)” on page 159.)



Important: If both you and your trading partner are using CCL, CCL V1.2 adds an interesting new SNI connectivity alternative called IP TGs. With IP TGs, CCL-to-CCL traffic is encapsulated in TCP/IP frames for transport over IP networks. Additionally, SNA COS-based prioritization can be enabled through the use of configurable TCP ports or TOS specification.

CCL IP TG support is discussed further in 16.3.4, “CCL V1.2 IP TG support” on page 294.

Advantages of a CCL-based SNI solution include:

- ▶ CCL provides the *most transparent* migration path from IBM 3745 Communication Controllers.
- ▶ *CCL minimizes change* to the current network management environment.
- ▶ The CCL solution allows you to migrate SNI connectivity from IBM 3745 Communication Controller hardware without having to fundamentally change your SNI environment (minimizing the requirement to coordinate changes with SNI partners).
- ▶ CCL, in conjunction with DLSw or IP TG support, allows you to leverage strategic IP-based connectivity and enables the incremental migration of applications (such as file transfer) from SNA to IP.

Considerations of a CCL-based SNI solution include:

- ▶ Care must be taken with regard to the use of MLTGs and SNA class of service in a CCL solution because the required use of router technology may interfere with those functions. (If both you and your trading partner are using CCL, IP TG support can eliminate this concern.)
- ▶ CCL runs in a Linux for zSeries environment and may require additional skills and mainframe system resources (such as logical partitions and CPU cycles) to support it.

Note: CCL CPU cycles may be less expensive than other host CPU cycles because Linux partitions can use Integrated Facility for Linux (IFL) processors.

SNI adjacent network connectivity

In some cases, you may be able to simply remove your GWNCP and connect directly to your SNI partners' GWNCP. If that SNI partner is a value-added network (VAN) service provider, they may even be able to provide connections to multiple other SNI partners for you. As shown in the example in Figure 5-14, routers can be used to transport the SNA messages between NET A and NET B (using bridging or DLSw) while NET B provides the SNI gateway functions.

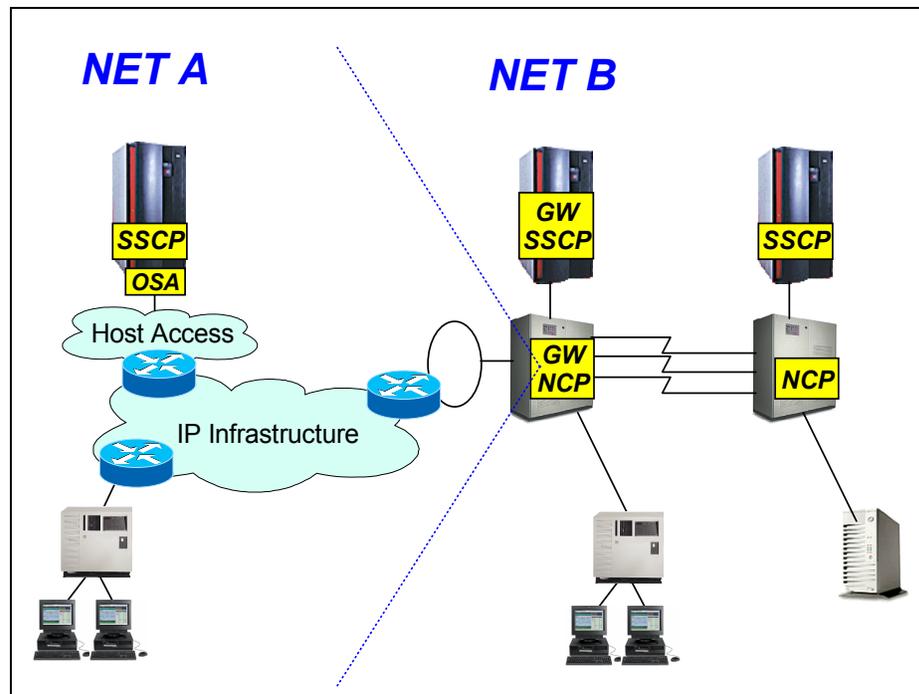


Figure 5-14 SNI adjacent network connectivity

SNI adjacent network connectivity does, indeed, allow you to migrate from your communication controller SNI function but it has some potentially significant limitations:

- ▶ It requires that your trading partner (or VAN) be able and willing to provide the SNI function for you.
- ▶ You must be willing to give up the privacy and security provided by your gateway NCPs and trust the security provided by your SNI partner.
- ▶ There will be an increase in the level of coordinated system definitions between you and your SNI partner. Specifically, to change from the SNI configuration shown in Figure 5-12 on page 167 to that shown in Figure 5-14 on page 172, you would need to:
 - a. Allocate a new subarea in the NET B GWNCP for the NET A network.
 - b. Add PATH definitions in all of the remaining NET A subarea nodes (VTAMs and NCPs) to this new subarea.
 - c. Add PATH definitions in the NET B GWNCP to all of the remaining NET A subareas.

Border Node (BN)

Border Node (BN) is the most direct *replacement* for SNI providing any-to-any SNA connectivity between different SNA networks. While SNI was developed for network-to-network connectivity using the subarea SNA model, BN is based upon APPN. There are a number of technologies associated with the BN architecture which, together, make BN a very significant and important alternative to SNI:

- ▶ Extended Border Node (EBN), a BN that supports an extended subnetwork boundary, a connection between two border nodes over which both nodes realize that a non-native network node exists on the other side, allowing for additional subnetwork hops and increased control over routing. EBN was implemented in VTAM V4R2 and first shipped in 1994. The key benefit of EBN is that it enables multiple network hops for example through a VAN.

The implementation of the VTAM EBN does not mean that your whole network must be migrated to APPN, which has been a concern in the past. The VTAM APPN implementation provides an interface between subarea networking and APPN. This makes it possible to retain a subarea SNA network, with only the SNI links migrated to extended border node links.

Important: VTAM search algorithms are completely different in subarea and APPN networks and the combination of them in a mixed network is quite complex. An understanding of that combination and careful planning are essential to achieve a trouble-free integration of subarea and APPN networking.

- ▶ Enterprise Extender (EE), a technology that enables SNA messages to flow over an IP-based network infrastructure (potentially even using Internet VPNs). EE is discussed in detail in Chapter 19, “Enterprise Extender (EE)” on page 331. Enterprise Extender together with EBN provide a way to migrate to an IP infrastructure and replace SNI function.

Some SNI connections are used for very large transfers of information, such as sending billing statements to an outsourcing company that prints and mails them. While the multi-link transmission group capability of the communication controller gives it a degree of scalability, some companies have pushed that technology to its limits, thus necessitating the need for an even faster communications path between companies. EE can provide that path by moving the traffic from your older communication controller lines to your (typically high-bandwidth) router network connections without requiring changes to your application programs.

- ▶ Global Connection Network (GCN) allows the specification of a connection network that can include multiple SNA networks and that can be used for the session route without requiring the data path to traverse the EBNs, a previous limitation of EBN support. This allows a node within one APPN subnet to utilize the services of an EBN in setting up a session with a partner in another network, but yet be able to route the data on a direct path (via the EE Global Connection Network) to the partner. This removes potential performance bottlenecks at the EBNs while allowing the exploitation of a pervasive IP-based backbone network.



Restriction: While Global Connection Network (GCN) provides an attractive option for inter-organization communication, most organizations secure external communication using firewalls. Those firewalls employ Network Address Translation (NAT) technology to hide their internal network addressing scheme and avoid IP addressing conflicts with their partners. VTAM only supports GCN over Enterprise Extender (EE) links and EE GCNs only work across NAT boundaries with z/OS CS V1R5 or higher.

Figure 5-15 on page 175 shows EBN as an alternative to SNI.

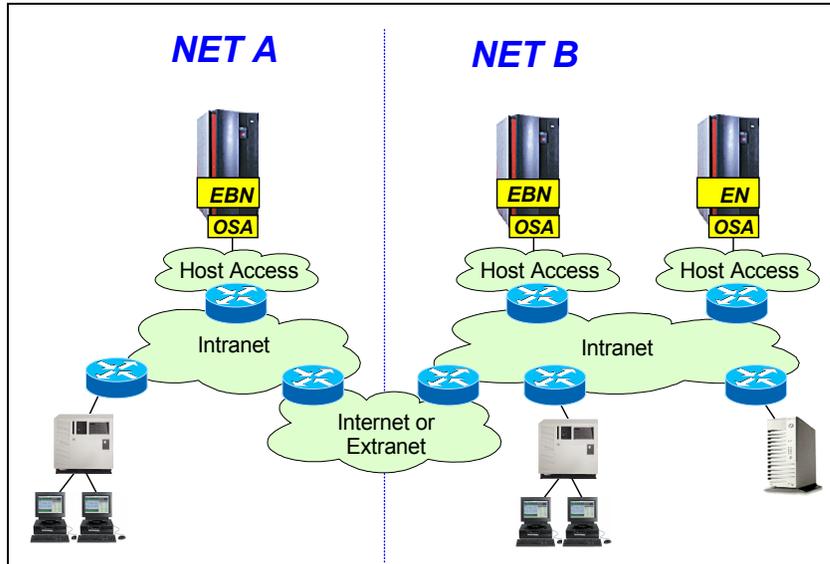


Figure 5-15 Extended Border Node alternative to SNI

Advantages of an EBN solution include:

- ▶ EBN offers the most complete *replacement* alternative for SNI.
- ▶ EBN, in conjunction with EE, enables you to migrate your SNI connectivity to a strategic IP-based infrastructure and potentially significantly improve the performance of cross-network applications.

Considerations of an EBN solution include:

- ▶ EBN is based on APPN technology and some organizations have not implemented APPN.



Restriction: VTAM for VM/ESA and VSE/ESA does not support Enterprise Extender.

TPF implements an APPN EN and, therefore, does not support EBN (a NN function). Also, TPF does not support EE. The best alternative for TPF users is to migrate their connections to IP using either MQSeries for TPF or MATIP.

Consolidating SNI

Consolidating SNI traffic into a few physical interfaces can help to reduce the overall costs of your communication controller environment. Here are some alternatives that could allow you to consolidate your SNI links:

- ▶ Consolidation onto frame relay links (with different virtual circuits terminating at different SNI partners) may allow reduction in the number of 3745s.
- ▶ Consolidation onto token-ring network interfaces with a routed network connecting the LANs together with a technology such as DLSw or bridging (Requires coordination of router *and* SNA internetworking technologies between companies).
- ▶ Consolidating SNI traffic by connecting through a VAN to trading partners.

Recommendation

Unlike any other SNI alternative, the CCL solution allows you to migrate SNI connectivity from IBM 3745 Communication Controller hardware without having to fundamentally change your SNI environment (minimizing the requirement to coordinate changes with SNI partners). Although the EBN solution requires a greater degree of coordinated change with your partners than does the CCL solution, it provides the most robust replacement alternative for SNI, including SNA class of service support.

Both the EBN solution (when used in conjunction with EE) and the CCL solution (when used in conjunction with DLSw or IP TG support) allow you to leverage strategic IP-based connectivity for SNA traffic and thereby enable the incremental migration of applications from SNA to IP.

5.3.5 APPN (and LEN) Composite Network Node (CNN)

As shown in Figure 5-16 on page 177, VTAM and NCP can work together to support communication from applications running in APPN (PU type 2.1) hosts to:

- ▶ Host-based subarea SNA applications or
- ▶ Applications running in other PU 2.1 nodes (which could be mainframe hosts, midrange systems such as AS/400s, or Intel or RISC-based servers)

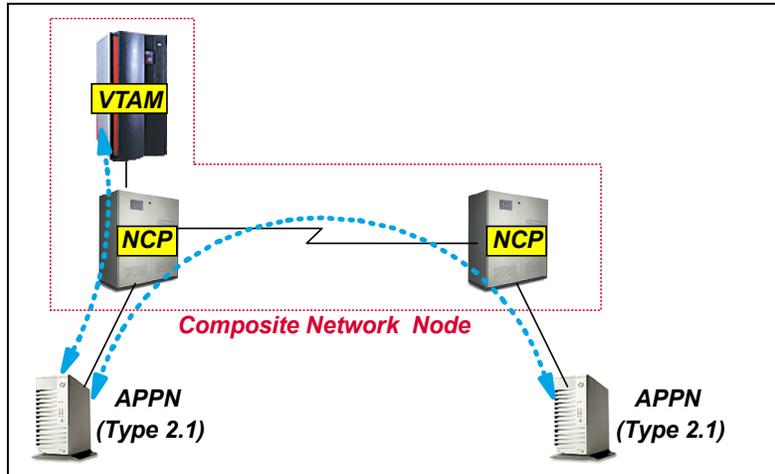


Figure 5-16 NCP APPN Composite Network Node support

This configuration is called Composite Network Node (CNN) because VTAM and NCP *together* form a single APPN Network Node. Originally, CNN support was important not only because it supported communication with PU 2.1 hosts but also because it allowed the use of the subarea SNA backbone network as a transport for communication between different PU 2.1 hosts. Much has changed in networking, however, since the VTAM and NCP support of PU 2.1 connectivity was first shipped in 1987 including:

- ▶ The emergence of IP as a strategic and mission-critical networking protocol. SNA messages are no longer the predominant traffic on the network; rather, SNA is just one of the protocols using the network.
- ▶ The development of multi-protocol routers capable of simultaneously routing IP traffic while also handling other protocols such as SNA.
- ▶ The ongoing enhancement of SNA networking protocols including, most importantly, Enterprise Extender (EE), which allows SNA traffic to use IP-based networking infrastructure. (EE is discussed in greater detail in Chapter 19, “Enterprise Extender (EE)” on page 331)

Today, you have two alternatives for migrating from the CNN function supported by your NCP:

- ▶ Using APPN or DLSw function in your routers to move the application traffic between APPN nodes in your network
- ▶ Using EE to allow your APPN host application traffic to use your IP infrastructure

Each of these alternatives are discussed below.

Note: Composite Network Node is also supported with NCP running on your mainframe using the Communication Controller for Linux on System z9 and zSeries (CCL) product; however, such a configuration does not appear to be useful in the context of supporting PU type 2.1 communications either to host applications or between external PU type 2.1 nodes.

Using APPN or DLSw function in your routers

Figure 5-17 shows the use of routers with APPN or DLSw function to carry traffic between APPN hosts.

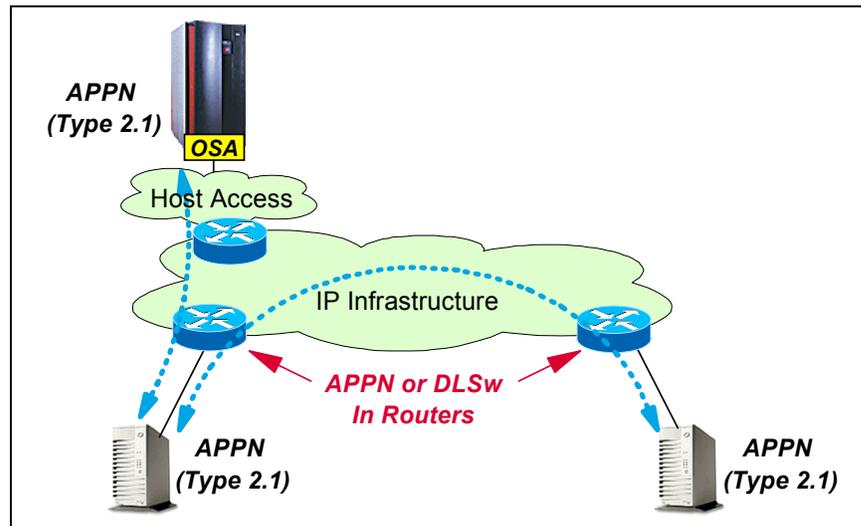


Figure 5-17 Using APPN or DLSw function in your routers

The CNN function provided by VTAM and NCP could be replaced by APPN support in your routers (for example Cisco's SNA Switching Services). Such a change would have minimal impact on your APPN hosts and would enable the consolidation of your SNA and IP wide-area links, perhaps allowing bandwidth upgrades (due to the cost savings) of your WAN and even enhancing the performance of your applications. Alternatively, you could use DLSw function in your routers to essentially bridge the SNA traffic over your IP network; however, it is important to realize that you would lose SNA class of service support for those links in the router network (this limitation is discussed in detail in 5.3.1, "SNA class of service (COS)" on page 159).

Using Enterprise Extender (EE) in your APPN hosts

Figure 5-18 shows the use of Enterprise Extender function in your APPN hosts to allow them to use an IP infrastructure as a transport for their application traffic.

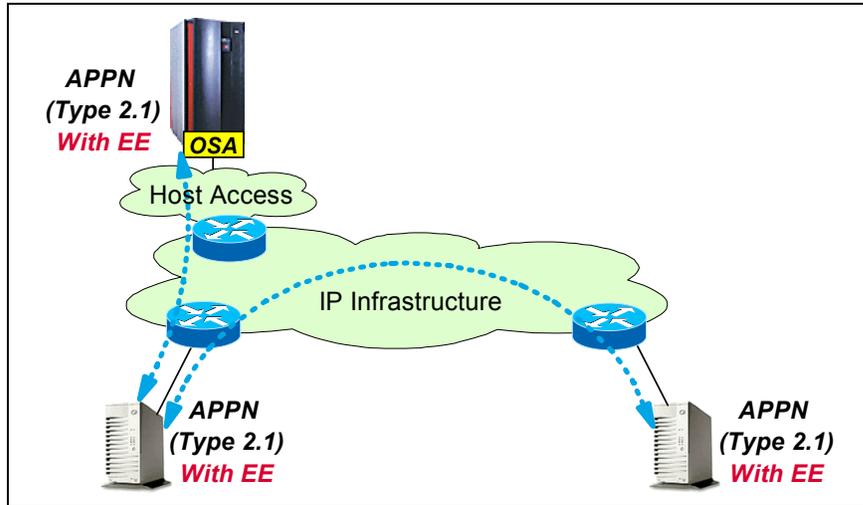


Figure 5-18 Using EE to in your APPN hosts

If your APPN hosts support EE, the CNN function provided by VTAM and NCP could also be replaced by using EE. The migration to EE function in the hosts might entail more impact to the hosts but it could completely avoid impact to your IP network other than enabling the consolidation of your SNA and IP wide-area links, perhaps allowing bandwidth upgrades (due to the cost savings) of your WAN and even enhancing the performance of your applications. Note that EE enables SNA class of service over the links in the router network by using specific UDP ports and setting the appropriate IP precedence bits for each SNA transmission priority, though you may need to implement a network-wide IP Quality of Service (QoS) policy if you have not already done so. (EE is discussed in detail in Chapter 19, “Enterprise Extender (EE)” on page 331.)



Restriction: VTAM for VM/ESA and VSE/ESA does not support Enterprise Extender.

Recommendation

If the APPN hosts in your network support EE function directly, use EE to allow them to use your IP infrastructure because that requires no special support in your routers. If, however, the APPN hosts in your network do not support EE, then use either APPN or DLSw function in your routers to carry the application traffic across your router network.

5.3.6 IP routing

NCP supports IP routing (independent of that provided by the 3746 Nways Multiprotocol Controller). If you are using NCP to route IP traffic into your host environment, migrate to the OSA-Express adapter (discussed in Chapter 15, “OSA-Express” on page 269). If you are using your NCP for other IP routing, consider migrating that function to routers or to layer-3 (L3) switches. L3 switches are switches that have very high-speed IP routing capabilities. In general, today’s routers and L3 switches can provide more robust IP routing (including such capabilities as IP differentiated services based prioritization) at a lower cost and with better performance.

5.3.7 Extended Recovery Facility (XRF)

Extended Recovery Facility (XRF) increases online availability of an IMS or CICS system by providing an alternate IMS or CICS system that monitors the active system so that it is ready to take over in the event that the primary system fails. Support for the XRF capability is provided by specific functions implemented across the following set of IBM software products:

- ▶ MVS™
- ▶ Data Facility Product (DFP)
- ▶ IMS
- ▶ CICS
- ▶ VTAM
- ▶ Network Control Program (NCP)

Figure 5-19 on page 181 illustrates the role of NCP in an XRF environment.

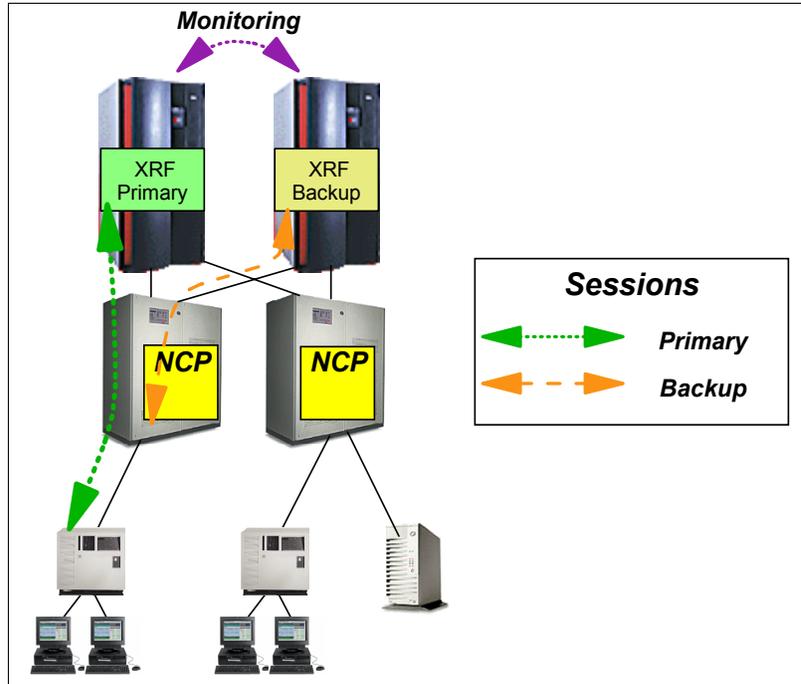


Figure 5-19 XRF environment

In the figure, sessions from the XRF primary system (such as the one shown to the left above) are active and supporting the communications needs of the application. Meanwhile, sessions from the backup system are pre-established but held by NCP. The devices in session with the XRF primary system do not see the sessions from the backup while they are being held. The XRF backup system monitors the status of the primary system through a “heartbeat” process. In the event of a primary XRF system failure, the backup system will send a command to NCP telling it to switch to the backup sessions, at which time the devices will very quickly be placed in session with the XRF backup system.

Alternatives for migrating from your 3745 communication controller support of XRF include:

- ▶ Continue to run NCP, supporting XRF, using CCL.
- ▶ Migrate from XRF to a Parallel Sysplex-based solution.

Each of these alternatives is discussed below.



Continue to run NCP using CCL to support XRF

NCP provides the only means of supporting XRF session switching capabilities. The IBM Communication Controller for Linux on System z9 and zSeries (CCL) product allows you to continue to run NCP, supporting your XRF environment, on your mainframe. CCL provides the *only* alternative that will allow you to migrate from your IBM 3745 Communication Controller while continuing to support XRF.

Figure 5-20 illustrates the use of CCL for an XRF environment.

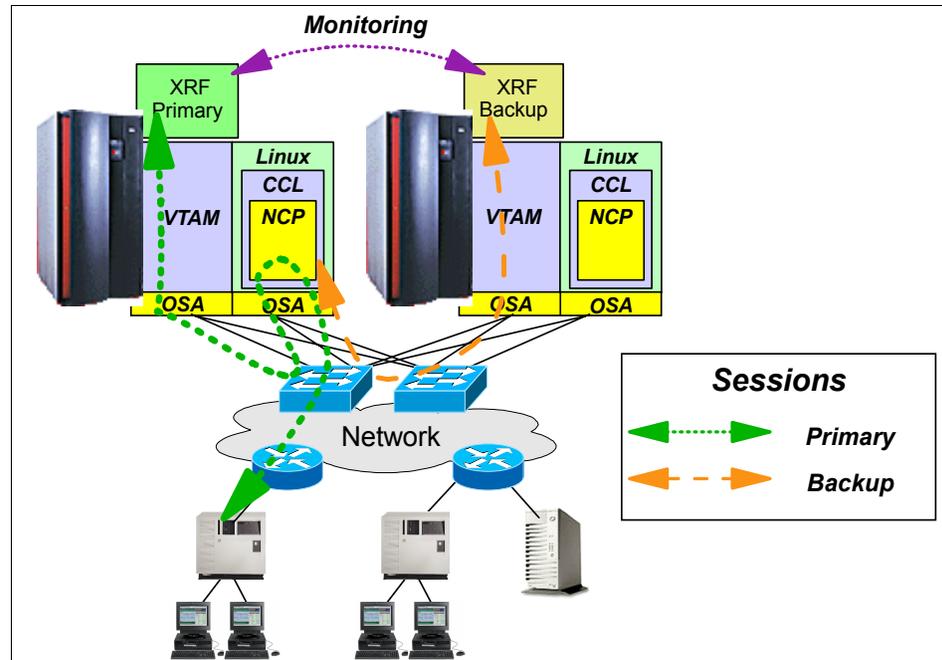


Figure 5-20 Using CCL in an XRF environment

The illustration shows a fully redundant local-area networking environment (CCLs, NCPs, OSA adapters, and LAN switches) comparable to the original 3745-based redundant environment. In order for XRF to work, the primary and backup sessions must go from their respective VTAMs through the LAN and to the CCL NCP that provides SNA boundary function (therefore, the peculiar down and up session paths in the illustration). From an XRF application and systems management perspective, however, the CCL NCP will work the same as running NCP in an IBM 3745 Communication Controller.

Advantages of using CCL to support XRF include:

- ▶ CCL provides the *most transparent* migration path from IBM 3745 Communication Controllers (and can be implemented with little or no coordination and change in the application environment).
- ▶ *CCL minimizes change* to the current network management environment.
- ▶ The CCL solution enables you to migrate XRF support from IBM 3745 Communication Controller hardware without having to fundamentally change your XRF environment.

Considerations of a CCL-based XRF solution include:

- ▶ Continued dependence on NCP connectivity for XRF prevents the use of Enterprise Extender technology (and Cisco SNA Switching Services) to provide non-disruptive rerouting of SNA sessions around network failures. (Enterprise Extender is discussed in detail in Chapter 19, “Enterprise Extender (EE)” on page 331.)
- ▶ CCL runs in a Linux for zSeries environment and may require additional skills and mainframe system resources (such as logical partitions and CPU cycles) to support it.

Note: CCL CPU cycles may be less expensive than other host CPU cycles because Linux partitions can use Integrated Facility for Linux (IFL) processors.

Migrate from XRF to Parallel Sysplex

IBM delivered XRF for IMS and CICS application high availability in the 1980s. Then, through the 1990s, IBM delivered Parallel Sysplex technology (to enable multiple mainframe systems to be combined for both high availability and scalability) as well as successive enhancements to SNA networking; most important, Enterprise Extender (EE) technology that enables non-disruptive rerouting of SNA sessions around network failures. Consequently, you may be able to deliver even higher end-to-end availability for your applications by leveraging Parallel Sysplex and advanced SNA technologies. Figure 5-21 on page 184 illustrates such a solution.

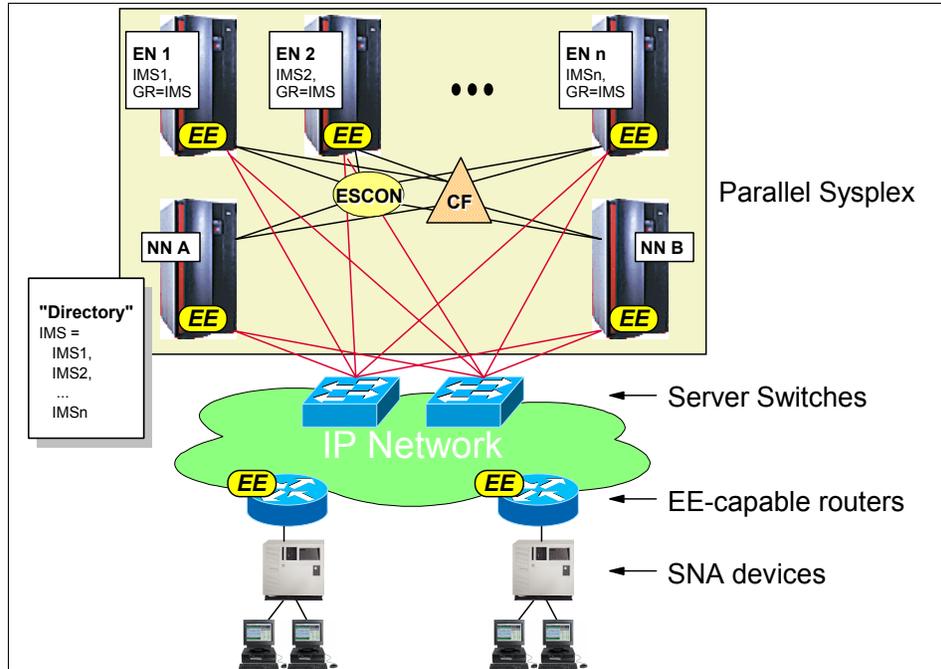


Figure 5-21 Replacing XRF with Parallel Sysplex technologies

The basic idea of XRF is to have a “primary” application host with an “alternate” application standing by in case of a failure. Users log on with a user application name variable, or USERVAR, which is assigned (via operator command) the actual application name of the current XRF primary application. In contrast, Parallel Sysplex technology is designed around having multiple concurrently active instances of a given application and distributing workload across them. Those multiple instances of the application together provide a single, virtualized, application referred to by a generic resource (GR) name.

Note: You should never have both a USERVAR and a Generic Resource with the same name concurrently active. Therefore, to migrate from using XRF to using sysplex generic resources, you can either:

- ▶ Migrate all at once by deactivating the USERVAR and then activating a Generic Resource with the same name
- or
- ▶ Use VTAM Interpret tables to migrate specific LUs, or groups of LUs, over a period of time.

In order to migrate incrementally, you must:

1. Define an Interpret table that translates session initiation requests for the USERVAR name to a new generic resource name
2. Assign the new Interpret table to specific LUs, which can be accomplished dynamically via operator command.

Migrating to generic resources from XRF/USERVARs using VTAM Interpret Tables and Autologon is discussed in details in z/OS Communications Server Technote 21220398, which can be found on the Web at:

<http://www.ibm.com/support/docview.wss?rs=852&uid=swg21220398>

In Figure 5-21 on page 184, the services of an IMS application are distributed across n different systems. (IMS is used here as an example; these could have been CICS applications.) Each system registers its IMS application instance as well as its GR name, “IMS,” with its serving Network Node (NN A in this example). Session requests for the GR name “IMS” will go to NN A, which will leverage sysplex policies (least busy, for example) to direct the session request to the optimal application instance. Thereafter, using Enterprise Extender (EE) technology, all session traffic will flow directly to the appropriate application host (not through a Network Node) and can be nondisruptively rerouted around failures in the network.

GR support works well for most application usage (for example, a terminal or automated teller machine) and, in some cases, can even ensure that subsequent parallel sessions go to the same application instance. Some application communications, however, require that sessions always go to a *specific* application instance. For example, some LU 6.2 applications require “persistent affinity” to an application instance (not allowing sessions if that specific instance is currently unavailable). To provide high availability for such persistent affinity applications (to quickly recover sessions and data from an application failure), you may be able to use Multi-Node Persistent Sessions (MNPS). MNPS continuously maintains session-state information in a Parallel Sysplex coupling facility for each

session (which may significantly increase CPU workload) and can switch sessions to a new instance if a given application instance fails. MNPS can only be used with applications that support MNPS session recovery, such as CICS.

Note: MNPS could also be used for application usage that has no affinity to any given application instance; however, GR recovery is faster than MNPS and requires far less overhead during normal operations.

Advantages of migrating from XRF to Parallel Sysplex support include:

- ▶ Leverages more modern, scalable, and strategic Parallel Sysplex technology.
- ▶ Provides a simpler and less expensive solution to support.
- ▶ Enables you to reduce the number of users affected by any given system failure by distributing users across more than two systems.
- ▶ Enables non-disruptive rerouting of SNA sessions around network failures through the use of Enterprise Extender technology (and Cisco SNA Switching Services). (Enterprise Extender is discussed in detail in Chapter 19, “Enterprise Extender (EE)” on page 331.)

Considerations of a Parallel Sysplex solution include:

- ▶ Requires potentially significant change to the application environment (and coordinated change across networking, systems, and application organizations).
- ▶ In the event of a system failure, it will likely take longer for the affected users to re-establish sessions than an XRF switch takes today.

Recommendation

CCL provides the most transparent alternative for migrating from your current IBM 3745 communication controllers and a CCL migration can, for the most part, be initiated and executed by just the networking organization. Leveraging sysplex support of generic resources is more strategic and scalable but requires significant changes to the application environment (and coordinated change across a much-broader group of organizations). Ultimately, the Parallel Sysplex alternative is likely to provide higher end-to-end availability as well as much better application scalability.



X.25 NCP Packet Switching Interface (NPSI)

IBM Product Number: 5688-035

This chapter discusses alternatives to the NPSI support of X.25 links in SNA networks and X.25 network access to SNA applications.



Note: NPSI can now be run on your mainframe, along with NCP, using the Communication Controller for Linux on System z9 and zSeries (CCL) V1.2 product, 5724-J38. In many cases, the CCL option will provide the most transparent migration from your IBM 3745 Communication Controllers.

See Chapter 16, “Communication Controller for Linux on System z9 and zSeries (CCL)” on page 273 for additional information about CCL.

Note: In addition to the technical solutions discussed in this chapter, you should consider the alternatives of leaving things as they are, eliminating the IT service in question, or replacing that service with a completely new solution. See “Important migration-strategy alternatives” on page 95 for a discussion of each of those alternatives.

6.1 NPSI product overview

NPSI enables access to Systems Network Architecture (SNA) application programs through an X.25 packet switched data network. As illustrated in Figure 6-1, NPSI is used to support:

- ▶ Certain SNA and non-SNA devices (such as terminals and printers) connecting by X.25 network facilities to a communication controller
- ▶ Remote communication controllers connecting by X.25 to data center controllers
- ▶ Data center to data center (including between companies) communication over X.25 networks

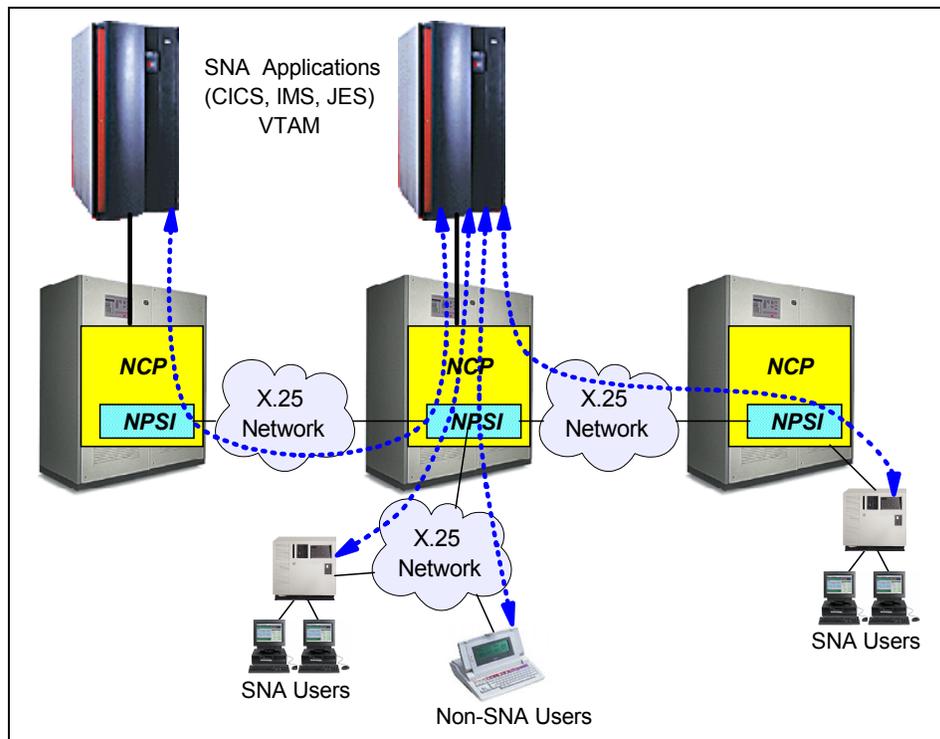


Figure 6-1 The role of NPSI

NPSI supports five LLC types. The LLC type used in any particular situation depends on the type of device with which NPSI communicates and the

characteristics of the connection. NPSI link-level protocol support can be considered in two groups:

- ▶ LLC types 2 and 3 are for SNA device to SNA application connections. Most SNA NPSI implementations use LLC type 3 which is also known as QLLC. The other SNA LLC type, LLC 2, is also called PSH.
- ▶ LLC types 0, 4, and 5 are for non-SNA device to SNA application connections and are also known as PCNE, GATE, and X.3 PAD respectively. Another key function to be aware of is DATE, which can be implemented for LLC types 0, 2, 3, and 5. These protocols tend to be used for very specific applications, sometimes involving custom-built terminals, and are used by many NPSI clients.

Note: X.25 is also supported via NCP, independently of NPSI, when used with microcode on the 3746-900. The microcode that drives serial lines on the 3746-900 is called ODLC. There is a licensed X.25 feature for ODLC that enables 3746-900 serial links to provide QLLC support for X.25 connections. If you have an LLC type 3, QLLC, connection attached to a port on a 3746-900, you can choose whether that port is supported by NPSI or by NCP and the X.25 feature microcode. However, LLC types 0, 2, 4, and 5 are only supported by NPSI.

For further information on 3746 microcode support of X.25, see 5.2.8, “X.25 (with licensed support feature on 3746)” on page 152.

For more information on NPSI, see:

- ▶ *X.25 NPSI Version 3 Licensed Program Specifications*, GC30-9605
- ▶ *X.25 NPSI Version 3 General Information*, GC30-3469
- ▶ *X.25 NPSI Version 3 Planning and Installation*, SC30-3470
- ▶ *X.25 NPSI Version 3 Host Programming*, SC30-3502

6.2 Functions and alternatives

The following sections provide additional detail about NPSI support and suggest alternative means of providing SNA and non-SNA device access over X.25 networks to SNA applications.

Important: In some cases, X.25 connections are used for interconnection between *different organizations*. Consequently, any attempt to change such an X.25 environment will require coordinating the change with those other organizations. The inter-organizational coordination of such changes often poses much more difficult challenges than do technical issues. Some approaches that may help include:

- ▶ Strive to work within international and/or industry standards such as message queuing (MQ), electronic data interchange (EDI), and extensible markup language (XML).
- ▶ Work closely with industry consortia and extranets. For example:
 - The automobile manufacturing industry has the Automotive Network Exchange.
 - The IVANS network provides similar interconnectivity for the insurance industry.
- ▶ Take advantage of value-added network services providers by letting them deal with the individual connectivity needs of your trading partners.
- ▶ Explore secure virtual private network (VPN) connectivity across an IP connection or across the Internet.

6.2.1 Supporting SNA (PSH and QLLC) communication over X.25

LLC type 2 supports NPSI-to-SNA peripheral node DTE where the peripheral node has PSH hardware (either integrated or as a separate device). PSH is a very old X.25 line driver that is far less common today than LLC type 3 (QLLC). PSH and QLLC are SNA end-to-end implementations. The SNA logical units are completely unaware of the X.25 component in the network. This means that CICS, TSO, IMS, printers, terminals, and intelligent devices all behave as if the connections were traditional SNA with channels, SDLC links, etc.

There are additional types of SNA X.25 end-to-end protocol support. For example, ELLC was designed for System/36-to-System/36 peer-to-peer SNA over X.25 connections. ELLC is not supported by NPSI. Most SNA peer-to-peer connections through NPSI are supported by means of PU2.1 peer-to-peer SNA protocols using QLLC for end-to-end control in the X.25 network. Also, QLLC has two options in NPSI. The boundary network node (BNN) option is used for the support of PU2.0 and PU2.1 SNA terminals. The intermediate network node (INN) option is used for the support of the PU4 peer SNA connections, such as between communication controllers, or PU4 to PU5 connections, such as between a communication controller and an ICA-attached mainframe host.



Important: When capacity planning for a migration from NPSI SNA support, be careful not to overlook the capacity requirements for SNA boundary function. Each terminal controller or other end device in an SNA network entails a certain amount of SNA boundary function workload that must be provided somewhere in the network. As devices are migrated from 3745 X.25 attachment, you must carefully consider where the boundary function will be provided for those devices (such as in a host-attached 3745 or in the host itself), as well as the impact of that increased boundary function workload. Alternatives for handling SNA boundary function workload are discussed in 5.3.3, “SNA subarea addressing, routing, and boundary function (BF)” on page 162.

Migrating to router X.25 support

The vast majority of SNA devices supported by NPSI are QLLC (LLC3). As illustrated in Figure 6-2 on page 192, some router implementations support QLLC.

For example, Cisco’s Web site at <http://www.cisco.com> depicts sample configurations in which an organization’s QLLC mainframe, mini, or control unit connects over X.25 to a router that supports QLLC. The router then connects to other SNA devices over an SNA link, LAN, or channel, thus providing full end-to-end communication for SNA over X.25 using QLLC without a communication controller. Some other vendors’ routers, including IBM, also support QLLC.

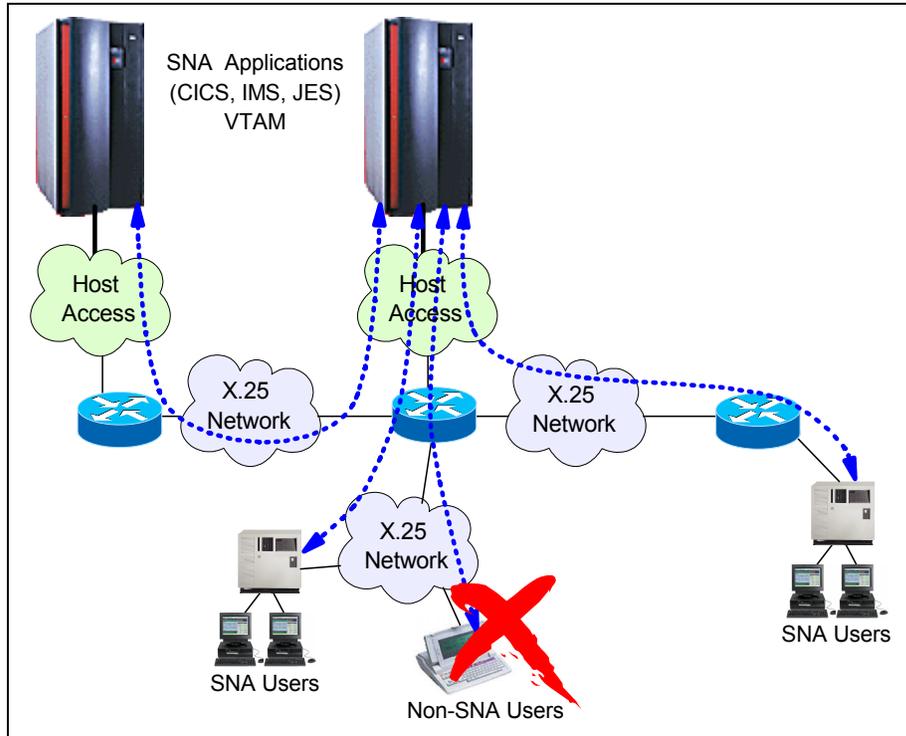


Figure 6-2 Router replacement of NPSI

Advantages of a router-based solution include:

- ▶ Your organization's X.25 end devices and network can remain in place.
- ▶ You have the flexibility of implementing native SNA or SNA over IP for the non-X.25 portion of the network without using a communication controller.

Considerations of a router-based solution include:

- ▶ Router-only solutions support QLLC, and will not handle devices using PSH, PCNE, PAD, or GATE without additional software. See "Host-based X.25 over TCP/IP (XOT)" on page 199 for an example of such software.
- ▶ Unless you implement DLUR function in the routers, router replacement of communication controllers for X.25 support will increase SNA boundary function requirements either in other communication controllers or in the host (see important note in 6.2.1, "Supporting SNA (PSH and QLLC) communication over X.25" on page 190).

Note: As you plan to migrate X.25 from your controllers, carefully explore *all* of your wide-area networking connectivity options. While routers may be the logical WAN connectivity devices for your locations, you may find newer service offerings such as frame relay services or virtual private network (VPN) offerings that are more cost-effective than your current network structure. Such service offerings can even include the provision and management of the routers you use to access the network.

Using routers to migrate from X.25

The routed network structure discussed in “Migrating to router X.25 support” on page 191 actually opens up yet another alternative: that of using the routed infrastructure to migrate away from your X.25 network. As illustrated in Figure 6-3, you may be better off migrating from X.25 as you implement your routed network. Note that you may be able to use router QLLC support (such as that provided by routers from Cisco Systems) to avoid requiring changes to your remote X.25 attached SNA devices.

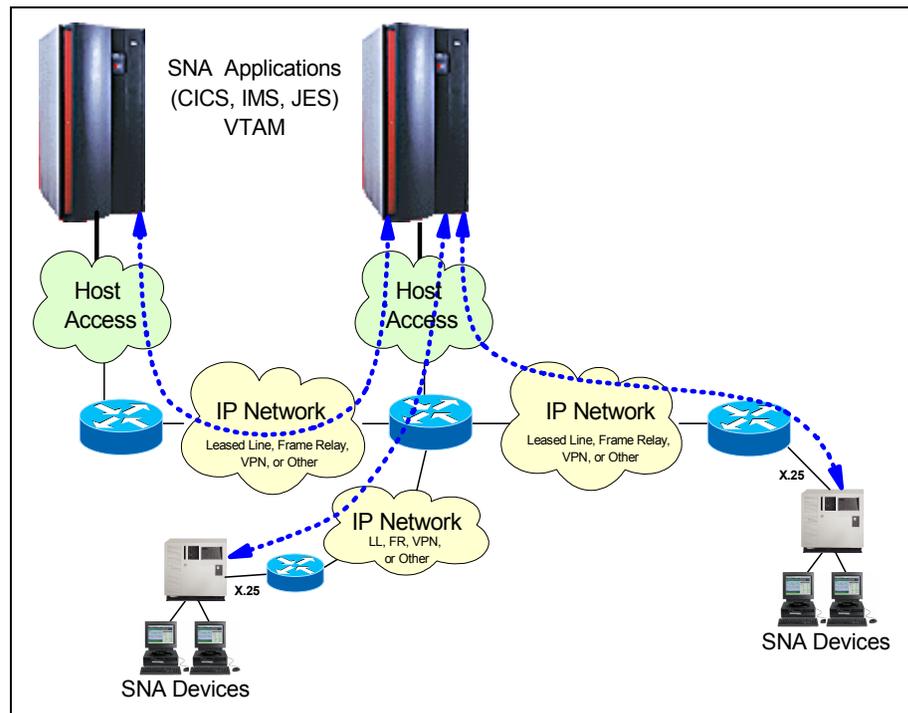


Figure 6-3 Using routers to migrate from X.25

The routers can use a technology such as DLSw or SNA Switching Services to deliver the SNA traffic to the host systems across the routed IP network infrastructure.

Advantages of a router-based migration from X.25 include:

- ▶ You can migrate toward a more strategic wide-area networking infrastructure with minimal impact to devices in the network.

Considerations of a router-based migration from X.25 include:

- ▶ Router-only solutions support QLLC, and will not handle devices using PSH, PCNE, PAD, or GATE without additional software. See “Host-based X.25 over TCP/IP (XOT)” on page 199 for an example of such software.
- ▶ Unless you implement DLUR function in the routers, router replacement of communication controllers will increase SNA boundary function requirements either in other communication controllers or in the host.

Connecting into an Integrated Communication Adapter (ICA)

Certain IBM mainframes, such as the 9221, support an integrated communications adapter (ICA) hardware feature. The ICA supports binary synchronous communication (BSC), synchronous data link control (SDLC), and X.25 link attachments directly through mainframe systems software, such as VTAM or BTAM. This means that the ICA supports X.25 links into the mainframe without NPSI or NCP (as shown in Figure 6-4 on page 195). The X.25 support is provided directly by VTAM under VM and VSE only and is for SNA X.25 QLLC only. For additional information, see the *VTAM Network Implementation Guide*, SC31-6494.

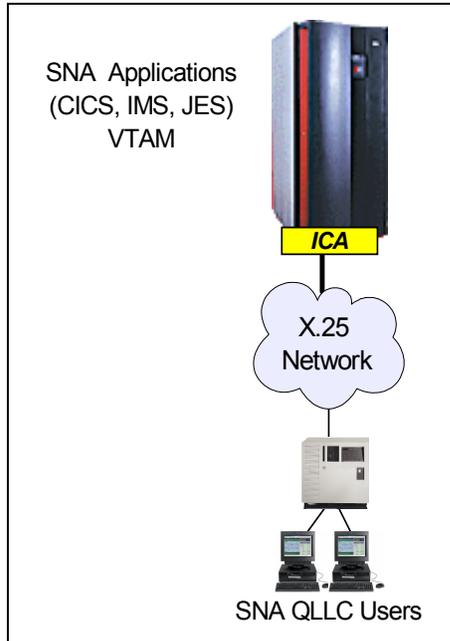


Figure 6-4 ICA X.25 support

Note: One of the key reasons for the development of communication controller products was to offload communication line handling from the host. If you are moving lines from a communication controller to an ICA, be careful to consider the possible increase in host CPU cycles it may require.

Advantages of an ICA-based solution include:

- ▶ Your organization's X.25 end devices and network can remain in place.
- ▶ It provides mainframe SNA application access with no need for a communication controller or a router.

Considerations of an ICA-based solution include:

- ▶ The last S/390 to support an ICA was the 9221, which is no longer marketed by IBM.
- ▶ The 9221 is not supported by the current VM operation system (z/VM). It is supported by VM/ESA; however, VM/ESA is no longer marketed by IBM.
- ▶ Only VM and VSE systems support the ICA, so if you are running applications on z/OS, you will need to ensure connectivity from the system with the ICA to the application LPAR.

- ▶ Because ICAs support limited numbers and speeds of links when compared with a communication controller, scalability may be an issue with an ICA solution.
- ▶ One of the key reasons for the development of communication controller products was to offload communication line handling from the host. By moving lines from a communication controller to an ICA, you may experience a substantial host CPU cycle increase.
- ▶ ICA replacement of communication controllers for X.25 support will increase SNA boundary function requirements in the host (see important note in 6.2.1, “Supporting SNA (PSH and QLLC) communication over X.25” on page 190).
- ▶ The ICA solution for migrating SNA devices off of NPSI covers QLLC implementations, but not PSH. We have found no alternative for PSH support short of replacing the devices.

Recommendation

A router-based solution will probably be your most cost-effective and strategic alternative. However, if you already have an ICA and the available ports and speeds are sufficient for your requirements, then the ICA may be an acceptable alternative.

6.2.2 Supporting non-SNA communication over X.25

The NPSI non-SNA device support includes the following types of connections:

- ▶ LLC type 0 connects to X.25 non-SNA DTEs using PCNE
- ▶ LLC type 4 connects to specific applications willing to control the X.25 Packet Level Protocol (PLP) that use GATE support
- ▶ LLC type 5 connects to non-SNA DTEs (such as start-stop terminals) through an ASCII PAD

NPSI supports communication from non-SNA devices by simulating a PU1 and LU1 on behalf of each device. The simulated PU/LU 1 is an IBM 3767, the same as with NTO. This simulated PU/LU1 exists only as a process within NPSI. This means that across the end-to-end connection for LLC types 0, 4, and 5, the VTAM-to-NPSI connection is SNA and the NPSI-to-device connection is non-SNA over X.25.

- ▶ For LLC type 0 and 5 (integrated PAD), the X.25 control packets (such as call and interrupt packets) are processed by NPSI.
- ▶ For LLC type 4 (GATE) and 5 (transparent PAD), as well as DATE, the control packets are processed by an application called Communication and Transmission Control Program (CTCP) that runs in the host.

Note: If your organization has the GATE keyword coded as GENERAL or DEDICATED in the NPSI generation statements, you should consider the user code controlling the end-to-end connection as well as CTCP relay code (such as Communication Subsystem For Interconnection (CSFI) for the GATE host-side support.

If the GATE keyword is coded as DEDICATED (DATE), all control packets are processed by CTCP for LLC types 0, 2, 3, and 5. The only way to migrate from 3745 support of DATE that we found is to run the NPSI program product in your mainframe using CCL V1.2.

In “Migrating to router X.25 support” on page 191, we described how QLLC support can be provided by routers. Such solutions address most SNA device-to-mainframe NPSI requirements. For non-SNA device-to-mainframe communication support (as well as for SNA PSH support), you could migrate X.25 support from your IBM 3745 Communication Controllers onto routers and, using router X.25 over TCP/IP (XOT) technology, transport the X.25 traffic over your TCP/IP network and into your mainframe hosts. In the mainframe, you can deliver the encapsulated X.25 traffic to either:

- ▶ NPSI running in CCL V1.2
- or
- ▶ A host-based X.25 over TCP/IP (XOT) product

Each of these alternatives are described in greater detail below.



Supporting NPSI in your mainframe using CCL V1.2

CCL V1.2 introduced NPSI program product support. Hence, you can now support *all of the X.25 capabilities* currently supported by NPSI on your 3745 communication controllers in a mainframe-based NPSI implementation. Figure 6-5 on page 198 illustrates how routers can be used to encapsulate X.25 packets in TCP/IP frames for transport into the host NPSI environment.

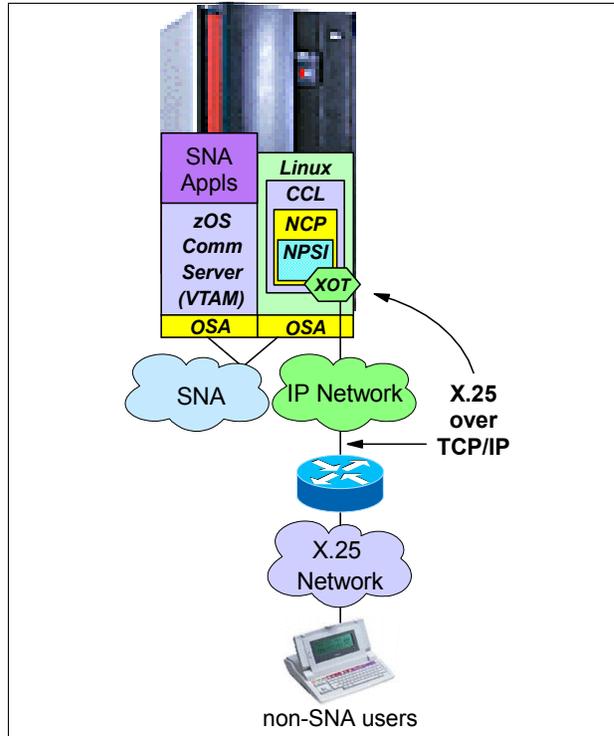


Figure 6-5 Supporting NPSI in your mainframe using CCL

In Figure 6-5, a router encapsulates inbound X.25 messages into TCP/IP messages and delivers them across the network where they are de-encapsulated by a corresponding XOT component running with the CCL. (Likewise, outbound X.25 messages are encapsulated by the mainframe XOT software, transported over the network, and de-encapsulated by the router.)

Note: The XOT protocol support in the mainframe must be provided by a separate, non-IBM, software product.

In the mainframe, the X.25 messages are delivered to NPSI (the same program product that runs on your 3745 communication controllers) and converted to SNA in exactly the same way as they would have been had they come in through an X.25 interface on a 3745. The SNA messages then flow, as they always have, through NCP and the SNA network to your SNA applications.

Advantages of supporting NPSI in your mainframe using CCL include:

- ▶ CCL provides comprehensive support of *all* NPSI functional capabilities

- ▶ CCL provides the *most transparent* migration path from IBM 3745 Communication Controllers.
- ▶ CCL *minimizes change* to the current network management environment.
- ▶ CCL leverages your highly reliable and increasingly cost-effective mainframe host environment.
- ▶ CCL requires no new host XOT software implementation effort.

Considerations of the CCL solution include:

- ▶ You must purchase and implement a separate, non-IBM product for support of XOT in your mainframe Linux environment.
- ▶ CCL runs in a Linux for zSeries environment and may require additional skills and mainframe system resources (such as logical partitions and CPU cycles) to support it.

Note: CCL CPU cycles may be less expensive than other mainframe CPU cycles because Linux partitions can use Integrated Facility for Linux (IFL) processors.

Host-based X.25 over TCP/IP (XOT)

More than likely, you are not using *all* of the varied NPSI capabilities in your environment. For certain non-SNA device-to-mainframe communication support, you may be able to use a host-based program that provides non-SNA support to complement the router X.25 support. As illustrated in Figure 6-6 on page 200, routers can encapsulate X.25 packets in TCP/IP frames for transport into the host environment. The “Host XOT” software can remove the X.25 packet from the TCP/IP “envelope” and provide support for various non-SNA LLC types.

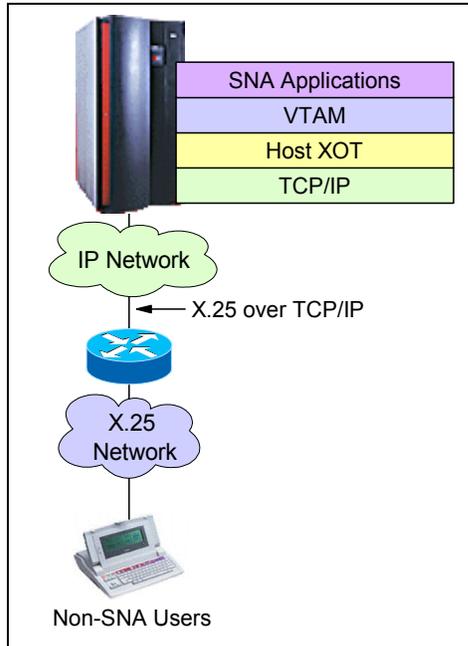


Figure 6-6 Host XOT solution

We found two such “Host XOT” offerings:

- ▶ Comm-Pro Associates, Inc. Host Network Access Support (HNAS)
- ▶ Computer Associates, Inc. Unicenter TCPaccess X.25 Server

Each of these offerings are discussed further below.

Comm-Pro Associates, Inc. Host Network Access Support (HNAS)

Comm-Pro Associates, Inc. (<http://www.comm-pro.com/>) X.25 Host Network Access Support (HNAS) product works with Cisco and IBM routers to avoid application changes by providing a wide-ranging suite of non-SNA NPSI-type support, including:

- ▶ PCNE (LLC0), including 3767 LU type 1
- ▶ Integrated PAD (LLC5) including host X.29 PAD control and PCNE 3767 function for start-stop terminals using X.3 PAD
- ▶ Transparent PAD (LLC5) for start-stop terminals using X.3 PAD
- ▶ GATE (LLC4) in which a host application uses a separate session to manage X.25 connections, disconnects, and data

In addition, HNAS provides support for GATE Fast Connect implementations. According to Comm-Pro (as well as some IBM customer experience with the HNAS product), no application changes are necessary, and sysgen requirements are replaced with small configuration data and VTAM application definition files. Comm-Pro's literature describes various additional HNAS security and performance features, such as the ability to associate individual devices with particular sessions, and the ability to map multichannel links on IBM routers to individual addresses (so that the CTCP application's view of the NPSI network can be preserved in the HNAS solution).

Note that there is no DATE support in HNAS. A list of applications that have been successfully tested with HNAS can be found at:

<http://evolution.sna.free.fr/hnase.html>

Comm-Pro HNAS product support is provided from Santa Clarita, California, USA, and Comm-Pro suggests that their international customers consider using one of their business partners in order to obtain local product support. For example, IBM Global Services, France, has a relationship with Comm-Pro to provide first-level support while Overlap (<http://www.overlap.fr>) provides second-level support.

Unicenter TCPAccess X.25 Server

Unicenter TCPAccess X.25 Server (formerly known as SOLVE:X.25) is a product from Computer Associates:

<http://ca.com/>

It provides GATE, PCNE, GATE Fast Connect, and DATE functionality. Unicenter TCPAccess X.25 Server runs on the mainframe under z/OS and works together with Cisco routers. Unicenter TCPAccess X.25 Server can handle LLC type 0, LLC type 4, and LLC type 5 (both integrated and transparent PAD) connections, thereby providing a means of preserving your organization's user-written GATE applications as well as other application code.

Recommendation

If you require NPSI capabilities that are not supported in a Host XOT software package, then CCL NPSI support will be your only viable migration option from your 3745 communication controllers. The CCL NPSI approach is also likely to provide an easier migration from your 3745 because it continues the use of the same software products that you are using today; however, the CCL approach will require the implementation of Linux in your mainframe environment if you do not already have it. Hence, in some situations, the use of Host XOT software may prove more straightforward.

The Comm-Pro HNAS product and the Computer Associates Unicenter TCPaccess X.25 Server product both handle a rich variety of LLC type 0 (PCNE) and LLC type 4 (GATE) functions. You should confirm that the particular functions that you require are provided by the product that you select, and check for availability of these products and their support in your country.



MERVA Extended Connectivity

IBM Product Number: 5655-110

This chapter discusses alternatives to MERVA Extended Connectivity support for X.25 connections to the S.W.I.F.T. network.

Note: In addition to the technical solutions discussed in this chapter, you should also consider the alternatives of eliminating the IT service in question or replacing that service with a completely new solution. See “Important migration-strategy alternatives” on page 95 for a discussion of each of those alternatives.

7.1 MERVA product overview

MERVA is a financial messaging system that helps institutions communicate with each other in the exchange of payments, settlements, and other financial transactions. For the last 20 years, the MERVA product family has been used by financial institutions to:

- ▶ Access the S.W.I.F.T. network
- ▶ Create service bureaus for S.W.I.F.T. users
- ▶ Interconnect branches within an enterprise
- ▶ Set up a real-time gross settlement network interconnecting financial institutions
- ▶ Synchronize remote multi-purpose applications

Hundreds of organizations worldwide have based their mission-critical applications on the MERVA messaging infrastructure. As shown in Figure 7-1, MERVA Extended Connectivity runs in IBM communication controllers to support MERVA connectivity to the S.W.I.F.T. network.

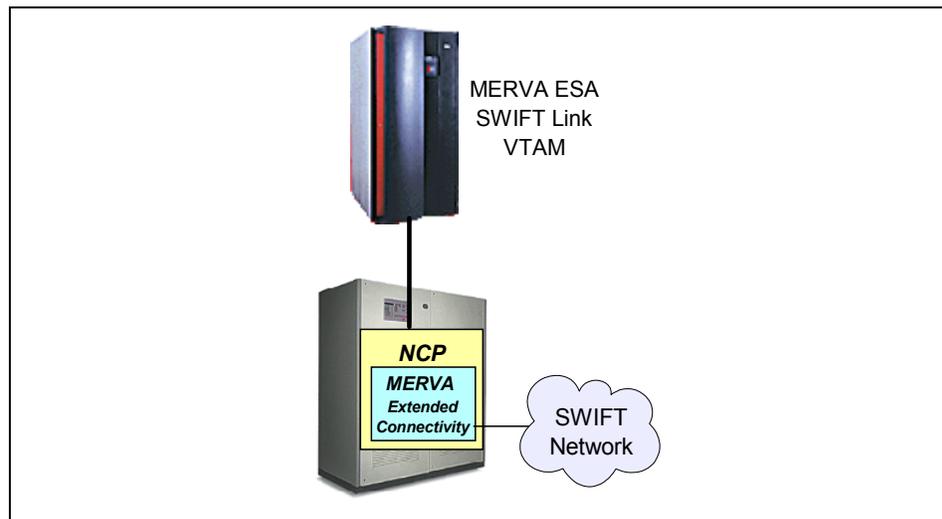


Figure 7-1 The role of MERVA Extended Connectivity

As of April, 2005, according to the Society for Worldwide Interbank Financial Telecommunication (S.W.I.F.T.) Web site (<http://www.swift.com/>), the S.W.I.F.T. network provides a shared worldwide data processing and communications link and a common language for international financial transactions. S.W.I.F.T. is an industry-owned co-operative supplying secure messaging services and interface software to over 7,650 financial institutions in 200 countries. S.W.I.F.T. carries nearly 10 million messages each day. S.W.I.F.T. provides messaging services to

banks, broker-dealers, and investment managers, as well as to market infrastructure providers in payments, treasury, securities, and trade. These services help S.W.I.F.T. customers reduce costs, improve automation, and manage risk.

For more information on the MERVA Extended Connectivity product, see: *MERVA Extended Connectivity for ESA Installation and User's Guide*, SH12-6157-01, available at:

<http://publibfi.boulder.ibm.com/epubs/pdf/cmvu0m01.pdf>

7.2 Functions and alternatives

The MERVA Extended Connectivity product provided communication controller based connectivity to the S.W.I.F.T. network. On April 2, 2005, the S.W.I.F.T. network X.25 connectivity was finally decommissioned. SWIFTNet is S.W.I.F.T.'s advanced IP-based messaging solution and alternative to the X.25 connectivity.

On October 1, 2002, IBM announced the WebSphere Business Integration for Financial Networks (WBI-FN) product to help financial services organizations take advantage of new network capabilities from SWIFT, simplify financial messaging and reduce IT costs. WBI-FN provides an integration hub delivering connectivity to multiple external financial networks. This product solution does not require the use of MERVA Extended Connectivity. For more information, see:

<http://www.ibm.com/software/integration/wbifn/>

Recommendation

Leverage IBM strategic product: WebSphere Business Integration for Financial Networks for connectivity to the S.W.I.F.T. Network.



X.25 SNA Interconnection (XI) and Network Supervisory Function (NSF)

IBM Product Numbers: 5685-035 and 5685-003

This chapter discusses alternatives to the XI support of X.25 packet forwarding across SNA networks and the NSF management of XI functions.

Note: In addition to the technical solutions discussed in this chapter, you should also consider the alternatives of eliminating the IT service in question or replacing that service with a completely new solution. See “Important migration-strategy alternatives” on page 95 for a discussion of each of those alternatives.

8.1 XI and NSF product overview

XI resides in one or more communication controllers with NCP in an SNA network. As shown in Figure 8-1, with XI, the SNA network can transport X.25 packets between compatible X.25 Data Terminal Equipment (DTEs). XI resources are managed by a co-requisite program, X.25 SNA Network Supervisory Function (NSF). NSF runs as a subtask of NetView and is designed to control the XI resources and the X.25 interface of the communication controllers running XI, and to collect X.25 traffic and accounting information.



Figure 8-1 The roles of XI and NSF

An extensive set of XI configuration parameters allows you to control routing, quality of service, flow control, call acceptance, security, usage and chargeback determinants, accounting, diagnostics, and timer and retry values.

For more information about XI and NSF, see *X.25 SNA Interconnection X.25 SNA Supervisory Function Planning*, GH11-3033-4.

8.2 Functions and alternatives

The basic purpose of XI is to enable your SNA network to transport X.25 traffic between compatible X.25 DTEs. Unlike NPSI, which provides protocol conversion from X.25 to SNA that allows X.25 devices (such as terminals and printers) to communicate with SNA applications, XI does not perform protocol conversion. Thus, aside from accounting and management functions that are ancillary to user function, there are no application considerations in XI migrations unless you migrate from X.25. This is because the SNA network is transparent to the X.25 devices and/or networks it connects. Some XI migration alternatives follow.

8.2.1 Transporting X.25 traffic

The function of the XI product is a little different from most communication controller functions because it is concerned only with transporting X.25 frames irrespective of the devices or applications that are the endpoints of the communication. There are various means of transporting X.25 traffic which, depending upon your unique circumstances, may or may not be viable.

Migrating end devices from X.25

Migrating end devices to use a new networking technology, such as frame relay, presents an opportunity to invest in newer and faster technologies, with potentially increased function. This, of course depends upon the availability of frame relay or some other suitable replacement service in your locations. It is important, however, that you check your X.25 configurations to determine what functions you are using. For example, X.25 contains a rich set of dial connection control functions. If any users now have switched X.25 network connections, and if they plan to use switched network connections in the new environment, you should make sure that the new environment can support any call optimization, accounting, and security functions currently implemented that will be required in the new environment.

Migrating to an X.25 service

Since XI provides a fully standard X.25 DCE interface, you could replace the SNA-based connections with circuits in an X.25 packet switched data network (PSDN) service. As with the migration to a new networking service mentioned above, the alternative of migrating to an X.25 service depends upon the availability of an X.25 service in your location.

Migrating to router transport of X.25

It may be possible to migrate from your SNA network transport of X.25 traffic to an IP router network transport of X.25 traffic. For example, according to the Cisco

Systems, Inc. Web site, their routers support transport of X.25 traffic over an IP router network.

You will need to determine whether your specific DTE and DCE connection requirements can be met, as well as whether performance, security, accounting, management, and other required functions can be met by such a solution.

Recommendation

Regardless of which of the above alternatives are viable in your environment, you should keep in mind that the list of parameters involved in XI customization is extensive. You should examine all explicitly coded parameter values as well as all default values to ensure that you understand the exact functionality of your current implementation. You should then determine whether all currently implemented functions are compatible with the capabilities of the new network. If not, you need to make sure that the changes required for the new configuration will meet your needs. For example, if you plan to replace the XI portion of the network with a router-based IP network, you must make sure that the routers can be configured for X.25 support in a manner that fulfills the requirements of your organization.

8.2.2 NSF-based chargeback for X.25 transport service

Some organizations make use of the substantial accounting information provided by NSF in order to facilitate chargeback for the network services utilized by another organization. While it may be possible to obtain limited information by querying router MIB variables, we have been unable to find any capability that comes close to that offered by NSF.



Emulation Program (EP)

IBM Product Number: 5735-XXB

Emulation Program (EP) software was developed almost 30 years ago to run on the IBM 3705 family of programmable communication controllers. EP on the 3705 provided a migration path from the hard-wired, non-programmable family of communication controllers (the IBM 2701 Data Adapter Unit and the IBM 2702 and 2703 Transmission Control Units) that preceded the development of programmable communication controllers. In this chapter, we discuss the migration options for hardware and software environments that depend upon EP.

Note: In addition to the technical solutions discussed in this chapter, you should also consider the alternatives of eliminating the IT service in question or replacing that service with a completely new solution. See “Important migration-strategy alternatives” on page 95 for a discussion of each of those alternatives.

9.1 EP product overview

EP is a licensed program that enables communication controllers to support the connection of link-attached binary synchronous communication (BSC) and start-stop devices (such as terminals and printers) to host applications over a byte channel. The devices typically supported by EP are:

- ▶ BSC 3270 devices that exchange information with non-SNA applications (such as BTAM applications)
- ▶ BSC remote job entry (RJE) devices such as 2780 and 3780
- ▶ Start-stop devices that exchange information with non-SNA applications (such as BTAM applications).

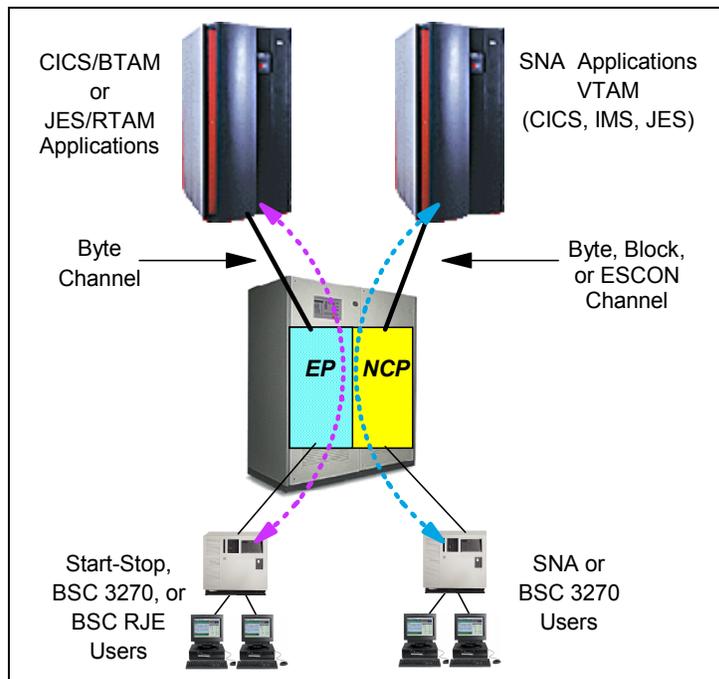


Figure 9-1 The role of EP

A communication controller running EP must be connected to the mainframe on a byte channel. Also, EP controlled lines are not supported in the 3746-900 frame. Each link under EP control is dedicated to a single application through its unique subchannel address.

You can use EP together with NCP in a single communication controller. This is called the partitioned emulation program (PEP) environment and is shown in Figure 9-1 on page 212. In PEP, the communication controller operates as a 270x device (non-SNA EP emulation mode) for some of the resources and as a communication controller under the control of NCP (SNA network control mode) for the other resources.

Over the years, organizations have generally migrated as many devices and applications as possible from EP for the following reasons:

- ▶ Cost: each EP network link is dedicated to a single application
- ▶ Manageability: EP code does not include modern control block access points for management control and analysis tools
- ▶ Recovery: EP does not contain flexible retry sequences that allow recovery from transient link errors of medium or long duration

If you still have applications or devices that depend on EP, and have not migrated them to more modern technologies such as SNA or TCP/IP, it is likely due to one or more of the following reasons:

- ▶ The required application changes would be costly, risky, complex, or simply impossible (due, for example, to unavailability of application source code)
- ▶ The devices involved are specialized and would be difficult or costly to replace
- ▶ The device-to-application network connectivity crosses organization lines, meaning that changes to the network and application environment would require the cooperation of both parties, which may be difficult to obtain or coordinate

For more information about EP and migrations, see:

- ▶ *VTAM Network Implementation Guide*, SC31-6494-01, for migrations involving the Integrated Communications Adapter)
- ▶ *NCP, SSP, EP Resource Definition Guide*, SC31-6223-09
- ▶ *NCP, SSP, EP Resource Definition Reference*, SC31-6224-09

9.2 Functions and alternatives

EP provides serial link attached devices with channel access to host applications. EP attachment is non-SNA, and typical device types supported are BSC 3270, BSC 2780/3780, and start-stop. We cover each of these below.

9.2.1 BSC 3270 terminal connection to non-SNA applications

It is easy to be confused about BSC 3270 support in the communication controller. This is because both NCP and EP can support BSC 3270 devices, which means that they may or may not be part of the SNA network. BSC 3270 devices connected to communication controller ports that are under the control of EP (and BTAM) are not part of the SNA network (referred to as “non-SNA”), and it is this type of connection that is covered in this section. For NCP-supported BSC 3270 devices, see 5.2.2, “BSC 3270 terminal connection to SNA applications” on page 146.

Important: The BTAM product was withdrawn from marketing and program services were discontinued March 31, 2002. Consequently, before pursuing the technical solutions discussed below, carefully consider eliminating the IT service in question or replacing that service with a completely new solution (discussed in “Important migration-strategy alternatives” on page 95).

One way to migrate BSC 3270 traffic off of EP would be to convert it from non-SNA to SNA. It is important to realize, however, that converting from non-SNA to SNA is much more complex than simply disconnecting from an EP port and connecting to an NCP port because the conversion would also require significant application changes.

Projects involving complex changes to applications written many years ago are risky. Therefore, the functional alternatives for BSC 3270 below provide for migration from EP without application changes. In our first alternative, both the applications and devices remain essentially unchanged. In our second alternative, the application is preserved but the user devices may be changed, if you wish, to workstations running SNA 3270 or TN3270.

Preserving BSC 3270 using an ICA or Hydra 3000

If you want to migrate from EP without changing your user devices or applications, you have two “box replacement” possibilities:

- ▶ If you have an Integrated Communications Adapter (ICA) and sufficient available ports on it, you may be able to migrate your BSC lines from your communication controllers to the ICA. Note that the last S/390 to support an ICA was the 9221, which is no longer marketed by IBM and that the 9221 is not supported by the current VM operation system (z/VM). It is supported by VM/ESA; however, VM/ESA is no longer marketed by IBM. Consequently, the ICA is a reasonable option for you *only* if you happen to already have one and your BTAM application is running in a VSE environment or under VM (including as an MVS guest).

Note: One of the key reasons for the development of communication controller products was to offload communication line handling from the host. If you are moving lines from a communication controller to an ICA, be careful to consider the possible increase in host CPU cycles it may require.

- ▶ You may be able to migrate your BSC lines from your communication controllers to an external device such as a Hydra Systems Hydra 3000. For more information, see:

<http://www.hydrasystems.com>

According to Hydra Systems, the Hydra 3000 works with any size mainframe and with VM, VSE, MVS, and OS/390®; however, while their technical support was confident that the Hydra 3000 should work for BSC 3270 traffic, they did not know of any organizations that are currently using the Hydra 3000 for BSC 3270.

Figure 9-2 illustrates the “box replacement” option.

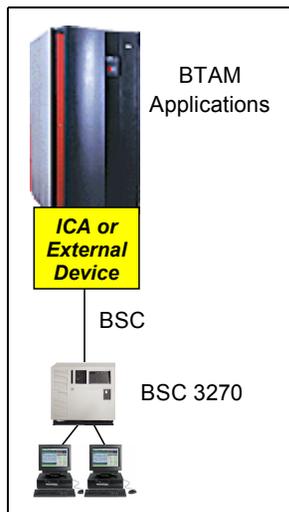


Figure 9-2 ICA or external device “box replacement” for BSC 3270

Preserving BTAM applications using a VM console

Interestingly, any 3270 device (including terminals or workstations running TN3270) can access a VM console and, through it, connect into a particular virtual machine as though it were a locally attached 3270 controller. With this capability, it may be possible to preserve your BTAM application environment by

moving it into an MVS guest system under VM. This approach, illustrated in Figure 9-3, would allow the use of newer access technologies (for example, SNA 3270, TN3270, or Web browsers) for your BTAM application and free you from dependence on old hardware devices.

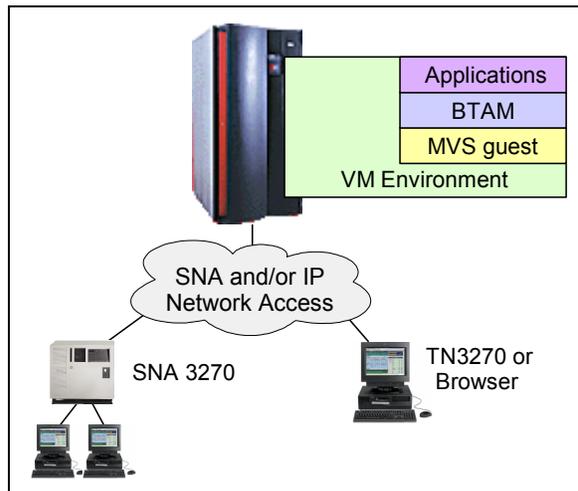


Figure 9-3 Preserving BTAM applications using VM console

In this example we suppose that the BTAM application is running under MVS. This solution would also apply to BTAM applications running under VSE or VM.

In order to preserve your BTAM application using VM, you must:

- ▶ Migrate your BTAM environment into an MVS guest system under VM.
- ▶ Change the BTAM definitions so that they point to local non-SNA 3270 devices instead of remote BSC devices.
- ▶ Add SPECIAL 3270 statements to the directory entry of the BTAM MVS guest and have the appropriate local 3270 device addresses generated in your MVS system.

You can also set up those devices so that they are automatically dialed to the virtual local 3270 device address specified on the SPECIAL statement.

You can then upgrade your BSC devices to modern workstations with no change to your BTAM applications because any SNA 3270, or TN3270, devices can appear as locally attached non-SNA devices to the MVS guest and access the BTAM application.

Note: Since this VM console solution is based upon 3270 device protocols, this solution works for BSC 3270 but will not work for BSC RJE.

Preserving BSC 3270 devices using protocol conversion

If you have BSC 3270 devices that you need to continue to support, you may be able to use a BSC 3270 protocol conversion solution. According to JBM Electronics Co., their Gateway products can provide BSC 3270 to TN3270 conversion:

<http://www.jbmelectronics.com>

Alternatively, according to INETCO Systems, their INETCO Connect product can provide BSC 3270 to SNA 3270 conversion:

<http://www.inetco.com>

Recommendations

While the “box replacement” alternative is clearly the simplest choice, using VM to preserve your BTAM application may be a more strategic option. The VM-based solution could facilitate new application possibilities including Web-enabling BTAM applications. It also enables access to BTAM applications from workstations with SNA 3270, TN3270, or Web browsers eliminating support requirements for aging BSC terminals and freeing up user desk space by no longer requiring workstations dedicated to a single application.

9.2.2 BSC RJE connection to non-SNA applications

BSC RJE is most often associated with IBM 2780 and 3780 device protocols. While the original 2780 and 3780 devices were introduced in the late 1960s to support remote job entry applications (such as submitting punched-card batch programs for execution), the communication protocols that IBM developed for them became the de facto standard for file transfer for much of the computing industry and is still widely used today, almost 40 years later. BSC RJE protocols are also used in some bank teller machines and store cash registers.

Important: In many cases, BSC RJE is used for interconnection between *different organizations*. Consequently, any attempt to change such an environment will require coordinating the change with those other organizations. The inter-organizational coordination of such changes often poses much more difficult challenges than do technical issues. Some approaches that may help include:

- ▶ Strive to work within international and/or industry standards such as message queuing (MQ), electronic data interchange (EDI), and extensible markup language (XML).
- ▶ Work closely with industry consortia and extranets. For example:
 - The automobile manufacturing industry has the Automotive Network Exchange.
 - The IVANS network provides similar interconnectivity for the insurance industry.
- ▶ Take advantage of value-added network services providers by letting them deal with the individual connectivity needs of your trading partners.

Explore secure virtual private network (VPN) connectivity across an IP connection or across the Internet.

BSC RJE to SNA RJE protocol conversion

If you have BSC RJE access requirements that you need to continue supporting, you may be able to use BSC RJE to SNA RJE protocol conversion. As illustrated in Figure 9-4 on page 219, the protocol conversion device would convert BSC 2780 or 3780 protocol into SNA 3770 protocol thereby allowing RJE access over any SNA path into your host environment and removing the dependence on EP-based access.

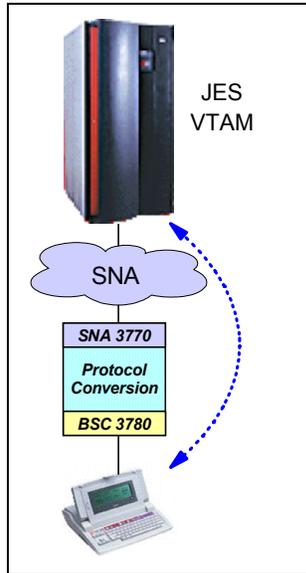


Figure 9-4 BSC RJE to SNA RJE protocol conversion

Two companies that indicated that their products provide such protocol conversion capabilities are:

- ▶ JBM Electronics Co. via their Gateway products at:
<http://www.jbmelectronics.com>
- ▶ INETCO Systems via their INETCO Connect product at:
<http://www.inetco.com>

This approach could be especially useful if you do not have the option of migrating the devices (for example, if the devices belong to organizations other than your own); however, you will need to investigate whether application changes are necessary.

Preserving BSC RJE connections using an ICA or Hydra 3000

If you wish to migrate BSC RJE from EP without changing your user devices or applications, you have two “box replacement” possibilities:

- ▶ If you have an Integrated Communications Adapter (ICA) and sufficient available ports on it, you may be able to migrate your BSC lines from your communication controllers to the ICA. Note that the last S/390 to support an ICA was the 9221, which is no longer marketed by IBM and that the 9221 is not supported by the current VM operating system (z/VM). It is supported by VM/ESA; however, VM/ESA is no longer marketed by IBM. Consequently, the ICA is a reasonable option for you *only* if you happen to already have one. As

illustrated in Figure 9-5, once a 3780-type RJE station has sent a job to a VM/RSCS or VSE/Power job entry subsystem, that job can be sent by NJE to an MVS system and that the NJE connection can be SNA even if the RJE jobs being shipped are non-SNA.

Note: One of the key reasons for the development of communication controller products was to offload communication line handling from the host. If you are moving lines from a communication controller to an ICA, be careful to consider the possible increase in host CPU cycles it may require.

- ▶ You can migrate your BSC lines from your communication controllers to an external device such as a Hydra Systems Hydra 3000. For more information, see:

<http://www.hydrasystems.com>

For RJE host access, the Hydra 3000 attaches to the mainframe over a byte channel thereby providing access to the job entry system. According to Hydra Systems, the Hydra 3000 works with any size mainframe and with VM, VSE, MVS, and OS/390.

Figure 9-5 illustrates the “box replacement” option. The option of connecting the MVS system directly to an external device such as a Hydra Systems Hydra 3000 is not depicted.

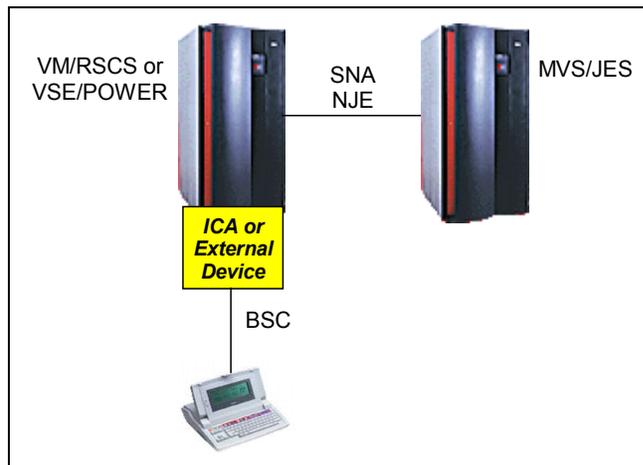


Figure 9-5 ICA or external device “box replacement” for BSC RJE

Migrating from RJE to message queuing (MQ)

There is a natural affinity between RJE and message queuing (MQ) applications. Both are typically used for file-transfer-type data movement between systems. While RJE is closely tied to the networking infrastructure (such as BSC or SNA), MQ provides independence from the networking infrastructure and, therefore, can enable your migration to a higher-performance and more strategic IP-based network.

Because it involves impact to *both* ends of your BSC RJE connections, migrating from RJE to MQ seems almost the same as the migration strategy of completely replacing the IT service in question. However, depending upon your IT environment, it may be possible to migrate to MQ with only small changes, or even potentially no change, at each end. In addition, a migration to MQ may leverage an investment that your organization has already made in MQ technology and may be strategically appropriate for the business process currently being supported by RJE.

Figure 9-6 illustrates an MQ solution that might be able to replace your use of RJE.

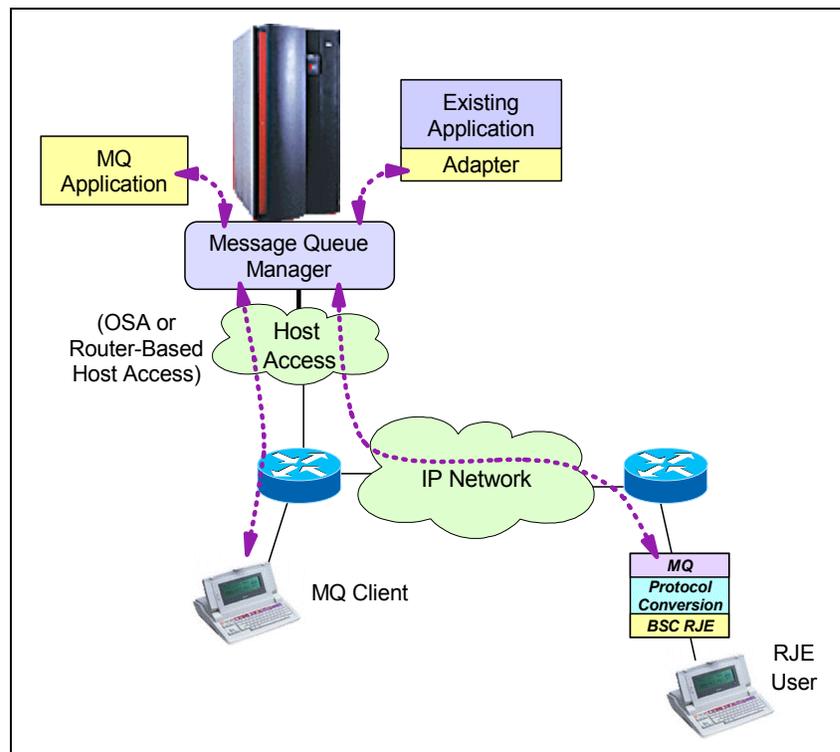


Figure 9-6 MQ alternative to RJE

Although the figure above shows a mainframe host-based MQ server, both the WebSphere MQ (formerly known as MQSeries) queue server and client software are available on a wide variety of platforms.

Using MQ would require implementation and use of MQ client interfaces at the communicating devices as well as at the host application. The left side of the illustration shows an application and a device that have been modified to use the MQ application program interface (MQI) directly. Such changes may or may not be possible depending upon the kinds of devices and applications you need to support. If you cannot change a device, it may be possible to use an intermediate device (such as a protocol converter) to provide access to the MQ environment. Likewise, if you cannot change an application, you may be able to develop an intermediate application to provide transparent access to the MQ environment.

To learn about services available from IBM that could help you with a migration to MQ, visit:

<http://www.ibm.com/software/integration/websphere/services/>

Advantages of migrating from RJE to MQ include:

- ▶ Higher availability and better management tools by migrating to a widely used commercially available program product such as the IBM WebSphere MQ products.
- ▶ Simplicity and scalability: Depending upon your environment, this solution could be implemented without any special, external, dedicated hardware platforms.
- ▶ Provides the option of leveraging your mainframe host environment or using an outboard platform as your business needs dictate.
- ▶ Higher performance by enabling migration of your RJE requirements onto your typically higher-bandwidth IP infrastructure.
- ▶ MQ flexibility eliminates client protocol dependencies and enables incremental network migrations.

A consideration of migrating from RJE to MQ is:

- ▶ Having to implement and support MQ interfaces for devices and applications

Recommendation

If message queuing is, indeed, a strategic technology for your IT organization, explore the possibility of migrating from RJE to MQ and pursue the MQ alternative if it seems feasible. Otherwise, a “box replacement” approach (using an ICA or an external device such as the Hydra 3000) may be your easiest

migration option. The protocol conversion alternative would require a bit more effort in considering the upstream SNA infrastructure requirements but would allow you to migrate your RJE traffic away from the current byte-channel and toward more flexible and strategic host connectivity options.

9.2.3 Start-stop terminal connection to non-SNA applications

The only alternative we can find for migrating start-stop devices communicating with mainframe non-SNA applications through EP is to change the application, the devices, or both.



Network Terminal Option (NTO)

IBM Product Number: 5735-XX7

This chapter discusses alternatives to the NTO support for communication between certain devices that use non-SNA (start-stop or BSC) protocols and SNA applications.

Note: In addition to the technical solutions discussed in this chapter, you should also consider the alternatives of eliminating the IT service in question or replacing that service with a completely new solution. See “Important migration-strategy alternatives” on page 95 for a discussion of each of those alternatives.

10.1 NTO product overview

NTO is an IBM-licensed program that resides in a communication controller. NTO runs under NCP and extends NCP capabilities by providing protocol conversion for a select group of non-SNA devices so that they can access SNA applications. Specifically NTO supports the following BSC and start-stop devices:

- ▶ BSC
 - IBM 2780/3780 Data Communications Terminals
 - IBM PABX Switching System Models 1750 and 3750 (emulating the 3780)

Note: Any 3780 or 3780 emulation through NTO requires a user-written application on the host; there is no host subsystem support for these devices.

- ▶ Start-stop
 - IBM 2740 Communications Terminal Model 1
 - IBM 2741 Communications Terminal
 - Teletypewriter Exchange (TWX) Service Models 33, 33ASR, and 35
 - World Trade Teletypewriter Terminals (WTTY)
 - IBM 3767 (when set in start-stop mode) terminals
 - IBM 1980/2980 financial terminals

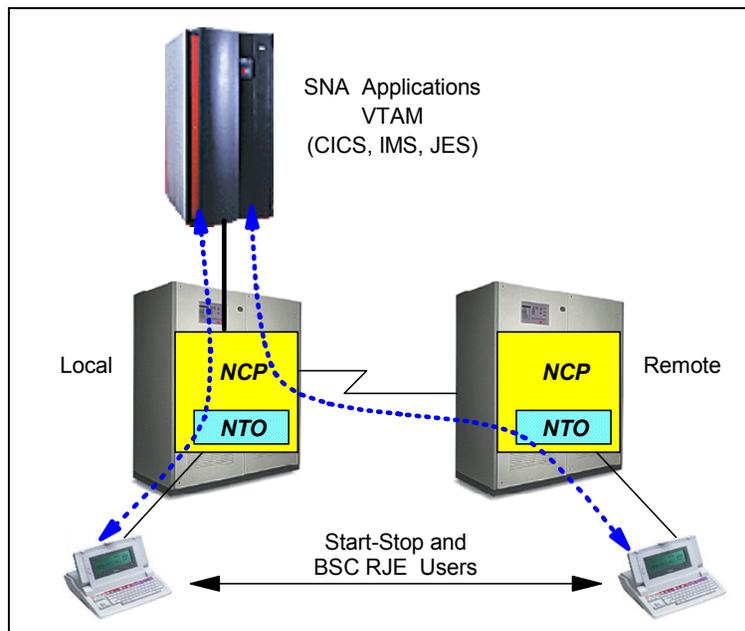


Figure 10-1 The role of NTO

As shown in Figure 10-1 on page 226, NTO running under NCP can take non-SNA start-stop and non-SNA BSC input and convert it to a line-by-line mode SNA stream. NTO support also includes ASCII-to-EBCDIC conversion. This enables the non-SNA devices mentioned above to communicate with VTAM host applications such as CICS and TSO. NTO is also used for SNA device connections using NRF (discussed further in Chapter 11, “Network Routing Facility (NRF)” on page 237). The SNA LU presented by NTO to the SNA application on behalf of the non-SNA device is an LU Type 1 (IBM 3767). The 3767 is a keyboard and printer-type device. It does not have a display, so the SNA application will be sending and receiving printer-like blocks of (often) 132 characters representing one line of print. In other words, the application treats the terminal much more like a printer than, say, a 3270-type display.

For additional technical information about NTO function and setup, see *Network Terminal Option Planning, Migration, and Resource Definition (Release 11)*, SC30-3347-08.

10.2 Functions and alternatives

NTO usage generally falls into one of the following three functional categories:

1. Connecting start-stop terminals to SNA mainframe applications
2. Connecting computer systems using non-SNA BSC RJE protocols to SNA mainframe applications
3. Connecting non-SNA computer systems peer-to-peer using NTO with NRF

We discuss alternatives for each of these functions below.

10.2.1 Start-stop terminal connection to SNA host applications

Following are the alternatives for the NTO function that supports connecting start-stop terminals to SNA mainframe applications.

Protocol conversion

As shown in Figure 10-2 on page 228, start-stop terminals can be supported by a special-purpose hardware device that converts start-stop to SNA 3767 protocol. The protocol conversion can also include ASCII-to-EBCDIC data stream conversion. For example, the INETCO Connect product from INETCO Systems provides this type of protocol conversion. See:

<http://www.inetco.com>

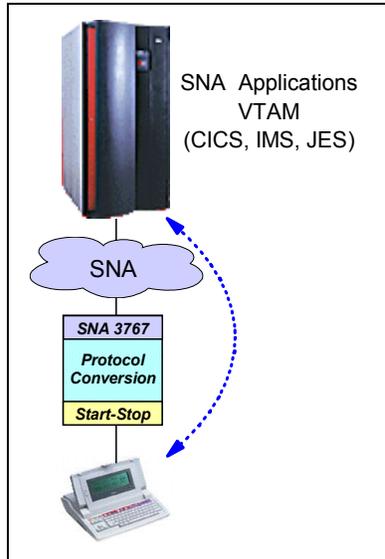


Figure 10-2 Start-stop terminal support

Using X.25-specific protocol conversion functions

Historically, X.25 support for start-stop terminal access to SNA host applications has included PCNE functionality, which converts start-stop over X.25 into SNA 3767 (see also Chapter 6, “X.25 NCP Packet Switching Interface (NPSI)” on page 187). This function is very similar to NTO conversion of start-stop directly to SNA 3767. It is therefore possible that you could preserve your current NTO-supported start-stop terminal to SNA host application environment by adding an X.25 component to your network and replacing NTO with a non-communication-controller-based program that includes PCNE support.

For example, both Comm-Pro Associates, Inc. (<http://www.comm-pro.com/>) and Computer Associates (<http://ca.com/>) appear to provide host-based programs that include PCNE support. As illustrated in Figure 10-3 on page 229, you could use an X.3 PAD function (available in RS/6000® software, in special-purpose hardware devices, and in some routers) and connect start-stop terminals to a router, such as those from Cisco Systems, Inc., that supports X.25 over TCP/IP (XOT). The encapsulated start-stop traffic can then be transported across your IP network and delivered through your host IP access infrastructure where the Comm-Pro HNAS product, or the Computer Associates Unicenter TCPaccess X.25 Server (formerly known as Solve:X.25) product, will provide PCNE support to deliver the start-stop terminal traffic to VTAM applications as SNA 3767 LU1 devices (the same format used by NTO).

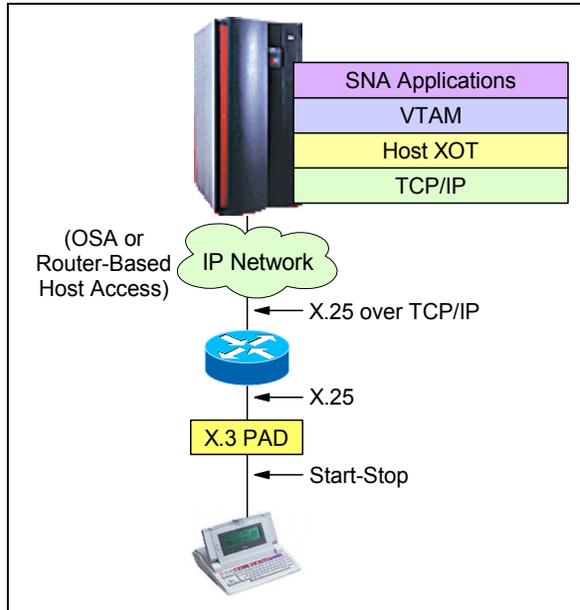


Figure 10-3 Host XOT solution

If you are using Cisco routers, according to Cisco Systems, you can connect the start-stop terminals directly into an asynchronous port on the router and the router software can perform both the X.3 PAD function and the XOT encapsulation.

Recommendation

Both the hardware-based protocol conversion and the X.25 host software protocol conversion can preserve your existing user devices and applications. The hardware solution is probably the easiest. However, you might select the X.25 solution if X.25 networking of your devices provides your organization with cost or connectivity benefits.

10.2.2 BSC RJE connection to SNA host applications

BSC RJE is most often associated with IBM 2780 and 3780 device protocols. While the original 2780 and 3780 devices were introduced in the late 1960s to support remote job entry applications (such as submitting punched-card batch programs for execution), the communication protocols that IBM developed for them became the de facto standard for file transfer for much of the computing industry and is still widely used today, almost 40 years later. BSC RJE protocols are also used in some bank teller machines and store cash registers. Following

are the alternatives for the NTO function that supports connecting non-SNA BSC 2780/3780-type devices to SNA host applications.

Important: In many cases, BSC RJE is used for interconnection between *different organizations*. Consequently, any attempt to change such a BSC RJE environment may require coordinating the change with those other organizations. The inter-organizational coordination of such changes often poses much more difficult challenges than do technical issues. Some approaches that may help include:

- ▶ Strive to work within international and/or industry standards such as message queuing (MQ), electronic data interchange (EDI), and extensible markup language (XML).
- ▶ Work closely with industry consortia and extranets. For example:
 - The automobile manufacturing industry has the Automotive Network Exchange
 - The IVANS network provides similar interconnectivity for the insurance industry
- ▶ Take advantage of value-added network services providers by letting them deal with the individual connectivity needs of your trading partners.
- ▶ Explore secure virtual private network (VPN) connectivity across an IP connection or across the Internet.

BSC RJE to SNA RJE protocol conversion

If you have BSC RJE access requirements that you need to continue supporting, you may be able to use BSC RJE-to-SNA RJE protocol conversion. As is illustrated in Figure 10-4 on page 231, the protocol conversion device would convert BSC 2780 or 3780 protocol into SNA 3770 protocol, thereby allowing RJE access over any SNA path into your host environment and removing the dependence on NTO-based access.

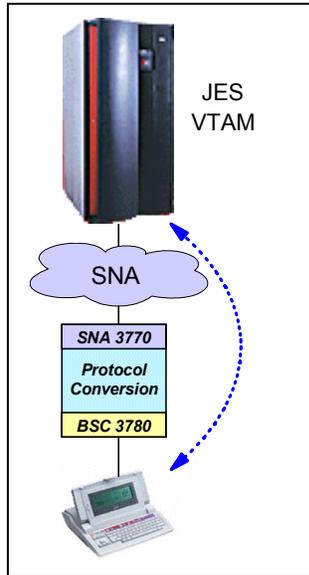


Figure 10-4 BSC RJE to SNA RJE protocol conversion

Two companies that indicated that their products provide such protocol conversion capabilities are:

- ▶ JBM Electronics Co. via their Gateway products (the host interface can be TCP/IP or serial) at:

<http://www.jbmelectronics.com>

- ▶ INETCO Systems via their INETCO Connect product at:

<http://www.inetco.com>

This approach could be especially useful if you do not have the option of migrating the devices (for example, if they are computer systems that belong to organizations other than your own); however, you will need to investigate whether application changes are required.

Migrating to SNA RJE

Change the BSC RJE protocol being used to SNA 3770 RJE. Determine what, if any, application changes may be required. SNA RJE emulation function is available on the AS/400®, RS/6000, and PC. For an example of an RS/6000 3770 SNA RJE package, go to the Web site of TPS Systems, Inc. and select Products, then Legacy/SNA Solutions, and look at the TPS/RJE product:

<http://www.tps.com>

For an example of a PC Windows-based 3770 SNA RJE package, visit the Web site of CQ Computer Communications Inc. and select software solutions and then the CQ-3770 Windows product:

<http://www.cq-comm.com/>

Preserving BSC RJE connections using an ICA or Hydra 3000

- ▶ If you have an Integrated Communications Adapter (ICA) and sufficient available ports on it, you may be able to migrate your BSC lines from your communication controllers to the ICA as shown in Figure 10-5. Once a 3780-type RJE station has sent a job to a VM/RSCS or VSE/Power job entry subsystem, that job can be sent by NJE to an MVS system. The NJE connection can be SNA even if the RJE jobs being shipped are non-SNA. The NTO conversion from non-SNA RJE to SNA RJE is not provided by this alternative. Consequently, JCL and application changes will be required.

Note: One of the key reasons for the development of communication controller products was to offload communication line handling from the host. If you are moving lines from a communication controller to an ICA, be careful to consider the possible increase in host CPU cycles it may require.

Note that the last S/390 to support an ICA was the 9221, which is no longer marketed by IBM and that the 9221 is not supported by the current VM operation system (z/VM). It is supported by VM/ESA; however, VM/ESA is no longer marketed by IBM. Consequently, the ICA is a reasonable option for you *only* if you happen to already have one.

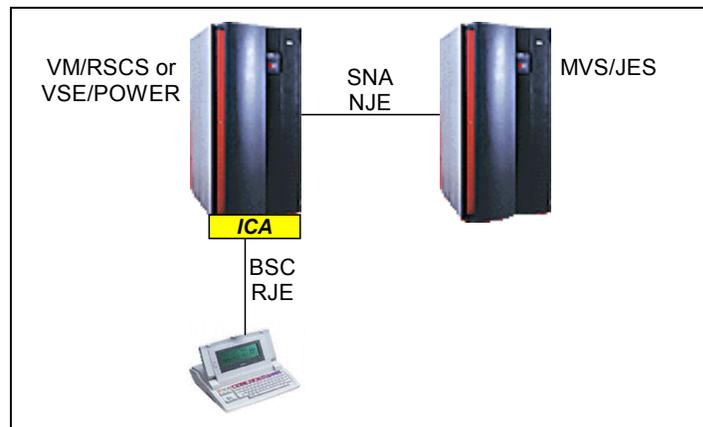


Figure 10-5 Preserving non-SNA BSC 2780/3780

The Hydra Systems, Inc. Hydra 3000 product could be an alternative to the ICA. The Hydra 3000 byte channel connects to VM, VSE, and MVS systems. For more information on the Hydra 3000, see:

<http://www.hydrasystems.com/>

Migrating from RJE to message queuing (MQ)

There is a natural affinity between RJE and message queuing (MQ) applications. Both are typically used for file-transfer-type data movement between systems. While RJE is closely tied to the networking infrastructure (such as BSC or SNA), MQ provides independence from the networking infrastructure and, therefore, can enable your migration to a higher-performance and more strategic IP-based network.

Because it involves impact to *both* ends of your BSC RJE connections, migrating from RJE to MQ seems almost the same as the migration strategy of completely replacing the IT service in question. However, depending upon your IT environment, it may be possible to migrate to MQ with only small changes, or even potentially no change, at each end. In addition, a migration to MQ may leverage an investment that your organization has already made in MQ technology and may be strategically appropriate for the business process currently being supported by RJE.

Figure 10-6 on page 234 illustrates an MQ solution that might be able to replace your use of RJE.

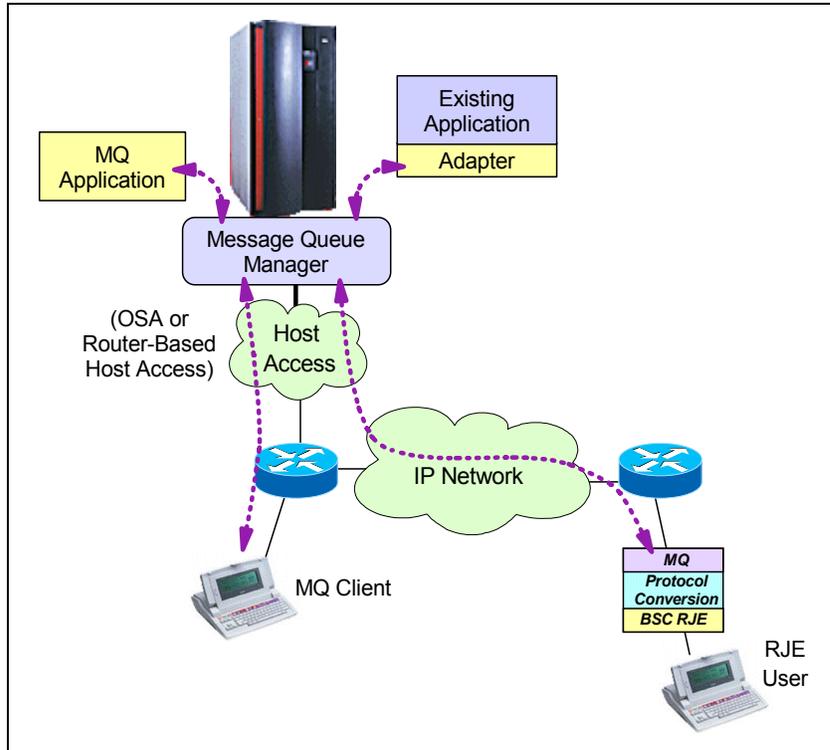


Figure 10-6 MQ alternative to RJE

Although the figure above shows a mainframe host-based MQ server, both the WebSphere MQ (formerly known as MQSeries) queue server and client software are available on a wide variety of platforms.

Using MQ would require implementation and use of MQ client interfaces at the communicating devices as well as at the host application. The left side of the illustration shows an application and a device that have been modified to use the MQ application program interface (MQI) directly. Such changes may or may not be possible depending upon the kinds of devices and applications you need to support. If you cannot change a device, it may be possible to use an intermediate device (such as a protocol converter) to provide access to the MQ environment. Likewise, if you cannot change an application, you may be able to develop an intermediate application to provide transparent access to the MQ environment.

To learn about services available from IBM that could help you with a migration to MQ, visit:

<http://www.ibm.com/software/integration/websphere/services/>

Advantages of migrating from RJE to MQ include:

- ▶ Higher availability and better management tools by migrating to a widely used commercially available program product such as the IBM WebSphere MQ product.
- ▶ Simplicity and scalability: Depending on your environment, this solution could be implemented with no special external dedicated hardware platforms.
- ▶ Provides the option of leveraging your mainframe host environment or using an outboard platform as your business needs dictate.
- ▶ Higher performance by enabling migration of your RJE requirements onto your typically higher-bandwidth IP infrastructure.
- ▶ MQ flexibility eliminates client protocol dependencies and enables incremental network migrations.

A consideration of migrating from RJE to MQ is:

- ▶ Having to implement and support MQ interfaces for devices and applications

Recommendation

If message queuing is, indeed, a strategic technology for your IT organization, explore the possibility of migrating from RJE to MQ and pursue the MQ alternative if it seems feasible. Otherwise, either migrating to SNA RJE or using protocol conversion would involve the least risk of change to the application. Attempting to preserve the BSC RJE using an ICA or an external gateway device (such as the Hydra 3000) would require changing a current SNA application to pre-SNA, which seems like a poor investment of time and resource to arrive at a difficult-to-manage solution.

10.2.3 Peer-to-peer connection of non-SNA devices

NTO is often used in conjunction with the NRF product to support peer-to-peer communication between non-SNA devices. If you are using NTO in conjunction with NRF, determine whether tunneling through a router connection will achieve the peer-to-peer connectivity you require. If NTO protocol conversion is required, you should examine the alternatives for NTO described in this chapter and fit them with the NRF alternatives discussed in 11.2.2, “Peer-to-peer connections involving non-SNA devices” on page 245.



Network Routing Facility (NRF)

IBM Product Number: 5668-963

This chapter discusses alternatives to the NRF support for communication between SNA-dependent LUs where neither LU may be located in a VTAM host. In addition, we cover environments where NRF is used with the NTO product to transport non-SNA traffic over your SNA network.



Note: NRF can now be run on your mainframe, along with NCP, using the Communication Controller for Linux on System z9 and zSeries (CCL) product, 5724-J38. In many cases, the CCL option will provide the most transparent migration from your IBM 3745 Communication Controllers.

See Chapter 16, “Communication Controller for Linux on System z9 and zSeries (CCL)” on page 283 for additional information about CCL.

Note: In addition to the technical solutions discussed in this chapter, consider the alternatives of eliminating the IT service in question or replacing that service with a completely new solution. See “Important migration-strategy alternatives” on page 95 for a discussion of each of those alternatives.

11.1 NRF product overview

NRF is an IBM-licensed program that resides in a communication controller together with NCP. As illustrated in Figure 11-1, NRF provides an SNA message routing function directly between external SNA devices (such as automated teller machines, retail store controllers, and non-VTAM application systems) without requiring the messages to go through a host processor or between an external device and application programs running in the host processor. While VTAM is involved in all NRF session establishments, the actual traffic flows through NRF once the sessions are established. NRF can maintain its terminal sessions and continue routing functions even in the event of a failure of the owning VTAM system. Essentially, NRF enables peer-to-peer communication between SNA secondary LUs. Without NRF, secondary LUs can only communicate with primary LUs, such as applications under the control of VTAM on the mainframe host.

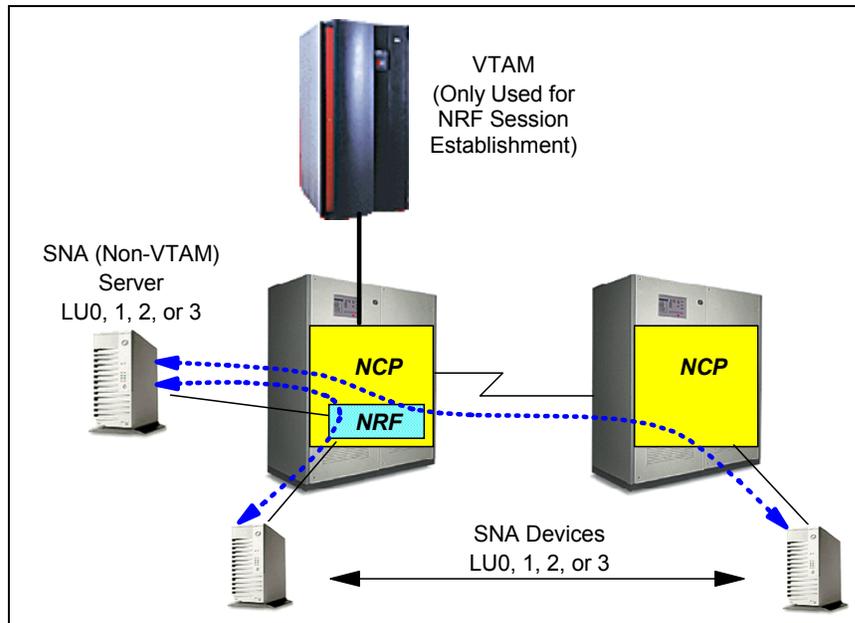


Figure 11-1 The role of NRF

NRF provides a primary LU image that functions as a point of service in the network to which secondary LUs can connect and have their messages routed directly to other secondary LUs. NRF provides a very flexible message routing capability including the ability to “spray” messages from a device to multiple recipients.

For support of non-SNA devices, NTO implements an SNA secondary LU that can connect to NRF. As described in Chapter 10, “Network Terminal Option (NTO)” on page 225, NTO can convert non-SNA 3780 or TWX device protocols into IBM 3767 SNA LU type 1 traffic. In this manner, an SNA Store Controller (PU Type 2 or 2.1) using one or more of its LU Type 1 or 2 LUs can connect to NRF which, in turn, connects to NTO which provides 3767 LU emulation for non-SNA RJE or start-stop terminal devices; thereby allowing message flows between secondary LUs and non-SNA end points. In addition, NRF (in conjunction with NTO) is used in some environments to enable the use of an SNA network for the transport of non-SNA peer to peer traffic.

For more information about NRF see:

- ▶ *NRF General Information*, GC27-0594-06
- ▶ *NRF Licensed Program Specifications*, GC27-0595-12
- ▶ *NRF Planning*, SC27-0593-10
- ▶ *NRF Migration, Resource Definition, and Customization*, SC31-6203-05

11.2 Functions and alternatives

NRF usage generally falls into one of the following two functional categories:

- ▶ Peer-to-peer SNA communication (between pre-APPN devices)
- ▶ Peer-to-peer communication where at least one of the devices is non-SNA

We discuss alternatives for each of these functions below.

11.2.1 Peer-to-peer connection of SNA devices (before PU type 2.1)

NRF supports peer-to-peer communication across an SNA network where neither peer is in a VTAM host. In such environments, the primary function that needs to be replaced is the NRF SNA message routing function. Alternatives for migrating from running NRF on your IBM 3745 Communication Controllers (which we discuss in the following sections) include:

- ▶ Running NCP and NRF using Communication Controller for Linux on System z9 and zSeries (CCL)
- ▶ Host-based SNA message routing
- ▶ Using a server for message routing
- ▶ Message queuing (MQ)



Running NCP and NRF using CCL

Figure 11-2 on page 240 illustrates the Communication Controller for Linux on System z9 and zSeries (CCL) solution. Essentially, your NCP and NRF software

can be migrated from your IBM 3745 Communication Controller into CCL on your mainframe with only minor configuration changes. While IBM 3745 Communication Controller supported a wide variety of physical interfaces, the LAN-only physical interface support of the CCL can be accommodated with the use of routers, connecting through your campus and wide area network, for connectivity to the SNA devices. Many organizations have historically used routers for such connectivity because the physical interfaces on routers are usually less expensive than comparable interfaces on IBM 3745 Communication Controllers.

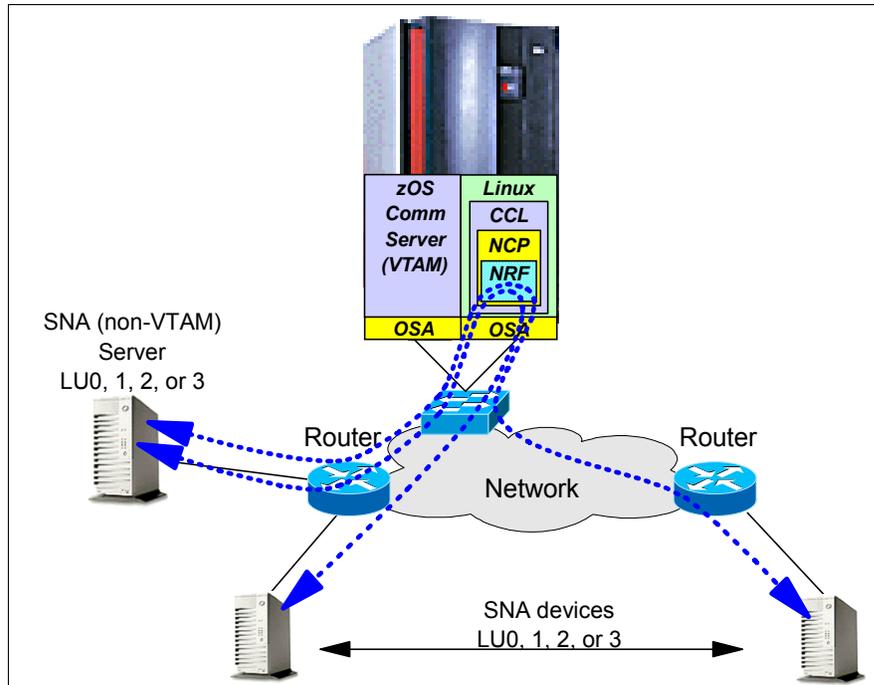


Figure 11-2 Running NCP and NRF using CCL

Advantages of the CCL solution include:

- ▶ CCL provides the *most transparent* migration path from IBM 3745 Communication Controllers.
- ▶ *CCL minimizes change* to the current network management environment.
- ▶ CCL leverages your highly reliable and increasingly cost-effective mainframe host environment.
- ▶ Simplicity and scalability: Depending upon your environment, a CCL solution could be implemented with no special external dedicated hardware platforms.
- ▶ CCL requires no new software development effort.

Considerations of the CCL solution include:

- ▶ CCL does not support NTO; therefore, this solution will not directly support peer-to-peer communication if any of the devices are non-SNA (see 11.2.2, “Peer-to-peer connections involving non-SNA devices” on page 245 for alternatives for addressing non-SNA requirements).
- ▶ CCL runs in a Linux for zSeries environment and may require additional skills and mainframe system resources (such as logical partitions and CPU cycles) to support it.

Note: CCL CPU cycles may be less expensive than other mainframe CPU cycles because Linux partitions can use Integrated Facility for Linux (IFL) processors.

Host-based SNA message routing

Figure 11-3 illustrates the flow of messages if you were to develop a host-based message routing function. The “host NRF replacement” software could be a VTAM application that establishes SNA sessions with each of the outboard SNA devices such as ATMs, store systems, or servers. In much the same way as NRF does it today, VTAM LOGAPPL commands could be used to drive session establishments and the application could implement “routing tables” that indicate where messages need to be forwarded.

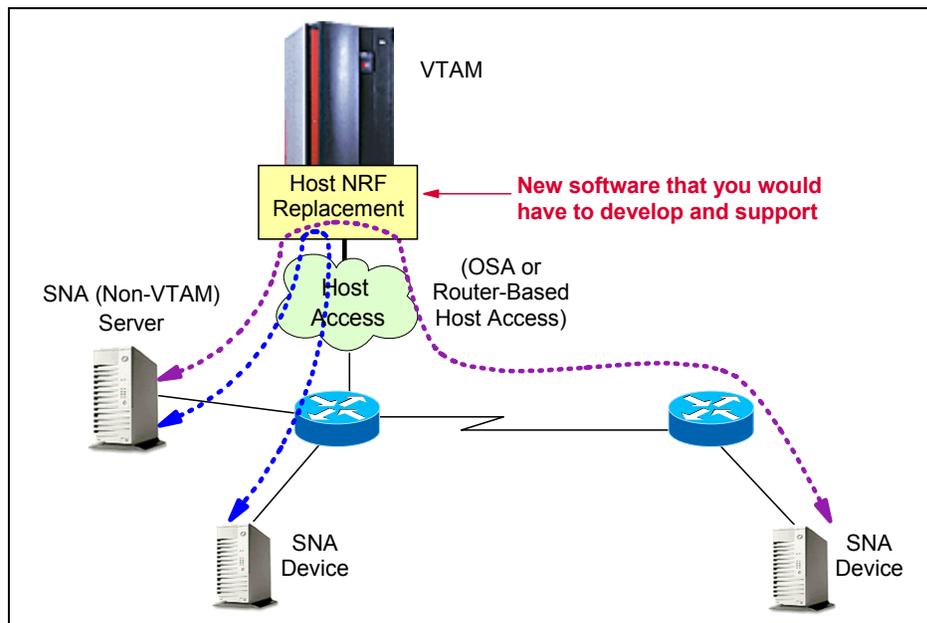


Figure 11-3 Host-based NRF replacement

Advantages of the host-based NRF replacement solution include:

- ▶ Simplicity and scalability: Depending upon your environment, this solution could be implemented with no special external dedicated hardware platforms.
- ▶ Leveraging your highly reliable and increasingly cost-effective mainframe host environment.
- ▶ Good host-based network management (NetView) visibility of the devices and sessions.

Considerations of the host-based NRF replacement solution include:

- ▶ Having to write and support new software.
- ▶ Using additional host cycles for message routing.
- ▶ This solution offers no synergy with protocol conversion devices you may require to support diversity of protocols used by the communicating devices.

Using a server for message routing

Figure 11-4 illustrates the flow of messages if you were to develop an RS/6000 or PC-based message routing function. The “NRF replacement” software could be an application, running in an SNA PU type 2.1 node, which has SNA sessions with each of the outboard SNA devices such as ATMs, store systems, or servers. In much the same way as NRF does it today, the application could implement “routing tables” that indicate where messages need to be forwarded.

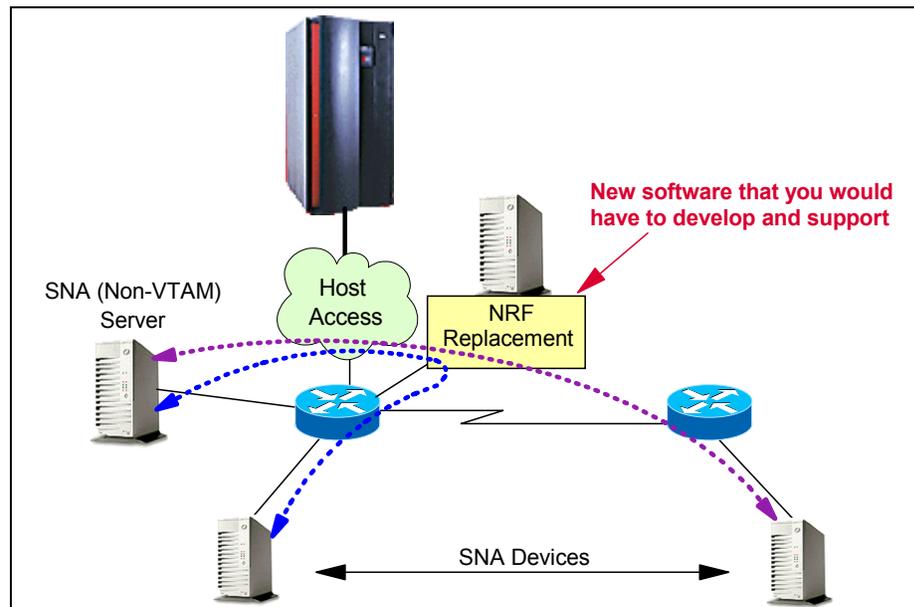


Figure 11-4 RS/6000 or PC-based NRF replacement

Advantages of the outboard NRF replacement solution include:

- ▶ Does not require host cycles for message routing.
- ▶ If your organization's strategic program development platform is the RS/6000 or PC environment rather than the mainframe, this alternative may allow use of your preferred environment.

Considerations of the outboard NRF replacement solution include:

- ▶ Having to write and support new software.
- ▶ Scalability: For large environments, you may have to have a number of systems working together to handle message routing requirements.
- ▶ Availability: Reliability of RS/6000 or PC-based solutions cannot match that of the mainframe environment and may be complicated by having to have several systems to handle the workload.
- ▶ Host-based network management (NetView) visibility of the devices and their sessions may be limited.

Message queuing (MQ) to replace NRF

Message queuing has become a very popular means of communication among diverse systems. As shown in Figure 11-5, by using software such as IBM WebSphere MQ (formerly known as MQSeries), you can avoid the expense, time, and risk involved in developing your own message routing capability.

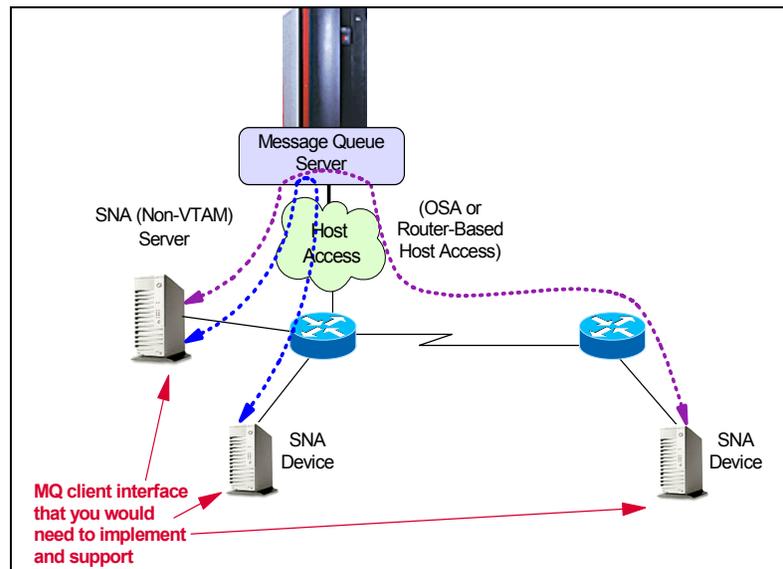


Figure 11-5 Using message queuing to replace NRF

Although the figure above shows a mainframe host-based MQ server, both the WebSphere MQ server and client software are available on a wide variety of platforms. While you would not need to develop the message routing software as was required in some of the alternatives discussed earlier, using MQ would require implementation and use of MQ client interfaces at the communicating devices. Such changes may or may not be possible, depending upon the kinds of devices you need to support. If you cannot change an end device, it may be possible to use an intermediate device (such as a protocol converter) to provide access to the MQ environment (see “NRF replacement with protocol conversion” on page 246).

Advantages of the MQ-based NRF replacement solution include:

- ▶ Potentially lower risk, lower cost, and faster implementation by using a commercially available program product such as IBM WebSphere MQ.
- ▶ Higher availability and better management tools by using a commercially available program product such as IBM WebSphere MQ.
- ▶ Simplicity and scalability: Depending upon your environment, this solution could be implemented with no special external dedicated hardware platforms.
- ▶ Provides the option of leveraging your mainframe host environment or using an outboard platform as your business needs dictate.
- ▶ Provides for migration of your message routing requirements onto your strategic IP infrastructure.
- ▶ MQ flexibility eliminates client protocol dependencies and enables incremental network migrations.

A consideration of the MQ-based NRF replacement solution is:

- ▶ Having to implement and support new client interfaces at all communicating devices

Recommendation

Historically, mainframe-based message switching was undesirable due to the higher cost of host CPU cycles. That was one reason why NRF was developed to run on the communication controller platform. However, the cost of host CPU cycles has declined substantially over the years and host message switching has gained in popularity because of mainframe reliability, scalability, security, and manageability. Consequently, you should consider a host-based solution.

Running NCP and NRF using Communication Controller for Linux on System z9 and zSeries (CCL) will provide by far the easiest, most transparent migration path from IBM 3745 Communication Controllers. In addition, CCL CPU cycles may be less expensive than other mainframe CPU cycles because Linux partitions can use Integrated Facility for Linux (IFL) processors. In most cases,

message queuing solutions, such as the IBM WebSphere MQ products, will provide a more strategic solution for your organization. In fact, you may be able to leverage an MQ server already in use in your IT environment.

11.2.2 Peer-to-peer connections involving non-SNA devices

As shown in Figure 11-6, using NTO and NRF together, communication controllers can support peer-to-peer communications where one or both of the devices is non-SNA. This is accomplished by connecting non-SNA devices to NTO, connecting NTO to NRF, and then using NRF SNA message routing capabilities. Picture this as an onion with NRF in the core, NTO in the middle, and the communicating devices at the edge.

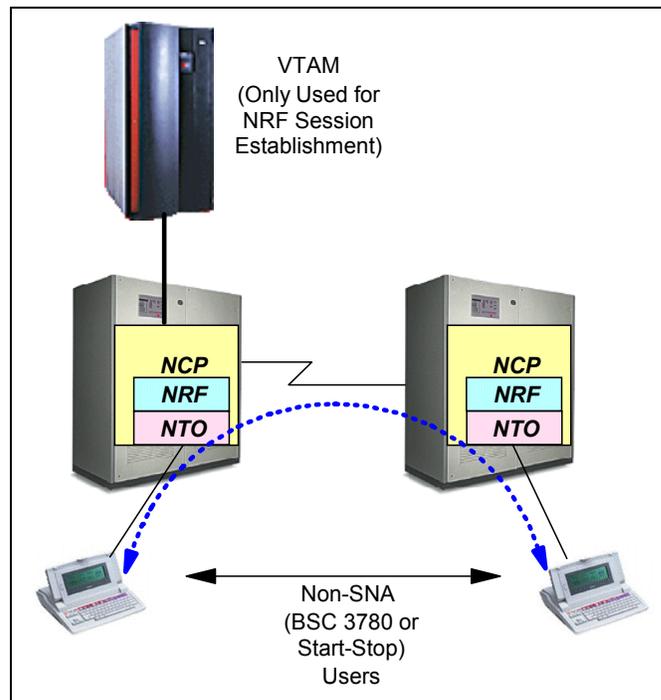


Figure 11-6 Using NRF and NTO together

The following alternatives support such connectivity without NRF and NTO.

Router tunneling

If the devices that are being interconnected are compatible with one another, a multiprotocol router function called *tunneling* (Figure 11-7 on page 246) may be sufficient to replace the combination of NRF and NTO. Cisco Systems provides such a capability, which they call *Block Serial Tunneling (BSTUN)*.

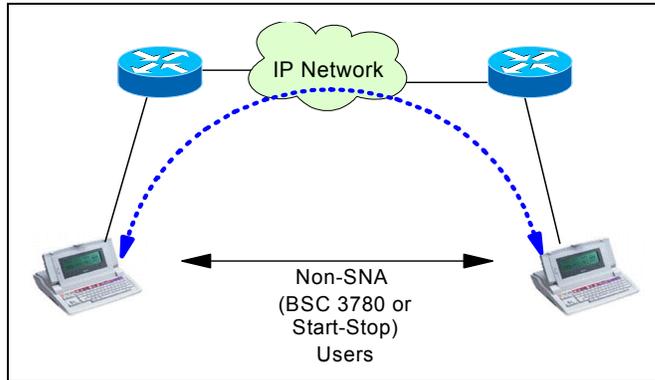


Figure 11-7 Router tunneling

The most significant advantages of this solution are its simplicity and the fact that the solution leverages your IP infrastructure.

You may be able to supplement the router tunneling solution by using protocol conversion equipment to resolve any device incompatibilities. Following are the Web sites for companies that appear to have products that might help in addressing such protocol conversion requirements:

- ▶ TPS Systems, Inc. at
<http://www.tps.com>
- ▶ CQ Computer Communications Inc. at:
<http://www.cq-comm.com/>

NRF replacement with protocol conversion

One approach to replacing the functions provided by NRF and NTO together would be to first address the challenge of replacing NRF function (as discussed in 11.2.1, “Peer-to-peer connection of SNA devices (before PU type 2.1)” on page 239), and then supplement it with protocol conversion as necessary to address the needs of the communicating devices. In such a solution, the protocol converter would need to either convert the BSC or start-stop traffic into:

- ▶ SNA (to use an SNA -based message routing program) or
- ▶ An MQ-compatible client format (to use an MQ-based solution)

Figure 11-8 on page 247 shows an MQ-based solution using software such as IBM WebSphere MQ. With MQ, while you would not need to develop your own message routing function, you would need to implement and use MQ client interfaces at the communicating devices. Such changes may or may not be possible depending upon the kinds of devices you need to support. If you cannot

change an end device, it may be possible to use an intermediate device (such as a protocol converter) to provide access to the MQ environment.

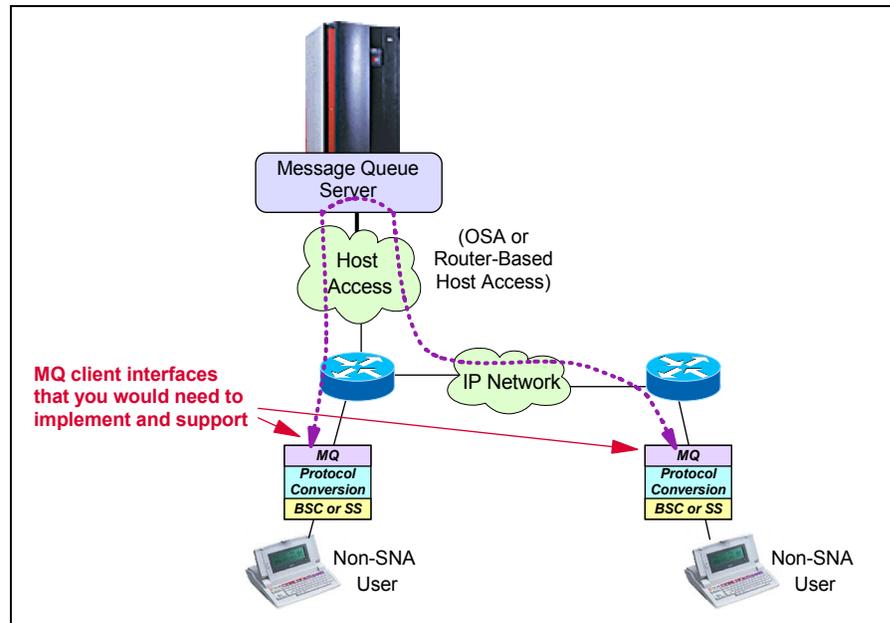


Figure 11-8 Using message queuing to replace NRF/NTO

Although the figure above shows a mainframe host-based MQ server, both the WebSphere MQ server and client software are available on a wide variety of platforms. Likewise, you have the option of writing your own SNA-based host or outboard message routing software.

While we found no protocol conversion devices that come prepackaged with MQ client interfaces, the IBM MQSeries family includes clients that are compatible with the base platforms of each of the following protocol conversion offerings:

- ▶ JBM Electronics Co.:
<http://www.jbmelectronics.com>
- ▶ INETCO Systems:
<http://www.inetco.com>

It may be possible to work with these companies, or others like them, to customize a solution to meet your needs.

Advantages of a NRF replacement with protocol conversion solution include:

- ▶ It enables migration toward MQ, a potentially more strategic means of accomplishing the particular business function being supported.

- ▶ Simplicity and scalability: Depending upon your environment, it may be possible to implement this solution with no special external dedicated hardware platforms.
- ▶ Provides the option of leveraging your mainframe host environment or using an outboard platform as your business needs dictate.
- ▶ Provides for migration of your message routing requirements onto your strategic IP infrastructure.

A consideration of a NRF replacement with protocol conversion solution is:

- ▶ Having to implement and support new client interfaces

Using a protocol conversion platform for message routing

In an NRF/NTO environment, some form of protocol conversion is required. In our investigation, we found two companies that offer products that provide diverse protocol conversion capabilities that may also be able to provide capability similar to NRF message routing. Customization or coding of some sort would be required to implement the message routing. Following are the Web sites for those companies:

- ▶ TPS Systems, Inc.:

<http://www.tps.com/>

- ▶ INETCO Systems:

<http://www.inetco.com>

Figure 11-9 illustrates the flow of messages if you were to use a protocol conversion platform for your message routing function.

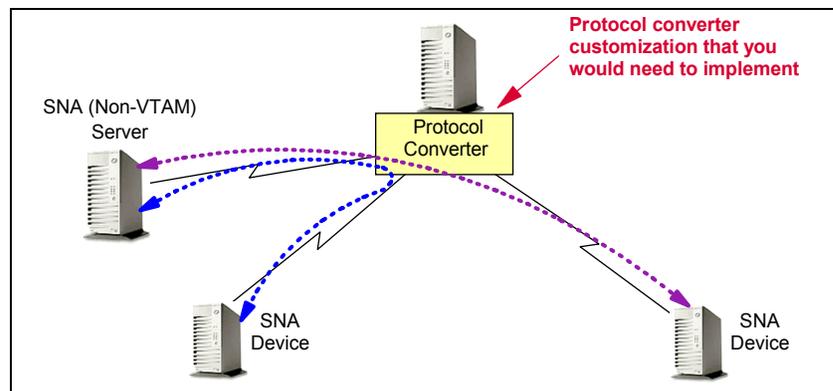


Figure 11-9 Protocol conversion platform-based NRF/NTO replacement

An advantage of the protocol conversion platform-based NRF/NTO replacement alternative is:

- ▶ For environments having a diversity of devices and protocols and, therefore, requiring protocol conversion, you may be able to leverage the protocol conversion platform for the NRF and NTO replacement function.

Considerations of the outboard NRF/NTO replacement alternative include:

- ▶ May require a significant amount of customization of the protocol conversion platform to achieve the desired solution.
- ▶ All of the communicating devices would have to be connected into a central protocol conversion platform, making such a solution impractical for large or distributed environments.
- ▶ Availability: Reliability of RS/6000 or PC-based solutions cannot match that of the mainframe environment and may be complicated by having to have several systems to handle the workload.

Recommendation

If the devices that are being interconnected are compatible with one another, router tunneling may be your most straightforward option. However, you should carefully consider the particular business function being supported because it may be the case that your organization is better served by pursuing a more strategic solution, such as MQ, to the business problem than simply tunneling legacy protocols.



Non-SNA Interconnection (NSI)

IBM Product Number: 5668-951

This chapter discusses alternatives to the NSI support for communication between binary synchronous communication (BSC) remote job entry (RJE) devices and host applications across an SNA network.

Note: In addition to the technical solutions discussed in this chapter, you should also consider the alternatives of eliminating the IT service in question or replacing that service with a completely new solution. See “Important migration-strategy alternatives” on page 95 for a discussion of each of those alternatives.

12.1 NSI product overview

NSI supports communications over the SNA backbone network between non-SNA host subsystems (RSCS, JES, Power) and:

- ▶ BSC RJE devices
- ▶ Other non-SNA host subsystems

Such BSC devices are dedicated to a specific application through a communication controller on a byte channel.

For those facilities that are controlled by NSI, the BSC data and control characters are encapsulated in SNA and transported through the network. When the data exits NSI and is sent to the terminal (via a BSC RJE line) or to the subsystem (via an emulation subchannel), the SNA envelope is removed and the original BSC data stream is presented to the receiving facility.

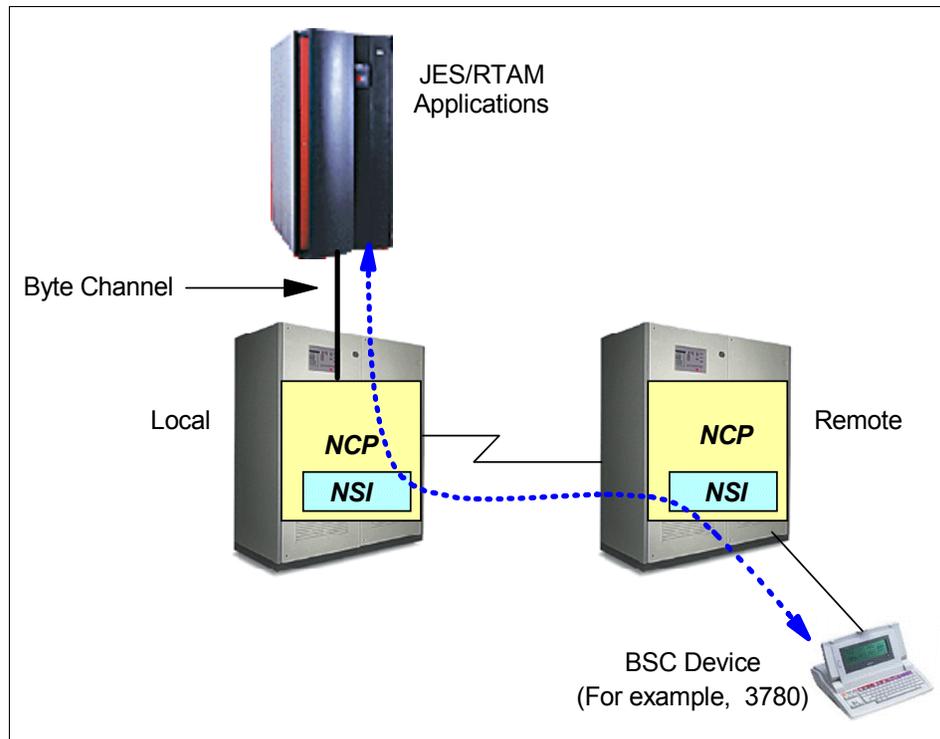


Figure 12-1 The role of NSI

12.2 Functions and alternatives

The NSI product uses encapsulation to provide two functions:

- ▶ Supporting non-SNA host NJE traffic over an SNA network
- ▶ Transporting traffic between BSC RJE devices and their host application across an SNA network

12.2.1 Non-SNA NJE to NJE connections between hosts

Historically, NSI supported an SNA network connection between non-SNA NJE subsystems. While this was originally an important capability to allow network migrations to SNA, most companies have since migrated their non-SNA NJE application to SNA NJE without issues. If your organization is using NSI for such non-SNA NJE applications, consider migrating those applications to SNA NJE.

12.2.2 Transporting BSC traffic across your network

If you are using NSI to carry BSC traffic over your SNA wide area network, a multiprotocol router function called “*tunneling*” may be sufficient to replace the NSI transport of BSC traffic across your network. Instead of the BSC 3780 traffic entering the network at a communication controller NSI/NCP point of presence, a router capable of tunneling BSC within IP can be used. Cisco Systems provides such a capability, which they call *Block Serial Tunneling (BSTUN)*. This alternative is shown in Figure 12-2 on page 254. For access to the host, the BSC traffic from the routers would require an ICA or Hydra 3000. The ICA and Hydra 3000 for BSC RJE access to the host are discussed further in 9.2.2, “BSC RJE connection to non-SNA applications” on page 217.

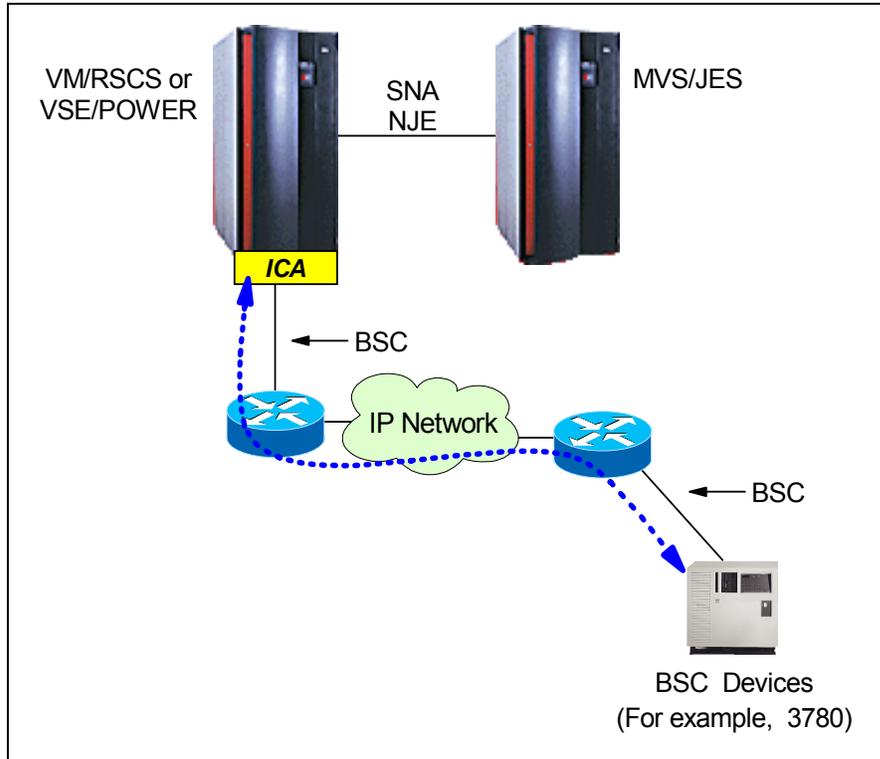


Figure 12-2 Router tunneling

The most significant advantages of this solution are its simplicity and the fact that it can leverage your IP infrastructure.



Teleprocessing Network Simulator (TPNS)

IBM Product Number: 5688-121

TPNS enables you to test your applications and their supporting networks by simulating them so that you can determine the impact of current and projected transaction loads on your host processor and network configuration. This chapter explores alternatives for the set of functions provided by TPNS.

Note: In addition to the technical solutions discussed in this chapter, you should also consider the alternatives of eliminating the IT service in question or replacing that service with a completely new solution. See “Important migration-strategy alternatives” on page 95 for a discussion of each of those alternatives.

13.1 Teleprocessing Network Simulator (TPNS)

TPNS is a collection of programs that generates transactions based upon user scripts in order to test your IT infrastructure. TPNS provides code that can be run in mainframe host computers (with either TCP/IP or SNA) or in IBM communication controllers. When TPNS is run in a mainframe host, it can generate transactions for any host in the network. When TPNS is run in an IBM communication controller (as shown in Figure 13-1), it can be used to simulate your SNA network. TPNS is therefore valuable for load testing of new applications (*function testing*), load testing of modifications to existing applications (*regression testing*), and for testing the effect of potential changes in load on currently existing applications (*performance testing*).

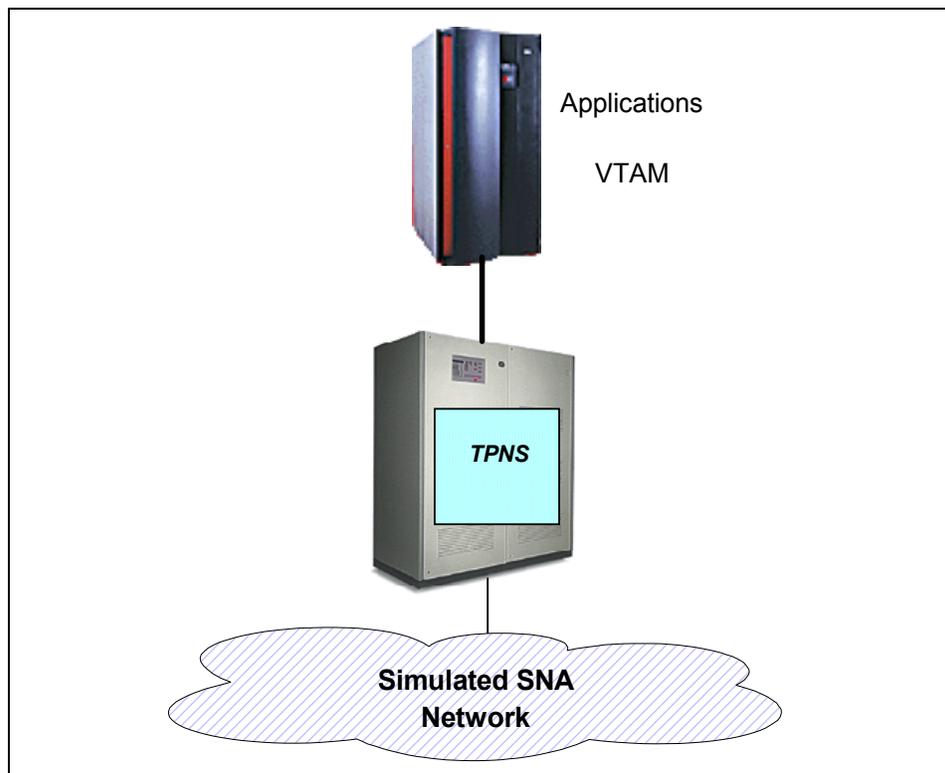


Figure 13-1 The role of TPNS in a communication controller

For more information about TPNS, see *TPNS Teleprocessing Network Simulator General Information, Version 3 Release 5, GH20-2487-08*.

13.2 Functions and alternatives

TPNS runs in a communication controller in order to enable simulation of SNA subarea and peripheral traffic.

13.2.1 Simulating subarea traffic

If you are migrating from your IBM communication controller platform, you can continue to run TPNS in your host and have it generate transactions for any other host. This will stress test both the application host and the connecting network. For example, the network connection between the TPNS host and the application host could be an IP network in which SNA is transported over DLSw or Enterprise Extender. Similarly, the network connection could be a simple channel-to-channel (CTC).

It should be noted that TPNS runs stand-alone (without NCP). There is no alternative for the functions that TPNS provides on a communication controller; however, since TPNS does not run regular production work, it would probably entail little risk to keep and run communication controller-based TPNS even in an unsupported mode, to perform testing as required to drive SNA traffic across a router-based network, or to test mainframe loading without requiring TPNS on a second host.



Part 3

Strategic solution technologies

In this part we provide more in-depth discussions of the most strategically important alternative solution technologies mentioned in the functional alternatives part of the book. It includes a chapter on each of the following key solutions:

- ▶ The Net390 architecture
- ▶ The Open Systems Adapter
- ▶ Communication Controller for Linux on System z9 and zSeries (CCL)
- ▶ TCP/IP
- ▶ TN3270
- ▶ Enterprise Extender (EE)
- ▶ Message queuing (MQ)



Net390 architecture

This chapter discusses the Net390 architecture, a best-practices network and systems design architectural framework developed by IBM Global Services through engagements with numerous clients. Because of its strategic focus, the Net390 architecture establishes a context for the other strategic solution technologies discussed in this book.

14.1 What is Net390?

Net390 is an architectural framework developed by IBM Global Services. The Net390 architecture was developed and successively refined through numerous consulting engagements with large organizations. These engagements focused on developing IT strategies and data center architectures and designs capable of delivering robust SNA and IP services:

- ▶ Providing for smooth scalability and continuous availability
- ▶ Improving responsiveness to e-business requirements
- ▶ Taking advantage of new technologies

The term “Net390” emphasizes that it is not the network alone, nor the mainframe host environment alone, but rather the synergy that can be achieved from looking at both together that results in an optimal IT infrastructure solution.

14.2 Why is Net390 strategic?

Your IT infrastructure is the set of components required to deliver the IT services needed by your organization. These components include:

- ▶ Servers (including middleware, applications, and storage)
- ▶ The network
- ▶ Client desktops
- ▶ Tools for network and systems management and administration

Because of the central role your network and servers play, it is particularly important that they are designed to provide high availability and smooth scalability. Such availability and scalability are key strengths of the IBM S/390 (or zSeries) servers. The Net390 architecture offers a best-practices structure for your mainframe environment and its network connectivity that will position you for the future while enabling you to incrementally migrate from your communication controllers as your business needs dictate. Being strategically focused, Net390 also provides a context for understanding the applicability of strategic technologies, such as those discussed in subsequent chapters of this book.

14.3 What you need to understand about Net390

Net390 ties the infrastructure portion of your IT strategy to the set of IT services your organization must deliver, and the service levels that must be met, in order to support the current and strategic needs of your business. We call this service-oriented focus the “service delivery model”.

The service delivery model focuses your organization on the delivery of a business-driven portfolio of services that facilitate the creation of information and the delivery of that information to clients. By focusing on the set of services offered by IT, your organization will be better positioned to:

- ▶ Optimize the infrastructure by reducing the costs associated with duplicate services, unnecessary services, or poorly delivered services.
- ▶ Harden the infrastructure to reduce business risk by focusing on the end-to-end delivery of key services.
- ▶ Enable new business initiatives faster because underlying, common services will already be in place with well-known capability, service levels, and cost structure.

Once you understand the set of services that must be supported by the IT organization, those services can be categorized as:

- ▶ *Application-unique*: Those that provide content creation and management. Examples of application-unique services include:
 - Enterprise Resource Planning
 - Supply Chain Management
 - Customer Relationship Management
 - Business Intelligence
- ▶ *Application support*: Those that facilitate content access. Examples of application support services include:
 - TN3270 servers
 - Web and host access
 - Directory, security, and mobility servers
 - File and print management

While an outage of an application-unique service may have a substantial impact on business operations, a loss of an application support service may be catastrophic, resulting in the outage of several or even all of the application-unique services. Consequently, the Net390 model, shown in Figure 14-1 on page 264, isolates application support services into a separate set of host logical partitions (LPARs) where they can be configured for the highest possible availability.

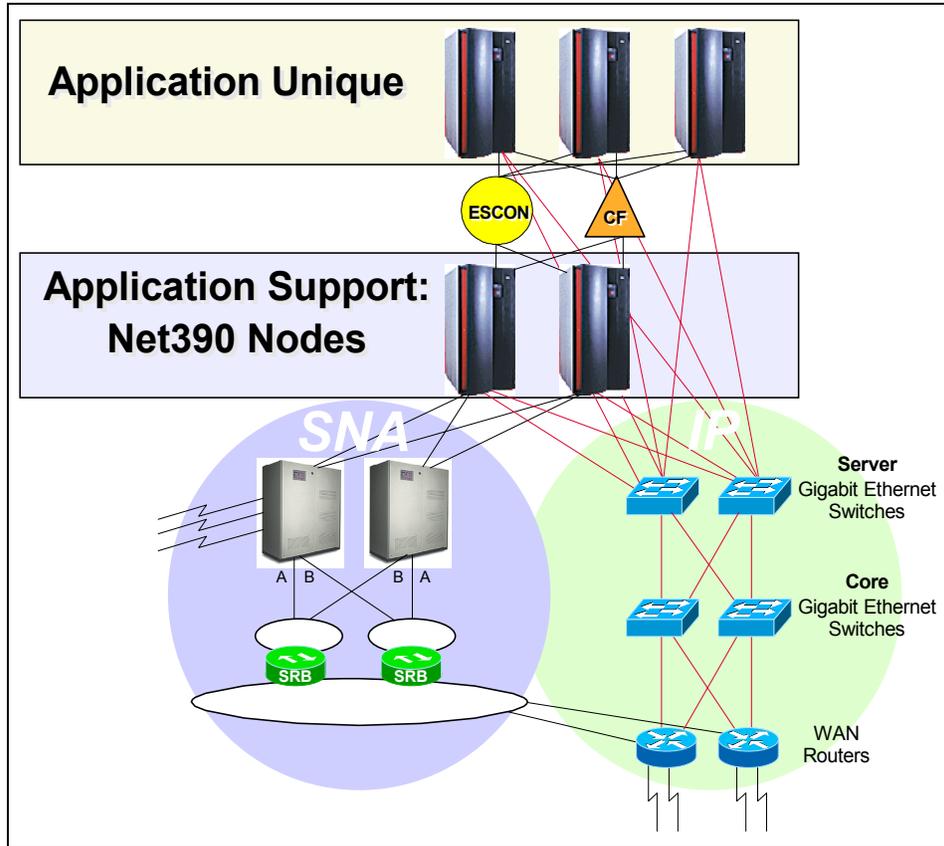


Figure 14-1 The Net390 model

The value of this approach is that the application-support nodes, which should be continuously available, can be isolated from the applications that run on the application-unique logical partitions. This minimizes any potential availability impact that applications might have on application-support services. It also enables an organization to deploy new services on the S/390 and zSeries servers with minimal impact on existing applications, thereby reducing the risk to the business as new S/390 applications are added or scaled.

Additionally, hardware and software upgrades to the Net390 nodes can be performed independently of the applications and, consequently, new support capabilities and enhancements can be implemented more quickly. Historically, in SNA environments, such logical partitions have been known as Communications Management Configuration (CMC) nodes with well-understood and proven benefits. In effect, Net390 extends the benefits of the CMC design to modern SNA and IP environments.

Key features of the Net390 architecture include:

- ▶ Net390 provides a framework that supports and exploits the scalability and availability strengths of the Parallel Sysplex. Within the host complex, all host nodes are interconnected via ESCON and Cross-System Coupling Facility (XCF). Host-to-network connectivity is provided by high-speed LAN interconnection such as gigabit Ethernet switches through OSA-Express interfaces (discussed in Chapter 15, “OSA-Express” on page 269). While it is not required for the application-unique nodes to connect to the network via the high-speed LAN infrastructure, in our experience such connectivity is usually inexpensive and worthwhile.
- ▶ ESCON attachments support communication controller connectivity to all hosts in the complex. By configuring SNA subarea connectivity from the communication controllers to application support (Net390) nodes, and defining APPN connectivity to application-unique nodes, application-unique nodes can be configured for SNA as APPN end nodes, thereby substantially simplifying the configuration of the SNA environment. This does not mean that SNA traffic must transit the Net390 nodes to reach the application-unique nodes, rather:
 - SNA traffic carried via Enterprise Extender technology (discussed in Chapter 19, “Enterprise Extender (EE)” on page 331) can exploit your high-performance and highly available IP-based infrastructure through the gigabit Ethernet switches and OSA-Express interfaces into the destination application-unique host.
 - Other SNA traffic, supported by the communication controllers, can be carried across the ESCON infrastructure into the destination host.
- ▶ In some situations, such as with BSC 3270 traffic, the application host requires direct subarea connectivity to the communication controllers. Such applications should be isolated into a “migration” LPAR (not shown in the illustration).
- ▶ Examples of the application-support functions that should be placed in Net390 nodes include:
 - Communication Management Configuration (CMC): SNA “ownership” of the network providing SSCP services (including both SSCP-PU and LU sessions)

Note: The CMC provides two important functions in an SNA network:

- ▶ It provides SNA ownership of network resources, in particular, communication controllers.
- ▶ By virtue of the CMC ownership of the SNA network, it provides an ideal location for your SNA network management focal point (typically implemented with the NetView product).

Because the CMC is specific to SNA, its importance will diminish as you migrate from your communication controllers and evolve your network from SNA to IP.

- Dependent LU Server (DLUS): working in conjunction with Dependent LU Requesters (DLURs) in the network to provide SNA boundary function for dependent LUs across an APPN (and Enterprise Extender) network
- Network Node (NN) server: providing topology, directory, and route selection services for the application-unique nodes in the complex as well as support for generic resources in a Parallel Sysplex
- Interchange Node (ICN): providing logical interconnectivity between the existing SNA environment and the APPN-based application-unique nodes
- SNI Gateway SSCP: supporting SNI connectivity from the communication controllers
- TN3270E servers: supporting IP-based access to the SNA application environment for 3270-type users (see Chapter 18, “TN3270” on page 311)
- Session Managers: used by many organizations for security as well as user convenience by managing SNA user application access (usually with a screen listing the applications that users are authorized to access)
- Host environment print services: for management of host print to both SNA and IP printers
- Network and systems management and automation of the host environment

The Net390 structure enables the S/390 and the zSeries servers in the Parallel Sysplex configuration to be the premiere service-delivery platforms in the industry. Leveraging the synergies between host systems and the network, Net390 combines the best attributes of SNA with the flexibility and reach of TCP/IP and the Web. The benefits of Net390 include:

- ▶ Continuous availability

Net390 enables continuous availability for users by integrating dynamic network routing protocols, for example, OSPF and High-Performance Routing (HPR), into the continuous operation characteristics of the Parallel Sysplex server.

- ▶ Scalability

By building on z/OS features such as generic resources, Workload Manager (WLM), and Service Policy Agent, and by extending these features into the network with Enterprise Extender, TN3270 Server, and Sysplex Distributor, Net390 provides scalability and enhances performance.

- ▶ Security

Net390 provides secure access to mainframe applications and data by leveraging the security features of the S/390 server (including world-class cryptographic capabilities) and the network.

- ▶ Investment protection

Net390 protects current investments in existing applications, routed networks and SNA equipment (including communication controllers) by enabling z/OS access from the intranet, Internet, and extranets. Connectors between WebSphere software on the S/390 server and back-end systems (CICS, IMS, and DB2® software) make existing data and application logic available to Web users.

- ▶ Predictable response times

Net390 enables the preservation of SNA service levels, and leverages new networking technologies that enable IP service levels through the use of bandwidth management, application content-specific prioritization, and workload distribution to achieve system-level response-time management.

- ▶ Flexibility

Net390 provides a flexible structure that is based on many existing and emerging technologies. As a best-practices approach, only those Net390 components that provide the greatest benefit, given current business requirements and time frames, need be implemented.

14.4 Strategies for exploiting Net390

The Net390 architecture positions your organization for incremental strategic migration by using the Net390 nodes as the integration point between your current IT infrastructure environment (including your IBM communication controllers) and your strategic future environment based upon OSA-Express, TCP/IP, TN3270E, and Enterprise Extender technologies.

IGS has developed a Net390 project based upon work with numerous clients. A Net390 project helps you to make the appropriate strategy, infrastructure, product, design, system management, and implementation decisions required to enable your IT infrastructure to support planned communication controller migration initiatives as well as leverage S/390 and zSeries servers into your e-business infrastructure. The main components of a Net390 project are:

1. Review current and proposed applications that involve the S/390 and zSeries servers and develop a service-delivery architecture capable of supporting your current and strategic business needs.
2. Make appropriate refinements to the Net390 architecture to support your business model, your investment in the S/390 and other servers, and your network.
3. Architect your data center network structure, including intra-data center, data center-to-data center, and disaster recovery center communications, as appropriate, to optimally support IP and SNA access to S/390 servers.
4. Design your Internet access network to provide the integration between your S/390 servers and the Internet.
5. Develop a detailed plan that identifies the steps to enable you to build your target environment.



OSA-Express

This chapter discusses OSA-Express, the IBM strategic family of mainframe host network adapters.

15.1 What is OSA-Express?

In short, OSA-Express is a family of adapters that provides direct mainframe-to-network attachment. The following sections discuss OSA-Express from an historical perspective and from an architectural perspective, and also provide a summary of available features.

For additional information about OSA-Express, visit the IBM zSeries Web site:

<http://www.ibm.com/servers/eserver/zseries/networking/>

15.1.1 An historical perspective

For 26 years, the parallel channel was the only mainframe attachment interface. Channel-attached IBM 3745 and 3174 controllers provided mainframe network connectivity.

As the IBM System/370™ parallel I/O infrastructure reached the limits in its ability to support large system images and performance, ESCON was introduced as a new I/O interconnection architecture supporting a new topology for high-speed, long-distance data exchange. At that time, the connectivity options to access mainframe resources from the network included:

- ▶ IBM 3745 Communication Controller for SNA
- ▶ Interconnect Controllers for TCP/IP
- ▶ Other similar devices with parallel or ESCON channel attachment interfaces.

The IBM zSeries organization believed that providing an industry-standard open interface would simplify the data center topology and lead to lower numbers of devices between the mainframe and the end users, while reducing complexity and the total cost of computing. Thus, the genesis of the Open Systems Adapter (OSA) began. With OSA, the mainframe could be connected directly to the data center network. The expense and complexity associated with the use of parallel or ESCON-attached Interconnect Controllers or routers, as well as communication controllers, could begin to be reduced.

Today, the current generation of OSA, called OSA-Express and OSA-Express2 (Gigabit Ethernet and 10 Gigabit Ethernet), provides the strategic mainframe interfaces to the network.

Terminology: If not specifically stated, the term OSA-Express applies to both the OSA-Express and the OSA-Express2 features throughout this chapter. Also, when a reference is made to Gigabit Ethernet it may also be applicable for 10 Gigabit Ethernet.

15.1.2 An architectural perspective

The 17 MBps Channel Request Handler (CRH) bus in S/390 CMOS machines attaches via Parallel channel, ESCON channel, or OSA-2 to the network. LAN bandwidth trends indicate that the CRH-based ESCON infrastructure is insufficient to meet future needs.

The IBM goal of is to offer balanced system performance in the mainframe. The internal bus structure and the I/O subsystem must keep up as the architecture is enhanced. OSA-Express uses the new Self-Timed Interconnect (STI) bus to overcome the speed limitations of the CRH bus. OSA-Express adapters connect directly to the STI bus.

- ▶ STI bus = 500 MBps - Used by OSA-Express
- ▶ STI bus = 1 GBps - Used by OSA-Express2

All of the OSA-Express features occupy one slot in an I/O cage.

- ▶ Each z990 I/O cage supports 20 OSA-Express features. The features have two ports each, with the exception of one port for the OSA-Express2 10 GbE LR feature. A maximum of 24 features can be installed in a z990 server.
- ▶ The z890 supports a maximum of 20 OSA-Express features except the smallest sub-uniprocessor (2086-110 allows 12). Each feature has two ports with the exception of one port for the OSA-Express2 10 GbE LR feature.
- ▶ The z800 and z900 I/O cages support 12 OSA-Express features. The features have two ports each. A maximum of 12 features (24 ports) can be installed in a z900 or z800 server. In a z900, the total number of OSA-Express and OSA-2 features must not exceed 12.

Note: OSA-Express features are not interchangeable between 9672 servers and zSeries servers.

See Table 15-1 on page 272 for details on the maximum number of OSA-Express ports supported by each zSeries server, based on feature.

The maximum combined number of OSA-Express2, OSA-Express, FICON, FICON Express, Crypto Express2, PCIXCC, PCICC, and PCICA features supported by each zSeries server family are:

- ▶ 60 features for the z990 server, 48 for the model 2084-A08
- ▶ 20 features for the z890 server, 12 for the model 2086-110
- ▶ 48 features for the z900 server
- ▶ 16 features for the z800 server

Not all features are supported by all servers.

15.1.3 Summary of features

Table 15-1 lists the OSA-Express feature support on all zSeries servers. Details for each feature follow the table.

Table 15-1 zSeries server support for OSA-Express

Feature name	Feature code	Maximum ports ^a				Connector type	Cable type	Maximum unrepeated distance ^b
		z800	z900	z890	z990			
OSA-Express ATM SM	2362	24	24	n/a	24 ^c	SC Duplex	SM 9 μ	20 km
OSA-Express ATM MM	2363	24	24	n/a	24 ^c	SC Duplex	MM 50 μ	2 km
							MM 62.5 μ	2 km
OSA-Express GbE LX ^d	1364	24	24	40	48	LC Duplex	SM 9 μ	5 km
							MCP ^f	550 m (500)
OSA-Express GbE LX	2364	24	24	24 ^e	24 ^e	SC Duplex	SM 9 μ	5 km
							MCP ^f	550 m (500)
OSA-Express GbE SX ^g	1365	24	24	40	48	LC Duplex	MM 62.5 μ	220 m (166) 275 m (200)
							MM 50 μ	550 m (500)
OSA-Express GbE SX	2365	24	24	24 ^e	24 ^e	SC Duplex	MM 62.5 μ	220 m (166) 275 m (200)
							MM 50 μ	550 m (500)
OSA-Express 1000BASE-T	1366	n/a	n/a	40	48	RJ 45	UTP Cat5	100 m
OSA-Express Fast Ethernet ^h	2366	24	24	24	24	RJ 45	UTP Cat5	100 m
OSA-Express Token Ring	2367	24	24	40	48	RJ 45	STP	100 m
							UTP	45 m
OSA-Express2 GbE LX	3364	n/a	n/a	40	48	LC Duplex	SM 9 μ	5 km
							MCP ^f	550 m (500)
OSA Express2 GbE SX	3365	n/a	n/a	40	48	LC Duplex	MM 62.5 μ	220 m (166) 275 m (200)
							MM 50 μ	550 m (500)
OSA-Express2 10 GbE LR	3368	n/a	n/a	20	24	SC Duplex	SM 9 μ	10 km

a. Maximum number of ports are 24 for model 2086-110 and 48 for model 2084-A08

- b. Minimum fiber bandwidth in MHz/km for multimode fiber links are included in parentheses where applicable.
- c. With RPQ 8P2258 this feature can be carried forward from a z900 when upgraded to a z990 in non-QDIO mode (OSE CHPID type) only.
- d. Feature 1364 is no longer orderable on z890 and z990. It has been replaced by feature 3364.
- e. Can be brought forward on an upgrade from z900 or z800
- f. Mode Conditioning Patch (MCP) cables enables the 1 Gbps single mode cards to connect to multimode fiber.
- g. Feature 1365 is no longer orderable on z890 and z990. It has been replaced by feature 3365.
- h. Feature 2366 is no longer orderable on z890 and z990. It has been replaced by feature 1366.

Multiple Image Facility (MIF)

MIF enables OSA-Express channels installed on zSeries servers to be shared among logical partitions.

Spanned channels

Spanning is the ability to configure channels to multiple Logical Channel Subsystems. When defined that way, the channels can be transparently shared by any or all of the configured logical partitions, regardless of the Logical Channel Subsystem to which the logical partition is configured.

15.2 Why is OSA-Express strategic?

OSA-Express is the fastest, simplest, and least expensive way of connecting mainframe hosts to networks. In particular, it offers the optimal host-to-network connectivity for TCP/IP which, itself, is a strategic technology for most organizations.

As data center networking evolves, the data center LAN infrastructure is increasingly moving to high-speed core switches: predominantly Gigabit Ethernet. To optimize connection to the campus backbone, the S/390 and zSeries hosts implemented the Self-Timed Interconnect bus structure and OSA-Express. Additionally, the e-business infrastructure is leaning heavily on the use of Internet technologies. These factors are the underpinnings of the IBM commitment to OSA-Express, Queued Direct I/O (QDIO), and TCP/IP as strategic technologies for IBM mainframes, with Gigabit Ethernet being today's optimal networking connection to zSeries and S/390 G5 and G6. Optimizations in z/OS, z/VM, TPF, and Linux continue to be made to enhance this TCP/IP-based mainframe networking strategy.

15.3 What you need to understand about OSA-Express

OSA-Express can be used to support both TCP/IP and SNA environments. Each are discussed below as well as the management tools that you can use to support your OSA-Express environment. For more information on OSA-Express, see the *zSeries Connectivity Handbook*, SG24-5444.

15.3.1 TCP/IP support

A networking I/O design called Queued Direct I/O (QDIO) was introduced specifically for TCP/IP communication. QDIO is supported by z/OS, z/VM, z/VSE™, TPF, and Linux on zSeries. It incorporates:

- ▶ “IP Assist” which performs the following functions:
 - Media Access Control (MAC) handling: Formatting datagrams for LAN-specific data
 - Address Resolution Protocol (ARP) function: Identifying the physical addresses
 - Packet filtering: Screening and discarding broadcast LAN packets instead of interrupting host processing
 - Building and maintaining a table of IP addresses to be used for packet routing
- ▶ IP Multicast: Sending one copy of data or message to a group, reducing network congestion. Increasing the number of recipients in the group does not require a corresponding increase in resources or bandwidth.
- ▶ LPAR-to-LPAR Communication, to share an adapter on the same server for traffic between two LPARs without sending the traffic out onto the network and back into the mainframe.
- ▶ Direct Memory Access (DMA) to allow the OSA-Express microprocessor to communicate directly with CS for z/OS to move data directly from the microprocessor to host memory.
- ▶ Ability to dynamically define the feature; receiving the definitions from the host TCP/IP stack. Dynamic updating of the OSA address table reduces configuration and setup time, eliminates duplicate data entry, and reduces the chance of data entry errors and incompatible definitions.
- ▶ Priority Queuing (unique to the Policy Server in z/OS): Data is placed in one of four queues for outbound traffic and one queue for inbound traffic. TCP/IP packet prioritization via Type of Service and Differentiated Services bits are checked by OSA-Express and queued in priority sequence for outbound transmission.

QDIO is optimized for TCP/IP traffic transport and has immediate and significant benefit potential including lower System Assist Processor utilization, fewer CPU cycles used to handle networking I/O, and improved response time potential.

15.3.2 SNA support

But how does OSA-Express provide benefit to the SNA traffic that communication controller customers need to handle? In cases where you want to transport SNA traffic over OSA-Express, in order to leverage the benefits of QDIO, IBM provides two technologies to integrate SNA-based traffic onto your TCP/IP network:

- ▶ Enterprise Extender (discussed in Chapter 19, “Enterprise Extender (EE)” on page 331) can be used to carry SNA traffic over TCP/IP directly into the mainframe using SNA Advanced Peer-to-Peer Networking (APPN) technology. Some limited Enterprise Extender performance information can be obtained from the Enterprise Extender white paper at:
<http://www.ibm.com/servers/eserver/zseries/networking/pdf/ee-snasw.PDF>
- ▶ TN3270 server (discussed in Chapter 18, “TN3270” on page 311) provided by CS for z/OS supports TN3270 client and browser access to SNA applications running on the mainframe. In most environments, this can effectively support movement of LU2 (3270) traffic from communication controllers onto the TCP/IP transport.

These SNA to IP integration technologies benefit from advancements in the TCP/IP application environment introduced as part of OSA-Express and CS for z/OS-based TCP/IP strategy. CS for z/OS TCP/IP stack capabilities based on the Virtual IP Address (VIPA), VIPA takeover, VIPA takeback, and the load-balancing Sysplex Distributor can help you to build highly available access to both SNA and TCP/IP applications on the Parallel Sysplex mainframe.

In its non-QDIO operational mode, an OSA-Express port can support native SNA, APPN, and HPR (using Link Services Architecture protocol) as well as TCP/IP (using LAN Channel Station protocol). In this mode, the OSA-Express protocol support is similar to the prior OSA-2 adapters, but at an increased speed due to the improved OSA-Express hardware and STI connection. The SNA support in non-QDIO mode uses an XCA connection just like an ESCON attached router and does *not* provide SNA boundary function services.



Important: Communication controllers save host CPU cycles by offloading certain network functions, including SNA device support (also known as the SNA boundary function). Removing communication controllers and migrating to OSA non-QDIO native SNA support will consequently result in increased host CPU cycles and storage requirements. The amount of increase will depend on:

- ▶ The number of devices in the network and their transaction rates, and
- ▶ Whether a dependent LU requester (DLUR) is implemented in the network to support SNA devices.

Alternatives for handling SNA boundary function workload are discussed in 5.3.3, “SNA subarea addressing, routing, and boundary function (BF)” on page 162.



Important: Large networks have encountered problems such as SNA session setup failures or the inability to activate additional resources as the result of running out of SNA element addresses in a VTAM subarea. Migrating devices from your communication controllers to OSA non-QDIO native SNA support can add to this problem, because the device support (boundary function) ends up moving from the 3745 subarea into the VTAM subarea.

Alternatives for handling the SNA subarea function are discussed in 5.3.3, “SNA subarea addressing, routing, and boundary function (BF)” on page 162.

While the OSA-Express Gigabit Ethernet only supports QDIO mode, each port on the OSA-Express Fast Ethernet, 1000BASE-T Ethernet, 155 ATM (Ethernet LAN Emulation), and token ring can be independently defined to use either QDIO or non-QDIO. Performance tests have shown that the OSA-Express Fast Ethernet in non-QDIO mode is capable of 11 MBps running full duplex with an SNA streams benchmark. This compares favorably in terms of performance with both ESCON-attached routers and controller, assuming you do not require communication-controller-unique functionality.

15.3.3 Layer 2 support

OSA-Express2 and OSA-Express Ethernet features on z890, z990, and System z9 can support two transport modes: Layer 2 (Link Layer) and Layer 3 (Network or IP Layer). In Layer 2 mode, the OSA can support IP (IPv4, IPv6) and non-IP (AppleTalk, DECnet, IPX™, NetBIOS, or SNA) traffic. Layer 2 support can help

facilitate server consolidation. Layer 2 support also enables the Communication Controller for Linux on System z9 and zSeries (CCL) to support:

- ▶ Fiber OSA connectivity (rather than just the copper OSA features originally supported by the CCL)
- ▶ Multiple MAC addresses per OSA port (reducing OSA port requirements when migrating traffic from multiple 3745 communication controller token-ring interfaces)

CCL Layer 2 support is discussed further in 16.3.1, “Layer 2 support” on page 290.

When using Layer 2, packet forwarding decisions are based on Media Access Control (MAC) addresses instead of Layer 3 information (such as SNA routing information or IP addresses). z/VM guests no longer have to share a single MAC address with OSA because each operating system attached to the Layer 2 interface has its own MAC address. Complexity can be reduced and LAN administrators can configure and maintain the mainframe environment the same as they do non-mainframe environments.

An OSA-Express2 or OSA-Express Ethernet feature can filter inbound datagrams by Virtual Local Area Network identification (VLAN ID, IEEE 802.1q), the Ethernet destination MAC address, or both. Filtering can reduce the amount of inbound traffic being processed by the operating system, helping to reduce CPU utilization. OSA port sharing is supported only between virtual switches that are using the same transport mode (Layer 2 with Layer 2 and Layer 3 with Layer 3). A Layer 2 guest cannot communicate directly with a Layer 3 guest.

15.3.4 Management

There are multiple ways at different levels to monitor OSA-Express:

- ▶ The OSA/Support Facility (OSA/SF) provides both GUI interface and VM REXX access to OSA-Express settings, traffic statistics, and performance information (zSeries only). This program is integrated in z/OS, z/VM, and z/VSE.

Note: OSA/Support Facility is no longer required to manage SNMP data for the OSA-Express features. An SNMP subagent exists on an OSA-Express feature, which is part of a direct path between the z/OS or Linux master agent (TCP/IP stacks) and an OSA-Express Management Information Base (MIB).

- ▶ OSA-Express can also be managed using TCP/IP SNMP protocol support and associated SNMP management products such as SNMP MIB

(Management Information Base). MIB information can be found on Resource Link™ at:

<http://www.ibm.com/servers/resourceLink>

- Tivoli® NetView provides basic SNMP MIB support and also provides a network performance monitor for IP which can assist with the collection of performance information.
- Internetwork Status Monitor, from Cisco Systems, can run as a NetView application. ISM can provide OSA-Express MIB device information.
- Other vendor SNMP management products can be used to support OSA-Express.

Note: Simple Network Management Protocol (SNMP) is supported for all of the OSA-Express features when configured in QDIO mode. SNMP support for the non-QDIO LAN Channel Station (LCS) mode became available in May 2004 and applies to all of the OSA-Express features supported on z890 and z990 in conjunction with TCP/IP applications only. It supports the same SNMP requests and alerts offered in QDIO mode (Get, GetNext, Trap, and Set) and is exclusive to z/OS V1.6 and later environments.

- ▶ S/390 Resource Measurement Facility (RMF™) can be set to issue reports about performance events as they occur. The RMF Channel Path Activity report aids in the performance analysis of the multi-path channel connection used by OSA-Express, which consists of the logical READ path for moving data from the network to the server, and a WRITE path for moving data from the server to the network. This same performance information is provided by OSA/SF and via SNMP on a zSeries system.

Note: RMF enables you to capture useful performance data for the OSA-Express features:

- ▶ Microprocessor utilization (per LPAR image if it applies)
 - ▶ Physical PCI (Peripheral Component Interconnect) bus utilization
 - ▶ Bandwidth per port (both read and write directions), per LPAR image
- ▶ Basic hardware information can be obtained using the standard hardware console as well.

15.3.5 TPF support

TPF supports the OSA-Express Gigabit features on zSeries and on S/390 G5 and G6 servers. TPF also supports OSA-Express2 Gigabit Ethernet and 10 Gigabit Ethernet LR on z990 and z890. This support includes use of QDIO, and offers the highest performance network connectivity for TCP/IP. Support includes VIPA and movable VIPA, which enables TCP/IP connections to be balanced in a loosely coupled environment. TPF is also able to route around a failed OSA adapter onto a second adapter to provide high availability. Support for OSA-Express requires TPF Version 4.1 with PUT 13 (or higher) or requires TPF APR PJ27333.

Note: The Communication Controller for Linux on System z9 and zSeries V1.2 (CCL) can only connect to TPF *internally* (within a given mainframe system between CCL and TPF LPARs or Virtual Machines) and requires the use of a new OSA for NCP (OSN) CHPID type that provides CDLC channel connectivity. OSN-based CCL connectivity is discussed in 16.3.2, “CCL V1.2 OSA for NCP (OSN) CDLC channel support” on page 291.

15.4 Strategies for exploiting OSA-Express

OSA-Express gives you a remarkably simple and scalable means of connecting your mainframe host environment to your network. Yet it is still just one piece of your overall IT infrastructure. As mentioned in Chapter 14, “Net390 architecture” on page 261, it is not the network alone, nor the mainframe host environment alone, but rather the synergy that can be achieved from looking at both together that results in an optimal IT infrastructure solution.

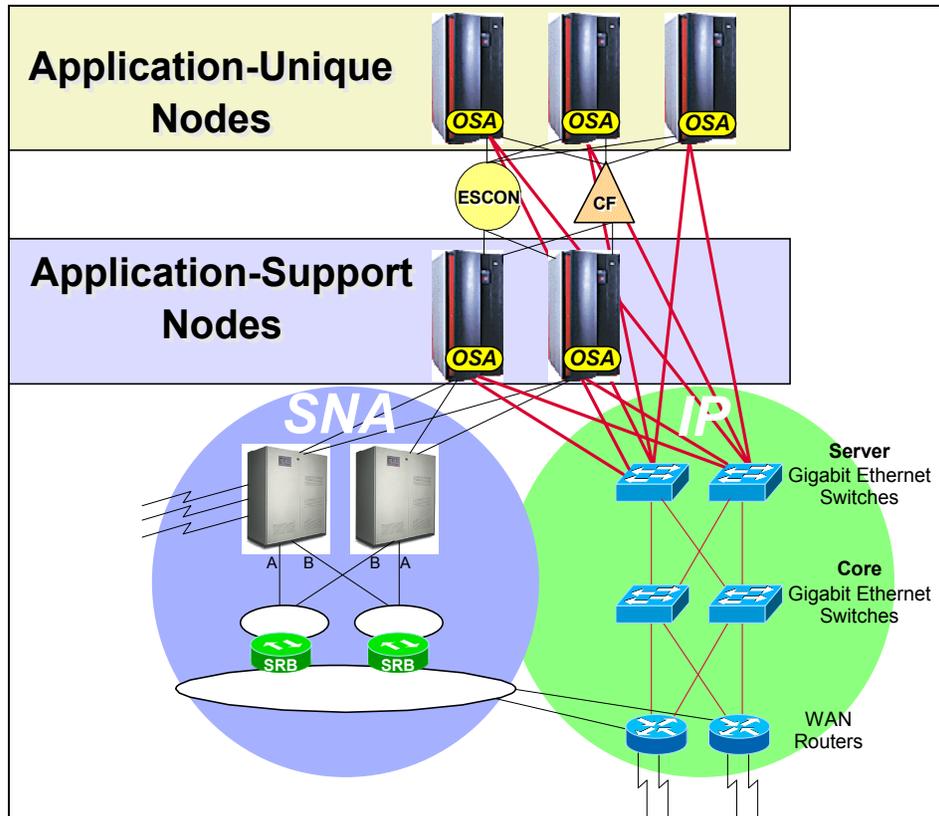


Figure 15-1 The Net390 model

As shown in Figure 15-1, OSA-Express can provide high-bandwidth connectivity to your strategic IP infrastructure. In a high-availability configuration with dual gigabit Ethernet switches, and taking advantage of IP high-availability capabilities provided by CS for z/OS such as dynamic routing and VIPA, you can build availability for your IP environment comparable to the availability that you have come to expect from your SNA environment. The key difference between the two environments will be that your OSA-Express and gigabit Ethernet-based IP infrastructure will provide orders-of-magnitude greater bandwidth than your SNA environment and at a much lower cost of ownership. Therefore, it will be desirable to migrate traffic from your SNA environment onto your IP infrastructure over time. As is described in 15.3.2, “SNA support” on page 275, TN3270 and EE technologies can provide powerful tools to help you in that migration.

Note: Even though OSA-Express adapters have two ports, prudent planning for redundancy indicates the use of redundant adapters and dual paths to redundant switches for high-availability configurations.



Communication Controller for Linux on System z9 and zSeries (CCL)



IBM Product Number: 5724-J38

IBM announced the IBM Communication Controller for Linux on zSeries V1.1 on February 15, 2005, and the Communication Controller for Linux on System z9 and zSeries V1.2 on October 25, 2005. Put simply, these products enable you to run the NCP, NPSI, and NRF program products (originally designed to be run on IBM communication controller hardware) on Linux for zSeries. Consequently, in many cases, CCL will provide the most transparent and functionally equivalent option for migrating from your current IBM 3745 Communication Controllers. The CCL option is discussed in greater detail, where applicable, in Part 2, “Functional alternatives reference” on page 93. This chapter discusses the CCL product and provides some strategies for taking advantage of it.

16.1 What is CCL?

With the CCL product, other program products that were originally designed to run on IBM 3745 Communication Controller hardware can be run on Linux for zSeries instead. Currently supported products include:

- ▶ Network Control Program (NCP), 5648-063
- ▶ X.25 NCP Packet Switching Interface (NPSI), 5688-035
- ▶ Network Routing Facility (NRF), 5668-963

CCL supports these products by emulating an IBM 3745 model 31A communication controller (with 16 MB of memory). In most cases, moving NCP, NPSI, and NRF from the 3745 hardware into Linux on zSeries requires only minimal definition changes.

Figure 16-1 shows a high-level diagram of the major CCL components.

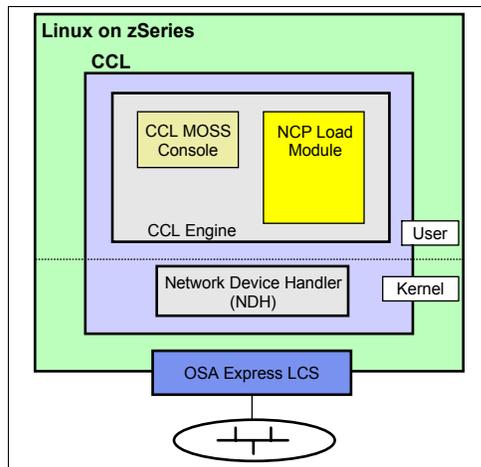


Figure 16-1 CCL components

The most fundamental of these components are:

- ▶ The Communication Controller for Linux on System z9 and zSeries Engine (CCL Engine) is a program that runs in the user space on the Linux operating system. The CCL Engine provides an environment that emulates many of the functions of the 3745 Communication Controller hardware, enabling your NCP load module to run on a zSeries machine without changes.
- ▶ The Communication Controller for Linux on System z9 and zSeries Network Device Handler (NDH) is software that runs in the Linux kernel. The NDH is logically positioned between the network interface card (NIC) device driver and one or more CCL Engines, and provides the communications path

between the CCL Engines and other devices in the network. Only a single instance of the NDH is loaded per Linux image, regardless of how many CCL Engines (NCPs) are running in that Linux image.

- ▶ The CCL MOSS console, accessible via a Web browser, is used to perform many of the functions typically performed using the MOSS console of the 3745 hardware.

For more information about CCL, see *Communication Controller for Linux on System z9 and zSeries Implementation and User's Guide*, SC31-6872. You can also find CCL Technotes, white papers, and other useful information on the CCL support Web pages at:

<http://www.ibm.com/software/network/ccl/support/>

16.2 Why is CCL important?

CCL is the first product other than IBM communication controller hardware to support 3745-based program products and their broad set of functional capabilities. In many cases, CCL provides the most transparent and functionally equivalent option for migration from your current IBM 3745 Communication Controllers. Of particular importance are:

- ▶ SNA Network Interconnection (SNI): CCL provides an attractive option because it eliminates the need for coordinated interorganizational changes that are required by every other alternative. (Discussed in 5.3.4, “SNA Network Interconnection (SNI)” on page 167.)
- ▶ Token-ring to Ethernet migration: The CCL can handle SNA traffic on Ethernet LAN interfaces, whereas the 3745 only supports SNA on token-ring LANs. (Discussed in “Migrating from token ring to Ethernet” on page 117.)
- ▶ Extended Recovery Facility (XRF): Prior to the introduction of CCL, no functional alternative existed other than to make significant application changes. (Discussed in 5.3.7, “Extended Recovery Facility (XRF)” on page 180.)
- ▶ Network Routing Facility (NRF): Prior to the introduction of CCL, no functional alternative existed other than to make significant application changes. (Discussed in 11.2.1, “Peer-to-peer connection of SNA devices (before PU type 2.1)” on page 239.)
- ▶ X.25 NCP Packet Switching Interface (NPSI): Using X.25 over TCP/IP (XOT) technology and CCL V1.2 NPSI program product support, your mainframe environment can support *all of the X.25 capabilities* currently supported by NPSI in your IBM 3745 communication controllers. Routers are used to support the actual X.25 physical interfaces, and XOT technology transports the X.25 traffic over your TCP/IP network into your mainframes. Advantages

of the CCL NPSI approach over the previously available host XOT software alternatives include comprehensive support of NPSI functional capabilities and easier migration. (Discussed in “Supporting NPSI in your mainframe using CCL” on page 200.)

16.3 What you need to understand about CCL

The IBM 3745 family of communication controllers runs a number of software program products and supports a variety of hardware communication interfaces. In contrast, the CCL product runs only the NCP, NPSI, and NRF software program products and only supports:

- ▶ OSA Ethernet or token-ring LAN interfaces
- ▶ OSA-based CDLC channel connectivity (using the new OSA for NCP CHPID type, discussed in 16.3.2, “CCL V1.2 OSA for NCP (OSN) CDLC channel support” on page 291)

In spite of these CCL software and physical interface limitations, CCL can provide significant value in many situations (as discussed in 16.2, “Why is CCL important?” on page 285); however, you must understand and accommodate those limitations in order implement a solution that meets your needs.

Part 2, “Functional alternatives reference” on page 93 provides a comprehensive listing of IBM 3745 Communication Controller functional capabilities and can serve as a guide to help you to identify the best alternatives (including CCL) for the specific controller functions that you use. Many of the most important controller functions are provided by the NCP software that is supported by CCL.

Table 16-1 CCL support summary

CCL supports	CCL does <i>not</i> support
<p>Software:</p> <ul style="list-style-type: none"> ▶ NCP (V7R5 and above) and compatible levels of NPSI, and NRF ▶ SSP, NTuneMON, NetView, and NPM continue to work as they have 	<ul style="list-style-type: none"> ▶ Other 3745 software: XI/NSF, EP, NTO, NSI, MERVA, or TPNS ▶ Functions provided by the 3746 NNP or MAE ▶ NCP-based IP routing

CCL supports	CCL does <i>not</i> support
<p>Physical interfaces:</p> <ul style="list-style-type: none"> ▶ OSA token-ring LAN ▶ OSA Ethernet LAN (supported on <i>all</i> Ethernet features through the use of OSA Layer 2 support, discussed further in 16.3.1, “Layer 2 support” on page 290) ▶ OSA-based CDLC channel connectivity (using the new OSA for NCP CHPID type, discussed in 16.3.2, “CCL V1.2 OSA for NCP (OSN) CDLC channel support” on page 291) <p>Note: Though <i>NCP software</i> only supports SNA over token ring, the CCL Network Device Handler transparently converts Ethernet traffic to token-ring format for the NCP.</p>	<ul style="list-style-type: none"> ▶ SDLC, Frame Relay, ISDN, Parallel Channel, BSC, ALC, Start-Stop, or X.25 <p>Note: Routers can be used to handle traffic from SDLC, Frame Relay, X.25, and ISDN WAN links and forward the traffic across a network to the CCL (NCP, NPSI, and NRF).</p>

The limited physical interface connectivity supported by the CCL can be complemented with the use of routers for serial link aggregation (as illustrated in Figure 16-2 on page 288). In fact, many organizations have historically used routers for such link aggregation because the physical interfaces on routers are usually less expensive than comparable interfaces on IBM 3745 Communication Controllers.



Important: Multi-link transmission groups (MLTGs) will not work across Data Link Switching (DLSw) networks because the local acknowledgement of logical link control, type 2 (LLC2) protocols interferes with MLTG sequence numbering. Therefore, if you are currently using multi-SDLC-link transmission groups between 3745s, then you cannot simply duplicate it using CCLs and routers. For additional information and alternatives, see 5.3.2, “Multi-link transmission group (MLTG)” on page 161.

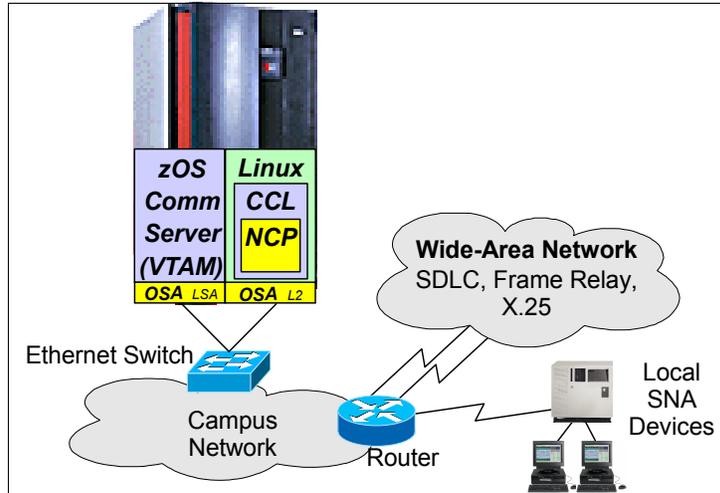


Figure 16-2 CCL network connectivity

CCL connects to the network through OSA ports. CCL Ethernet connectivity can exploit OSA Layer 2 support (discussed in 16.3.1, “Layer 2 support” on page 290) to enable:

- ▶ Efficient, QDIO-based, data transfer
- ▶ Support of *all* Ethernet OSA features (both copper and fiber)
- ▶ Assignment of multiple CCL MAC addresses to a single OSA port (which can be very useful when consolidating traffic from multiple 3745 TICs)

Although *NCP software* only supports SNA over token ring, the CCL Network Device Handler (NDH) transparently converts Ethernet traffic to token-ring format for the NCP.

Important: Certain OSA features support token ring; however, you should use Ethernet whenever possible because token-ring technology is nearly obsolete. *Token ring is not supported on IBM System z9 servers.*

The OSA adapter connects the CCL to a LAN and, through the LAN, to routers. The routers can be configured to take traffic from a variety of link types, convert it to Ethernet (or token-ring) frames, and pass it to the CCL. In essence, the CCL, the LAN, and the routers together form a *virtualized* communication controller.

Note: CCL can only connect to TPF *internally* (within a given mainframe system between CCL and TPF LPARs or Virtual Machines) and requires the use of a new OSA for NCP (OSN) CHPID type. OSN-based CCL connectivity is discussed in 16.3.2, “CCL V1.2 OSA for NCP (OSN) CDLC channel support” on page 291.

As with CCL, z/OS Communications Server (in particular, its VTAM subcomponent) connect to the network through an OSA interface. VTAM, however, uses an OSA adapter configured to provide a Link Service Architecture (LSA) interface. LSA provides SNA device attachment to VTAM hosts using External Communications Adapter (XCA) communications (the same interface as was used by the IBM 3172 Interconnect Controller product). LSA is supported only on the following (copper) Ethernet OSA features:

- ▶ G5/G6 Processors: OSA Express FC 2340 Fast Ethernet (10/100 Mb)
- ▶ z800 or z900 Processors: OSA Express FC 2366 Fast Ethernet (10/100 Mb)
- ▶ IBM System z9: OSA Express2 FC 3366 1000BaseT Ethernet



Note: VTAM and CCL can share an OSA port; however, the utility of that approach is very limited because they cannot use it for communication between each other.

A VTAM can connect to a CCL either within a given mainframe system (using the OSN as described in 16.3.2, “CCL V1.2 OSA for NCP (OSN) CDLC channel support” on page 291) or externally through a LAN. When connected internally, the CCL NCP is configured and managed as a local, channel-attached NCP. When LAN-connected, CCL NCPs must be defined to VTAM as XCA PUs, and such CCL NCPs must be operated as remote NCPs.

VTAM may be either the owning host or a data host to LAN-connected CCL NCPs. To support VTAM ownership of a “remote” (LAN-connected) CCL NCP, a new keyword has been added to the XCA PU statement (via VTAM PTF) to enable VTAM to activate and own CCL NCP resources over an XCA interface: ALLOWACT=NO/YES. The currently available VTAM APARs for ALLOWACT support are:

- ▶ OS/390 and z/OS VTAM: APAR OA10425
- ▶ VSE/ESA VTAM: APAR DY46311
- ▶ z/VM VTAM: APAR VM63677

The following sections discuss additional connectivity options that are available for the CCL.

16.3.1 Layer 2 support

OSA Layer 2 support is the ability for Ethernet OSA adapters, in QDIO mode, to pass data frames based on their Ethernet MAC addresses rather than Layer 3 information (such as a TCP/IP address). The operating system also must be able to support Ethernet-format frames (rather than just Layer 3 frames), and two do:

- ▶ z/VM
- ▶ Linux on System z9 and zSeries

(Notably absent at the time of the writing of this book is z/OS.)

Finally, programs running in those operating system environments (such as CCL) must also be able to handle Ethernet-format frames.

CCL V1.1 initially shipped without Layer 2 support. Consequently, CCL V1.1 could only use OSA features that supported native SNA, LAN Channel Station, (OSE CHPID type) connectivity, which limited CCL to only using copper Ethernet OSAs. However, Layer 2 support has subsequently been added to the CCL product (and ported back into the V1.1 code). So now, with Layer 2 support, CCL can use any of the OSA Ethernet features (including fiber connectivity) and does so in the more efficient QDIO mode.

An additional advantage of Layer 2 support is the ability to define multiple Ethernet MAC addresses for a single OSA port. A very useful technique for migrating from 3745 communication controller token ring interfaces is to:

1. Define a MAC address on the OSA port for a CCL NCP that is identical to a MAC address that is being used on a 3745 token-ring port.
2. Use router load balancing technology that, recognizing the existence of the duplicate addresses, is able to distribute sessions between them.
3. When you are comfortable with the operation of the new CCL infrastructure, you can disable the communication controller token-ring interfaces and let all connections go through the CCLs.

Prior to CCL Layer 2 support, if you needed to migrate several different 3745 token-ring MAC addresses, you had to have an equal number of OSA ports (because OSA LCS support allows only one MAC address per port). Using Layer 2 to support multiple MAC addresses on a single OSA port enables you to reduce the number of required OSA ports to just enough to meet the availability and scalability that you require.

Note: A CCL NCP can support up to eight MAC addresses (corresponding to the NCP support limitation of up to eight token-ring interfaces). Those eight MAC addresses could all be assigned to a single OSA port, distributed across eight different OSA ports, or any combination in between. For example, an optimal availability configuration for a single CCL NCP with eight MAC addresses might be to have four MAC addresses, on each of two OSA ports, on separate OSA adapters.

If, for migration reasons, you need to support even more than eight MAC addresses, you can run additional CCL NCPs (either in the same Linux image or in a separate one depending on your performance and scalability requirements), distributing those MAC addresses across however many OSA ports you require for availability (probably just a couple).

Note: Duplicate MAC addresses must be distributed across different Linux images. (If you set up duplicate MAC addresses within one Linux image, all traffic will be routed to the same CCL NCP TIC, defeating the point of duplicate MAC addressing.)

16.3.2 CCL V1.2 OSA for NCP (OSN) CDLC channel support

The OSA for NCP CHPID type (OSN) essentially provides OSA-based CDLC channel connectivity between LPARs or Virtual Machines in the same mainframe system. Support for OSN was added to CCL V1.2.

CCL V1.2 OSN support is particularly important to TPF users. TPF only supports channel connectivity to an NCP and can, therefore, only connect to a CCL NCP through OSN.

Communications Server (VTAM) can also leverage the OSN CHPID type (CDLC channel) for *internal* connectivity to CCLs. Where CCLs are running in the same mainframe system as a VTAM, OSN connectivity between that VTAM and the CCLs will be considerably more efficient than using VTAM XCA communications through OSA adapters and external LAN switches.

One additional advantage of using OSN for VTAM connectivity to CCLs is that the CCL NCP is configured and managed as a local, channel-attached NCP rather than a remote NCP. As a local NCP, the CCL NCP can be loaded and managed in the same way as local 3745 communication controllers, further reducing the changes required to migrate from 3745s to CCLs.

OSN CHPID type support requires *all* of the following:

- ▶ CCL V1.2
- ▶ An OSA-Express2 Gigabit Ethernet SX, Gigabit Ethernet LX, or 1000BASE-T Ethernet feature (#3364, #3365, and #3366 respectively)
- ▶ A z9-109 mainframe system

16.3.3 CCL DLSw support

IBM developed DLSw technology in the early 1990s to enable routers to reliably transport LAN-based SNA and NetBIOS traffic over TCP/IP networks. Because IBM shared the technology through working with open standards (RFC 1434, March 1993), DLSw was quickly adopted and implemented by all major router vendors.

One thing that can be confusing when discussing DLSw technology is that router vendors often refer to a number of different technologies, packaged together, as DLSw (or, in the Cisco Systems case, DLSw+), including:

- ▶ SDLC-to-LLC2 conversion
- ▶ X.25 QLLC-to-LLC2 conversion
- ▶ MAC-addressed-based load balancing

DLSw technology takes advantage of the fact that SNA and NetBIOS traffic use Logical Link Control Type 2 (LLC2) connection-oriented transport protocols on LANs. Consequently, much of the DLSw functionality is concerned with transparently supporting LAN-oriented LLC2 protocols over a TCP/IP-based wide-area network rather than the more obvious encapsulation and transport of the actual SNA (or NetBIOS) messages. Indeed, DLSw can require considerable router processor resources. Figure 16-3 on page 293 illustrates DLSw between routers.

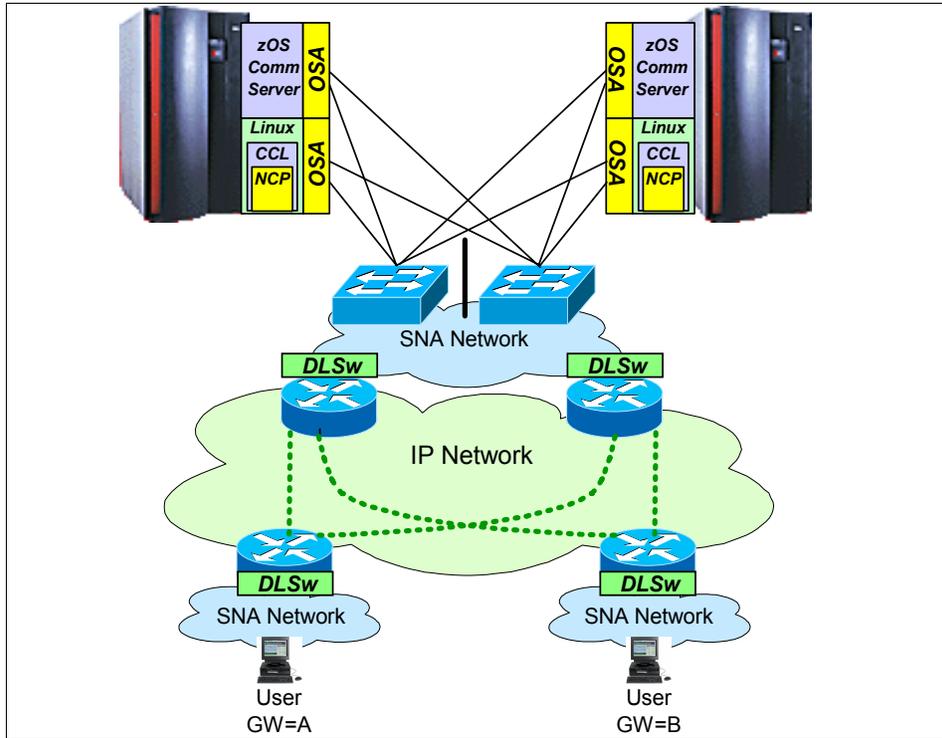


Figure 16-3 Data Link Switching (DLSw) between routers

In Figure 16-3, routers serving remote SNA users encapsulate SNA messages for transport over the IP network to the data center DLSw “peer” routers. Due to the processor workload required by DLSw support, large networks can often require several routers in the data center. In terms of SNA session availability, a failure of any DLSw router will result in the failure of all SNA sessions that go through it. The failure of a data center router will almost always have the biggest impact because data center routers handle the SNA traffic from many remote DLSw routers. Figure 16-4 on page 294 illustrates DLSw to CCLs.

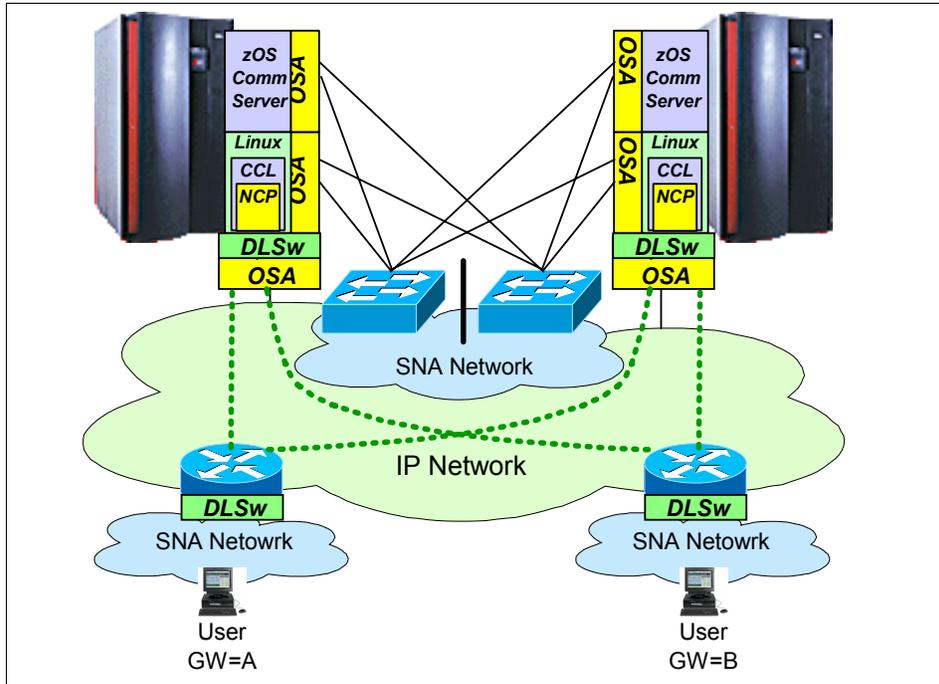


Figure 16-4 Data Link Switching (DLSw) to CCLs

Using CCL DLSw support, you may be able to eliminate the need for data center DLSw routers. If you are planning a new deployment of DLSw technology, you may find the CCL DLSw support to be more scalable and, potentially, less expensive relative to supporting the same amount of DLSw workload in routers. However, the most-compelling advantage of CCL DLSw support may come from improving availability through eliminating a single point of failure on SNA session paths (the data center DLSw routers).

16.3.4 CCL V1.2 IP TG support

SNA transmission groups (TGs) are logical connections between adjacent SNA nodes used to pass SNA session traffic. CCL V1.2 IP TG support provides the ability to encapsulate and transport SNA traffic over TCP/IP connections. Of course, all TCP/IP encapsulation technologies require compatible support in each of the two adjacent nodes. CCL V1.2 is the only product that has implemented IP TG support; therefore (at the time of the writing of this book) *IP TGs can only be implemented between CCLs.*

As with DLSw, IP TG support is a TCP/IP encapsulation technology; however, DLSw support is focused on extending *LAN-based SNA* (LLC2) traffic over IP

networks, which is much more complex (because of the LLC2 LAN protocols and timers) than just transporting traffic from an SNA TG. Consequently, IP TG support is much simpler and more efficient for the purposes of connecting CCLs together across an IP infrastructure. In addition, where DLSw cannot support SNA Class of Service (COS) because DLSw cannot differentiate between different types of SNA traffic, CCL IP TGs can be configured to support SNA COS-based prioritization through the use of configurable TCP ports or TOS specification.

So when do you need to connect CCLs together across an IP infrastructure?

IP TGs may not be very useful for intra-data-center connectivity. Within a data center (or, more specifically, within the range of local area networking), it generally makes sense to connect each VTAM to every CCL NCP. Between VTAMs and CCLs within a single mainframe system, the OSN (discussed in 16.3.2, “CCL V1.2 OSA for NCP (OSN) CDLC channel support” on page 291) will provide the most efficient connectivity. Between VTAMs and CCLs running on different mainframe systems, Ethernet and Layer 2 support (discussed in 16.3.1, “Layer 2 support” on page 290) provides the optimal high-performance and efficient connectivity.

Extending LAN protocols between data centers (across a wide-area network, using technologies such as DLSw) can be inefficient and costly. Consequently, CCL IP TGs could be quite useful for inter-data-center connectivity between CCLs, especially in situations where the CCL NCP is required for connectivity, such as:

- ▶ For SNI connectivity between trading partners when both partners have CCL NCPs (see 5.3.4, “SNA Network Interconnection (SNI)” on page 167)
- ▶ For SNA connectivity for environments where Enterprise Extender (EE) technology (discussed in Chapter 19, “Enterprise Extender (EE)” on page 331) cannot be used (such as for connectivity to VM or VSE environments).

16.4 Strategies for exploiting CCL

The CCL product fits very nicely into the Net390 architectural framework (described in Chapter 14, “Net390 architecture” on page 261) by providing an effective means of consolidating yet more support functions into network-support (Net390) nodes. CCL can help to simplify your data center network environment by giving you a solid means of eliminating your aging token-ring infrastructure and reducing dependence on ESCON channels. While it is clear from scanning through the pages of this book that there are often many alternative technology solutions that must be considered, few alternatives can be implemented as

nearly transparently as CCL. For that reason alone, the CCL may well play an important role in your future SNA networking environment.

For example, consider the options for migrating third-party business partner SNI links. (This is discussed in 5.3.4, “SNA Network Interconnection (SNI)” on page 167.) Migrating directly to the use of TCP/IP communications is almost certainly the most strategically attractive option. Short of a TCP/IP migration, implementing EE and EBN technology likely offers the best efficiency and performance. However, the sheer burden of trying to coordinate such changes with each of your partners will make the nearly transparent migration to CCL very attractive (providing the flexibility to make a strategic migration at a future time that works well for both you and your partners).

One unusual characteristic of a CCL solution is that, although certain aspects of it may seem new (for example, the use of Linux), most operational aspects of it will be *very similar* to your current operating environment. CCL NCP operations are supported through VTAM exactly as if the NCP were running in a 3745. Special 3745 operations that are performed today using the 3745 MOSS console can be performed by the CCL MOSS console through a Web browser interface. Similarly, the NetView, NPM, and NTuneMON interfaces to a CCL NCP will remain unchanged; however, certain parts of the presented information will be different, for example:

- ▶ Under the NTuneMON “3745 HARDWARE INFO,” the MICROCODE EC field will show the CCL version and release, such as CCLV1R1.
- ▶ With NPM, the CCL will not show CCU or TIC utilization.

Consider making your CCL implementation an early project in your IBM 3745 Communication Controller migration efforts. Perhaps start with a very simple configuration (as illustrated in Figure 16-2 on page 288) and use it in low-risk environments until you have gained sufficient operational experience. From there, if needed, you can build out a fully redundant high availability solution for mission-critical application environments (as illustrated in Figure 16-5).

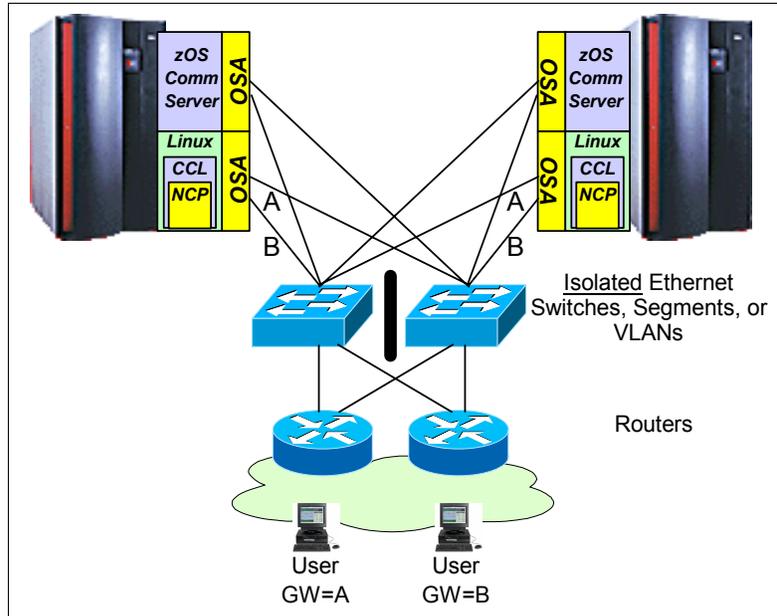


Figure 16-5 CCL Ethernet high availability

The Ethernet high availability configuration is discussed in greater detail in “Implementing CCL Ethernet high availability” on page 118.



TCP/IP

TCP/IP has become the de facto standard for networking in today's world of universal communication requirements. As a result, the trend of convergence toward a single ubiquitous network infrastructure leads to one based on the IP protocol. This chapter discusses the TCP/IP networking protocol suite and the increased importance of this architecture in today's networking infrastructures.

17.1 What is TCP/IP?

Although Transmission Control Protocol (TCP) and Internet Protocol (IP) are two building blocks of a complete protocol suite, their relative importance gives rise to the common name for this family of protocols: TCP/IP. TCP/IP follows the layering approach, whereby protocol layers provide functionality to layers above and make use of the functionality provided by layers below as illustrated in Figure 17-1.

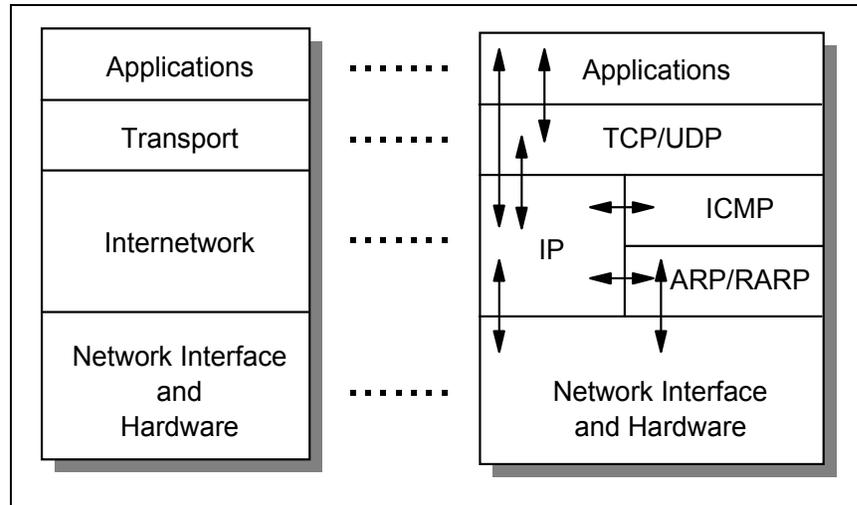


Figure 17-1 TCP/IP layering

In essence, the Internet Protocol allows for communication across multiple heterogeneous networks. For example, it may be desirable to connect an Ethernet network with an ATM network. Networks based on the TCP/IP suite are called internetworks (or internets), the largest and most famous of which is the Internet.

IP provides a best-effort service of delivering data from one machine in an internetwork to another. Transport functionality such as reliable data streaming, congestion control, and flow control must be provided at a higher level, such as TCP. Some applications leverage the reliable stream-oriented functionality provided by TCP. Others, however, make use of the User Datagram Protocol (UDP) which adds little functionality to basic IP besides the ability to demultiplex data to multiple applications. Using UDP, the application is responsible for ensuring reliability by perhaps retransmitting sent data or tolerating loss of data.

17.2 Why is TCP/IP strategic?

Because of the simplicity of TCP/IP, largely due to its simple layered architecture, it has become the single most important and widely used networking protocol suite. As a result, organizations have invested heavily in their IP infrastructures. The net effect of such investment has been increased development of advanced functionality with the TCP/IP protocol suite to accommodate the increased needs of e-business networks. Functions such as Quality of Service (QoS) and Voice over IP (VoIP) have emerged.

In addition, available applications on the TCP/IP platform have increased its importance and penetration into the marketplace. Of course, the single application that has led most to TCP/IP's success has been the World Wide Web. It has led to increased use by the general public and has been the impetus for the creation of numerous e-business applications.

17.3 What you need to understand about TCP/IP

The TCP/IP protocol suite is rich with applications and functionality. This section highlights some of these.

17.3.1 Routing infrastructure

The network infrastructure of a TCP/IP network is the collection of routers that are implemented. Routers are highly specialized machines that reside on a plurality of networks and are capable of forwarding (actually routing) data packets from one physical network to other (potentially heterogeneous) networks. This operation is illustrated in Figure 17-2 on page 302.

In this example, host C is resident on both network X and network Y. It provides the necessary functionality to allow hosts A and B to communicate. This involves receiving data packets on one network interface and transmitting them over the other.

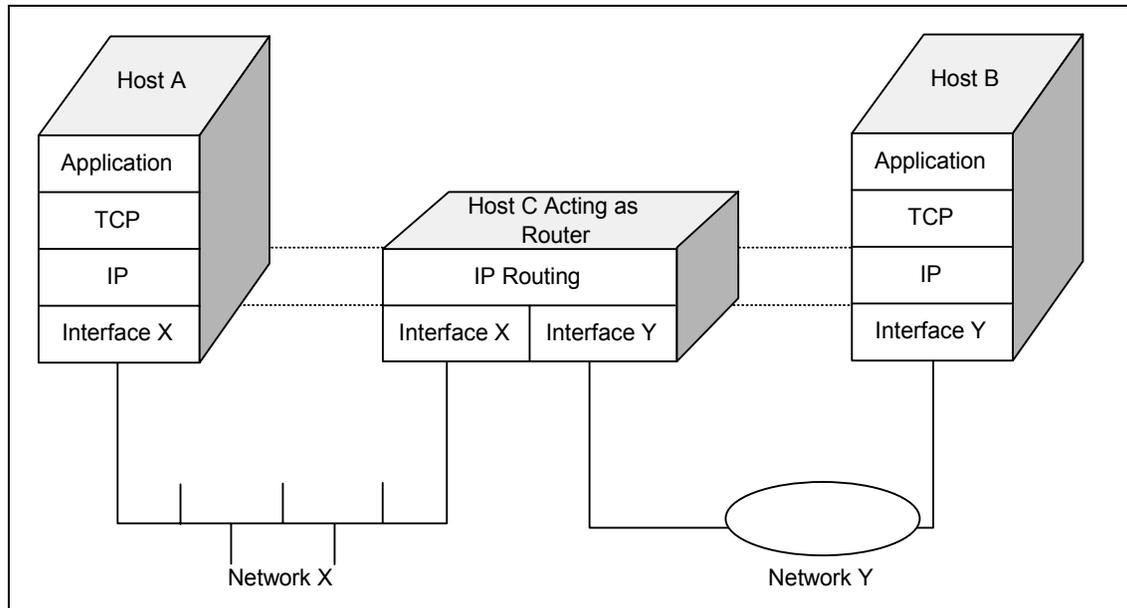


Figure 17-2 IP routing operation

Careful planning in the capabilities of these routers, their placement, and their configuration yields a highly available, high-performance network. It is critical to end-system perceived performance to appropriately design a routing blueprint for the network that leverages redundancy, routing hierarchy, and high-performance networks.

17.3.2 TCP/IP security

Originally, TCP/IP was based on a system of trust. That is, because of its roots in academia, the TCP/IP architecture made no attempt to prevent or circumvent malicious activity on the network. Today, it seems that nearly the entire civilized world has access to the global Internet. Malicious attacks are daily, sometimes hourly, events to many organizations on the Internet. This includes the use of computer viruses and extends into powerful Denial of Service attacks that are difficult to defend against.

In essence, protocols have been developed to aid in the continued usability and performance of networks. They allow for private communications ensuring data integrity and with mechanisms for authentication. The tools in this battle make use of encryption, authentication, network policy, and admission control.

The result is a number of protocols and functions that aid in securing the network. A Virtual Private Network (VPN), for example, provides secure tunnels

by encrypting data, authenticating communicating partners, and ensuring integrity of transmitted data. The leading VPN protocol is IPSec, although other tunneling protocols are also in use today. Secure Sockets Layer (SSL) provides much the same functionality, but at an application layer.

Firewalls allow for the specification of a network policy that limits the entrance of certain data into a protected network. For example, a service such as FTP (which uses specific TCP port numbers) that your organization does not allow between your network and the Internet can be blocked by the firewall to eliminate potential security implications of that service.

17.3.3 Quality of Service

Initially, IP was defined to provide only a best-effort delivery service. A network built using no additional QoS mechanisms is still very robust and services many disparate applications, as proven by the popularity and scale of the global Internet. However, as enterprises make heavier use of the Internet and build their own intranets, bandwidth will inevitably become constrained. No longer will the network be able to offer the desired services to each and every application.

Convergence beyond data networking adds new requirements for traffic delivery. Voice, video, and other digitally encoded streams have real-time transport service requirements that cannot be accommodated without supplemental protocols and policies. QoS protocols are necessary to allow intelligent packet forwarding decisions to be made based on the end-to-end service goals of the application.

There are three commonly referred to models of end-to-end QoS:

- ▶ Best-effort service

Best-effort service is the type of service provided by all general IP-based networks. The network will deliver data on a first-in, first-out basis as long as resources are available to do so. No guarantees or assurances are made with respect to delay, packet loss, or throughput. You could say a best-effort service lacks any QoS mechanisms.

- ▶ Differentiated Services

Differentiated Services involves the handling of individual packets or flows within a network node. Each packet is associated with a particular class of service. Each node along the network path handles packets in a cooperative manner according to a common set of rules resulting in end-to-end service classes.

- ▶ Integrated or Reserved Services

Integrated Services, also known as “*Reserved or Guaranteed Services*,” provides the bandwidth and delay characteristics as requested by the application or configured for specific types of traffic.

17.3.4 High availability and load distribution

The traditional view of a single server has been primarily that of a single machine with perhaps a few network interfaces (IP addresses). This tends to lead to many potential points of failure within the server: the machine itself (hardware), the operating system (including TCP/IP stack) kernel executing on the machine, or a network interface (and the IP address associated with it).

Clustering techniques that address the load balancing of connections requests also typically provide for some high availability. That is, these techniques dispatch connections to target servers and can exclude failed servers from the list of target servers that can receive connections. In this way, the dispatching function avoids routing connections and requests to a server incapable of satisfying such requests.

Load balancing is the ability for a cluster to spread workload evenly (or based on some policy) to target servers comprising the cluster. Usually, this load balancing is measured by some notion of perceived load on each of the target servers. By providing load balancing, clustering techniques also provide for other system requirements in addition to the dispatching of connections. These include the ability to advertise some single system-wide image or identity so that clients can uniquely and easily identify the service. Additionally, clustering techniques also provide for horizontal growth of the system and ease of management.

17.3.5 Web-based applications

In today’s Internet, World Wide Web traffic, which mostly uses the Hypertext Transfer Protocol (HTTP), greatly surpasses any other application protocol in using the most bandwidth. Modern computer operating systems provide Web browser applications by default, some even provide Web servers, thus making it ever easier for end users and businesses to explore and exploit the vast capabilities of worldwide networked computing.

The World Wide Web, now commonly referred to as simply the Web, has become a de facto standard for the development and deployment of business-critical applications. There has been a clear shift toward developing Web-friendly applications that allow for a consistent look and feel for users.

17.3.6 Print solutions

The line printer requester (LPR) allows access to printers on other computers running the line printer daemon (LPD) as though they were on your computer. The clients provided (LPR, LPQ, LPRM or LPRMON or LPPORTD) allow the user to send files or redirect printer output to a remote host running a remote print server (LPD). Some of these clients can also be used to query the status of a job as well as to delete a job.

17.4 Strategies for exploiting TCP/IP

Most organizations have TCP/IP-based networks currently implemented. We summarize some of the important aspects of TCP/IP networks and strategies that should be considered when building and maintaining these types of networks.

17.4.1 TCP/IP convergence

TCP/IP provides a single convergence point for all communications in your organization. By consolidating your network into a single networking infrastructure, you can avoid duplication of communication function and minimize communication costs. This includes not only the acquisition and maintenance of machinery that comprises the network fabric, but also the personnel required to maintain these.

Important: In Part 2, “Functional alternatives reference” on page 93, we identify *functional alternatives* to key IBM 3745 Communication Controller capabilities. For example, APPN Border Node (BN) provides an alternative to SNA Network Interconnection (SNI). However, wherever practical, you should try to migrate from the use of SNA altogether. In certain cases, such as with many currently available commercial file transfer software packages, the change from SNA to TCP/IP can be accomplished just by reconfiguring the software at each end of the communication.

When considering migrating applications to the use of TCP/IP, it is useful to consider three basic categories of SNA communications: file transfer, interactive, and application to application. The migration to the use of TCP/IP will vary depending upon the category of application as well as the specific application products in use.

File transfer

Today, most commercially available file transfer applications can be configured to support file transfer over either SNA or TCP/IP. Typically, the current use of SNA is a vestige of the fact that, for many organizations, SNA was used for communications long before TCP/IP came along. Such environments should now be migrated to the use of TCP/IP.

For older file transfer environments for which the software cannot be reconfigured to support TCP/IP, you must either continue to support them over SNA or migrate to more modern file transfer software. Where practical, upgrade to more modern file transfer software because:

- ▶ Native use of TCP/IP will be more efficient than accommodating SNA in the network, and
- ▶ Modern file transfer software packages include advanced capabilities which can make the communication more efficient, manageable, and secure.

Interactive

Interactive environments are more complex to migrate to the use of TCP/IP than file transfer environments because they involve many endpoints rather than just the two ends of a file transfer. Options for migrating interactive users include:

- ▶ Web-enabling the application using a technology such as the IBM Host Access Transformation Server (HATS), or
- ▶ Using TN3270 terminal emulation

Figure 17-3 illustrates the use of Web-enabling or TN3270 for SNA applications.

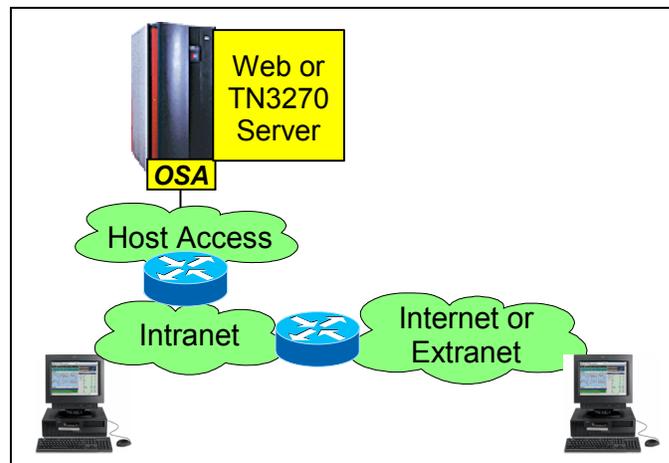


Figure 17-3 IP enabling with Web or TN3270

Given the desire of many organizations to migrate current applications to their intranet or even the World Wide Web, Web-enabling an application can be very strategically attractive; however, for any specific desktop to be migrated from SNA, all of the SNA applications accessed by that desktop have to be Web enabled. Using TN3270 is often more straightforward because it gives access from any TN3270 client to *all* mainframe SNA applications while allowing the migration of the desktop and intervening network to TCP/IP. For more information on TN3270, refer to Chapter 18, “TN3270” on page 311.

Application to application

Common SNA application to application environments involve communication between applications developed for IMS, CICS, or Message Queuing (MQ) subsystems. Like file transfer software, MQ systems can be configured to communicate equally well over SNA or TCP/IP “channels.” Migrating IMS and CICS subsystems to TCP/IP communication is more complex. The WebSphere MQ product (formerly known as MQSeries) comes with “bridge” software which enables it to integrate with IMS or CICS. Consequently, it may be possible in your environment to provide a high-performance, IP-capable connectivity for existing IMS or CICS applications using WebSphere MQ. For additional information about:

- ▶ The MQSeries - IMS Bridge
<http://www.ibm.com/software/ts/mqseries/platforms/imsbridge.html>
- ▶ MQSeries - CICS/ESA® Bridge
<http://www.ibm.com/software/ts/mqseries/txppacs/male.html>

Figure 17-4 shows how such a solution would work.

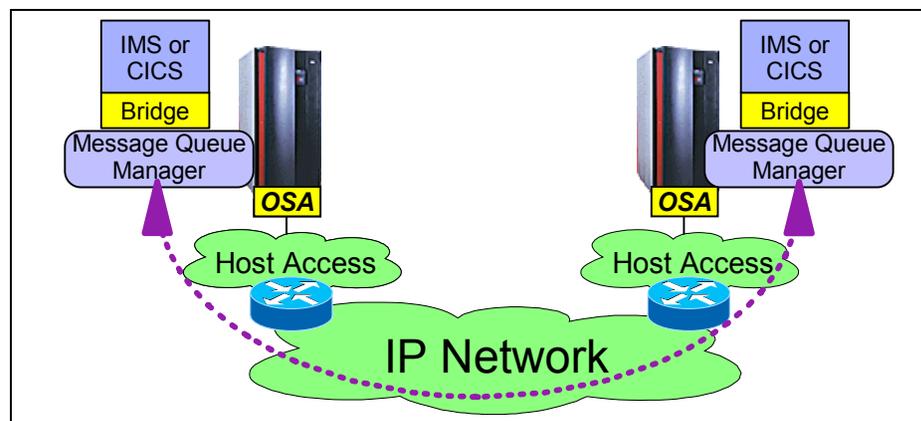


Figure 17-4 Using MQ for IMS or CICS

Advantages of migrating application-to-application message exchanges to MQ include:

- ▶ MQ is optimized for application-to-application message exchanges and provides many functions (including message-level reliability and error handling) that can enhance your application environment.
- ▶ Such a migration provides a first step in a potentially strategic migration of application communications toward a message queuing structure.
- ▶ Using MQ may improve performance by enabling migration of your application-to-application message exchanges onto your typically higher-bandwidth IP infrastructure.
- ▶ MQ flexibility eliminates your SNA network dependencies and enables incremental network migrations.

Considerations of migrating application-to-application message exchanges to MQ include:

- ▶ If your organization does not currently use MQ, you would need to implement and support the new MQ (and bridge) software environment.
- ▶ The IMS or CICS application traffic may not be well-suited to message queuing due to timing sensitivities.

To learn about services available from IBM that could help you with a migration to MQ, visit:

<http://www.ibm.com/software/integration/websphere/services/>

Alternatively, for SNA application to application environments, Enterprise Extender technology can be used to allow SNA data to be transported over an IP infrastructure. Quality of Service policies can be established to propagate the classes of service associated with SNA. For more on Enterprise Extender, see Chapter 19, “Enterprise Extender (EE)” on page 331.

17.4.2 Secure your network

Critical to today’s TCP/IP networks is ensuring the integrity, authenticity, and privacy of data. In addition, network services must be readily available to meet users’ demands as envisioned when deploying the application. All of these objectives are points of attack for potential malicious hackers in the Internet and in corporate intranets. This group of people includes everything from teenage kids playing with their first computer to high-tech corporate spies with black ties and pen pocket protectors. Therefore, it is imperative that your TCP/IP network be armed with proper security mechanisms.

A good strategy to employ when securing an IP-based network is the DTA motto: “Don’t Trust Anybody.” This mentality dictates security mechanisms that will make it as difficult as possible to gain unauthorized access to sensitive information and resources. It involves the quick revocation of departed employees’ user IDs as well as the physical security of corporate networks.

From a network perspective, the first line of defense against Internet-born attacks is a firewall that filters out unwanted or unexpected data directed toward your network. The use of SSL and VPNs protects the integrity, authenticity, and privacy of data exchanged with trusted peers. The use of Intrusion Detection Systems helps to detect Denial of Service attacks and circumvent them before they can cause severe damage. When combined, all of these tools can be used as part of a global security strategy that can avoid unwanted distractions or even worse, lost revenue due to an IP-based attack.

17.4.3 Quality of Service

Another interesting way of viewing QoS is illustrated in Figure 17-5. In essence, the most complete QoS is Integrated Services, which allows for the explicit reservation of network resources and ensures that these resources will be available to the application when needed. Unfortunately, Integrated Services is not without its cost, primarily in the form of the resources required to maintain per-flow state information throughout the network. That is, each application instance’s existence must be visible to all routers in a network. Such requirements can lead to serious scalability issues.

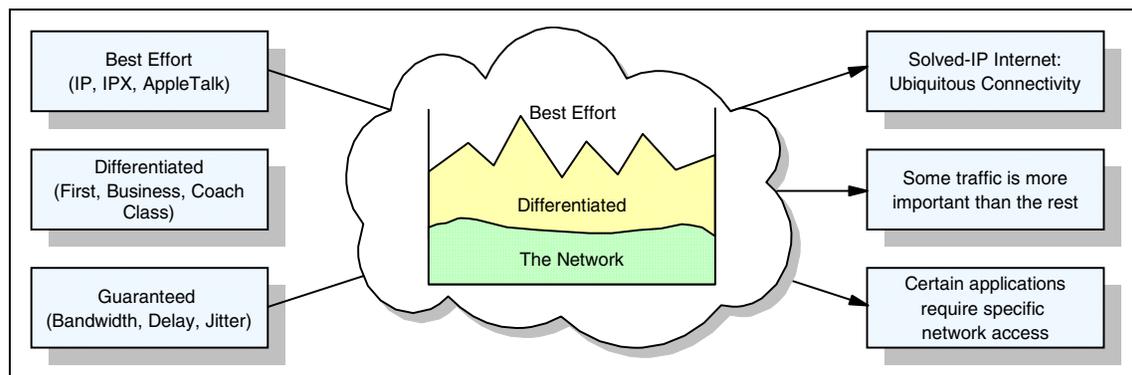


Figure 17-5 Quality of Service models

As a result, Differentiated Services is currently the most widely accepted Quality of Service mechanism used today.

17.4.4 High availability and load distribution

The key to high availability in a network is redundancy. Redundant network components and paths will ensure that access to critical network resources remain available even during individual network component outages. Likewise, redundant network interfaces aid in providing highly available access to servers. Finally, redundant servers can ensure that no single server outage will impact the overall availability of the service.

Network components such as routers and network paths should be redundant so that if a particular path in a network fails, another one may be used to ensure connectivity. For example, component redundancy can be implemented with the use of the Virtual Router Redundancy Protocol (VRRP) and by establishing multiple paths between pairs of routers. Dynamic routing protocols such as OSPF allow for the network to dynamically recover from failed network routes.

Additionally, servers may have multiple network interfaces. The use of Virtual IP Addresses (VIPAs) and dynamic routing allow for the continuous reachability to a server even if part of the network to which it attaches or some of its network interfaces fail. With VIPA, an IP address is associated with a specific service in addition to the network interfaces and dynamic routing protocols are used to identify all possible routes (network interfaces) to that service.

Finally, redundant servers can be used to provide the highest possible availability for a network service (such as Web service). Servers can be clustered together to represent one single server image to connecting clients. Such load-balancing clustering technologies include WebSphere Edge Server, Sysplex Distributor, and Cisco's MultiNode Load Balancing (MNLB). All of these have in common the use of a cluster address to identify the service. They distribute connections to servers based on current load and can avoid failed servers, thereby further contributing to the high availability goals of the service. Alternatively, dynamic VIPA can also be used to maintain a single image service while providing for a backup. With dynamic VIPA, should the primary server fail, the backup can dynamically assume the VIPA address of the primary to maintain service availability.

One important type of server in the z/OS environment that requires high levels of availability and load distribution is the TN3270 Server. This server usually serves a large number of clients and the expected availability on behalf of clients is quite high. Some clustering technology such as Sysplex Distributor coupled with the use of dynamic VIPA can provide for the high demands placed on this service.



TN3270

This chapter discusses Telnet 3270 (TN3270). TN3270 provides you with a fast, straightforward means of migrating interactive SNA, 3270-based application users to your strategic IP-based environment. A migration from SNA 3270 to TN3270 can go a long way toward achieving your goal of migrating users from your communication controller environment.

18.1 What is TN3270?

TN3270 is a terminal emulation protocol that allows you to log on to a remote SNA mainframe host via TCP/IP as though you were directly connected to it. Essentially, the TN3270 server provides a logical bridge between an IP network and a direct SNA connection as illustrated in Figure 18-1. The server receives data on the IP connection and transmits it over the SNA session, and vice versa.

A TN3270 server running on a mainframe host provides the most efficient means for communications with applications running on that host. Regardless of the location of the TN3270 server, it can establish sessions with multiple hosts.

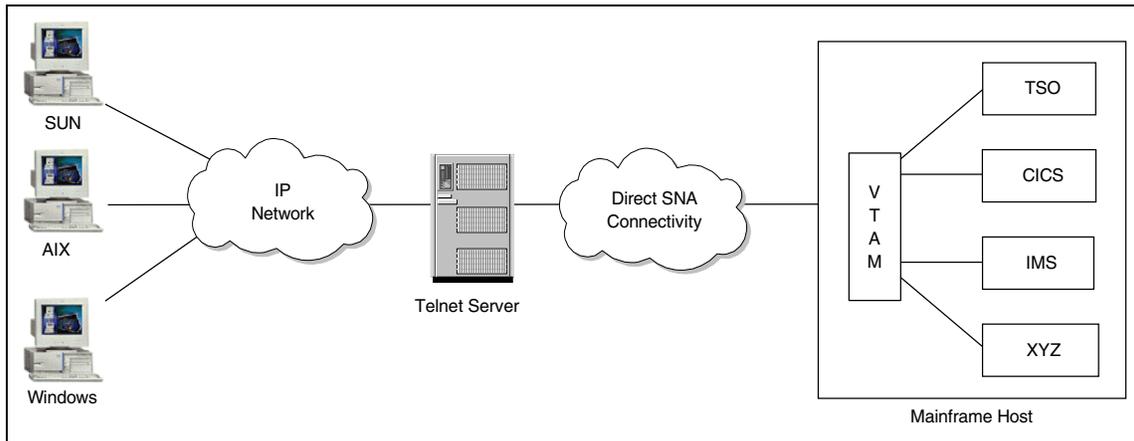


Figure 18-1 TN3270 operation

18.2 Why is TN3270 strategic?

A migration from SNA 3270 to TN3270 can go a long way toward achieving your goal of migrating users from your communication controller environment into your strategic IP-based environment. The TN3270 protocol provides an efficient means for transporting 3270 traffic over an IP network while enabling the migration of user access to TCP/IP-only connectivity. Because the TN3270 server uses specific knowledge about the data it is transporting, it can optimize its functionality to get the most out of the TCP/IP connection and the SNA session. For example, it exploits and optimizes for the low-bandwidth, delay-sensitive nature of the data stream.

18.3 What you need to understand about TN3270

In this section, we describe some of the interesting characteristics of TN3270. We make special reference to the z/OS TN3270 Server because of its strategic importance. This is further discussed in “TN3270 server placement” on page 323.

Telnet operation

The TN3270 protocols originally grew out of the basic Telnet protocol defined in RFC 854. This base Telnet protocol was first extended through RFC 1041 which introduced the concept of IBM 3270 terminal access using the basic Telnet protocol. Since then, TN3270 has been precisely defined and extended through the following RFC standards documents:

- ▶ RFC 1576 - TN3270 Current Practices
- ▶ RFC 1646 - TN3270 Extensions for LU Name and Printer Selection
- ▶ RFC 1647 - TN3270 Enhancements (TN3270E, now obsolete and replaced by RFC 2355)
- ▶ RFC 2355 - TN3270 Enhancements

RFC 2355 is the current definition of the enhanced TN3270 protocol that is generally referred to as TN3270E where implementations based on an RFC prior to RFC 1647 generally are referred to as TN3270. Most current literature uses the term TN3270 to refer to Telnet 3270 in general, including both the base TN3270 and the TN3270E protocol levels, unless there is a specific need to distinguish between the two protocol levels.

The original Telnet protocol defined in RFC 854 introduced the concept of a Network Virtual Terminal (NVT). An NVT is a virtual device that has basic limited characteristics common to a wide range of real terminals. An NVT must be supported by all the Telnet servers and clients, both basic Telnet and Telnet 3270 clients and servers.

The Telnet protocol allows servers and clients to negotiate their characteristics, because many hosts will wish to provide additional services to the NVT. Once a TCP connection has been established, both sides of the connection are capable of working on the minimum level that is implemented by the NVT. After this minimum understanding is achieved, they can negotiate additional options to extend the capabilities of the real hardware in use. Because of the symmetrical model used by Telnet, both the server and the client may propose additional options to be used.

The basics of TN3270

The TN3270 server is simply a hybrid TCP/IP and SNA application to which TN3270 clients connect using a TCP connection. When a TN3270 client connects to the TN3270 server, the server selects an SNA logical unit (LU) to represent the connected TN3270 client in the SNA network, and establishes an SNA session to an SNA 3270 application, such as CICS. The TN3270 server from then on simply relays the unmodified 3270 data stream between the TCP connection and the SNA session.

If the TN3270 server runs on an operating system platform that does not implement an SNA PU Type 5 (a VTAM) but instead implements a PU Type 2 or PU Type 2.1, the LU that is used by the TN3270 server is a traditional dependent secondary LU, and some form of an SNA network infrastructure must exist between the TN3270 server and the VTAM that acts as the SSCP for that secondary LU. If the TN3270 server runs on an operating system platform that does implement an SNA PU Type 5, such as z/OS, the LU that is used by the TN3270 server is not a traditional dependent LU, but an application minor node LU. SNA communication between the TN3270 server and VTAM is in that case optimized and does not involve any SNA network flows, but is carried out using VTAM programming interface interactions only. These flows are illustrated in Figure 18-2.

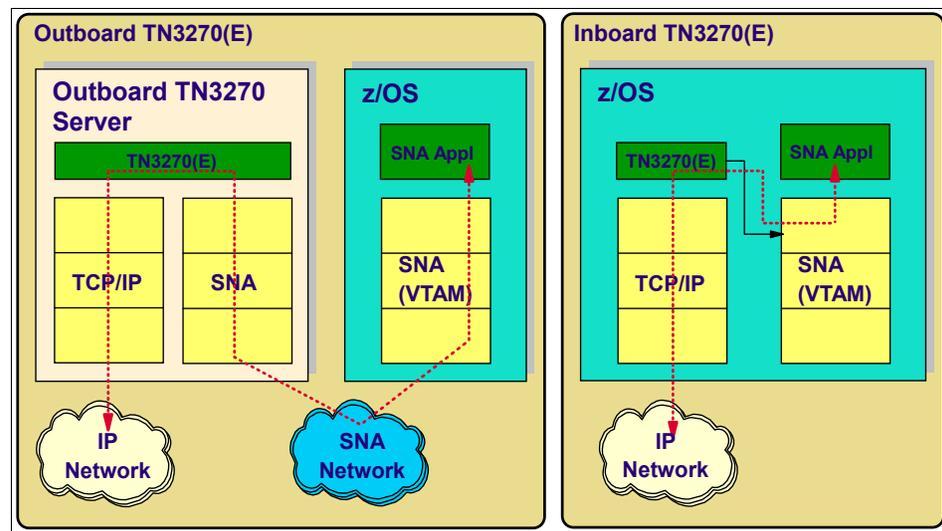


Figure 18-2 Outboard and inboard TN3270 server flows

You should note that on z/OS the TN3270 server is a separate application, not part of VTAM, although it uses VTAM.

As shown in Figure 18-3, a Telnet client that connects to a TN3270 server will always end up being associated with an SNA LU and that SNA LU will be used to establish an SNA session to SNA applications. A traditional Telnet client (non-TN3270 client) that connects to the traditional Telnet server on z/OS, the OTelnetD server, will not be involved in SNA communication at all, but will be associated with a UNIX® process using UNIX System Services interprocess communications functions. In that UNIX process, a UNIX shell program will be started and the Telnet client user will engage in a traditional UNIX shell dialog. A traditional Telnet client may be used to connect to a TN3270 server, but since a traditional Telnet client does not understand the 3270 data stream, the TN3270 server will in such a case establish a line-mode SNA session with an SNA application and only allow line-mode communication between the Telnet client and the SNA application.

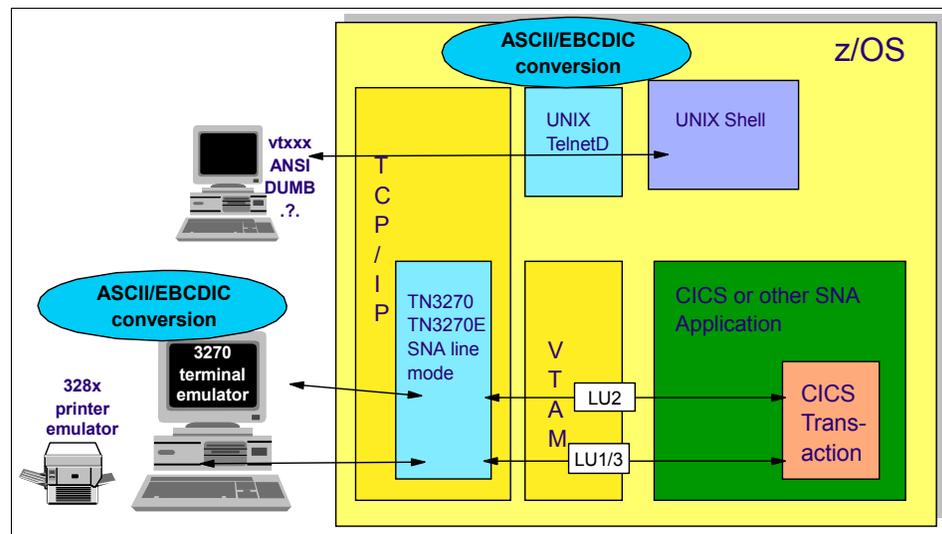


Figure 18-3 Traditional Telnet and TN3270

The main differences between traditional Telnet and TN3270 terminal emulation are:

- ▶ 3270 terminal emulation uses block mode where traditional Telnet uses line-mode or raw mode.
- ▶ The 3270 data stream between the TN3270 client and TN3270 server is a mix of binary 3270 control codes and EBCDIC character data where a traditional Telnet data stream between the Telnet client and Telnet server is based on ASCII character data. For TN3270 emulation, the TN3270 client performs conversion between ASCII and EBCDIC, whereas for traditional Telnet, the Telnet server performs conversion between ASCII and EBCDIC.

The TN3270 connection is accomplished by the negotiation of the following Telnet options:

- ▶ Terminal Type: Specifies the client “terminal type” (such as IBM-3278-2-E) to the host
- ▶ Binary Transmission: Specifies that the receiver should interpret characters received from the sender as 8-bit binary data (other than the special Interpret As Command (IAC) character and the Telnet command that follows it)
- ▶ End of Record: Because 3270 terminal emulation uses block mode, the length of the data may vary. Also, CR LF does not delineate a “new line” in binary transmission mode. Therefore, every command and its related data must be separated with the IAC End of Record (EOR) sequence.

TN3270 enhancements (TN3270E)

The TN3270 function was enhanced by RFC 1647 and RFC 2355.

The z/OS Telnet server implements RFC 2355 TN3270 enhancements plus some selected TN3270 draft RFC recommendations, in particular, the contention resolution enhancements (Send Data Indicator, Keyboard Restore Indicator, and the BID flows). TN3270E enhances the traditional TN3270 protocol as follows:

- ▶ Provides capability to emulate the 328x printers
- ▶ Enables the Telnet client to request a specific 3270 device name (an LU name)
- ▶ Supports ATTN and SYSREQ keys
- ▶ Adds support for SNA positive/negative responses

Telnet clients and servers negotiate the support of TN3270E. If either side does not support TN3270E, traditional TN3270 can be used. Once both sides have agreed on using TN3270E, they begin to negotiate the subset of TN3270E options. These options are device-type and a set of supported 3270 functions:

- ▶ Printer data stream type
- ▶ Device status information
- ▶ The passing of BIND information from server to client
- ▶ Positive/negative response exchange

TN3270E now supports:

- ▶ Client requests for a specific device name
- ▶ 3287 LU1 and LU3 printer emulation
- ▶ SYSREQ
- ▶ SNA definite, exception, and no-response requests
- ▶ The forwarding of BIND images to the client.

Support of SNA responses to the client helps synchronize data flow and provides a more accurate measurement of user response. For example, definite response can be used by various 3270 response time monitoring components, such as the built-in response timer monitor in the z/OS TN3270 server or NetView Performance Monitor to enable them to estimate user response times to ensure that application performance is meeting agreed-to service levels.

Based on the time stamps (A, B, and C), shown in Figure 18-4, and the assumption that the time it takes the initial request to get from the client to point A is roughly equivalent to the time it takes for the response to get from the client to point C, the following response times can be calculated:

Round-trip response time = Time C - Time A
IP response time = Time C - Time B
SNA response time = Round trip response time - IP response time

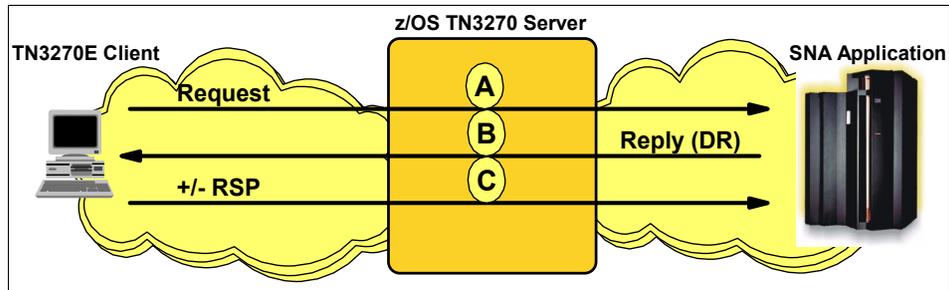


Figure 18-4 Estimating user response times

18.3.1 Secure TN3270 sessions

Many of the currently available TN3270 clients, such as Personal Communication (PCOM) or Host On Demand (HOD) from IBM, support secure sockets for the TCP connection to an SSL/TLS-enabled TN3270 server, such as the z/OS TN3270 server that supports both implicit SSL/TLS and TN3270-protocol negotiated SSL/TLS connections. Secure sockets provide secure data transmission between a secure sockets TN3270 server and a Secure Sockets TN3270 client. On an SSL/TLS-protected TN3270 connection, all data that is exchanged over the TCP connection between the TN3270 client and the TN3270 server is encrypted. Data received over the secure TCP connection from the TN3270 client is decrypted before the data is sent to other processes, such as VTAM and data coming from VTAM is encrypted before being sent over the secure TCP connection to the TN3270 client.

SSL/TLS is based on a public key infrastructure where, at a minimum, the TN3270 server has a key ring with a private and a public key, and a digital certificate. (Optionally, in addition, the client may have a key ring with a private

and a public key, and a digital certificate.) TN3270 clients can use the server's digital certificate to authenticate the server to which they are connecting. The SSL/TLS protocols can then, based on the server's key ring and the content of the digital certificate, generate the necessary session keys to encrypt the traffic between the TN3270 client and the TN3270 server. To conduct commercial business on the Internet, you might use a widely known Certificate Authority (CA), such as VeriSign, to get a high assurance server certificate. For a relatively small private network within your own enterprise or group, you can issue your own server certificates, called self-signed certificates, for your own use by using the GSKKYMAN utility shipped as part of z/OS System SSL or the RACF® RACDCERT command.

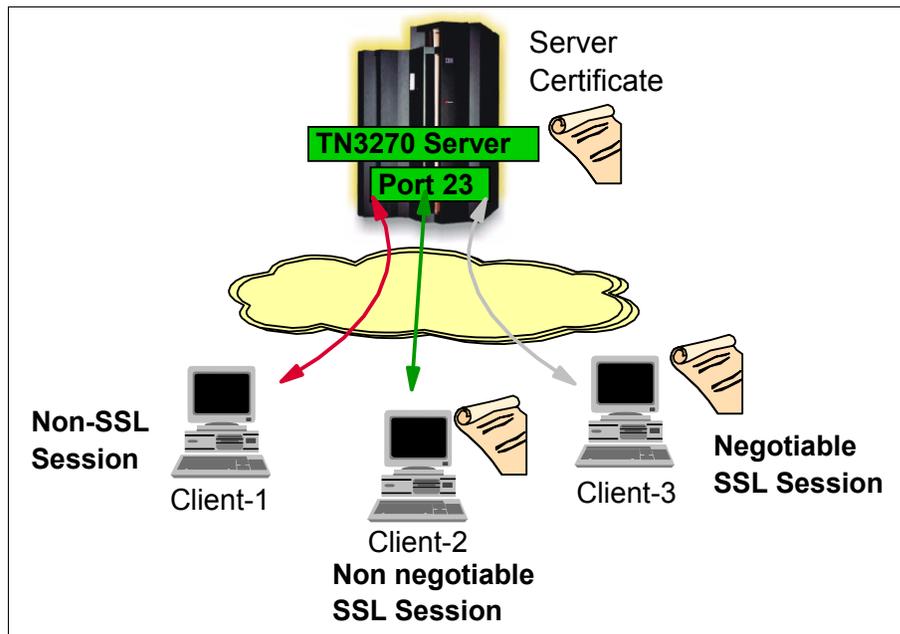


Figure 18-5 TN3270 SSL security

SSL/TLS can optionally use client key rings and digital certificates in addition to the server key ring and certificate. If client certificates are used, then the server during TN3270 server connection setup can authenticate who the client is based on the client's digital certificate. The TN3270 server on z/OS extends this capability to determine what the RACF user ID is that is associated with the client certificate that was presented during connection setup. This enables the TN3270 server to learn the RACF user ID of the user before any 3270 data has been exchanged over the TN3270 connection, which allows for TN3270 server port protection via RACF resource definitions so that unauthorized users can be rejected even before they see the traditional USS message 10 screen on their

terminal window. It also allows the z/OS TN3270 server to assign a TN3270 server resource, such as selecting an LU name, based on the user ID of the user who established the connection.

Client authentication is required if the Express Logon Feature of the z/OS TN3270 server is to be used. Express logon allows an user who has already been authenticated based on a client digital certificate to log on to various z/OS subsystems, such as TSO or CICS, without being prompted for a RACF password.

On z/OS, the key rings and the digital certificates can be managed by RACF and stored in the RACF database, or they can be managed by a utility known as GSKKYMAN that stores the key rings and certificates in HFS files.

If a TN3270 server outside z/OS is used, such as the TN3270 server in the IBM Communications Server for Linux on zSeries, SSL connections between TN3270 clients and such a TN3270 server are supported. However, remember that the SNA session between the TN3270 server and the SNA application is not protected by SSL. If the SNA session traverses insecure network segments, remember the SNA data stream may pass through such network segments in the clear.

Some TN3270 servers outside z/OS also support the Express Logon Feature if client certificates are used for the SSL connection from the TN3270 client. However, since the TN3270 in such a case does not run on z/OS and does not have a direct interface available to query RACF about userIDs and passtickets, a separate subsystem needs to be set up to provide such an interface. This subsystem is known as the Digital Certificate Authentication Server (DCAS), which is started on z/OS. The non-z/OS resident TN3270 server then establishes a separate protected SSL connection to the DCAS server for client authentication and to obtain the RACF userID and passticket that are needed for the Express Logon Feature.

18.3.2 Telnet printer support

As is illustrated in Figure 18-6 on page 320, most TN3270E client emulators can emulate both a 3270 display workstation, such as an IBM 3279 Model 3, and a 3270 printer workstation, such as an IBM 3287 Model 1 printer. Printer data streams may either be SNA Character Stream (SCS - LU Type 1) or the 3270 printer data stream (DCS - LU type 3).

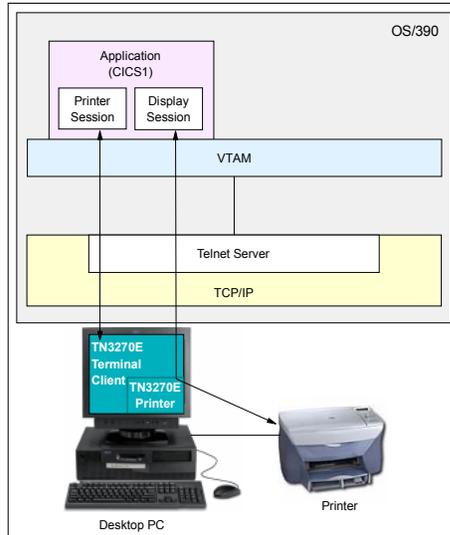


Figure 18-6 z/OS Telnet with TN3270E printer emulation

TN3270 printer emulation allows the Telnet administrator to use a single technology (that is, TN3270) to provide both 3270 display terminal and 3270 printer support.

A TN3270 server cannot “acquire” a TN3270 client emulator for a printer session. That is because a TN3270 server only accepts incoming TCP connections from TN3270 emulators. It does not establish outbound connections to the TN3270 clients. Therefore, a printer connection to the TN3270 server always must be established from the TN3270 emulator workstation to the TN3270 server. The TN3270 server may then do one of two things: It may open a secondary LU and wait for the SNA application subsystem to acquire the printer LU from the SNA side, or it may log the secondary LU on directly to the application subsystem. The last method is especially useful when the application subsystem supports auto-install of printer definitions, which is the case with both CICS and IMS.

Using TN3270 for 3270 print is very similar to using a traditional real 3270 printer device attached through a terminal controller. When the application subsystem, such as CICS, sends print to the TN3270 emulator, it is a synchronous process. If the paper in the printer jams, you have to be able to go back to the CICS application that generated the print and request the application to recreate the print. With TN3270 printing, there is, in general, no spool system in-between the SNA application that generates the print and the actual printer. For occasional print, this behavior is most likely acceptable. For printing large amounts of business documents, such as invoices or shipping information, however, this

behavior is likely to be unacceptable (just as when a real 3270 printer was used in the past).

For more-than-occasional print requirements, therefore, some form of spooling system is probably required to enable better management of print data, in particular, to handle situations where printers jam or malfunction. The IBM InfoPrint Server product provides a more robust IP-based print solution by collecting the print in a spool data set and providing an operator interface for management of the print data including the ability to reroute print to alternate physical printers and restart failed print operations.

Printer association

Once printer emulation is available, it is sometimes necessary to establish control over how display LU names and printer LU names for TN3270 connections from a single workstation are coordinated. A typical example is an installation where a CICS application assumes that if a display LU name is of the form xnnT then the printer LU name that is located right next to that display terminal is named xnnP. That was a safe assumption back when the display and printer device indeed were real 3270 devices that were connected via coax cables to specific ports in an IBM terminal control unit that was connected via a serial line into a specific interface on an IBM Communication Controller where the NCP definitions clearly defined the relationship between these two LU names. With TN3270, however, we cannot assume such a relationship unless we plan for it in our TN3270 infrastructure.

There are different ways a predetermined relationship between a display LU and a printer LU can be established. The most flexible method is to use a function within the TN3270E protocol that is known as printer association. Briefly, printer association works as follows:

1. The workstation connects to the TN3270 server and requests a display session. The TN3270 server determines which LU name to use for that display session based on whatever LU name assignment configuration has been defined on the TN3270 server.
2. The chosen LU name is returned to the workstation in a TN3270 protocol flow.
3. The workstation connects again to the TN3270 server and now requests a printer session. Along with this request, it sends the LU name it was given for the display session and asks the TN3270 server to assign a printer LU name that is associated with this display LU name. Configuration definitions on the TN3270 server will then specify which printer LU name to use with the associated display LU name and assign that LU name to the printer session.

Note: If you use printer association, you should in general also use a TN3270 server option to ensure that if one of the two sessions is terminated, then the associated session will be terminated too; otherwise, print might end up in the wrong location. The z/OS TN3270 server supports a configuration option called “DropAssocPrinter” to enforce such coordinated session termination.

Printer association is a TN3270E protocol function. Even if a TN3270 client supports TN3270E, it is not always the case that it also supports printer association. Verify with your TN3270 emulator product documentation whether it supports printer association or not before you base your implementation on that capability.

How to assign LU names to TN3270 clients

Assigning LU names to your TN3270 sessions can be very easy or it can be extremely complex. If your SNA applications do not have any dependency on LU names, you can simply define a pool of LU names in your TN3270 server that is large enough to support the highest number of concurrent sessions you expect, and let the server pick the next available LU name in that pool when a TN3270 client connects. In many cases, however, some SNA applications depend on a certain naming structure to determine some amount of user authorization based on LU name prefixes or whole LU names. A simple example is where all LU names in the accounting department start with ACCTnnnn and the SNA applications in CICS or IMS authorize use of accounting applications based on the terminal's LU name. In this case, we need to be able to assign an LU name that starts with ACCTnnnn to all TN3270 sessions that are established from TN3270 clients in the accounting department.

Such processing is generally referred to as “*LU nailing*” or “*LU mapping*.” Each TN3270 server supports one or more methods to perform such an LU name assignment. All TN3270 servers support assigning LU names based on client IP address or parts of the client IP address (a subnet). Other TN3270 servers, including the z/OS TN3270 server, have much more elaborate methods to determine which LU name to assign and can base the decision on the client host name, which local IP address the client connected to on the server, which port number it connected to on the server, which network interface the connection arrived over, and even the RACF user ID of the TN3270 user if secure connections are used in combination with client certificates. This last option is especially useful for traveling users, who may connect in from all over the world using completely unpredictable local IP addresses.

The z/OS TN3270 server has an extensive set of definitions to control assignment of LU names to TN3270 clients. It also, as a last resort, supports an installation exit routine that can be used to implement local custom methods for

assignment of LU names. In general, the TN3270 server uses its configuration definitions to pick and assign an LU name to a TN3270 client. This is referred to as “*generic LU name assignment*.”

The TN3270E protocol allows a TN3270E client during connection setup to request a specific LU name or request an LU name in a specific LU group identified by an 8-character name. This is referred to as “*specific LU name assignment*.” In such a case, the TN3270 server uses its LU name assignment configuration definitions to decide if it will allow or deny that specific client to be assigned the LU name it requests.

18.4 Strategies for exploiting TN3270

There are a number of issues that must be considered when planning for the use of a TN3270 server. Most notably, the placement of the server will influence performance and scalability. Additionally, high availability and load balancing of both the TN3270 server and the ultimate SNA applications should be considered.

18.4.1 TN3270 server placement

We have seen that a TN3270 server can give workstations access to SNA applications without the need to have an SNA protocol stack implemented on that workstation. The benefits of using TN3270 are clear, but perhaps not so clear is the decision of where to place the TN3270 server. Essentially, there are two choices:

- ▶ On a server in the network
- ▶ On the mainframe host

As TCP/IP networking emerged, it was adopted first by small and midrange computer systems and only later by mainframes. Early implementations of TCP/IP for mainframe hosts also had significant performance issues. Finally, mainframe host CPU cycles have historically been more expensive than small and midrange computer systems. Consequently, many organizations implemented TN3270 servers outside of the mainframe host (such as on Cisco routers or on Microsoft SNA Servers). Such outboard TN3270 server implementations became very popular because they enabled organizations to migrate users onto their strategic IP infrastructures. Unfortunately, in many instances, they fell prey to their own popularity encountering the following problems:

- ▶ Scalability: TN3270 to SNA 3270 conversion requires a considerable amount of resources on the server in terms of CPU capacity and memory for message buffers. As the number of users of PC-based servers grows, additional machines have to be added. Today, some organizations have many

PC-based servers with large staffs trying to manage and administer them, which inevitably leads to the second issue.

- ▶ **Availability:** Designing high-availability configurations for huge banks of PC-based servers is very complex. Often, organizations using outboard TN3270 server implementations encounter unacceptable availability for their SNA host application service level requirements.
- ▶ **Cost:** Finally, considering the floor space, management, administration, and high-availability configurations, the lower hardware costs of outboard TN3270 server implementations were not quite the bargain that they appeared to be on the surface.

As the TN3270 service has evolved from a workgroup service to an enterprise-wide service, an enterprise-class solution is now required. Over the past decade, the cost of host CPU cycles has declined, and IBM mainframe TN3270 server efficiency has improved, to the point where it is less expensive to run your TN3270 server in the mainframe host. When you consider the outstanding scalability and availability characteristics of mainframe hosts, we think that your TN3270 server function should be consolidated into your mainframe host environment.

Another aspect to consider when determining the optimal location of a TN3270 server is the general objective to reduce the complexity of the physical SNA network. The network protocols between a TN3270 client and the TN3270 server are TCP/IP-based and traverse an IP network infrastructure. The network protocols between a TN3270 server and the mainframe SNA application hosts are SNA-based and traverse an SNA network. If the TN3270 server is physically located far away from the mainframe data center, there either needs to be an SNA wide area network infrastructure maintained between that remote location and the data center, or a technology such as Data Link Switching (DLSw) or Enterprise Extender (EE) needs to be deployed to convert the SNA traffic so it flows over an IP network infrastructure between the remote location and the data center. To avoid this additional complexity of the overall enterprise network infrastructure, it is desirable to extend the IP-based communication between the TN3270 client and the TN3270 server all the way into the data center in order to minimize or completely remove the SNA network between the TN3270 server and the mainframe SNA applications.

A TN3270 server can be located on the mainframe in either the traditional mainframe operating system, such as z/OS, or in Linux on zSeries running the Communications Server for Linux on zSeries product (CS Linux). When the TN3270 server runs in Linux, the SNA network between Linux and the other mainframe operating systems may be implemented as EE over any available IP connectivity (including HiperSockets™), a channel-to-channel connection, or a shared local area network to which both Linux and the traditional mainframe

operating systems are attached using OSA adapters. Choosing between the TN3270 server in CS Linux on zSeries and the TN3270 server in z/OS does require research and one should carefully consider areas such as:

- ▶ Performance
- ▶ Scalability
- ▶ Functions
- ▶ Availability
- ▶ Manageability
- ▶ Cost
- ▶ Security

The following table summarizes some of these areas and compares selected characteristics of the two implementations.

Table 18-1 TN3270 server in CS Linux on zSeries compared to TN3270 server in z/OS

Areas of interest	TN3270 server in CS Linux on zSeries	TN3270 server in z/OS
Secure TN3270 support	Yes (SSL only) including client authentication (signature verification only)	Yes (SSL and TLS) including client authentication (signature verification) with optional SAF authentication and port protection (SERVAUTH)
Express logon	Yes (TCP-SSL connection with a z/OS DCAS server)	Yes (direct SAF interaction)
Support for RFC2355E contention resolution (important for both HOD and PCOMM)	Yes	Yes
zSeries hardware crypto exploitation	No	Yes
LU name assignment (LU name nailing)	Client IP address, client host name	Client IP address (including ranges), client host name, MVS user ID, server IP address, server interface name
Real or placeholder LU name assigned	Placeholder LU name (the locally defined name)	The real LU name
Printer association support	Yes	Yes
Specific LU requests	Yes	Yes

Areas of interest	TN3270 server in CS Linux on zSeries	TN3270 server in z/OS
ANS=CONT support	Yes	No
USS table support	N/A (dependent LUs done by the VTAM SSCP and controlled via VTAMLST definitions)	Yes TN3270 server reuses VTAM USS table definitions (z/OS V1R6 adds support for SCS mode USS)
Selecting SNA application	N/A (dependent LUs; done by the VTAM SSCP and controlled via VTAMLST definitions)	Yes LOGAPPL and QINIT support
Definitions	LU definitions on Linux and in VTAMLST (one PU per 255 LUs)	LU definitions in z/OS TN3270 server and VTAMLST (ACBs, cloning supported)
Capacity	Testing with 20,000 concurrent sessions has been done	128,000 concurrent sessions have been successfully tested
Load balancing	Traditional connection balancing	Traditional connection balancing. Sysplex Distributor adds value in terms of real-time LPAR capacity and server availability.
zSeries processor	Can be an IFL processor	A traditional zSeries processor

In general, the TN3270 server on z/OS will provide more function, better scalability, better performance, and simplified management and operation than any other platform.

18.4.2 High availability and load balancing

Achieving high availability and load balancing is critical for today's high-performance servers. TN3270 servers benefit greatly from high-availability techniques because of the large numbers of clients that often use them. In general z/OS is the strategic platform for the TN3270 server and in this section, we describe high availability and load balancing techniques for TN3270 traffic. For additional information, refer to "High availability and load distribution" on page 304.

Sysplex Distributor

The strategic load balancing (or more accurately, load distribution) technique on the z/OS platform, Sysplex Distributor, can be used to distribute TCP connections to multiple z/OS images in a sysplex. A distributing TCP/IP protocol stack advertises reachability to a distributed VIPA and accepts connections for a particular service, such as TN3270. The distributing TCP/IP protocol stack forwards individual connections to target TCP/IP protocol stacks within the sysplex based on a predefined policy and the current workload of those servers, as seen by WorkLoad Manager (WLM).

A failed target stack will be avoided by the distributing stack, thereby reducing the adverse effect of failed servers. As a result, higher levels of availability can be attained. Further, because the IP address of the service is a Virtual IP Address (VIPA), the service is highly available with respect to network interface failures as well.

External load balancers

Various load balancing technologies can also be deployed in the switch infrastructure that connects the z/OS Sysplex to the underlying IP network infrastructure. If such external load balancing techniques are used, it is recommended they are combined with the z/OS Load Balancing Advisor technology that allows them to learn about server and TCP/IP stack availability in the z/OS Sysplex and also current workload characteristics of the various LPARs in the Sysplex as seen by WorkLoad Manager. As an example, Cisco's Content Switching Module (CSM) in the Catalyst 6500 switch currently supports the z/OS Load Balancing Advisor.

LU name assignment when load balancing connections

In general, you must design your TN3270 server cluster in such a way that two TN3270 servers in the Sysplex will never try to assign the same LU name. The TN3270 servers do not coordinate LU name assignment, so it is up to your configuration to ensure they never assign the same LU name. How you do that depends on how you assign LU names:

- ▶ Generic request: Server decides entirely which LU name to assign.
 - Each TN3270 server must have its own group of LU names (generic pools)
 - There can be no overlap in the LU names each of the TN3270 servers uses
- ▶ Specific request with pool name: Server validates and decides which LU name in the named pool to assign.
 - Pool names can be the same in more TN3270 servers, but each server must have its own group of LU names assigned to those pool names.

- There can be no overlap in the LU names in the groups each of the TN3270 servers uses.
- ▶ Specific request with LU name: Server validates and assigns requested LU name.
 - Generally TN3270 servers can allow assigning the same LU names since theoretically a client workstation will only request an LU name once.
 - Theory does not always work in practice, so be somewhat careful with this approach. Some emulators support a “run the same” function that will start a second instance of the current emulator session and the second session will request the same LU name a second time. If that second request ended up in another TN3270 server, it might be honored a second time. To avoid this from happening, use timer-based affinity in the load balancer, so it will send all connection requests from the same client to the same server instance for a defined period of time.

Other considerations when load balancing connections

There are two TN3270 functions that require additional planning in a load-balanced TN3270 server environment. The first function is printer association. The display connection and the printer connection must end up in the same TN3270 server instance. If the display connection and the printer connection are load-balanced to different TN3270 servers, the printer association request will fail since the TN3270 server where the printer connection ended up has no knowledge of the display LU name, which was assigned by another TN3270 server. The second function to address is reconnect processing. A reconnect connection must go to the same TN3270 server to which the original connection was established. If the reconnect request ends up in another TN3270 server, it will fail.

The solution to both issues is to use a timer-based client IP address affinity that will send multiple connections from the same client IP address to the same server instance within a configurable timer window. Such timer-based affinity is supported by most load-balancers, including Sysplex Distributor.

SNA Generic resources

Another important aspect of high availability and load distribution is that of the eventual SNA application. Consider Figure 18-7 on page 329 which shows two Telnet servers (TelnetA and TelnetB) receiving connections from Telnet clients (via Sysplex Distributor). At TCP connection time, the Telnet client will be connected to the z/OS TN3270 server that is best suited for the incoming TCP connection. Once this connection has been established, the client chooses an application name. By using a generic resource (GR) application name (for example, CICS), VTAM will choose an optimal application instance for the client. This will work nicely if there is no GR instance in the same z/OS where the Telnet

server resides. But if there is an application that is a part of GR, local VTAM will always choose a zero-hop instance, that is CICS1. This is because Telnet terminals are applications for VTAM.

OS/390 V2R5 and above provide new options to bypass VTAM's bias to choose the nearest instance when resolving a GR name. It allows you to change the coding of the GR resolution exit, ISTEEXGR. VTAM checks the setting of the flag that indicates if VTAM should prefer zero-hop routes for applications or locally attached resources when the exit is first invoked. For more information, refer to *SNA in a Parallel Sysplex Environment*, SG24-2113.

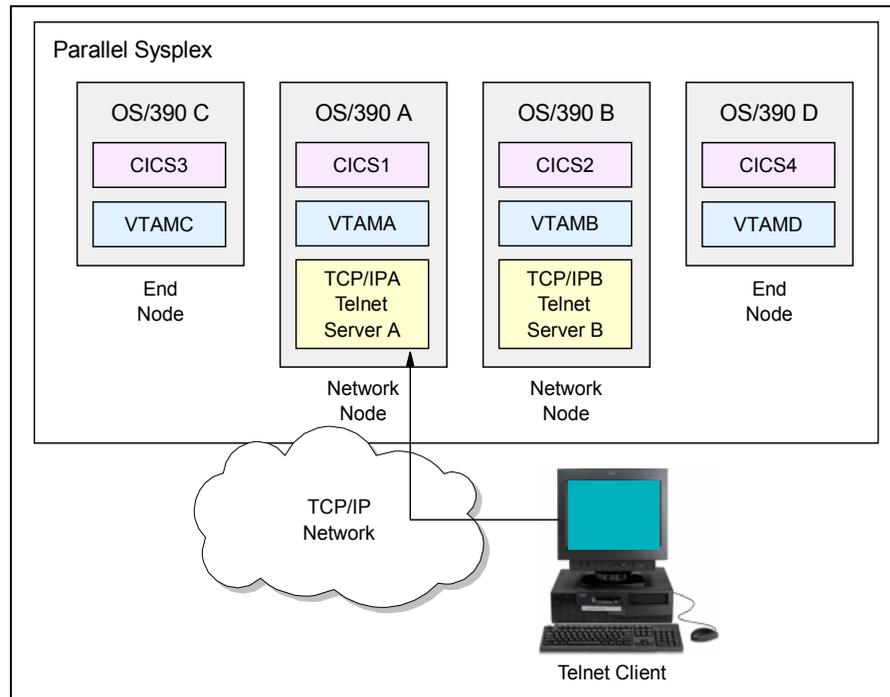


Figure 18-7 z/OS Telnet with VTAM generic resources in a sysplex environment



Enterprise Extender (EE)

This chapter introduces and discusses the importance of Enterprise Extender technology. In a nutshell, Enterprise Extender allows for the use of SNA transport protocols (namely APPN and HPR) over an Internet Protocol (IP) network. It enables the leveraging of IP-based infrastructural network components for use in delivering SNA traffic. This leads to the consolidation of your network infrastructure into one strategic IP-based network.



Restriction: VTAM for VM/ESA and VSE/ESA do not support Enterprise Extender.

19.1 What is Enterprise Extender (EE)?

Enterprise Extender architecture carries SNA (HPR) traffic of any LU type over an IP infrastructure without requiring changes to that infrastructure. It essentially treats the IP network as a particular type of SNA logical connection, in much the same way as an ATM or frame relay network is treated. In this manner, these SNA protocols act as transport protocols on top of IP, as does any other transport protocol such as Transmission Control Protocol (TCP).

19.2 Why is EE strategic?

Enterprise Extender combines features of SNA and IP to offer the best of both worlds when running SNA traffic over an IP backbone. Because of its design, Enterprise Extender is extremely flexible. It can be used in all networks from the smallest to the largest, and provides you with a choice of where the SNA/IP boundary is placed.

In this section we summarize some of the benefits of Enterprise Extender that make this solution strategic.

19.2.1 Cost-effectiveness and resource convergence

The use of Enterprise Extender allows you to avoid costly application rewrites by providing means of IP-enabling SNA applications that cannot be migrated to IP using technologies such as TN3270. That is, EE allows the continued use of native SNA applications over a different network: the IP network. This allows you to eliminate the SNA infrastructure altogether. The ultimate result is the convergence to a single network infrastructure that carries both IP and SNA application data.

Enterprise Extender has been designed to run over existing IP networks without requiring any change to applications or IP routers. SNA applications see the same SNA network interfaces as before, whereas IP routers continue to see the same IP (UDP) packets. Only the nodes at the edges of the IP network (potentially just the host systems) need to be aware of Enterprise Extender.

19.2.2 Flow and congestion control

TCP/IP and HPR both provide their own unique, network-specific mechanisms for flow and congestion control. TCP uses a windowed technique, whereas HPR uses a technique based on data rate. Enterprise Extender introduced a new variant of the HPR flow control method known as *responsive mode adaptive rate-based (ARB) flow control*. Responsive mode ARB, like basic mode ARB, is

designed to prevent network congestion; however, unlike basic mode ARB, it can also ensure a fair division of network capacity between the four SNA priorities and native IP traffic.

19.2.3 Class of service

One of the biggest issues facing those who wish to transport SNA over an IP network is the question of maintaining SNA's class of service. In SNA, the class of service specified for a particular session is used to determine both the route taken by the session and the transmission priority allotted to it.

With an IP backbone the route is essentially unpredictable because of IP's connectionless property. However, IP provides for a transmission priority using the precedence bits in the IP header. Many routers now support the use of these bits, but in the past they have tended to use the TCP or UDP port number as a means of assigning priorities to packets.

Enterprise Extender supports the use of both the precedence bits and the port numbers to inform the IP network of the transmission priority. Use of the precedence bits is recommended because the UDP or TCP port numbers are carried inside the IP datagram, whereas the precedence bits are in the IP header. Thus encrypted packets have unreadable port numbers and fragmented packets have no port numbers after the first fragment. For such encrypted or fragmented packets, intermediate routers cannot determine the appropriate priority.

19.2.4 Exploits ubiquitous Internet connectivity

Enterprise Extender enables remote branches or workstations to be connected to the SNA backbone using the Internet, with no application changes required, while maintaining SNA connectivity from end to end. Dependent LU sessions can be carried on an Enterprise Extender connection as easily as any others, and by utilizing the dependent LU requester function (available on all current IBM workstation and router platforms) they can take advantage of the Enterprise Extender technology all the way into the most remote locations.

19.2.5 Session availability

TCP/IP has always had the ability to reroute packets around failing components, without disrupting the connection, by means of the connectionless property of IP. More recently SNA has implemented the same function, albeit in a rather different fashion. The high-performance routing (HPR) extension to SNA is connection-oriented as SNA has always been, but when it detects a failure it will move an existing connection around a failing component. The use of HPR

transport over an IP network provides non-disruptive rerouting around failed network components using either IP or HPR methods, depending on the location of the failure.

19.2.6 Why not Data Link Switching (DLSw)?

Like Enterprise Extender (EE), Data Link Switching (DLSw) is a standard that was developed by IBM to enable the transport of SNA traffic across IP networks. However, there are some very important differences between EE and DLSw:

- ▶ DLSw was developed for use in multiprotocol routers and is essentially a more efficient extension of token-ring source route bridging (SRB) for wide area networks. EE can be implemented in routers (such as with Cisco Systems, Inc. “SNA Switching Services”) but can also be used in end systems such as hosts and intelligent workstations.
- ▶ DLSw runs only at the edges of the IP network and can, therefore, benefit from IP dynamic rerouting around failed network components without SNA session disruption; however, a failure of DLSw endpoint router will disrupt all SNA sessions that cross that router. EE, too, runs only at the edges of the IP network and benefits from IP dynamic rerouting around failed network components without SNA session disruption. However, because EE can be implemented in end systems (particularly your mainframe hosts), non-disruptive rerouting of SNA sessions around failed links or nodes in the network does not end at the data center router (as with DLSw). EE can non-disruptively reroute around outages of branch routers, telephone lines, and data center routers.
- ▶ Because DLSw runs only in routers, and because it entails substantial workload in the router, most implementations require multiple large, DLSw “peer” routers in each data center. By running EE in your mainframe hosts, you can avoid the cost and complexity of having pools of large data center routers. This can be of particular benefit if you have multiple data centers and/or a disaster recovery center.
- ▶ While DLSw is an industry standard, vendors such as Cisco Systems, Inc. have proprietary extensions that work only with their products.
- ▶ Because the DLSw standard is based on the concept of source-route bridging of LLC type 2 traffic, DLSw does not support HPR (which uses LLC type 1). However, some proprietary vendor extensions to DLSw support HPR connections in certain circumstances.
- ▶ DLSw, by itself, cannot enable the support of SNA class of service over the IP network because it can only distinguish between SNA and other protocols: not between different types of SNA traffic. Because EE integrates SNA APPN technology with your IP infrastructure, it can enable appropriate prioritization

for each class of SNA traffic. (For a more detailed discussion of SNA class of service support, see 5.3.1, “SNA class of service (COS)” on page 159.)

19.3 What you need to understand about EE

This section describes the mechanism used to carry SNA traffic over an IP network. It includes a discussion of interesting topics that arise when doing so, including the maintenance of Class of Service and TCP-friendly congestion and flow control.

19.3.1 Transporting SNA over IP

The designers of Enterprise Extender had the task of architecting the way in which SNA and IP-based protocols would be layered in order to transport SNA data over the IP network. They essentially had three choices of encapsulation of SNA data units:

- ▶ Raw IP datagrams. Datagrams are completely compatible with the HPR principles, as they flow through the network with minimal overhead and provide no error recovery of any sort. However, raw IP provides no means of multiplexing, particularly with no IETF designated protocol value for HPR. Using a non-designated protocol value could lead to inconsistencies with security measures that filter IP packets based on this value. Additionally, although raw IP allows priority and type of service to be specified, in practice not all networks or routers are (or can be) configured to support this.
- ▶ UDP packets provide the multiplexing required because they contain UDP port numbers. This allows Enterprise Extender packets to be distinguished from other IP packets. It also permits a priority scheme to be implemented independent of the type of service bits, since many routers are able to prioritize traffic based on the received port number. UDP, in addition, has low overhead since it does not concern itself with error recovery or flow control.
- ▶ A TCP connection also provides multiplexing by means of port numbers, but it incurs a significantly higher overhead than raw IP or UDP. A TCP connection handles error recovery, retransmission and flow control; none of these is required for an HPR connection because the RTP endpoints are responsible for all of them. Moreover, if the Enterprise Extender connection is only part of the HPR path, then one end or the other will be an ANR node; the additional burden of a TCP connection is unacceptable for an ANR node, which should do as little SNA processing as possible.

UDP was therefore the method chosen for Enterprise Extender, as Figure 19-1 on page 336 illustrates. It shows how SNA data is transported over the IP cloud in UDP frames.

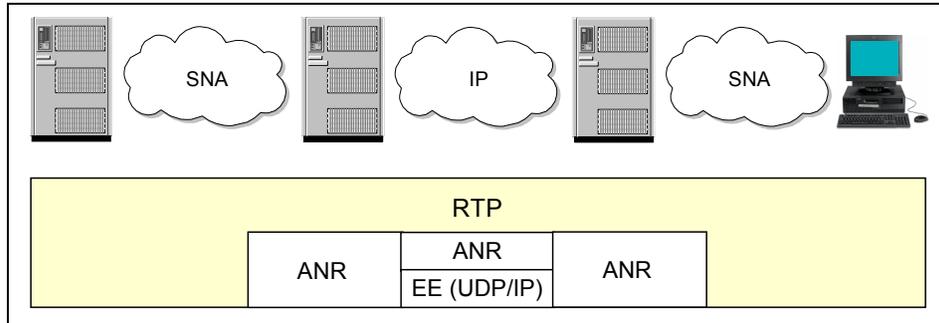


Figure 19-1 Enterprise Extender layering

19.3.2 Enterprise Extender architecture

Enterprise Extender is very similar in concept to the way native SNA is implemented over ATM:

- ▶ The underlying transport network appears as an APPN TG but uses logical data link control (LDLC) to exchange XIDs and NLPs. LDLC is a subset of LLC2 that eliminates much of the error handling and acknowledging that RTP makes unnecessary at link level. It is similar in concept to the qualified logical link control (QLLC) used to transport SNA traffic over an X.25 network, but there are some major differences that LDLC has to allow for:
 - LLC2 requires the use of several fixed timers, which are by their nature incompatible with a variable-route variable-delay IP network.
 - LLC2 performs error recovery, which is not necessary in HPR.
 - LLC2 requires in-order delivery, which cannot be guaranteed on an IP network.

LDLC, used also for native SNA over ATM, includes only the XID, TEST, DISC, DM and UI frame types. These are sufficient to establish the connection (XID), send data (UI), terminate the connection (DISC), and respond in the negative to a previous frame (DM). The TEST frames are used to check whether a connection is still active, a function required by HPR over IP.

- ▶ The UDP port number identifies the destination of the datagram as being the partner IP host's ANR routing function. Several UDP ports (12000-12004) have been registered with the Internet Assigned Number Authority (IANA) for this purpose. Each of these default ports is mapped to one of the APPN transmission priority values, with the fifth (12000) being used for XID exchange. An Enterprise Extender implementation may choose to alter these port numbers, but by using the registered defaults you can be reasonably sure that no other application will conflict with Enterprise Extender. ANR labels are mapped to the partner's IP address.

- ▶ Because there is no link-level error recovery and no guarantee that packets will arrive in order on an Enterprise Extender connection, only HPR NLPs can be transported once the XID flows are completed. Therefore, both partner nodes must support control flows over RTP.
- ▶ A Connection Network name can be defined on each EE node that connects to the IP network. This can dramatically simplify configuration tasks by allowing the EE nodes to predefine links to only a few other nodes and to dynamically bring up links with others for session traffic. Predefining all such logical links between each pair of EE nodes using the destination's host name or IP address could be an unpleasant task, just as predefining partners' MAC addresses is on a LAN.
- ▶ The SNA transmission priority is mapped to the UDP port number, which is why five UDP ports have been registered for Enterprise Extender use. The main reason for this is that many IP routers can be configured to prioritize traffic based on the port number. The Enterprise Extender architecture also permits the use of the precedence bits in the IP header for the same purpose. These bits are reserved in the TCP/IP architecture for exactly this usage, but not all routers take account of them. LDLC commands use the same precedence bit setting (the highest) as network priority NLPs. Table 19-1 shows the correspondence between APPN priorities, IP precedence bits, and UDP port numbers.

Table 19-1 Use of UDP/IP for APPN priorities

APPN Priority	IP Precedence	UDP Port
N/A (LDLC commands)	B'110'	12000
Network	B'110'	12001
High	B'100'	12002
Medium	B'010'	12003
Low	B'001'	12004

19.3.3 Flow and congestion control

The original ARB algorithm introduced with HPR works very well with SNA traffic alone, but is less efficient in the Enterprise Extender environment when SNA and IP traffic must coexist. Upon detection of a lost packet (a common occurrence in IP networks), the original ARB would immediately reduce its sending rate by a significant amount, thus impacting performance.

An enhanced algorithm known as Responsive Mode ARB was introduced with Enterprise Extender, and is now an option for RTP nodes whether or not their

HPR connection includes an Enterprise Extender link. Nodes that support Responsive Mode ARB can negotiate their level of ARB support during route setup exchange, and fall back to the original ARB if their partner does not support Responsive Mode. Responsive Mode ARB provides these features:

- ▶ It competes fairly with TCP congestion control.
- ▶ It can tolerate a certain level of lost data without significant degradation.
- ▶ It gives priority to short transmissions.
- ▶ It allocates a fair bandwidth to sustained transmissions, independent of propagation delays.
- ▶ It can ramp up its transmission rate faster at startup.

19.3.4 Enterprise Extender implementation

At the time this redbook was written, the following products implemented Enterprise Extender:

- ▶ Communications Server for OS/390
- ▶ z/OS Communications Server
- ▶ Communications Server for Windows
- ▶ Communications Server for Linux
- ▶ Communications Server for AIX
- ▶ Cisco routers with SNA Switching Services (SNASw)
- ▶ Microsoft Host Integration Server 2004

19.4 Strategies for exploiting EE

There are many ways to implement an Enterprise Extender network with many implications or restrictions that are not visible at first glance. As a result, we will discuss the available options and a general strategy for migrating to an EE network. We make general recommendations based on typical subarea configurations and issues.

19.4.1 The IP backbone as an APPN Connection Network

One extremely important aspect of Enterprise Extender is the ability to view the IP network as an APPN Connection Network. In this case, the benefit comes from the ability to dynamically establish a single one-hop HPR link to any host to which IP connectivity is enabled, provided that the host implements Enterprise Extender. In general, this allows the routing function to be handled entirely within IP. IP routers serve as the only routing nodes (hosts) in the network and no ANR routing is done. This minimizes the duplication of routing function by IP and HPR.

Usually, the migration to Enterprise Extender will not require change in the IP routers already in place in your network fabric except perhaps for the implementation of a organization-wide Quality of Service policy if you do not already have one.

19.4.2 EE within the sysplex

Typically, installations have at least one mainframe host, which may or may not participate in a sysplex. If you have a sysplex, you can view the group of hosts comprising the sysplex as an island of closely coupled machines. Participating in a sysplex gives these machines the ability to communicate over the media provided by the sysplex, using both SNA (APPN) and TCP/IP, natively. As a result, when migrating to an Enterprise Extender network, enable the EE function in all hosts as shown in Figure 19-2 on page 340, even data hosts in the sysplex.

Typically the APPN links in such a scenario would be:

- ▶ XCF links, which activate automatically between each LPAR in the sysplex
- ▶ MPC+ links, which need to be defined and enabled on every host

In fact, an APPN installation in a sysplex could include both types of links. In either case, the links form a network mesh, providing connectivity between any two hosts in the sysplex. In the case of MPC+, however, adding a new LPAR in the sysplex requires a change to every other host in the sysplex to enable a new MPC+ connection between the new host and the other hosts. On the other hand, MPC+ provides higher throughput rates than XCF links, making them a more viable option, until now.

With Enterprise Extender, SNA can take advantage of high-performance connectivity options such as:

- ▶ Gigabit Ethernet connectivity with OSA-Express
- ▶ Hipersockets (only within a single zSeries server)

Each host in the sysplex can be given a physical interface to the IP network (for example OSA) then EE links can be used to transport SNA host-to-host traffic. Such a configuration is illustrated in Figure 19-2 on page 340. If OSA-Express or Hipersockets is being used to connect the LPARs, performance will be even better than through ESCON.

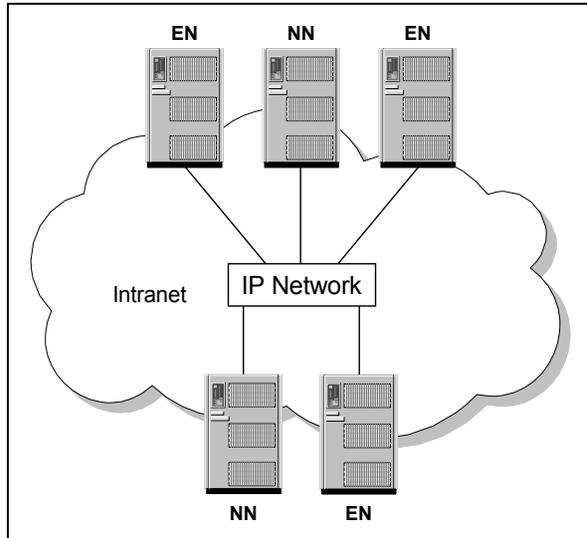


Figure 19-2 Enabling Enterprise Extender links within a sysplex

19.4.3 Connecting branches to hosts

There are many options for connecting branches to the data center. The single most important motivational factor in using Enterprise Extender to the branch is the ability to use existing TCP/IP network connectivity (perhaps even the Internet) instead of costly dedicated lines. Among the issues involved in the migration to Enterprise Extender in branches is the decision of using either branch extender network nodes or network nodes, the placement of EE and DLUR functions, and whether to use the branch IP router as the serving APPN network node (BrNN or NN).

BrNN vs. NN

Branch extender is an architecture used to reduce the number of NNs in the APPN topology. A BrNN acts as an EN to the upstream network while it presents the appearance of a NN to downstream nodes. Therefore, it is not part of the APPN topology and will not receive topology database updates (TDUs). If it registers its LUs to its serving NN (typically a VTAM), it will not receive any APPN directory search broadcasts, which can be quite extensive if the network is large and no measures are taken to reduce them.

Essentially, BrNN provides a powerful tool to simplify the logical structures of APPN networks by providing a hierarchical “proxy” of sorts that represents a summarization of those nodes that are configured to be downstream. Therefore, an APPN network of 40,000 nodes can be hierarchically structured as a network

of 1,000 BrNNs, each transparently summarizing the resources of the 40 nodes below it. Because the BrNN creates a hierarchical logical structure on the network and summarizes information from its downstream nodes, its use imposes certain limitations:

- ▶ No node downstream of a BrNN can provide the DLUR function. The BrNN must provide that function for all downstream nodes. Note that this is a logical configuration limitation, not a physical configuration limitation. For example, if you have a branch location for which you would like to have an additional DLUR (perhaps for high availability), you simply need to define another node in that branch as a separate BrNN, with its own DLUR, and all of the BrNNs in that branch can connect back to the data center using EE and the IP-based connection network.
- ▶ NNs cannot be downstream of a BrNN and all traffic from nodes downstream of a BrNN must transit the BrNN (even if there is a direct IP-based connection from the workstation to the data center). Again, this is a logical configuration limitation, not a physical configuration limitation. If your networking needs, for some reason, require a NN on the same remote LAN as a BrNN, use EE and configure the NN to be on the IP-based connection network. The limitation is that the NN cannot route through the BrNN.

Due to these limitations, the BrNN can pose a potential single point of failure and performance bottleneck that must be considered in your network design. Pure NNs are subject to APPN broadcasts and TDU flows, which can seriously impact the scalability of an APPN network. Consequently, use BrNN whenever possible.

Dependent LU Requester (DLUR) placement

In order to support older SNA devices such as IBM 3174s, 3274s, or SNA gateways (referred to in the SNA architecture as dependent LUs), the Dependent LU Requester (DLUR) function provides subarea SNA boundary function across APPN networks. Essentially, DLUR provides an LU6.2-based session over which the old subarea SNA SSCP-PU and SSCP-LU control session traffic (management and session establishment/termination) is passed to the host Dependent LU Server (DLUS). The actual LU-LU session traffic is sent from the PLU host to the DLUR using APPN routing, across the most appropriate path in the network (based upon the requested class of service). Prior to APPN and DLUR, if the mainframe host upon which a device depended for its session establishment services were to become unavailable, the device would no longer be able to support new requests for sessions unless taken over via operator command by another host. In such a situation with DLUR, a backup DLUS can be specified to be automatically contacted for SNA host session services.

Note: DLUR only supports dependent SLUs (not dependent PLUs). Specifically, the AS/400 SNA Primary LU Support (SPLS) function does not work if the AS/400 is attached via DLUR.

Essentially, DLUR can replace the SNA boundary function currently provided by NCP on your communication controller. Therein, however, lies a common misconception: Many network designers mistakenly conclude that, because they are replacing the boundary function provided by a communication controller, they should locate the DLUR function where their controllers are: in the data center.

The SNA boundary function requires a great deal of resources, both in terms of memory and CPU cycles. When distributed to branches, such as when using DLUR in routers or branch servers, the overall boundary function workload in the network is divided into smaller amounts which can more easily be handled by the routers or servers. Additionally, network performance and availability are enhanced in three important ways:

- ▶ Message segmentation is performed close to the dependent LU rather than in the data center and, consequently, fewer, larger messages can be sent across the network rather than sending the multiple smaller message segments.
- ▶ Since all dependent LU traffic must first flow through its SNA boundary function, when the boundary function is located close to the dependent LU, the network route between the dependent LU and the application can be more efficient. This is especially important when you have multiple data centers (or a disaster recovery center).
- ▶ Finally, locating the DLUR as close as possible to the dependent LUs extends the benefits of HPR (with its adaptive rate-based congestion control and non-disruptive rerouting around failed links or nodes in the network) over the typically least reliable and most congested part of your networking infrastructure: your wide-area communications links.

When DLUR is consolidated into the data center, each DLUR must serve a larger number of dependent LUs. Depending upon the number of dependent LUs in your network, you may require many servers or routers to provide the necessary DLUR function. The centralized DLUR message segmentation can result in less efficient use of your network and, if you have multiple data centers, your SNA session path will first have to go through the data center where your DLUR is located before going on to the other data center. Finally, if you use a disaster recovery center, you will need to duplicate your DLUR servers or routers across both centers thereby increasing the cost of having centralized DLURs.

Consequently, distribute your DLUR function whenever possible.

Supporting SNA in the branch

As discussed so far, there are many compelling reasons to extend the “AHDBEE” set of functions (APPN, HPR, DLUR, BrNN, and EE) to your remote locations. So where should you run them, in your routers or in branch office servers?

One option is to place the AHDBEE functions on the branch router itself. Generally speaking, a failure of this node would cause problems in both the APPN network and the IP network, which can be recovered from if alternate IP connectivity exists and a network node server backup is present on the branch network. However, the difficulty in recovering from such an error may be a reason to attempt to avoid such a configuration. Also, your router may not be able to support the needed function without costly hardware and software upgrades.

Alternatively, you may choose to run the AHDBEE functions in servers located in your remote locations. It is very common in certain industries to use a utility server in a remote location, such as a “branch platform” in banking or an “in-store processor” in retail. If you have such a server, it may be considerably less expensive to support your SNA from the server than from the router, leaving the router to do what it does best: Route IP packets.

The best place to support your SNA requirements for your remote locations will depend upon the unique requirements and decision criteria of your organization.

19.4.4 Securing Enterprise Extender links with IPSec

Enterprise Extender links use IP and are therefore subject to the same security concerns of any IP traffic. This may not be much of an issue within an intranet, but suddenly becomes a potential nightmare when the IP network is actually the Internet. In general, this applies to any insecure IP connection over which an EE link flows. The specific type of link, the sensitivity of the data flowing over the link, and your commitment to securing your data determine the definition of insecure. An example of such a link could be a link that has replaced a leased line of an SNI connection with an EE link between the two networks using the Internet.

IPSec is an optional protocol of the TCP/IP protocol suite that provides mechanisms for authenticating and encrypting data to transport over a series of links, one of which is considered insecure. Whether to authenticate, encrypt, or do both is a decision based on the security needs of your network. We recommend the use of the IPSEC Encapsulating Security Payload (ESP) protocol within IPSec for both encrypting and authenticating data over any insecure link. Furthermore, this encryption and authentication need not be done at an APPN node, and probably is more suitable at existing firewalls (especially if they implement IPSec already). In this case, the IPSec must be used in tunnel mode, since data must be forwarded beyond the IPSec end point.



Message queuing (MQ)

Message queuing (MQ) technology can be applied to address a wide variety of business problems, from the integration of business processes within an organization to the movement of information between organizations. This chapter discusses the application of MQ technology to the migration of functions from your communication controller environment.

20.1 What is MQ?

Message queuing is a method of program-to-program communication. Programs communicate by writing and retrieving application-specific data (messages) to and from queues, without having a private, dedicated, logical connection to link them. In addition, the MQ asynchronous message delivery mechanisms support multiple levels of delivery assurance and recovery. Using assured information delivery, a message queue manager will never lose a message.

Messaging means that programs communicate with each other by sending data in messages and not by calling each other directly. The messages can be data, programs, or other content being passed from one program to another, or from one system to another.

Queuing means that programs communicate through queues. Programs communicating through queues need not be executed concurrently; therefore, the primary service delivered by MQ is asynchronous message delivery.

The IBM WebSphere MQ family (formerly known as MQSeries) is award-winning middleware for commercial messaging and queuing. It is used by thousands of clients in every major industry in many countries around the world. The WebSphere MQ products enable programs to communicate with each other across a network of unlike components, such as processors, subsystems, operating systems and communication protocols. WebSphere MQ programs use a consistent application program interface (API) across a wide variety of platforms.

20.2 Why is MQ strategic?

MQ provides a widely accepted and increasingly popular mechanism for reliable asynchronous communication across diverse environments. The MQ application programming interface (MQI) removes the need for application programmers to write code to deal with the details of communications. It also enhances the network manager's control of the network topology by allowing him to reconfigure network connections and move queues from one place to another without affecting the applications using those queues. IBM views the WebSphere MQ product family as strategic and has devoted extensive resources to further enhancing it for their 7,000+ customers. Also, over 350 independent software vendors offer services and products for MQ environments.

Because of its robust support of asynchronous communications, MQ can provide your organization with a well-suited, strategic, successor technology to legacy business-to-business transfers of information such as RJE. Additionally, MQ provides a powerful set of tools for integrating diverse computing environments. Rather than developing a different solution for each inter-organizational

communication need, MQ provides you with a consistent and scalable means of exchanging information. WebSphere MQ software is available on more than 35 different platforms and presents a consistent application programming interface across them all. Finally, MQ provides reliability:

- ▶ Message-level reliability that ensures each message is delivered “once and only once.”
- ▶ Transaction-level reliability with built-in XA-compliant Resource Manager capability. (The X/Open XA interface is a specification that describes the protocol for transaction coordination, commitment, and recovery between a Transaction Manager and one or more Resource Managers.)

20.3 What you need to understand about MQ

The discussion of MQ in this book is intended to give you a fundamental basis for understanding the strategic role MQ could have in your IT infrastructure and how MQ might be applied in your environment to aid in your communication controller migration efforts.

While MQ is primarily thought of as an application development topic, its adoption can free applications from underlying infrastructure dependencies such as requirements for ongoing support of BSC lines, SNA networks, and SNI interconnections.

There is much more information about MQ that we cannot possibly cover in this book. For more about the IBM WebSphere MQ family of products, see:

- ▶ The IBM WebSphere MQ Family Web site
<http://www.ibm.com/software/ts/mqseries/>
- ▶ IBM Redpaper *MQSeries Primer*
<http://www.redbooks.ibm.com/redpapers/pdfs/redp0021.pdf>

We now discuss two MQ concepts, asynchronous messaging and adapters, that may be key in your communication controller migration efforts.

20.3.1 Asynchronous messaging

Fundamentally, MQ provides an asynchronous messaging capability. With MQ, the sending program proceeds with its processing without waiting for a reply to its message. In contrast, synchronous messaging waits for the reply before it resumes processing. This asynchronous mode of operation offers a number of advantages:

- ▶ **Reliability:** As discussed in 20.2, “Why is MQ strategic?” on page 346, MQ can provide reliability both at the message and the transaction levels. This reliability also saves your application programmers from having to develop code to handle myriad exception conditions.
- ▶ **Fault tolerance:** Messages will not be lost in spite of IT infrastructure failures such as network outages and, as long as a route is available (or becomes available) the messages get delivered.
- ▶ **Scalability:** Because programs are allowed to proceed with additional work without having to wait for message replies, greater efficiencies and scalability can be achieved.
- ▶ **Performance:** When converting existing applications to the asynchronous messaging style of MQ, performance is often enhanced by enabling more work to be done in parallel. When migrating from existing application environments such as BSC RJE toward MQ, performance may be enhanced because of the MQ support for a diverse set of networking transports, the most important of which is TCP/IP. For example, in the case of BSC RJE, a migration to MQ could allow migration of the traffic from 19.2 Kbps BSC links to high-speed (1.5 Mbps and above) TCP/IP network links.

20.3.2 Adapters

An adapter (also called bridge, link, or connector) is a program that moves data between a message on a queue and an application or environment. Adapters handle data inbound to and outbound from the application or environment.

Adapters can be very helpful in a migration to MQ. For example, in an RJE environment, with an adapter for your existing host application, you could support having some of your remote users send their messages using MQ while others continue to use the current RJE-based mechanism. Conversely, after you have migrated your host application to the MQ application programming interface (MQI), you could use an adapter to continue to support RJE-based users by implementing an adapter from JES.

20.4 Strategies for exploiting MQ

If your organization has MQ today, you may be able to leverage that existing investment by using it to enable simplification and optimization of your underlying networking infrastructure. If you do not have MQ today, you may be able to justify the initial investment in MQ with your infrastructure optimization efforts, while that investment pays dividends into the future through enabling your organization's application integration efforts.



Part 4

Appendixes



Physical inventory worksheets

This appendix provides you with a set of worksheets that can be used to help you to perform a physical inventory of your communication controller environment.

Instructions

This appendix contains configuration sheets that can be used to determine the exact physical specifications of a 3745 and/or 3746. They are provided as a tool for you to use and may be photocopied. Fill out a set for each 3745 or 3746-950.

Note: Even for older controller models (such as IBM 3705, 3720, and 3725 controllers), it is important to carefully explore and clearly understand the equipment that you have installed as well as how it is being used. However, a physical inventory is not necessary for these devices because IBM no longer supports them and you will almost always be better off migrating from them. For older model controllers, a logical and functional inventory will be sufficient.

These physical inventory sheets are intended to provide an organized means for identifying your installed communication controller hardware components. You need only fill out those sheets that apply. The physical inventory sheets are structured as follows:

3745 and attached frames

- ▶ Part 1: 3745 machine overview: Covers the basic components of the 3745 (such as memory and console specifications)
- ▶ Part 2: 3745 base machine configurations
 - Section A: 3745-x1x Model (for example, a 3745-210 or a 3745-61A)
 - Section B: 3745-130 Model
 - Section C: 3745-150 Model
 - Section D: 3745-160 Model
 - Section E: 3745-17x Model (for example, a 3745-170 or a 3745-17A)
- ▶ Part 3: 3746 adapter expansion configurations (3746-A11 and 3746-A12)
- ▶ Part 4: 3746 line expansion configurations (3746-L13, 3746-L14 and 3746-L15)
- ▶ Part 5: 3746-900 configuration

3746-950

- ▶ Part 1: 3746-950 machine overview
- ▶ Part 2: 3746-950 configuration

Fill out all inventory sheets that pertain to your particular controller configuration. For example, if you have a 3745-17A with a 3746-900 installed, fill out:

- ▶ Part 1: 3745 Machine Overview
- ▶ Part 2, Section E: 3745-17x Model
- ▶ Part 5: 3746-900 Configuration

Or, if you have a 3745-610 with 3746-A11, 3746-A12 and 3746-L13 installed, fill out:

- ▶ Part 1: 3745 Machine Overview
- ▶ Part 2, Section A: 3745-x1x Model
- ▶ Part 3: 3746 Adapter Expansion Configurations (both 3746-A11 and 3746-A12 sheets)
- ▶ Part 4: 3746 Line Expansion Configurations (only 3746-L13 sheet)

The following resources may be helpful in your efforts to inventory your communication controller environment:

- ▶ Your IBM Customer Engineer (CE)
- ▶ The Configuration Definition File-Extended (CDF-E) from the MOSS console

Note that IBM may charge you a fee for this service.

3745 and attached frames physical inventory

Part 1: 3745 Machine Overview

3745 Model _____ Serial Number _____

Customer Designation (Ex. NCP Name) _____

Indicate all attached 3746 expansion frames.

3746-A11 SN _____

3746-A12 SN _____

3746-L13 SN _____

3746-L14 SN _____

3746-L15 SN _____

3746-900 SN _____

Controller Expansion (Rack) Qty _____ (0,1 or 2)

(Please inventory each attached frame separately on 3746 Model Sheets.)

How much memory (per CCU if model 41x or 61x)?

4M _____

8M _____

16M _____

Part 1: 3745 Machine Overview

3745 Model _____ Serial Number _____

Console Information

3151 or Equivalent _____

Service Processor _____

Type:

9577 _____

9585 _____

3172 P/N 41H7520 _____

3172 P/N 55H7630 _____

7585 _____

6275 _____

6563 _____

Indicate serial numbers of other 37XXs that share this console.

Part 2: 3745 Base Machine Configuration

Section A: 3745-x1x Model (210/A, 310/A, 410/A, or 610/A)

3745 Model _____ Serial Number _____

Channel Board

Bus Group 1	Pos. 1 _____	Pos. 2 _____	Pos. 3 _____	Pos. 4 _____
Bus Group 2	Pos. 5 _____	Pos. 6 _____	Pos. 7 _____	Pos. 8 _____

Legend:

CADS = Channel Adapter Data Streaming

BCCA = Buffer Chaining Channel Adapter

TPS = Two Processor Switch

Part 2: 3745 Base Machine Configuration

Section A: 3745 Models x1x

3745 Model _____ Serial Number _____

TSS Board

Bus Group 1	Pos. 1 Adap: _____		Pos. 2 Adap: _____		Pos. 3 Adap: _____		Pos. 4 Adap: _____	
	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____
Bus Group 2	Pos. 5 Adap: _____		Pos. 6 Adap: _____		Pos. 7 Adap: _____		Pos. 8 Adap: _____	
	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____

Legend:

Adapters

HSS = High Speed Scanner

LSS = Low Speed Scanner

ELA = Ethernet/IEEE 802.3 Adapter

TRA1 = Token-Ring Adapter Type 1

TRA2 = Token-Ring Adapter Type 2

Port Cards

TIC = Token-Ring Interface Coupler

LAN = Ethernet LAN Attached

#x fff = Where 'x' is the Line Unit Area driven by the installed LSS and 'fff' is the model number of frame where the Line Unit Area resides

(Example: #2 L13 - would indicate

Line Unit Area 2 in the 3746-L13)

(Example: #3 210 - would indicate Line

Unit Area 3 in the 3745-210)

Part 2: 3745 Base Machine Configuration

Section A: 3745 Models x1x

3745 Model _____ Serial Number _____

Line Unit

Area 1				Area 2			
LIC Type _____							

Additional Line Units

Area 3				Area 4			
LIC Type _____							

Part 2: 3745 Base Machine Configuration

Section A: 3745 Models x1x

3745 Model _____ Serial Number _____

Line Unit

Area 5				Area 6			
LIC Type _____							

Additional Line Units

Area 7				Area 8			
LIC Type _____							

Part 2: 3745 Base Machine Configuration

Section B: 3745-130 Model

3745 Model 130 Serial Number _____

Channel Board

Pos. 8 _____	Pos. 7 _____	Pos. 6 _____	Pos. 5 _____
-----------------	-----------------	-----------------	-----------------

Legend:

CADS = Channel Adapter Data Streaming
BCCA = Buffer Chaining Channel Adapter
TPS = Two Processor Switch

Part 2: 3745 Base Machine Configuration

Section B: 3745-130 Model

3745 Model 130 Serial Number _____

TR Board

Pos. 1 _____		Pos. 2 _____	
Port _____	Port _____	Port _____	Port _____

Legend:

Adapters

TRA1 = Token-Ring Adapter Type 1

TRA2 = Token-Ring Adapter Type 2

Port Cards

TIC = Token-Ring Interface Coupler

TSS Board

Pos. 4 _____		Pos. 3 _____	
Port _____	Port _____	Port _____	Port _____

Legend:

Adapters

HSS = High Speed Scanner

LSS = Low Speed Scanner

Port Cards

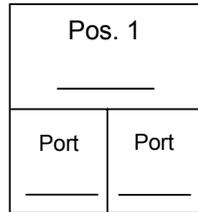
LAN = Ethernet LAN Attached

Part 2: 3745 Base Machine Configuration

Section C: 3745-150 Model

3745 Model 150 Serial Number _____

TR Board



Legend:

Adapters

TRA1 = Token-Ring Adapter Type 1

TRA2 = Token-Ring Adapter Type 2

Port Cards

TIC = Token-Ring Interface Coupler

Part 2: 3745 Base Machine Configuration

Section C: 3745-150 Model

3745 Model 150 Serial Number _____

TSS Board

Pos. 12 Adap: <u>NA</u>		Pos. 11 Adap: <u>NA</u>		Pos. 10 Adap: <u>NA</u>		Pos. 9 Adap: _____		Pos. 4 Adap: <u>NA</u>		Pos. 3 Adap: _____	
Port <u>NA</u>	Port <u>NA</u>	Port <u>NA</u>	Port <u>NA</u>	Port <u>NA</u>	Port <u>NA</u>	Port _____	Port _____	Port <u>NA</u>	Port <u>NA</u>	Port _____	Port _____

Legend:

Adapters

HSS = High Speed Scanner

LSS = Low Speed Scanner

ELA = Ethernet/IEEE 802.3 Adapter

Port Cards

LAN = Ethernet LAN Attached

#x = Where 'x' is the Line Unit Area

driven by the installed LSS

Line Unit

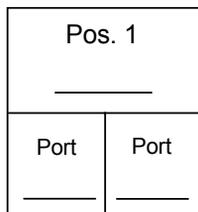
Area 1				Area 2			
LIC Type _____							

Part 2: 3745 Base Machine Configuration

Section D: 3745-160 Model

3745 Model 160 Serial Number _____

TR Board



Legend:

Adapters

TRA1 = Token-Ring Adapter Type 1

TRA2 = Token-Ring Adapter Type 2

Port Cards

TIC = Token-Ring Interface Coupler

Part 2: 3745 Base Machine Configuration

Section D: 3745-160 Model

3745 Model 160 Serial Number _____

TSS Board

Pos. 12 Adap: _____		Pos. 11 Adap: _____		Pos. 10 Adap: _____		Pos. 9 Adap: _____		Pos. 4 Adap: _____		Pos. 3 Adap: _____	
Port _____	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____

Legend:

Adapters

HSS = High Speed Scanner

LSS = Low Speed Scanner

ELA = Ethernet/IEEE 802.3 Adapter

Port Cards

LAN = Ethernet LAN Attached

#x = Where 'x' is the Line Unit Area

driven by the installed LSS

Line Unit

Area 1				Area 2			
LIC Type _____							

Part 2: 3745 Base Machine Configuration

Section E: 3745-17x Model (3745-170 or 3745-17A)

3745 Model 17 Serial Number _____

Channel Board

Pos. 8	Pos. 7	Pos. 6	Pos. 5
_____	_____	_____	_____

Legend:

CADS = Channel Adapter Data Streaming
BCCA = Buffer Chaining Channel Adapter
TPS = Two Processor Switch

Part 2: 3745 Base Machine Configuration

Section E: 3745-17x Model (3745-170 or 3745-17A)

3745 Model 17 Serial Number _____

TR Board

Pos. 1 _____	
Port _____	Port _____

Legend:

Adapters

TRA1 = Token-Ring Adapter Type 1

TRA2 = Token-Ring Adapter Type 2

Port Cards

TIC = Token-Ring Interface Coupler

Part 2: 3745 Base Machine Configuration

Section E: 3745-17x Model (3745-170 or 3745-17A)

3745 Model 17 Serial Number _____

TSS Board

Pos. 12 Adap: _____		Pos. 11 Adap: _____		Pos. 10 Adap: _____		Pos. 9 Adap: _____		Pos. 4 Adap: _____		Pos. 3 Adap: _____	
Port _____	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____

Legend:

Adapters

HSS = High Speed Scanner

LSS = Low Speed Scanner

ELA = Ethernet/IEEE 802.3 Adapter

Port Cards

LAN = Ethernet LAN Attached

#x = Where 'x' is the Line Unit Area

driven by the installed LSS

Part 2: 3745 Base Machine Configuration

Section E: 3745-17x Model (3745-170 or 3745-17A)

3745 Model 17 Serial Number _____

Line Units

Area 1				Area 2			
LIC Type _____							

Additional Line Units

Area 3				Area 4			
LIC Type _____							

Part 2: 3745 Base Machine Configuration

Section E: 3745-17x Model (3745-170 or 3745-17A)

3745 Model 17 Serial Number _____

Additional Line Units

Area 5				Area 6			
LIC Type							
_____	_____	_____	_____	_____	_____	_____	_____

Additional Line Units

Note: Areas 7 and 8 are for LIC types 5 and 6 only.

Area 7				Area 8			
LIC Type							
_____	_____	_____	_____	_____	_____	_____	_____

Part 3: 3746 Adapter Expansion Configurations (3746-A11 and 3746-A12)

3746 Model A11 Serial Number

Channel Board

3746 Models A11 Only

Bus Group 1	Pos. 9 _____	Pos. 10 _____	Pos. 11 _____	Pos. 12 _____
Bus Group 2	Pos. 13 _____	Pos. 14 _____	Pos. 15 _____	Pos. 16 _____

Legend:

CADS = Channel Adapter Data Streaming
 BCCA = Buffer Chaining Channel Adapter
 TPS = Two Processor Switch

Part 3: 3746 Adapter Expansion Configurations (3746-A11 and 3746-A12)

3746 Model A11 Serial Number

TSS Board

3746 Models A11 Only

Bus Group 1	Pos. 9 Adap: _____	Pos. 10 Adap: _____	Pos. 11 Adap: _____	Pos. 12 Adap: _____
	LIC area # _____			
Bus Group 2	Pos. 13 Adap: _____	Pos. 14 Adap: _____	Pos. 15 Adap: _____	Pos. 16 Adap: _____
	LIC area # _____			
Bus Group 1	Pos. 17 Adap: _____	Pos. 18 Adap: _____	Pos. 19 Adap: _____	Pos. 20 Adap: _____
	LIC area # _____			
Bus Group 2	Pos. 21 Adap: _____	Pos. 22 Adap: _____	Pos. 23 Adap: _____	Pos. 24 Adap: _____
	LIC area # _____			

Legend:

Adapters

LSS = Low Speed Scanner
ALC = Airline Line Control Scanner

Port Cards

#x fff = Where 'x' is the Line Unit Area driven by the installed LSS and 'fff' is the model number of frame where the Line Unit Area resides
(Example: #2 L13 - would indicate Line Unit Area 2 in the 3746-L13)
(Example: #3 210 - would indicate Line Unit Area 3 in the 3745-210)

Part 3: 3746 Adapter Expansion Configurations (3746-A11 and 3746-A12)

3746 Model A12 Serial Number

TSS Board

3746 Models A12 Only

Bus Group 1	Pos. 25 Adap: _____	Pos. 26 Adap: _____	Pos. 27 Adap: _____	Pos. 28 Adap: _____
	LIC area # _____			
Bus Group 2	Pos. 29 Adap: _____	Pos. 30 Adap: _____	Pos. 31 Adap: _____	Pos. 32 Adap: _____
	LIC area # _____			

Bus Group 1	Pos. 33 Adap: _____	Pos. 34 Adap: _____	Pos. 35 Adap: _____	Pos. 36 Adap: _____
	LIC area # _____			
Bus Group 2	Pos. 37 Adap: _____	Pos. 38 Adap: _____	Pos. 39 Adap: _____	Pos. 40 Adap: _____
	LIC area # _____			

Legend:

Adapters

LSS = Low Speed Scanner
ALC = Airline Line Control Scanner

Port Cards

#x fff = Where 'x' is the Line Unit Area driven by the installed LSS and 'fff' is the model number of frame where the Line Unit Area resides
(Example: #2 L13 - would indicate Line Unit Area 2 in the 3746-L13)
(Example: #3 210 - would indicate Line Unit Area 3 in the 3745-210)

Part 4: 3746 Line Expansion Configurations (3746-L13, L14 and L15)

3746 Model L14 Serial Number

Line Unit

Area 1				Area 2			
LIC Type _____							

Additional Line Units

Area 3				Area 4			
LIC Type _____							

Area 5				Area 6			
LIC Type _____							

Area 7				Area 8			
LIC Type _____							

Part 5: 3746-900 Configuration

3746 Model 900 Serial Number _____

Overview

Extended Microcode Options:

Extended Functions 1 (5800) _____
Extended Functions 2 (5802) _____
Extended Functions 3 (5801) _____
Extended Functions 4 (5810) _____
Extended Functions 5 (5812) _____
Extended Functions 6 (5813) _____
X.25 (5030) _____
IP (5033) _____

Multiaccess Enclosure (MAE)
present?

Yes _____ No _____

MAE Microcode Options: _____

Extended Functions 1 (5804) _____
Extended Functions 2 (5805) _____
Extended Functions 3 (5807) _____
TN3270 Server (5806)

Network Node Processor (NNP)

Qty _____ (0,1 or 2)

NNP Type:

Type 1 (3172) _____
Type 2 (7585) _____
Type 3 (6275) _____
Type 4 (6563) _____
Type 5 (6578) _____

Part 5: 3746-900 Configuration

3746 Model 900 Serial Number

Base Enclosure Top View

P	N	M	L	K	J	H	G	F	E	D	C	Rear Side
Port _____	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____	Port TIC3	Port CBC	Port N/A	Port N/A	
Processor _____		Processor _____		Processor _____		Processor _____		Processor CBSP Type _____		Processor Power Supply		
Slot 6 (P)		Slot 5 (M)		Slot 4 (K)		Slot 3 (H)		Slot 2 (F)		Slot 1 (D)		Front Side

Notes: 1. Indicate CBSP Type in Slot 2 (F)

2. For 3745 Models 41A and 61A, Slot 3 (H) must be a TRPx and Port G must be a CBC

Processors:

TRP = Token-Ring Processor Type 1
 TRP2 = Token-Ring Processor Type 2
 TRP3 = Token-Ring Processor Type 3
 ESCP = ESCON Processor Type 1
 ESCP2 = ESCON Processor Type 2
 ESCP3 = ESCON Processor Type 3
 CLP = Communication Line Processor
 CLP3 = Communication Line Processor Type 3
 SIE = Switch Interface Extension (MAE connection)

Ports:

TIC3 = Token-Ring Coupler Type 3
 ETH = Ethernet Port/Ethernet-TR Bridge
 ESCC = ESCON Coupler Type 1
 ESCC2 = ESCON Coupler Type 2
 LIC11x = Line Interface Coupler Type 11 where 'x' is
 the Line Connection Box (LCB) ID where the
 ARCS are installed that correspond to this LIC
 LIC12 = Line Interface Coupler Type 12

Part 5: 3746-900 Configuration

3746 Model 900 Serial Number _____

Expansion Enclosure 1 Top View

P	N	M	L	K	J	H	G	F	E	D	C	Rear Side
Port	Port	Port	Port	Port	Port	Port	Port	Port	Port	Port	Port	
_____	_____	_____	_____	_____	_____	_____	_____	_____	_____	_____	_____	
Processor		Processor		Processor		Processor		Processor		Processor		
_____		_____		_____		_____		_____		_____		
Slot 12 (P)		Slot 11 (M)		Slot 10 (K)		Slot 9 (H)		Slot 8 (F)		Slot 7 (D)		Front Side

Processors:

TRP = Token-Ring Processor Type 1
 TRP2 = Token-Ring Processor Type 2
 TRP3 = Token-Ring Processor Type 3
 ESCP = ESCON Processor Type 1
 ESCP2 = ESCON Processor Type 2
 ESCP3 = ESCON Processor Type 3
 CLP = Communication Line Processor
 CLP3 = Communication Line Processor Type 3
 SIE = Switch Interface Extension (MAE connection)

Ports:

TIC3 = Token-Ring Coupler Type 3
 ETH = Ethernet Port/Ethernet-TR Bridge
 ESCC = ESCON Coupler Type 1
 ESCC2 = ESCON Coupler Type 2
 LIC11x = Line Interface Coupler Type 11 where 'x' is
 the Line Connection Box (LCB) ID where the
 ARCS are installed that correspond to this LIC
 LIC12 = Line Interface Coupler Type 12

Part 5: 3746-900 Configuration

3746 Model 900 Serial Number _____

Expansion Enclosure 2 Top View

P	N	M	L	K	J	H	G	F	E	D	C	Rear Side
Port _____	Port _____											
Processor _____		Processor _____		Processor _____		Processor _____		Processor _____		Processor _____		Front Side
Slot 18 (P)		Slot 17 (M)		Slot 16 (K)		Slot 15 (H)		Slot 14 (F)		Slot 13 (D)		

Processors:

- TRP = Token-Ring Processor Type 1
- TRP2 = Token-Ring Processor Type 2
- TRP3 = Token-Ring Processor Type 3
- ESCP = ESCON Processor Type 1
- ESCP2 = ESCON Processor Type 2
- ESCP3 = ESCON Processor Type 3
- CLP = Communication Line Processor
- CLP3 = Communication Line Processor Type 3
- SIE = Switch Interface Extension (MAE connection)

Ports:

- TIC3 = Token-Ring Coupler Type 3
- ETH = Ethernet Port/Ethernet-TR Bridge
- ESCC = ESCON Coupler Type 1
- ESCC2 = ESCON Coupler Type 2
- LIC11x = Line Interface Coupler Type 11 where 'x' is the Line Connection Box (LCB) ID where the ARCS are installed that correspond to this LIC
- LIC12 = Line Interface Coupler Type 12

Part 5: 3746-900 Configuration

3746 Model 900 Serial Number

LCB ID Line Connection Box Layout

	0	+1	+2	+3	+4	+5	+6	+7	+8	+9	+10	+11	+12	+13	+14	
ARC Type																To LCBE
Attach																
Length																

Line Connection Box Expansion Layout

	+16	+17	+18	+19	+20	+21	+22	+23	+24	+25	+26	+27	+28	+29	+30	
ARC Type																To LCB
Attach																
Length																

Length:

ARC Types:

V.35
V.24
X.21

Attach:

DCE
DTE

.6 = .6 Meters
1.2 = 1.2 Meters
2.4 = 2.4 Meters
5 = 5 Meters
10 = 10 Meters
12 = 12 Meters
15 = 15 Meters
ST = Stub cable to 3745 cable

Notes:

Fill out one sheet for each
LCB/LCBE combo installed

Note: A 3746-900 may have as many as 32 LCB/LCBE pairs. Fill out one of these sheets for each one.

Part 5: 3746-900 Configuration

3746 Model 900 Serial Number _____

Multiaccess Enclosure (MAE)

How is this MAE connected to the 3746-900?

Token-Ring _____

Directly _____

Base Power	PCMCIA Card	Slot 1 LIC _____	Slot 2 LIC _____	Slot 3 LIC _____	Slot 4 LIC _____
Second Power	RS232	Slot 5 LIC _____	Slot 6 LIC _____	Slot 7 LIC _____	Slot 8 LIC _____

LIC Types:

280 = 2-Port Token-Ring

281 = 2 Port Ethernet

282 = 8 Port V.24/EIA-232

283 = 1 Port ISDN PRI-T1/J1

284 = 1 Port Multi-Mode ATM

286 = 1 Port Multi-Mode FDDI

287 = 1 Port ESCON

288 = 1 Port Fast Ethernet

289 = 1 Port HSSI

290 = 6 Port V.35/V.36

291 = 8 Port X.21

293 = 1 Port Single-Mode ATM

294 = 1 Port 155 Multi-Mode ATM

295 = 1 Port 155 Single-Mode ATM

297 = 4 Port ISDN T1/J1

297+ = 4 Port ISDN T1J1 plus 4-P Daughter

299 = 1 Port Parallel Channel

3746-950 physical inventory

Part 1: 3746-950 Configuration

3746 Model 950 Serial Number _____

Overview

Extended Microcode Options:

Extended Functions 1 (5800) _____
Extended Functions 2 (5802) _____
Extended Functions 3 (5801) _____
Extended Functions 4 (5810) _____
Extended Functions 5 (5812) _____
Extended Functions 6 (5813) _____
X.25 (5030) _____
IP (5033) _____

Multiaccess Enclosure (MAE)
present?

Yes _____ No _____

MAE Microcode Options:

Extended Functions 1 (5804) _____
Extended Functions 2 (5805) _____
Extended Functions 3 (5807) _____
TN3270 Server (5806) _____

Network Node Processor (NNP)

Qty _____ (1 or 2)
NNP Type:
Type 1 (3172) _____
Type 2 (7585) _____
Type 3 (6275) _____
Type 4 (6563) _____
Type 5 (6578) _____

Service Processor Type:

9577 _____
9585 _____
3172 P/N 41H7520 _____
3172 P/N 55H7630 _____
7585 _____
6275 _____
6563 _____

Part 2: 3746-950 Configuration

3746 Model 950 Serial Number _____

Base Enclosure Top View

P	N	M	L	K	J	H	G	F	E	D	C	Rear Side
Port _____	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____	Port _____	Port N/A	Port N/A	
Processor _____		Processor _____		Processor _____		Processor _____		Processor CBSP Type _____		Processor Power Supply		Front Side
Slot 6 (P)		Slot 5 (M)		Slot 4 (K)		Slot 3 (H)		Slot 2 (F)		Slot 1 (D)		

Note: Indicate CBSP Type in Slot 2 (F)

Processors:

TRP = Token-Ring Processor Type 1
 TRP2 = Token-Ring Processor Type 2
 TRP3 = Token-Ring Processor Type 3
 ESCP = ESCON Processor Type 1
 ESCP2 = ESCON Processor Type 2
 ESCP3 = ESCON Processor Type 3
 CLP = Communication Line Processor
 CLP3 = Communication Line Processor Type 3
 SIE = Switch Interface Extension (MAE connection)

Ports:

TIC3 = Token-Ring Coupler Type 3
 ETH = Ethernet Port/Ethernet-TR Bridge
 ESCC = ESCON Coupler Type 1
 ESCC2 = ESCON Coupler Type 2
 LIC11x = Line Interface Coupler Type 11 where 'x' is
 the Line Connection Box (LCB) ID where the
 ARCS are installed that correspond to this LIC
 LIC12 = Line Interface Coupler Type 12

Part 2: 3746-950 Configuration

3746 Model 950 Serial Number _____

Expansion Enclosure 1 Top View

P	N	M	L	K	J	H	G	F	E	D	C	Rear Side
Port	Port	Port	Port	Port	Port	Port	Port	Port	Port	Port	Port	
_____	_____	_____	_____	_____	_____	_____	_____	_____	_____	_____	_____	
Processor		Processor		Processor		Processor		Processor		Processor		
_____		_____		_____		_____		_____		_____		
Slot 12 (P)		Slot 11 (M)		Slot 10 (K)		Slot 9 (H)		Slot 8 (F)		Slot 7 (D)		Front Side

Processors:

TRP = Token-Ring Processor Type 1
 TRP2 = Token-Ring Processor Type 2
 TRP3 = Token-Ring Processor Type 3
 ESCP = ESCON Processor Type 1
 ESCP2 = ESCON Processor Type 2
 ESCP3 = ESCON Processor Type 3
 CLP = Communication Line Processor
 CLP3 = Communication Line Processor Type 3
 SIE = Switch Interface Extension (MAE connection)

Ports:

TIC3 = Token-Ring Coupler Type 3
 ETH = Ethernet Port/Ethernet-TR Bridge
 ESCC = ESCON Coupler Type 1
 ESCC2 = ESCON Coupler Type 2
 LIC11x = Line Interface Coupler Type 11 where 'x' is
 the Line Connection Box (LCB) ID where the
 ARCS are installed that correspond to this LIC
 LIC12 = Line Interface Coupler Type 12

Part 2: 3746-950 Configuration

3746 Model 950 Serial Number

Expansion Enclosure 2 Top View

P	N	M	L	K	J	H	G	F	E	D	C	Rear Side
Port	Port											
_____	_____	_____	_____	_____	_____	_____	_____	_____	_____	_____	_____	
Processor		Processor		Processor		Processor		Processor		Processor		
_____		_____		_____		_____		_____		_____		
Slot 18 (P)		Slot 17 (M)		Slot 16 (K)		Slot 15 (H)		Slot 14 (F)		Slot 13 (D)		Front Side

Processors:

TRP = Token-Ring Processor Type 1
 TRP2 = Token-Ring Processor Type 2
 TRP3 = Token-Ring Processor Type 3
 ESCP = ESCON Processor Type 1
 ESCP2 = ESCON Processor Type 2
 ESCP3 = ESCON Processor Type 3
 CLP = Communication Line Processor
 CLP3 = Communication Line Processor Type 3
 SIE = Switch Interface Extension (MAE connection)

Ports:

TIC3 = Token-Ring Coupler Type 3
 ETH = Ethernet Port/Ethernet-TR Bridge
 ESCC = ESCON Coupler Type 1
 ESCC2 = ESCON Coupler Type 2
 LIC11x = Line Interface Coupler Type 11 where 'x' is
 the Line Connection Box (LCB) ID where the
 ARCS are installed that correspond to this LIC
 LIC12 = Line Interface Coupler Type 12

Part 2: 3746-950 Configuration

3746 Model 950 Serial Number _____

LCB ID _____ Line Connection Box Layout

	0	+1	+2	+3	+4	+5	+6	+7	+8	+9	+10	+11	+12	+13	+14	
ARC Type																To LCBE
Attach																
Length																

Line Connection Box Expansion Layout

	+16	+17	+18	+19	+20	+21	+22	+23	+24	+25	+26	+27	+28	+29	+30	
ARC Type																To LCB
Attach																
Length																

Length:

- | | | |
|---|------------------------------|---|
| ARC Types:
V.35
V.24
X.21 | Attach:
DCE
DTE | .6 = .6 Meters
1.2 = 1.2 Meters
2.4 = 2.4 Meters
5 = 5 Meters
10 = 10 Meters
12 = 12 Meters
15 = 15 Meters
ST = Stub cable to 3745 cable |
|---|------------------------------|---|

Notes:

Fill out one sheet for each LCB/LCBE combo installed

Note: A 3746-950 may have as many as 32 LCB/LCBE pairs. Fill out one of these sheets for each one.

Part 2: 3746-950 Configuration

3746 Model 950 Serial Number _____

Multiaccess Enclosure (MAE)

How is this MAE connected to the 3746-950?

Token-Ring _____

Directly _____

Base Power	PCMCIA Card	Slot 1 LIC _____	Slot 2 LIC _____	Slot 3 LIC _____	Slot 4 LIC _____
Second Power	RS232	Slot 5 LIC _____	Slot 6 LIC _____	Slot 7 LIC _____	Slot 8 LIC _____

LIC Types:

280 = 2-Port Token-Ring

281 = 2 Port Ethernet

282 = 8 Port V.24/EIA-232

283 = 1 Port ISDN PRI-T1/J1

284 = 1 Port Multi-Mode ATM

286 = 1 Port Multi-Mode FDDI

287 = 1 Port ESCON

288 = 1 Port Fast Ethernet

289 = 1 Port HSSI

290 = 6 Port V.35/V.36

291 = 8 Port X.21

293 = 1 Port Single-Mode ATM

294 = 1 Port 155 Multi-Mode ATM

295 = 1 Port 155 Single-Mode ATM

297 = 4 Port ISDN T1/J1

297+ = 4 Port ISDN T1J1 plus 4-P Daughter

299 = 1 Port Parallel Channel



Logical and functional inventory worksheets

The purpose of the logical and functional inventory is to provide a framework for you to review what resources are currently active and what functions your current controller environment is providing. When working through this exercise, it is important to include *only* those resources that you determine are still used. In the past, we have experienced many situations where resources still existed in communication controller configurations but were actually redundant or no longer necessary.

Instructions

This appendix contains configuration sheets that can be filled out to determine the logical and functional characteristics of a communication controller. These configuration sheets may be photocopied for your use.

Fill out a set for each 3745 or 3746-950. You need only fill out those sheets that apply to your particular controller environment. These sheets are structured as follows:

3745 and attached frames

- Part 1: General NCP Information - the basic components
- Part 2: NCP Owned Resource Section - the details on the line, LAN, and channel resources
- Part 3: NNP Owned Resource Section - the APPN or IP resources that are owned by the NNP

Note: If you have a 3745-410/A or 610/A running in Twin-Dual mode, fill out a Logical and Functional Inventory Worksheet for both NCPs running on the machine.

3746-950 frames

- Part 1: General Information - information about the operating system to which the 3746-950 is attached
- Part 2: Resource Information - the details on the APPN and or IP resources on the 3746-950

Fill out all sheets that pertain to your particular controller configuration.

For example, if you have a 3745-170 fill out:

- ▶ Parts 1 and 2 of the 3745 Logical and Functional Worksheet

Or, if you have a 3745-61A running in twin-dual mode and a 3746-900 with a Network Node Processor, fill out:

- ▶ Parts 1, 2, and 3 of the worksheet for the NCP running on CCU-A
- ▶ Parts 1 and 2 of the worksheet for the NCP running on CCU-B

To assist in filling out the Logical and Functional Inventory Worksheet, it will be helpful to review your NCP generation statements. NCP generation statements are structured as follows:

- ▶ Start-stop PEP line groups

- ▶ Start-stop NCP line groups
- ▶ BSC PEP line groups
- ▶ BSC NCP line groups
- ▶ Line groups defined as SDLC, including the following (these resources can be defined in any order):
 - SDLC telecommunication links
 - Network Terminal Option (NTO) resources
 - NetView Performance Monitor (NPM) resources, with definitions for both
 - NPM and Network Session Accounting (NSA)
 - Network Routing Facility (NRF) resources
 - X.25 resources
 - X.25 SNA interconnection (XI) resources
 - X.21 resources
 - Frame-relay resources
 - ISDN resources
- ▶ 3746 Model 900 and NTRI Token-Ring resources
- ▶ Ethernet-type LAN resources
- ▶ 370 I/O and ESCON channel adapter line groups
- ▶ User line groups
- ▶ SNA network interconnection (SNI) non-native resources.

Additionally, tools such as NTuneMON, NetView, and NPM will be helpful in determining whether resources are still active.

For users of the Network Node Processor for APPN and/or IP resources, on the 3746-900 and 3746-950, the Controller Configuration and Management (CCM) tool will also be useful.

Logical and Functional Inventory Worksheet for 3745 (All Models) and 3746-900

Part 1: General NCP Section

Please indicate the number of resources currently needed.

NCP Name _____

3745 Serial Number _____

If 3745 model 410/A or 610/A, indicate how NCP is defined:

Twin-Dual _____

Twin-Standby _____

Twin-Backup _____

Which operating system(s) are you using?

MVS (OS/390) _____

VM _____

VSE _____

TPF _____

Specify any IBM special products you are using. Check all that apply.

EP _____

XI _____

NTO _____

MERVA _____

NRF _____

NSI _____

NPSI _____

Other (Please specify) _____

Specify any user provided products you are using. _____

Access method(s) your NCP communicates with:

VTAM _____

BTAM _____

Other _____

Do you currently utilize transmission groups?

Yes _____

No _____

Part 2: NCP Owned Resource Section

Token-Ring

Downstream PU Count (DSPU)	Logical Unit Count (LUDRPOOL)	TICs In Use
<i>Example: 25,000</i>	<i>50,000</i>	<i>4</i>
Note: Counts in this table are for all Token-Ring resources within this NCP.		

Are you currently using duplicate TIC addresses to balance and backup your Token-Ring SNA traffic?

Yes _____

No _____

Ethernet LAN (NCP Owned Resources)

What type of traffic is on your Ethernet LAN connection?

SNA _____

IP _____

If you are running IP, how many routes are you supporting? _____

Part 2: NCP Owned Resource Section Channels

Channel Type	LPAR Count
Bus and Tag	
ESCON	
<p>Note: If you are currently not using ESCON, is your mainframe ESCON capable?</p> <p>Yes _____</p> <p>No _____</p>	

Part 3: Network Node Processor Resources (3746-900 with NNP Owned Resources)

APPN

APPN: Indicate the number of APPN sessions.

Functions	Count
PU1/PU2/LENs connected?	
ENs/NNs connected?	
SSCP-LU (control) sessions activated by the 3746 DLUR?	
LU-LU sessions (dependent) activated by the 3746 DLUR?	
LU-LU sessions (independent) activated by the 3746 DLUR?	
LU-LU sessions established by other NNs through the 3746 NN?	

Is the 3746 operating as a Branch Extender Node:

Yes _____

No _____

Part 3: Network Node Processor Resources (3746-900 with NNP Owned Resources)

IP

IP: Indicate the number of IP sessions.

Routing Protocol	Number of Routes
OSPF	
BGP	
RIP	

Are you currently using TN3270e server functions on the MAE?

Yes _____

No _____

Logical and Functional Inventory Worksheet for 3746-950

Please indicate the number of resources currently needed.

Part 1: General Information Section

3746 Model 950 Serial Number _____

Which operating system(s) are you using?

MVS (OS/390) _____

VM _____

VSE _____

TPF _____

Part 2: Resource Section

APPN

APPN: Indicate the number of APPN sessions.

Functions	Count
PU1/PU2/LENs connected?	
ENs/NNs connected?	
SSCP-LU (control) sessions activated by the 3746 DLUR?	
LU-LU sessions (dependent) activated by the 3746 DLUR?	
LU-LU sessions (independent) activated by the 3746 DLUR?	
LU-LU sessions established by other NNs through the 3746 NN?	

Is the 3746 operating as a Branch Extender Node:

Yes _____

No _____

Part 2: Resource Section

IP

IP: Indicate the number of IP sessions.

Routing Protocol	Number of Routes
OSPF	
BGP	
RIP	

Are you currently using TN3270e server functions on the MAE?

Yes _____

No _____



Avoiding VTAM network addressing problems

In recent years, a growing number of organizations have been implementing SNA session manager applications and host-based TN3270 servers to provide user access to their SNA applications. At the same time, many organizations are migrating users away from communication controller attachment (such as 3745 token-ring gateways) using OSA or router attachment instead. Both of these changes can significantly increase the number of SNA network element addresses that must be allocated from the VTAM network address pool. As a result, organizations that implement changes such as these without planning for the increased use of network addresses often encounter problems, such as SNA session setup failures or the inability to activate additional resources, which result from depletion of the VTAM element address pool.

This appendix discusses VTAM enhanced addressing, a function that was developed to allow VTAM to address more than the architected element address limit of 65,536 SNA resources (from 0 to 65,535). By using a portion of the subarea number field as a logical extension of the element address, enhanced addressing allows VTAM to address several million SNA resources, thereby supporting a much greater number of SNA devices and sessions. However, because other subarea products (like NCP) do not support or recognize element addresses greater than 65,535 (commonly referred to as “high-order” addresses), there are still many cases where VTAM must continue to use “low-order” element addresses (in the 0 to 65,535 range).

The following topics are discussed in this appendix:

- ▶ “Understanding VTAM network addressing” on page 405 provides an overview of how VTAM assigns network addresses to various types of LU-LU session-capable and non-session-capable resources.
- ▶ “VTAM enhanced addressing” on page 410 provides an overview of VTAM’s enhanced addressing philosophy and describes the enhanced addressing improvements provided in various releases of VTAM.
- ▶ “Maximizing VTAM use of enhanced addressing” on page 422 provides suggestions on how clients can configure their VTAMs to maximize the use of high-order network addresses.

With VTAM enhanced addressing, you can support greater numbers of:

- ▶ TN3270 users
- ▶ Session manager users
- ▶ VTAM-managed SNA devices

If you are curious about how many element addresses are in use in your various VTAMs, you may use the **DISPLAY STATS,TYPE=VTAM** command. Near the end of the statistics displayed by this command are messages that show both the number of low-order addresses (called *element addresses*, meaning in the range 0-65,535) and the number of high-order addresses (called *extended element addresses*, meaning in the range 65,536 to several million) that are in use, as well as the high-water mark (the most element addresses assigned at any one time) for each range.

Important: Most organizations should turn on enhanced addressing by specifying the VTAM Start Option ENHADDR=YES. (See “VTAM enhanced addressing” on page 410 for more information.)



Note: The IBM Communication Controller for Linux on System z9 and zSeries (CCL) product enables you to run NCP on your mainframe, which gives you a way to migrate devices from your IBM 3745 Communication Controller hardware without having to move device addresses from your NCP subarea into a VTAM subarea. For additional information about the CCL product, see Chapter 16, “Communication Controller for Linux on System z9 and zSeries (CCL)” on page 283.

Understanding VTAM network addressing

Network addresses are the architected means for identifying specific resources in an SNA network, referred to as Network Addressable Units or NAUs. Network addresses are assigned to SNA resources in a variety of ways (like pre-definition or dynamic assignment), and are also used in a variety of ways. In general, most network addresses are assigned to SNA resources that participate in LU-LU sessions, such as SNA applications and devices. But other types of devices, like PUs and LINEs, must also be assigned network addresses which are used for SSCP-PU sessions or as a general means of identification.

SNA sessions

Every LU-LU session that is established must be represented by a unique pair of network addresses: one network address that is assigned to the primary LU (PLU) of the session and one that is assigned to the secondary LU (SLU). Because many SNA application programs support “parallel sessions” (two or more concurrently active sessions between the same two session partners), assigning only a single network address to each session-capable SNA resource is not always sufficient to guarantee network address pair uniqueness for all sessions. Therefore, it is often necessary to assign more than one network address to some SNA resources.

Depending on where a resource resides in the network (or attaches to the network), network addresses may be assigned from any subarea node (typically NCP or VTAM). However, it is the SSCP component of VTAM that manages the assignment of network addresses during LU-LU session establishment, even if one or both of the network addresses being assigned comes from an NCP network address pool. To optimize the use (and reuse) of network addresses for LU-LU sessions and to simplify the algorithms used to assign network addresses to SNA resources, VTAM follows these general guidelines:

- ▶ Each session-capable SNA resource that is capable of being the SLU for a session is assigned one network address that is used for every session in which this resource acts as the SLU. For resources that support only one active session at a time (like dependent LUs) and applications that do not support parallel sessions, the SLU network address may also be used for sessions in which the resource acts as the PLU, since network address pair uniqueness is still guaranteed.
- ▶ Each session-capable SNA resource that is capable of acting as the PLU for a session may also be assigned one or more network addresses (as needed) that is only used for sessions in which this resource is acting as the PLU. For SNA resources that establish parallel sessions, a new PLU network address is assigned for each parallel session, since the SLU typically uses the same network address for each parallel session. To conserve network addresses,

these PLU network addresses can be reused for parallel sessions with other session partners.

Network addresses assigned to LU-LU session-capable resources can be classified as pre-assigned network addresses or dynamically assigned network addresses. Pre-assigned network addresses are permanently assigned to an SNA resource when the resource is first activated by VTAM. Dynamically assigned network addresses are assigned to an SNA resource on an as-needed basis, and are unassigned (returned to the network address pool) when they are no longer being used. This distinction is important because early iterations of the VTAM enhanced addressing function only assigned high-order network addresses for dynamically assigned network addresses.

Applications (APPLs)

VTAM application programs are typically pre-assigned two network addresses when they are activated by VTAM: one network address that is always used when this APPL is acting as the SLU of a session, and one that is only used when this APPL is acting as the PLU of a session. The only exceptions to this rule are APPLs defined as not supporting parallel sessions (PARSESS=NO). If parallel sessions are not supported, then the session partner for every concurrent session must be unique, so the session partner network address must also be unique. This special case allows PARSESS=NO APPLs to use a single pre-assigned network address for every session regardless of whether it is acting as the PLU or SLU.

To ensure network address pair uniqueness for APPLs that support parallel sessions, VTAM dynamically assigns a new PLU network address to an APPL for each additional parallel session in which that APPL is acting as the PLU.

Note: Pre-assignment of network addresses to applications when they are activated does not depend on the application program actually starting (opening its ACB). For example, if you activate an application major node with 2,000 applications and only 20 of them are active and running at any given time, you are still using at least 2,000 element addresses.

Dependent PUs and LUs

Network addressing for dependent physical units (PUs) and logical units (LUs) is handled differently depending on how the dependent PU attaches to the network. In traditional subarea networks, dependent PUs must be (or must appear to be) directly attached to a subarea node (NCP or VTAM). This physical adjacency is required in order to be able to route dependent PU and LU flows (like ACTPU,

ACTLU and INITSELF) through a subarea network between the owning SSCP and the dependent PU or LU.

Although APPN networks also support subarea-attached dependent LUs, the APPN architecture provides a function called Dependent LU Requester/Dependent LU Server (DLUR/DLUS) that allows dependent PUs more flexibility in the way they attach to the network. DLUR nodes are APPN nodes that support the attachment of dependent PUs and dependent SLUs to an APPN network. But these dependent PUs and SLUs still require the services of an owning SSCP and the owning SSCP function is provided by an APPN DLUS node.

VTAM is the only product available that supports the DLUS function. The DLUR function is provided by a variety of distributed communications products including:

- ▶ Communications Server for Windows
- ▶ Communications Server for Linux
- ▶ Communications Server for AIX
- ▶ Cisco routers with SNA Switching Services (SNASw)
- ▶ Microsoft Host Integration Server 2004

DLUR nodes and DLUS nodes communicate using a pair of LU 6.2 sessions referred to as “CPSVRMGR sessions” (because they use the CPSVRMGR logmode name). The CPSVRMGR sessions allow the DLUS and DLUR to send dependent PU and LU flows (like ACTPU, ACTLU and INITSELF) through an APPN network between the owning DLUS and the DLUR-served PUs and LUs. This allows dependent PUs to attach to an APPN network from remote locations without the need for an adjacent subarea node.

Subarea-attached dependent PUs and LUs

Subarea-attached dependent PUs and LUs are pre-assigned a network address from the network address pool of the subarea to which they attach when they are activated. This pre-assigned network address is used for the SSCP-PU or SSCP-LU session. The pre-assigned network address for each dependent LU is also logically associated with its higher-level dependent PU.

Because all dependent LUs have a single session limit, the network address that is pre-assigned to a subarea-attached dependent LU is also used for every LU-LU session with that dependent LU; that is, no dynamically assigned network addresses are ever needed. This is possible because the session path of every LU-LU session to a subarea-attached dependent LU must traverse the attaching subarea node through the higher-level dependent PU (and the pre-assigned LU network address is already logically associated with the higher-level PU).

Note: Subarea-attached dependent PUs and LUs are always pre-assigned only low-order addresses. This is true regardless of whether these dependent PUs and LUs attach to an NCP or a VTAM subarea, and regardless of the VTAM release. That is, none of the enhanced addressing functions in any release of VTAM allow subarea-attached dependent PUs and LUs to use high-order addresses.

DLUR-attached dependent PUs and LUs

DLUR-served dependent PUs and LUs are pre-assigned a network address from the network address pool of their owning DLUS (VTAM) when they are activated. This pre-assigned network address is used for the SSCP-PU or SSCP-LU session. As with subarea-attached LUs, the pre-assigned network address for each DLUR-served LU is logically associated with its higher-level dependent PU.

For LU-LU sessions involving DLUR-served LUs, the only APPN node that is required to be on the session path is the DLUR node itself: The owning DLUS node is not required to be on the session path for the actual LU-LU sessions. This means that the DLUS node is not always required to assign network addresses to DLUR-served LUs for every session; rather, network addresses are only assigned by a DLUS node to DLUR-served LUs for a session when the session path traverses an APPN link station that is also owned by the DLUS node.

If the path of an LU-LU session involving a DLUR-served LU happens to pass through the owning DLUS node, the network address that is pre-assigned to DLUR-served LU (for the SSCP-LU session) cannot be used for the LU-LU session. This is because LU-LU sessions that traverse APPN nodes require a network address that is logically associated with the APPN link station (PU) that is used for the session. (This is identical to the way network addresses are assigned to independent LUs for sessions that traverse APPN nodes.) Because the pre-assigned LU network address is always associated with the higher-level dependent PU, it cannot be used for LU-LU sessions that traverse APPN links. Instead, all network addresses assigned to DLUR-served LUs for LU-LU sessions are dynamically assigned as needed.

Note: Just as with LEN and APPN independent LUs, DLUR-served LUs can establish LU-LU sessions over different APPN link stations (PUs). Because every network address that is dynamically assigned to a DLUR-served LU for a session is associated with a specific APPN link station (PU), a DLUR-served LU will have a different dynamically assigned SLU address for each APPN link station that it uses for an LU-LU session. However, because DLUR-served LUs can have only one active session at a time, the number of network addresses that are dynamically assigned to a DLUR-served LU at any given time is usually very small. DLUR-served LUs do not require any PLU network addresses, because the DLUR/DLUS architecture does not support dependent PLUs.

Independent LUs (CDRSCs)

When VTAM first provided support for independent LUs, VTAM modeled them after dependent LUs. That is, independent LUs were defined as LOCADDR=0 LUs under a PU. Just as for dependent LUs, VTAM is responsible for assigning network addresses to independent LUs when they participate in LU-LU sessions over link stations that are owned by VTAM. However, because SSCP-LU sessions are never established with independent LUs, VTAM never pre-assigns network addresses to independent LUs (therefore, the LOCADDR=0 operand on the LU definition); instead, all network addresses assigned to independent LUs are dynamically assigned.

Another significant difference between dependent and independent LUs is the fact that independent LUs do not have a session limit (they can have multiple sessions active at the same time) and can be either the PLU or the SLU of a session. This means that dynamically assigning only a single network address to an independent LU may not be sufficient to guarantee network address pair uniqueness for every LU-LU session. Instead, just as with applications, it may be necessary to dynamically assign one SLU address and one or more PLU addresses to each independent LU. For independent LUs, all network addresses are dynamically assigned as they are needed and they are returned to the pool of available network addresses when they are no longer being used for any sessions.

In later releases, VTAM expanded the independent LU support to provide the “multi-tail” function, which allows independent LUs to access a subarea network over multiple different link stations at the same time. With this support, VTAM discontinued representing independent LUs as “LUs” and, instead, represented them as cross-domain resources (CDRSCs). If you continue to define independent LUs as LOCADDR=0 LUs, VTAM converts them to CDRSCs when they are activated. These converted CDRSCs are placed in the ISTPDILU major node, which stands for “Pre-Defined Independent LUs.” This CDRSC-based

independent LU model also worked well for APPN independent LUs when VTAM provided APPN support in VTAM V4R1, because LEN and APPN independent LUs are virtually identical with respect to the types of network attachment and LU-LU sessions they support.

Note: In a typical subarea network, CDRSCs usually represent SNA resources that are owned by another subarea-attached (cross-domain) VTAM. In this case, any network addresses assigned to these “subarea CDRSCs” for LU-LU sessions are assigned by the owning VTAM (or an NCP in that VTAM’s domain). The enhanced network addressing information presented here pertains only to CDRSCs that represent independent LUs; that is, CDRSCs that use a LEN or APPN link station (PU) owned by the local VTAM to establish LU-LU sessions. Subarea CDRSCs are never candidates for enhanced addressing because cross-domain sessions always involve at least two subareas, which means low-order network addresses are always required.

Note: Because every network address assigned to an independent LU is associated with a specific LEN or APPN link station, independent LUs represented as CDRSCs follow the same guidelines with respect to network address assignment as when they were represented as LUs, but on a link station basis. That is, each LEN or APPN independent LU may be assigned one SLU network address and one or more PLU network addresses for each link station over which it establishes sessions.

VTAM enhanced addressing

The majority of a subarea node’s network addresses are assigned to session-capable resources (APPLs, LUs, and CDRSCs) for the purpose of establishing LU-LU sessions. As such, the original VTAM enhanced addressing support (in VTAM V4R2) focused on assigning high-order network addresses to these types of resources. The enhanced addressing improvements provided by subsequent releases of VTAM allow high-order addresses to be used more often for session-capable resources but they also allow high-order addresses to be assigned to some types of SNA resources that do not participate in LU-LU, such as LINEs and PUs.

Important: Most organizations should turn on enhanced addressing by specifying the VTAM Start Option ENHADDR=YES.

In general, high-order network addresses are assigned to eligible session-capable resources (APPLs, LUs, and CDRSCs) only when ENHADDR=YES is specified at VTAM startup. The default value for ENHADDR (NO) was chosen because the network addresses used for LU-LU sessions are provided to some user exit routines, and it was felt that there was some risk that these exit routines may not work correctly when high-order addresses are provided. However, in CS for OS/390 V2R7, VTAM was changed to no longer provide the “high-order” portion of the element address to user exits. Because there is no longer any risk to user exits, all clients using at least CS for OS/390 V2R7 are strongly encouraged to enable enhanced addressing for session-capable resources by specifying ENHADDR=YES when VTAM is started. (Note that the ENHADDR start option value cannot be changed using the MODIFY VTAMOPTS command.)

Note: High-order network addresses are assigned to eligible non-session-capable resources (like LINEs and PUs) regardless of the ENHADDR start option value. This is due to the fact that network addresses used for these types of resources are not provided to any user exit routines, so there was no risk involved with assigning high-order addresses for these types of resources even when ENHADDR=NO is specified or used by default.

To avoid compatibility issues with other subarea-capable products (like NCP), the general philosophy behind enhanced addressing for LU-LU sessions is as follows:

- ▶ If, for any given session, both the PLU and the SLU network addresses are assigned from the same VTAM subarea, then it is possible to use high-order network addresses for both the PLU and the SLU for that LU-LU session. But this does not always mean that high-order network addresses will be used (as described in subsequent sections).
- ▶ If, for any given session, the PLU and the SLU network addresses are assigned from different subareas (meaning that a virtual route is being used between subarea nodes), then low-order addresses must be used because it is impossible to tell if all of the subarea nodes along the virtual route support enhanced addressing.

Figure C-1 shows a sample network with a mix of APPN nodes and subarea nodes.

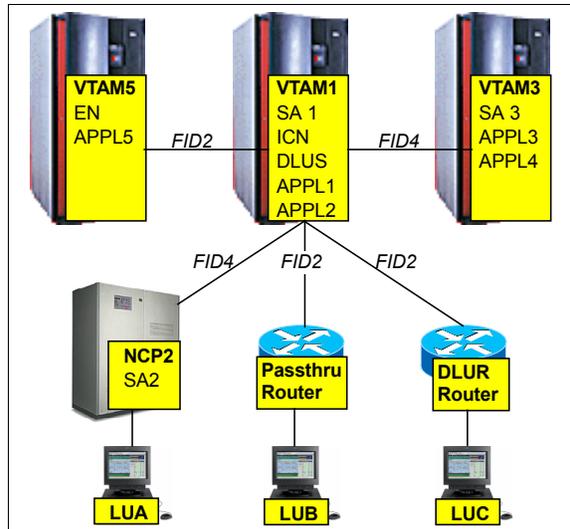


Figure C-1 Enhanced addressing in a network with a mix of APPN and subarea connections

Using Figure C-1 as a reference, the table in Figure C-2 demonstrates which nodes are allowed to use high-order network addresses for various sessions (which depends on the session path that is used). Columns 1 and 2 identify various primary PLU and SLU session partner combinations (one per row). Columns 3 through 6 indicate which nodes are responsible for assigning network addresses for one or both of the session partners, and whether each resource is eligible to use a high-order network address (HI) or is required to use a low-order network address (LO) for this session.

Primary LU	Secondary LU	VTAM1		NCP2		VTAM3		VTAM5	
		PLU	SLU	PLU	SLU	PLU	SLU	PLU	SLU
APPL1	APPL2	HI	HI						
APPL1	APPL3	LO				LO			
APPL1	APPL5	HI	HI					HI	HI
APPL3	APPL4					HI	HI		
APPL1	LUA	LO			LO				
APPL1	LUB	HI	LO						
APPL1	LUC	HI	HI						
APPL3	LUA				LO	LO			
APPL3	LUB		LO			LO			
APPL3	LUC		LO			LO			
APPL5	LUA	LO			LO			HI	HI
APPL5	LUB	HI	LO					HI	HI
APPL5	LUC	HI	HI					HI	HI

Figure C-2 Enhanced addressing example address table

Notice that low-order addresses must be used whenever the PLU and SLU network addresses are assigned from different subareas. High-order addresses are only allowed when the same VTAM subarea assigns both the PLU and SLU network addresses. Also note that, in cases where a VTAM assigns two network addresses for a session even though only one session partner resides in that VTAM (like for all of the sessions with APPL5), the other session partner is represented as a CDRSC and is treated like an independent LU.

Using this table, it is clear to see that pure APPN nodes (nodes with no subarea links at all, like VTAM5) stand to benefit the most from enhanced addressing because all “virtual routes” begin and end in the same VTAM subarea. However, even nodes that are part APPN and part subarea (like VTAM1) or pure subarea (like VTAM3) can still benefit from enhanced addressing to some extent.

Note: Just because high-order addresses are allowed to be used for a specific session does not mean that high-order addresses will be used for that session. The decision to use a high-order address or a low-order address depends on several factors, including which release of VTAM you are using and the fact that VTAM tries to reuse existing network addresses that are already assigned to a given session partner whenever possible.

Figure C-3 on page 414 provides a summary of VTAM enhanced addressing changes by release.

Type of resource	VTAM V4R2	CS for OS/390 V2R5	CS for OS/390 V2R7	Z/OS CS V1R4	Z/OS CS V1R6
Applications (APPLs)	D*		D* P*		
Independent LUs (CDRSCs)	D*		D*		
Subarea-attached LUs					
DLUR-served LUs	D*	P	D*		
DLUR-served PUs					P
EE logical LINEs and PUs				P	
RTP PUs					P

Key:

D: Enhancement affects dynamically-assigned addresses for this type of resource

P: Enhancement affects pre-assigned addresses for this type of resource

*****: Enhancement requires ENHADDR=YES to be specified at VTAM startup

Figure C-3 Summary of VTAM enhanced addressing changes by release



Restriction: VTAM for VM/ESA and VSE/ESA enhanced addressing support includes the VTAM V4R2-level capabilities described in this chapter; however, it does not include any of the enhancements included in subsequent VTAM releases.

The following sections provide more details about the enhanced addressing functions provided in various releases of VTAM or Communications Server for OS/390 (CS for OS/390) or z/OS Communications Server (z/OS CS). If the enhanced addressing improvements provided by a specific release affect more than one type of resource, then subsections further describe the benefits provided by these improvements for each type of resource.

VTAM V4R2

Because the majority of a subarea node's network addresses are assigned to session-capable resources (APPLs, LUs, and CDRSCs) for the purpose of establishing LU-LU sessions, the original VTAM enhanced addressing support provided in VTAM V4R2 focused on assigning high-order network addresses to these types of resources. Furthermore, the VTAM V4R2 enhanced addressing

function only attempts to use high-order network addresses for dynamically assigned network addresses; pre-assigned network addresses for APPLs and LUs are always low-order network addresses.

As mentioned earlier, high-order network addresses can only be used for sessions in which the network addresses of both session partners (PLU and SLU) are assigned by the same VTAM subarea. For sessions where this is true, the enhanced addressing support provided by VTAM V4R2 attempts to use high-order network addresses for the PLU, SLU, or both resources. However, this does not mean that high-order network addresses are always used whenever they are allowed. This is true for a couple of reasons:

- ▶ If an existing low-order network address assigned to either session partner can be used for this session, the existing low-order network address is used for the session rather than unnecessarily assigning another network address to that resource. This is true regardless of whether the existing low-order network address was pre-assigned or dynamically assigned.
- ▶ There are many cases where a resource that is eligible to use high-order network addresses is assigned its network address before the session partner's network address (or subarea number) is known. When this occurs, VTAM V4R2 takes the conservative approach of assigning a low-order network address to the resource for this session. This ensures that this network address can be used for this session regardless of which subarea eventually assigns the session partner's network address.

The following sections describe how the VTAM V4R2 enhanced addressing function is exploited by various types of session-capable resources.

Applications

The introduction of enhanced addressing in VTAM V4R2 affected the way dynamically assigned network addresses are allocated to applications that support parallel sessions. When additional PLU network addresses are dynamically assigned to an application for parallel sessions, VTAM dynamically assigns high-order network addresses when allowed (that is, when the network address of the session partner is assigned from the same VTAM subarea). If a high-order network address is not allowed for a particular session, then VTAM dynamically assigns a new low-order address as before. But for application programs that establish many parallel sessions to same-domain or APPN resources, this change can save a significant number of low-order addresses.

Note: High-order network addresses are dynamically assigned to APPLs only when ENHADDR=YES is specified at VTAM startup.

DLUR-attached dependent LUs

VTAM DLUS support was first made available with VTAM V4R2, the same release that introduced the enhanced addressing function. But because the VTAM V4R2-enhanced addressing function only attempts to use high-order network addresses for dynamically assigned network addresses, the network address that is pre-assigned to a DLUR-served LU and used for the SSCP-LU session is still a low-order address (just as with subarea-attached dependent LUs).

As described earlier, however, the network address that is pre-assigned to DLUR-served LUs cannot be used for LU-LU sessions; rather, network addresses assigned to DLUR-served LUs for LU-LU sessions are always dynamically assigned. As a result, DLUR-served LUs are eligible to use high-order network addresses for sessions (when the session partner's network address is assigned from the same VTAM subarea).

This means that DLUR-served LUs are always assigned at least as many low-order network addresses as subarea-attached LUs. And in some cases, DLUR-served LUs may use more low-order network addresses than subarea-attached LUs (if any of their session partners require a low-order address).

Note: High-order network addresses are dynamically assigned to DLUR-served LUs for use in LU-LU sessions only when ENHADDR=YES is specified at VTAM startup.

Independent LUs (CDRSCs)

Because all network addresses assigned to independent LUs (CDRSCs) are dynamically assigned, independent LUs are eligible to use high-order network addresses for every session (if the session partner's network address is assigned from the same VTAM subarea).

Note: High-order network addresses are dynamically assigned to CDRSCs only when ENHADDR=YES is specified at VTAM startup.

CS for OS/390 V2R5 DLUR-attached dependent LUs

When the first large clients began migrating thousands of users (dependent LUs) from NCP subarea attachment to APPN DLUR/DLUS attachment, it quickly became apparent that requiring VTAM to pre-assign a low-order network address for each DLUR-served LU (for the SSCP-LU session) would rapidly deplete the pool of available low-order addresses. To improve the scalability of the DLUR/DLUS function in large APPN networks, the VTAM enhanced addressing

function was extended in CS for OS/390 V2R5 to support pre-assigning high-order network addresses to DLUR-served LUs. (This support was also made available on prior releases of VTAM via APAR. APAR OW31455 provided this function for VTAM V4R4 and VTAM V4R4.1; APAR OW32705 provided this function for VTAM V4R3.) This CS for OS/390 V2R5 enhancement, combined with the VTAM V4R2 enhanced addressing function for dynamically assigned network addresses (which DLUR-served LUs use for LU-LU sessions), allows VTAM to avoid assigning any low-order network addresses to DLUR-served LUs unless a low-order address is required for a specific session.

Keep in mind that the high-order network addresses that are pre-assigned to DLUR-served LUs for SSCP-LU sessions are long term in nature. Dynamically assigned network addresses (which are used for LU-LU sessions, and which are still required to be low-order when the session traverses a subarea link) tend to be relatively short term in nature, and are returned to the pool of available network addresses when they are no longer being used for any sessions.

Note: Although high-order network addresses are dynamically assigned to DLUR-served LUs only when ENHADDR=YES is specified at VTAM startup, high-order network addresses are always pre-assigned to DLUR-served LUs for the SSCP-LU session regardless of whether ENHADDR=YES is specified at VTAM startup.

Note: The CS for OS/390 V2R5 enhanced addressing function still pre-assigns low-order network addresses to DLUR-served PUs. Pre-assigning high-order network addresses for DLUR-served PUs was implemented in z/OS CS V1R6.

CS for OS/390 V2R7

The enhanced addressing improvements provided by CS for OS/390 V2R7 were designed to extend the VTAM V4R2 enhanced addressing function to more fully exploit high-order network addresses. Prior to these enhancements, VTAM only attempted to use high-order network addresses for dynamically assigned network addresses (but not for pre-assigned network addresses), and only when the subarea number of the session partner was known to be the same VTAM subarea.

One of the enhanced addressing improvements provided by CS for OS/390 V2R7 allows VTAM to pre-assign high-order network addresses to some or all applications depending on how VTAM is configured. See “Applications” on page 418 for more details.

The second enhanced addressing improvement provided by CS for OS/390 V2R7 changes the way VTAM assigns the first network address for an LU-LU session. In cases where the first network address to be assigned is for a resource that is eligible to use high-order network addresses but the subarea number of the session partner is not yet known, VTAM now assigns a high-order network address to the resource for the LU-LU session. (Contrast this with VTAM V4R2 enhanced addressing where VTAM takes the conservative approach of assigning a low-order network address to the resource for this session.)

This second improvement may cause a small amount of additional overhead for some sessions. If a high-order network address is used for one resource and the network address of the session partner turns out to be assigned from another subarea, then this high-order address must be unassigned and a low-order address assigned instead. But it is generally felt that conserving low-order network addresses for when they are really needed is worth a small amount of overhead during session establishment.

Applications

The enhanced addressing improvements provided by CS for OS/390 V2R7 include several changes that affect addressing for VTAM applications. One of the enhanced addressing improvements provided by CS for OS/390 V2R7 affects the way VTAM pre-assigns network addresses for APPLs when VTAM is configured as a “pure NN” or “pure EN” (which means that they do not support any subarea connections at all, not even to NCPs). For pure NN and pure EN VTAMs, the network addresses assigned to both session partners for every session are always assigned by the same VTAM, so every LU-LU session is eligible to use high-order network addresses. Consequently, VTAM now always pre-assigns high-order network addresses to applications if VTAM is configured as a pure NN or pure EN.

For VTAMs that are not configured as pure NNs or pure ENs, CS for OS/390 V2R7 can still improve the exploitation of high-order network addresses for applications that do not support parallel sessions (PARSESS=NO). As described earlier, PARSESS=NO applications are pre-assigned a single network address that is used for every session. CS for OS/390 V2R7 causes VTAM to pre-assign a high-order network address for PARSESS=NO applications. But this pre-assigned high-order network address cannot be used for sessions that traverse a subarea link (where the session partner’s network address is assigned from a different subarea). So if and when a PARSESS=NO application establishes a session that traverses a subarea link, VTAM dynamically assigns a low-order network address for that session (and reuses this address for any subsequent sessions that also require a low-order network address). As with other dynamically assigned network addresses, when the low-order network address is no longer being used for any sessions, it is returned to the network address pool.

Finally, as described above, VTAM now dynamically assigns high-order network addresses to eligible applications even when the network address (or subarea number) of the session partner is not yet known. If it is eventually determined that the network address of the session partner will be assigned from a different subarea, then the high-order network address is unassociated with this session (and returned to the pool of available network addresses, if it is no longer in use) and a low-order network address is used instead.

Note: Because PARSESS=NO is the default value for TN3270 APPLs, starting with CS for OS/390 V2R7 TN3270 APPLs are typically pre-assigned only one (high-order) address. If a low-order address is required for a specific session, then it is dynamically assigned as needed and freed when it is no longer being used.

Note: High-order network addresses are assigned to APPLs only when ENHADDR=YES is specified at VTAM startup.

DLUR-attached dependent LUs

As described earlier, network addresses that are pre-assigned to DLUR-served dependent LUs are only used for the SSCP-LU session. For LU-LU sessions, additional network addresses are dynamically assigned to DLUR-served LUs (as needed) in much the same way as for independent LUs. As a result, the enhanced addressing improvements provided by CS for OS/390 V2R7 cause VTAM to dynamically assign high-order network addresses to eligible DLUR-served dependent LUs for LU-LU sessions even when the network address (or subarea number) of the session partner is not yet known. If it is eventually determined that the network address of the session partner will be assigned from a different subarea, then the high-order network address is unassociated with this session (and returned to the pool of available network addresses, if it is no longer in use) and a low-order network address is used instead.

Note: High-order network addresses are dynamically assigned to DLUR-served LUs for use in LU-LU sessions only when ENHADDR=YES is specified at VTAM startup.

Independent LUs (CDRSCs)

The enhanced addressing improvements provided by CS for OS/390 V2R7 cause VTAM to dynamically assign high-order network addresses to eligible CDRSCs that represent independent LUs even when the network address (or subarea number) of the session partner is not yet known. If it is eventually

determined that the network address of the session partner will be assigned from a different subarea, then the high-order network address is unassociated with this session (and returned to the pool of available network addresses, if it is no longer in use) and a low-order network address is used instead.

Note: High-order network addresses are dynamically assigned to CDRSCs only when ENHADDR=YES is specified at VTAM startup.

z/OS CS V1R4 EE logical LINES and PUs

One important reason for migrating from communication controllers is to simplify the wide area network infrastructure by converging on a single network protocol, namely the internet protocol (IP). As a result, Enterprise Extender (EE), the strategic SNA/IP integration technology from IBM, has emerged as a popular option. However, moving wide area network LINE attachment from communication controller hardware to mainframes (using EE) resulted in the same scalability problems as migrating users from communication controller attachment to OSA or router attachment. Specifically, the network addresses assigned to LINES and PUs, which previously were allocated from an NCP network address pool, are now allocated from VTAM's network address pool. To prevent this increase in VTAM-attached LINES and PUs from depleting VTAM's pool of low-order network addresses, enhanced addressing was expanded in z/OS CS V1R4 to allow VTAM to pre-assign high-order network addresses to all EE logical LINES and PUs.

Note: High-order network addresses are always pre-assigned to EE logical LINES and PUs regardless of whether ENHADDR=YES is specified at VTAM startup.

z/OS CS V1R6

The enhanced addressing functions provided in VTAM V4R2 and CS for OS/390 V2R7 allow VTAM to use high-order network addresses for most eligible LU-LU sessions. And the enhanced addressing functions provided by z/OS CS V1R4 allow VTAM to assign high-order network addresses for some common non-session-capable resources, namely EE logical LINES and PUs. But EE logical LINES and PUs are not the only non-session-capable resources that occur in significant numbers in large APPN networks. The enhanced addressing improvements provided by z/OS CS V1R6 allow VTAM to pre-assign high-order network addresses for other types of non-session-capable resources, targeting DLUR-served PUs and RTP PUs which occur in great numbers in today's large APPN networks.

DLUR-attached dependent PUs

When enhanced addressing was implemented for DLUR-served LUs in CS for OS/390 V2R5, the corresponding DLUR-served PUs were still assigned low-order network addresses. At the time, this was not considered a problem because there are typically many more DLUR-served LUs than PUs. It was felt that the additional low-order addresses that would be saved by developing logic to assign high-order network addresses for DLUR-served PUs would not justify the cost of that development effort. Since then, several organizations implementing extensive Enterprise Extender networks ran into scalability problems because their network design included defining 20,000 or more DLUR-served PUs. To prevent this growth in the number of DLUR-served PUs from depleting VTAM's pool of low-order network addresses, one of the VTAM enhanced addressing improvements provided by z/OS CS V1R6 allows VTAM to pre-assign high-order network addresses to all DLUR-served PUs.

Note: High-order network addresses are always pre-assigned to DLUR-served PUs regardless of whether ENHADDR=YES is specified at VTAM startup or not.

RTP PUs

Enterprise Extender (EE) technology requires the use of APPN and HPR. As HPR networks continue to grow in size, the number of HPR connections or "RTP pipes" that are established also increases substantially. Although a single RTP pipe can carry many sessions, each RTP pipe describes a specific path through the network to a specific partner RTP node for a specific SNA Class of Service. If a new session is established to a different partner RTP node, or over a different path to the same partner RTP node, or uses a different Class of Service, a new RTP pipe is created for that session (if an existing RTP pipe that meets all of the same criteria does not already exist). This commonly results in multiple RTP pipes being created between any given pair of partner RTP nodes. In a typical banking network with thousands (or even tens of thousands) of branch locations, the number of RTP pipes that are created as part of normal operation can be quite surprising.

Because RTP pipes are represented in VTAM as switched PUs (commonly referred to as RTP PUs), each RTP PU is assigned a network address from VTAM's address pool. In large HPR networks, scalability problems with respect to network addresses are almost a certainty if each RTP PU is assigned a low-order address. One of the VTAM-enhanced addressing improvements provided by z/OS CS V1R6 allows VTAM to pre-assign high-order network addresses to all RTP PUs as they are created.

Note: High-order network addresses are always pre-assigned to RTP PUs regardless of whether ENHADDR=YES is specified at VTAM startup.

Maximizing VTAM use of enhanced addressing

The first step toward ensuring that network address depletion does not occur in your network is to being aware of how many network element addresses VTAM is currently using. To determine this, you may use the **DISPLAY STATS,TYPE=VTAM** command. Near the end of the statistics displayed by this command are messages that show both the number of low-order addresses (called *element addresses*, meaning in the range 0-65,535) and the number of high-order addresses (called *extended element addresses*, meaning in the range 65,536 to several million) that are in use, as well as the high-water mark (the most element addresses assigned at any one time) for each range.

If VTAM is close to reaching the low-order element address limit (even if this only occurs during peak hours of operation), then it is time to take steps to ensure that low-order element address depletion does not occur. The following sections describe strategies that you can use to conserve low-order network addresses or maximize VTAM's use of high-order network addresses. Use this list to decide which of these strategies are most likely to produce the best results for your network.

Enable VTAM enhanced addressing for sessions

As stated earlier, high-order network addresses are generally assigned to eligible session-capable resources (APPLs, LUs, and CDRSCs) only when ENHADDR=YES is specified at VTAM startup. The default value for ENHADDR (NO) was chosen because the network addresses used for LU-LU sessions are provided to some user exit routines, and it was thought that there was some risk that these exit routines may not work correctly when high-order addresses are provided. However, in CS for OS/390 V2R7, VTAM was changed to no longer provide the "high-order" portion of the element address to user exits. So there is no longer any risk that user exits will stop working correctly when high-order addresses are used for sessions. Because of this, it is strongly recommended that all clients running at least CS for OS/390 V2R7 enable enhanced addressing for session-capable resources by specifying ENHADDR=YES when VTAM is started.

Note: The ENHADDR start option value can only be set when VTAM is started; it cannot be changed dynamically using the MODIFY VTAMOPTS command.

Use VTAM application cloning

Pre-assignment of network addresses to applications occurs when they are activated and does not depend on the application program actually starting (opening its ACB). If you predefine a large number of application programs that are not active at the same time (like TSO and TN3270 applications), it can consume a significant number of low-order network addresses. To avoid this, many organizations use a function introduced in VTAM V4R3 called “model applications” or “application cloning.”

With application cloning, the application minor node is dynamically created by VTAM from a model definition at the time the application opens its ACB. This process ensures that network addresses and other resources are allocated to application clones only when they are actually being used. Additionally, the addresses are also freed when the close ACB occurs putting the addresses back in the available pool for reallocation to another resource. Savings can be substantial for TSO, TN3270, and session manager applications because network addresses are allocated when a new user logs on (or when a new TN3270 user connects to a TN3270 server) and freed up when the user logs off.

Note: Most of the enhancements to VTAM addressing discussed in this appendix decrease the number of low-order addresses used by assigning more high-order addresses. Application cloning, however, is yet more efficient because it actually reduces the total number of network addresses (both high-order and low-order) as well as other VTAM resources that are used for applications by allocating them only when an application actually starts (opens its ACB). Therefore, although the enhanced addressing functions provided by CS for OS/390 V2R7 enable VTAM to pre-assign high-order network addresses to some applications (depending on how VTAM is configured), there are still significant advantages to using application cloning even in these environments.

Migrate your network from subarea to APPN

As shown in the example in Figure C-1 on page 412, APPN nodes stand to benefit the most from enhanced addressing. If your network is still mostly or completely subarea-based, then low-order network addresses will likely be required for most LU-LU sessions. (The only time high-order addresses are used by pure subarea nodes is for some same-domain LU-LU sessions.)

In most cases, migrating an existing subarea network to APPN significantly improves VTAM's ability to use high-order network addresses and this is especially true with more recent releases of VTAM. For VTAM releases prior to z/OS CS V1R6, however, there is a trade-off between the low-order network

addresses that are saved by migrating to APPN and the additional low-order addresses that are used for HPR connections (RTP PUs). Since there are typically many more LU-LU sessions than RTP PUs, this tradeoff generally results in fewer low-order network addresses being used after migrating from subarea to APPN.

Note: Subarea-to-APPN migration can be very complicated, particularly for large, complex SNA networks. Careful planning is required to avoid running into problems during the migration. Useful insights into planning such a migration can be found in:

- ▶ *Subarea to APPN Migration: VTAM and APPN Implementation*, SG24-4656
- ▶ *Subarea to APPN Migration: HPR and DLUR Implementation*, SG24-5204
- ▶ *Inside APPN - The Essential Guide to the Next-Generation SNA*, SG24-3669

These books are only available online at:

<http://www.redbooks.ibm.com/>

Migrate dependent LUs to DLUR-attachment

In the past, DLUR-served dependent PUs and LUs consumed at least as many and in some cases even more low-order network addresses as subarea-attached dependent PUs and LUs. However, CS for OS/390 V2R5 provided a substantial improvement because it allows DLUR-served LUs to be pre-assigned high-order network addresses instead. In any case, some low-order network addresses are typically still required for LU-LU sessions with DLUR-served LUs (depending on where the session partner is located). It is also important to remember that low-order network addresses are still pre-assigned to all DLUR-served PUs.

With CS for OS/390 V2R7, DLUR-served LUs typically use even fewer low-order addresses for LU-LU sessions, because high-order network addresses can more often be dynamically assigned for LU-LU sessions. The use of low-order addresses for DLUR-served LUs can be completely eliminated if the DLUS VTAM (CS for OS/390 V2R7 or above) is configured as a pure NN. The use of low-order addresses for DLUR-served PUs can be completely eliminated if VTAM is at least z/OS CS V1R6.

The single greatest inhibitor to migrating subarea-attached dependent PUs and LUs to DLUR-served PUs and LUs is the fact that at least some portion of the network must be migrated to APPN. Additionally, migrating end users from subarea-attached dependent LUs to DLUR-served LUs usually requires new hardware and software (for the DLUR nodes).

Migrate your APPN network to EE

If your network is already using APPN but is not yet using Enterprise Extender (EE), then migrating part or all of your APPN network to EE can conserve even more low-order network addresses if VTAM is at least z/OS CS V1R4. This is due to the fact that EE logical LINES and PUs use high-order network addresses, while other types of LINES and PUs use low-order network addresses.

Glossary

ACF. A group of IBM licensed programs, principally VTAM, TCAM, NCP, and SSP, that use the concepts of Systems Network Architecture (SNA), including distribution of function and resource sharing.

ACF/NCP. See *NCP*.

ACF/VTAM See *VTAM*.

ACP/TPF. See *TPF*.

Advanced Communications Function. See *ACF*.

Advanced Peer-to-Peer Networking. See *APPN*.

Advanced Program-to-Program Communication. See *APPC*.

API. A software interface that enables applications to communicate with each other. An API is the set of programming language constructs or statements that can be coded in an application program to obtain the specific functions and services provided by an underlying operating system or service program.

APPC. An advanced SNA API whose base function is to provide application programs with commands that coordinate and control the sending and receiving of data among networked devices. Some systems implement optional APPC functions, such as commands controlling database commits and synchpointing or “waking up” a partner program across a network to enable

communication. LU 6.2 is an implementation of APPC, and although not exactly technically correct, the terms APPC and LU 6.2 are nearly always used synonymously.

APPN. An extension to SNA featuring:

- ▶ More distributed network control that avoids critical hierarchical dependencies, thereby isolating the effects of single points of failure
- ▶ Dynamic exchange of network topology information to foster ease of connection, reconfiguration, and adaptive route selection
- ▶ Dynamic definition of network resources
- ▶ Automated resource registration and directory lookup

APPN extends the LU 6.2 peer orientation for user services to network control and can support multiple LU types, including LU 2 (see DLUR), LU 3, and LU 6.2.

asynchronous. Pertaining to two or more processes that do not depend upon the occurrence of specific events such as common timing signals. Without regular time relationship.

asynchronous communication. A method of communication supported by the operating system that allows an exchange of data with a remote device, using either a start-stop line or an X.25 line. Some common names for asynchronous transmission devices include: async Ascii, TTY, TWX, WTTY, IBM 3151, IBM 3767, and various Digital Equipment Corporation devices such as DEC VT100.

ATM. Asynchronous Transfer Mode is a networking technology in which information to be transmitted is divided into 53 byte cells. It is asynchronous in the sense that the recurrence of cells containing information from an individual user is not necessarily periodic. This means that ATM is architected to handle both voice and data. ATM is specified in international standards such as ATM Forum UNI 3.1.

B-channel. In the Integrated Services Digital Network (ISDN), a 64 Kbps channel for the transport of speech or data between the ISDN service provider and user.

backbone. A set of nodes and their interconnecting links that form a central, high-speed network interconnecting other, typically lower-speed networks or client nodes. In a local area network multiple-bridge ring configuration, a high-speed link to which the rings are connected by means of bridges or routers. A backbone may be configured as a bus or as a ring.

BAN. Boundary Access Node is a type of frame relay connectivity, typically used by routers, in which the routers use frame relay RFC 1490 bridged frame format to forward LAN-based SNA traffic to communication controllers. When the communication controller receives the frame, it treats it as if it were received on a token-ring interface. BAN enables remote routers to communicate with a mainframe directly through a communication controller, without using a router at the mainframe site.

bandwidth. The measure of the speed at which devices are using a communications facility. For serial links (such as leased or dial lines) and LANs, the speed is governed by modems, CSUs/DSUs, or NIC cards and is measured in bits per second (bps). In

asynchronous transfer mode (ATM) networks, bandwidth can be expressed as the capacity of a virtual channel, expressed in terms of peak cell rate (PCR), sustainable cell rate (SCR), and maximum burst size (MBS). In frame relay networks, bandwidth may refer to the network access link speed in bps, or the portions of that access link speed available to the user as committed information rate or burst rate.

binary synchronous communication. See *BSC*.

BNN. Boundary Network Node refers to a peripheral SNA device such as an IBM 3174. In a subarea SNA network, a BNN link is a connection between a peripheral device and an NCP or a VTAM. In an APPN-type SNA network, BNN-like function is implemented in the End Node (EN). Like BNNs, ENs are devices which obtain network-level directory and routing services from an INN-type node called a Network Node (NN). (See also INN.)

boundary function. In SNA, a capability of a subarea node to provide protocol support for attached peripheral nodes, such as:

- ▶ Interconnecting subarea path control and peripheral path control elements
- ▶ Performing session sequence numbering for low-function peripheral nodes
- ▶ Providing session-level pacing support

boundary node. In SNA, a subarea node with boundary function. A subarea node may be a boundary node, an intermediate routing node, both, or neither, depending on how it is used in the network.

branch extender. See *BX*.

bridge. A functional unit that interconnects two local area networks that use the same logical link control protocol but may use different media access control protocols. A bridge forwards a frame to another bridge or an end station based on the media access control (MAC) address.

BSC. One of the first data-link control technologies. In BSC, synchronization of characters is controlled by timing signals sent at the beginning of message blocks rather than within each character. BSC was also one of the first DLCs to support multi-point connectivity. Typical implementations of BSC are IBM 3270 (terminal and print interactive traffic) and IBM 2780 or 3780 (remote job entry batch traffic).

BX. Branch extender simplifies APPN network topology and supports the implementation of very large APPN networks. A branch extender network node acts as an EN to the upstream network while it presents the appearance of a NN to downstream nodes. Therefore, it is not part of the APPN topology and will not receive topology database updates (TDUs). If it registers its LUs to its serving NN (typically a VTAM), it will not receive any APPN directory search broadcasts further reducing APPN network broadcast traffic.

CCL. See Communication Controller for Linux on System z9 and zSeries.

CEPT. A type of data link used in Europe similar to a T1 line.

channel. A path along which signals can be sent, for example, data channel, output channel. In data communication, a means of one-way transmission. A functional unit, controlled by the processor, that handles the transfer of data between processor storage and local peripheral equipment.

channel-attached. Pertaining to the attachment of devices directly by input/output channels to a host processor. Pertaining to devices attached to a controlling unit by cables, rather than by telecommunication lines. IBM 3745 base frames channel connect to mainframes using copper parallel channels. IBM 3746-900 frames channel connect to mainframes using fiber optic ESCON channels.

CIP. 1. Classical IP, the implementation of IP over ATM. The procedures for encapsulating IP datagrams and ATM ARP traffic over ATM defined in RFC 1577. 2. Channel Interface Processor, the name for channel adapters for Cisco routers.

class of service. See *COS*.

client. A computer system or process that requests a service of another computer system or process that is typically referred to as a server. Multiple clients may share access to a common server.

client/server. In communications, the model of interaction in distributed data processing in which a program at one site sends a request to a program at another site and awaits a response. The requesting program is called a client; the answering program is called a server.

CLP. In a 3746-900 communication controller, the processor that manages SNA, X.25, or IP serial telecommunication lines. Supported interfaces are RS232-C, V.35, and X.21.

cluster. A station that consists of a control unit (a cluster controller) and the terminals attached to it. A group of APPN nodes that have the same network ID and the same topology database. A cluster is a subset of a NETID subnetwork. In high-availability cluster

multiprocessing (HACMP™), a set of independent systems (called nodes) that are organized into a network for the purpose of sharing resources and communicating with each other.

cluster controller. A device that can control the input/output operations of more than one device connected to it. A cluster controller may be controlled by a program stored and executed in the unit, for example the IBM 3601 Finance Communication Controller. Or, it may be entirely controlled by hardware, for example the IBM 3272 Control Unit.

CNN. A PU Type 5 node (VTAM) and its subordinate PU Type 4 nodes (NCPs) that support APPN network node protocols and appear to an attached APPN or LEN node as a single network node.

communication controller. A networking device, such as an IBM 3705, 3720, 3725, 3745, or 3746-900, that serves two important networking needs:

- ▶ As gateways for connecting mainframe host servers to networks offloading network functions such as line control and device support, thereby saving host CPU cycles
- ▶ As line concentration devices that could be placed in remote locations to save on communications line charges by consolidating the traffic from many remote devices into a few, usually higher speed, communication lines

Communication Controller for Linux on System z9 and zSeries. A program product that allows you to run other program products including NCP, NPSI, and NRF, originally designed to be run on IBM 3745 Communication Controller hardware, on Linux for zSeries.

connection. In data communication, an association established between functional units for conveying information. In Open Systems Interconnection architecture, an association established by a given layer between two or more entities of the next higher layer for the purpose of data transfer. In VTAM, synonym for physical connection. In SNA, the network path that links together two logical units (LUs) in different nodes to enable them to establish communications. In X.25 communication, a virtual circuit between two (pieces of) data terminal equipment (DTEs). A switched virtual circuit (SVC) connection lasts for the duration of a call; a permanent virtual circuit (PVC) is a permanent connection between the DTEs. In TCP/IP, the path between two protocol applications that provides reliable data stream delivery service. In the Internet, a connection extends from a TCP application on one system to a TCP application on another system.

COS. A set of requested connection characteristics (such as route security and transmission priority) used to select a route between SNA session partners and control queuing within the network. The “*class of service*” is derived from a mode name specified by the initiator of a session.

CPU. The part of a computer that includes the circuits that control the interpretation and execution of instructions. A CPU is the circuitry and storage that executes instructions.

cross-system coupling facility. A component of the MVS operating system that provides functions to support cooperation between authorized programs running within a sysplex.

data circuit-terminating equipment. See *DCE*.

data link control. See *DLC*.

data stream. All information (data and control commands) sent over a data link usually in a single read or write operation. A continuous stream of data elements being transmitted, or intended for transmission, in character or binary-digit form, using a defined format.

DATE. Dedicated access to the X.25 transport extension for IBM NPSI allows the CTCP to manage the virtual circuits to SNA and non-SNA X.25 DTEs.

DCE. In a data station, the equipment that provides the signal conversion and coding between the data terminal equipment (DTE) and the line. The DCE may be separate equipment or an integral part of the DTE or of the intermediate equipment. A DCE may perform other functions that are usually performed at the network end of the line.

dependent LU requester. See *DLUR*.

DHCP. A communications protocol defined by the Internet Engineering Task Force (IETF) that is used for dynamically assigning IP addresses to computers in a network. Using the Internet Protocol, each machine that can connect to the Internet needs a unique IP address. When an organization sets up its computer users with a connection to the Internet, an IP address must be assigned to each machine. Without DHCP, the IP address must be entered manually at each computer and, if computers move to another location in another part of the network, a new IP address must be entered. DHCP enables a network administrator to supervise and distribute IP addresses from a central point and automatically sends a new IP address when a computer is plugged into a different place in the network.

DiffServ. The Internet (both public ISPs and large intranets) needs scalable solutions for quality of service (QoS), so that different applications (interaction, browsing, e-commerce, video, telephony, etc.) can each get adequate service quality from common infrastructure. For smaller networks this will be achieved by RSVP and associated protocols. For the very large scale a simpler solution with coarser granularity is needed, very similar to the SNA Class of Service (COS) concept. This solution has been developed by the Internet industry and the IETF under the name Differentiated Services (DiffServ). In DiffServ, each IP packet contains a marker byte that selects a given grade of service at each router in the network. Packets are marked either by the originating computer or by a traffic classifier at the network boundary.

DLC. A set of rules used by nodes on a data link (such as an SDLC link or a token ring) to accomplish an orderly exchange of information.

DLSw. A commonly used method of transporting network protocols that use IEEE 802.2 logical link control (LLC) type 2 over an IP network. SNA and NetBIOS are examples of protocols that use LLC type 2.

DLUR. An APPN end node or network node that:

- ▶ Supports dependent LUs in its local node or in adjacently attached nodes, and
- ▶ Obtains SSCP services for these dependent LUs from a dependent LU server (DLUS) located elsewhere in an APPN network

The flows of SSCP services between DLUR and DLUS are carried over a special pair of LU 6.2 sessions. DLUR enables APPN-type SNA networks to carry IBM 3270-type traffic by providing an APPN connection for VTAM-to-3270 control sessions which are otherwise not supported in an APPN network.

DNS. In the Internet suite of protocols, the distributed database system used to map domain names to IP addresses.

The domain name system (DNS) is the way that Internet domain names are located and translated into Internet Protocol addresses. A domain name is a meaningful and easy-to-remember “handle” for an Internet address.

DTE. That part of a data station that serves as a data source, data sink, or both.

e-commerce. The subset of e-business that involves the exchange of money for goods or services purchased over an electronic medium such as the Internet.

EBN. Allows the connection of network nodes with different net IDs, and session establishment between LUs in different net 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 net ID subnetworks. An extended border node can also partition a single net ID subnetwork into two or more clusters or topology subnetworks with the same net ID, thus isolating one from the topology of the other. The purpose of EBN is to provide network isolation for APPN networks that is analogous to the network isolation that SNI provides for subarea SNA networks.

EDI. A standard format for exchanging business data. The standard is ANSI X12 and it was developed by the Data Interchange Standards Association. ANSI X12 is either closely coordinated with or is being merged with an international standard, EDIFACT.

An EDI message contains a string of data elements, each of which represents a singular fact, such as a price, product model number, and so forth, separated by delimiter. The entire string is called a data segment. One or more data segments framed by a header and trailer form a transaction set, which is the EDI unit of transmission (equivalent to a message). A transaction set often consists of what would usually be contained in a typical business document or form. The parties who exchange EDI transmissions are referred to as trading partners.

EDI messages can be encrypted. EDI is one form of e-commerce, which also includes e-mail and fax.

EE. Enterprise Extender provides a data link control (DLC) that supports High-Performance Routing (HPR) connections over Internet Protocol (IP) networks. To the HPR network, the IP backbone appears to be a logical link. SNA traffic is carried over the IP network in User Datagram Protocol (UDP) datagrams.

element address. In SNA, a value in the element address field of the network address identifying a specific resource within a subarea. On INN links, both the subarea and element portions of the address are present in the packet header. Since BNN links do not traverse subareas, packets on BNN links only contain an element address. NCPs and VTAMs are fully capable of routing BNN traffic. Even so, popular jargon often refers to SNA traffic as “unroutable” exactly because BNN packets do not contain subarea addresses.

emulation. The use of a data processing system to imitate another data processing system such that the imitating system accepts the same data, executes the same programs, and achieves the same results as the imitated system.

Emulation Program. An IBM control program that allows a channel-attached IBM communication controller to emulate the functions of an IBM 2701 Data Adapter Unit, an IBM 2702 Transmission Control, or an IBM 2703 Transmission Control.

enterprise extender. See *EE*.

Enterprise Systems Connection. See *ESCON*.

Extended Border Node. See *EBN*.

ESCON. A set of IBM and vendor products and services that dynamically interconnect S/390 computers with each other and with attached storage, locally attached workstations, and other devices using optical fibre technology and dynamically modifiable switches called ESCON Directors. In IBM mainframes, the local interconnection of hardware units is known as channel connection (and sometimes as local connection to distinguish it from remote or telecommunication connection). ESCON's Fibre Optic cabling can extend this local-to-the-mainframe network up to 60 kilometers (37.3 miles) with chained Directors. The data rate on the link itself is up to 200 Mbps (million bits per second) and somewhat less when adapted to the channel interface. Vendor enhancements may provide additional distance and higher amounts of throughput.

ESCON Director. A switch for ESCON optical channels that provides connectivity capability and control for attaching any two links to each other.

Ethernet. A 10 Mbps per second baseband local area network that allows multiple stations to access the transmission medium at will without prior coordination, avoids contention by using carrier sense and deference, and resolves contention by using collision detection and delayed retransmission. Ethernet uses carrier sense multiple access with collision detection (CSMA/CD). See also, Fast Ethernet and Gigabit Ethernet.

Fast Ethernet. A local area network (LAN) transmission standard that provides a data rate of 100 megabits per second (referred to as "100BASE-T"). Workstations with existing 10 megabit per second (10BASE-T) Ethernet card can be connected to a Fast Ethernet network. (The 100 megabits per second is a shared data rate; input to each workstation is constrained by the 10 Mbps card.)

FDDI. A standard for data transmission on fiber optic lines in a local area network (LAN) that can extend in range up to 200 Km (124 miles). The FDDI protocol is based on the token-ring protocol. In addition to being large geographically, an FDDI local area network can support thousands of users.

FRAD. Sometimes referred to as a "*frame relay access device*," this is a box that provides the interface between the user and a network that uses frame relay. The FRAD is usually included in a router.

frame relay. An interface standard describing the boundary between a user's equipment and a fast-packet network. In frame-relay systems, flawed frames are discarded; recovery is provided on an end-to-end basis rather than hop-by-hop. A technique derived from the Integrated

Services Digital Network (ISDN) D channel standard. It assumes that connections are reliable and dispenses with the overhead of error detection and control within the network.

frame relay access device. See *FRAD*.

frame-relay frame. The frame-relay frame structure defined by American National Standards Institute (ANSI) Standard T1.618.

frame-relay frame handler. See *FRFH*.

FRFH. The function in a frame-relay node that routes (or switches) frames along a permanent virtual circuit (PVC). A frame handler receives frames from an adjacent frame-relay node and uses the DLCI to forward them to the next node on the PVC. In NCP, the function that switches frames between communication controllers on a PVC segment. The NCP frame handler function can also switch frames to frame-relay terminating equipment.

GATE. NPSI's General Access to X.25 Transport Extension allows a host user application program, called the communication and transmission control program (CTCP), to monitor virtual circuits to non-SNA X.25 DTEs.

gateway. A functional unit that interconnects two computer networks with different network architectures. A gateway connects networks or systems of different architectures. A bridge interconnects networks or systems with the same or similar architectures. A functional unit that connects two networks or subnetworks having different characteristics, such as different protocols or different policies concerning security or transmission priority.

In SNI, the combination of machines and programs that provide address translation, name translation, and system services control point (SSCP) rerouting between independent

SNA networks to allow those networks to communicate. A gateway consists of one gateway NCP and at least one gateway SSCP. In an IBM Token-Ring Network, a device and its associated software that connect a local area network to a host.

gateway NCP. In SNI, an NCP that performs address translation to allow cross-network session traffic. The gateway NCP connects two or more independent SNA networks.

gateway SSCP. In SNI, an SSCP that is capable of cross-network session initiation, termination, take down, and session outage notification. A gateway SSCP is in session with the gateway NCP; it provides network name translation and assists the gateway NCP in setting up alias network addresses for cross-network sessions.

Gigabit Ethernet. A transmission technology based on the Ethernet frame format and protocol used in local area networks (LANs), providing a data rate of 1 billion bits per second (one gigabit). Gigabit Ethernet is defined in the IEEE 802.3 standard and is currently being used as the backbone in many enterprise networks.

Gigabit Ethernet is carried primarily on optical fiber (with very short distances possible on copper media). Existing Ethernet LANs with 10 and 100 Mbps cards can feed into a Gigabit Ethernet backbone.

HiperSockets. zSeries HiperSockets provides a "TCP/IP network in the system" that allows high-speed any-to-any connectivity among select Operating System (OS) images within a zSeries server, without any physical cabling or need for an external networking connection.

hop. In APPN, a portion of a route that has no intermediate nodes. It consists of only a single transmission group connecting adjacent nodes. To the routing layer, the logical distance between two nodes in a network.

host. A computer that is connected to a network (such as the Internet or an SNA network) and provides an access point to that network. Also, depending on the environment, the host may provide centralized control of the network. The host can be a client, a server, or both a client and a server simultaneously.

HPR. An addition to APPN that enhances data-routing performance and session reliability.

IEEE 802.2. An IEEE standard describing how data is formatted into frames for LAN transmission.

INN. Intermediate Network Node. In a subarea SNA network, an INN is a device that provides connectivity services (such as routing) to a peripheral node. NCP and VTAM devices are INNs. INN links connect NCPs and/or VTAMs to each other. Subarea SNA links are either BNN or INN.

Institute of Electrical and Electronics Engineers. A professional society accredited by the American National Standards Institute (ANSI) to issue standards for the electronics industry.

Internet Protocol. See *IP*.

intranet. A private network that integrates Internet standards and applications (such as Web browsers) with an organization's existing computer networking infrastructure.

IP. In the Internet suite of protocols, a connectionless protocol that routes data through a network or interconnected networks and acts as an intermediary between the higher protocol layers and the physical network.

ISDN. A digital end-to-end telecommunication network that supports multiple services including voice and data.

ISDN-BRI. The ISDN Basic Rate Interface defined as "2B+D" - 2 "bearer" channels of 64 Kbps each plus one 16 Kbps "digital signalling" channel.

ISDN-PRI. The ISDN Primary Rate Interface defined, in the US, as "23B+D" - 23 "bearer" channels of 64 Kbps each plus one 64 Kbps "digital signalling" channel. The Euro-ISDN standard specifies the primary rate interface as 30B+D.

ITU-T. The part of the International Telecommunication Union (ITU) that is responsible for developing recommendations for telecommunications.

JES. An IBM licensed program that receives jobs into the system and processes all output data produced by the jobs. The jobs processed through JES are often referred to as batch processing, and often consist of file transfers, print jobs, and other movements of data.

JES is a subsystem of the z/OS and MVS mainframe operating systems that manages jobs (units of work) that the system does. Each job is described to the operating system by system administrators or other users in job control language (JCL). The operating system then sends the job to the JES program. The JES program receives the job, performs the job based on priority, and then purges the job from the system.

There are two versions, JES2 and JES3. JES3 allows central control of the processing of jobs using a common work queue. Both z/OS and MVS provide an interactive menu for initiating and managing jobs.

LAN. A computer network located on a user's premises within a limited geographical area. Communication within a local area network is not subject to external regulations; however, communication across the LAN boundary may be subject to some form of regulation. A network in which a set of devices are connected to one another for communication and that can be connected to a larger network.

LAN Channel Station. See LCS.

LCS. The term commonly used to refer to the channel protocol introduced with the IBM 8232 LAN Channel Station product. LCS is the channel protocol supported by TCP/IP application in mainframe hosts. Each application defines a consecutive pair of subchannels, one for TCP/IP to read from the channel, and one for TCP/IP to write to the channel. The LCS interface allows LAN MAC frame to be transported over the channel, and provides a command interface to activate, deactivate, and query the LAN interfaces.

LEN. Low Entry Networking: An SNA capability of nodes to attach directly to one another, point-to-point, using basic peer-to-peer protocols to support multiple and parallel sessions between logical units.

link-attached. Pertaining to devices that are connected to a controlling unit by a data link.

link level. A part of Recommendation X.25 that defines the link protocol used to get data into and out of the network across the full-duplex link connecting the subscriber's machine to the network node. LAP and LAPB are the link access protocols recommended by the CCITT.

Link Services Architecture. See LSA.

Linux. A Unix-like, open source, operating system.

LLC. The data link control (DLC) LAN sublayer that provides two types of DLC operation for the orderly exchange of information. The first type is connectionless service (referred to as LLC1), which allows information to be sent and received without establishing a link. The LLC sublayer does not perform error recovery or flow control for connectionless service. The second type is connection-oriented service (referred to as LLC2), which requires establishing a link prior to the exchange of information. Connection-oriented service provides sequenced information transfer, flow control, and error recovery.

load balancing. Dividing the amount of work that a system has to do between two or more computers so that more work gets done in the same amount of time and, in general, all users get served faster. Load balancing can be implemented with hardware, software, or a combination of both. Typically, load balancing is the main reason for computer server clustering.

On the Internet, companies whose Web sites get a great deal of traffic usually use load balancing.

local area network. See *LAN*.

logical link. A pair of link stations, one in each of two adjacent nodes, and their underlying link connection, providing a single link-layer connection between the two nodes. Multiple logical links can be distinguished while they share the use of the same physical media connecting two nodes. Examples are 802.2 logical links used on local area network (LAN) facilities and LAP E logical links on the same point-to-point physical link between two nodes. The term logical link also includes the multiple X.25 logical channels that share the use of the access link from a DTE to an X.25 network.

logical link control. See *LLC*.

logical partition. See *LPAR*.

logical unit. See *LU*.

LPAR. The division of a mainframe computer's processors, memory, and storage into multiple sets of resources so that each set of resources can be operated independently with its own operating system instance and applications. The number of logical partitions that can be created depends on the system's processor model and resources available. Typically, partitions are used for different purposes such as database operation or client/server operation or to separate test and production environments. Each partition can communicate with the other partitions as if the other partition is in a separate machine.

LSA. The channel protocol used for SNA device attachment to VTAM hosts. LSA uses the same External Communications Adapter (XCA) support as the IBM 3172 with the Interconnect Control Program (ICP).

LU. A type of network accessible unit that enables users to gain access to network resources and communicate with each other.

LU 6.2. See "*advanced program-to-program communication*" (APPC).

MAC. In LANs, the sublayer of the data link control layer that supports medium-dependent functions and uses the services of the physical layer to provide services to the logical link control (LLC) sublayer. The MAC sublayer includes the method of determining when a device has access to the transmission medium.

MAE. The 3746 Multiaccess Enclosure is a 3746 hardware extension that expands the connectivity of the 3746 by offering eight adapter slots. These PCI-based adapters offer a broad range of high-speed LAN (ATM LAN emulation client and classical IP), increase traditional LAN interfaces (token-ring and Ethernet), and offer additional E1/T1 and low-speed lines. The MAE also brings ESCON and parallel channel adapters, ATM, FDDI, HSSI, Fast Ethernet, and MPC. It also provides specific functionality such as DLSw support, TN3270E Server, Network Dispatcher, Enterprise Extender (HPR over IP), and MPC+.

mainframe. A computer, usually in a computer center, with extensive capabilities and resources to which other computers may be connected so that they can share facilities.

medium access control. See *MAC*.

message queuing. See *MQ*.

MERVA ESA. An IBM licensed program designed to help financial institutions in the preparation and processing of financial messages.

MIPS. A measure of computer processing performance that is equal to one million instructions per second.

MLTG. In SNA, a multi-line transmission group is a group of links that are viewed as a single logical link. A “mixed-media multi-link transmission group (MMMLTG)” is one that contains links of different medium types (for example, token-ring, switched SDLC, non-switched SDLC, and frame-relay links).

MQ. A method of program-to-program communication. Programs communicate by writing and retrieving application-specific data (messages) to and from queues, without having a private, dedicated, logical connection to link them. In addition, the MQ asynchronous message delivery mechanisms support multiple levels of delivery assurance and recovery.

MQSeries. The former name of a WebSphere MQ, a family of IBM-licensed programs that provide message queuing services.

multileaving. A variation of BSC communication that allows several devices to communicate concurrently over a link without using station addresses.

multiplex. To interleave or simultaneously transmit two or more messages on a single channel.

MVS. The operating system from IBM that is installed on most of its mainframe computers. MVS has been said to be the operating system that keeps the world going. The payroll, accounts receivable, transaction processing, database management, and other programs critical to the world’s largest businesses are usually run on an MVS system. Although MVS tends to be associated with a monolithic, centrally controlled information system, IBM has in recent years repositioned it as a “large server” in a network-oriented distributed environment, using a 3-tier application model.

The latest version of MVS, z/OS, no longer bears the “MVS” in its name. Since MVS represents a certain epoch and culture in the history of computing and since many older MVS systems still operate, the term “MVS” will probably continue to be used for some time. Since z/OS also comes with UNIX user and programming interfaces built in, it can be used as both an MVS system and a UNIX system at the same time.

NCP. An IBM licensed program that provides communication controller support for single-domain, multiple-domain, and interconnected network capability.

NCPMON. A program that runs under the NetView program in a host attached to one or more NCPs. NCPMON monitors the resources of NCPs activated by VTAM in the host in which NCPMON is running.

NetView. See *Tivoli NetView for z/OS*.

Network Control Program. See *NCP*.

Network Dispatcher. The Interactive Network Dispatcher is load-balancing software that evenly distributes load among mirrored servers by routing TCP/IP session requests to different servers within a group of servers. It increases the performance of servers, and it balances the requests among all the servers by operating transparently to users and other applications. It is useful for applications such as e-mail servers, World Wide Web servers, distributed parallel database queries, and other TCP/IP applications.

network node. See *NN*.

Network Routing Facility. See *NRF*.

Net390. A strategy for enabling host systems for optimal support of both SNA and IP applications.

NJE. Network Job Entry, communication between host job entry subsystems (such as POWER™-to-JES-to-RSCS). Note: communications can be SNA or non-SNA

NN. In APPN networks, an NN (network node) is a node in the network that provides directory services (locates partner LUs) and routing services to other devices.

NNP. In order to provide the APPN network node function in the 3746-9X0, the network node processor (NNP) provides the necessary hardware resources and the licensed internal code.

node. In a network, a point at which one or more functional units connect channels or data circuits. In network topology, the point at an end of a branch. The representation of a state or an event by means of a point on a diagram. In a tree structure, a point at which subordinate items of data originate. Any device, attached to a network, that transmits and receives data. An endpoint of a link or a junction common to two or more links in a network. Nodes can be processors, communication controllers, cluster controllers, or terminals. Nodes can vary in routing and other functional capabilities. In VTAM, a point in a network defined by a symbolic name.

NPSI. An IBM licensed program that allows SNA users to communicate over packet switching data networks that have interfaces complying with CCITT Recommendation X.25. It allows SNA programs to communicate with SNA or non-SNA equipment over such networks.

NRF. An IBM licensed program that resides in NCP. NRF provides a path for routing messages between terminals and routes messages over this path without going through the host processor.

NTO. An IBM licensed program used in conjunction with NCP that allows certain non-SNA devices to participate in sessions with SNA application programs in the host processor. When data is sent from a non-SNA device to the host processor, NTO converts non-SNA protocol to SNA protocol; and when data is sent from the host processor to the non-SNA device, NTO converts SNA protocol to non-SNA protocol.

NTuneMON. A program that runs with Tivoli NetView for z/OS and monitors NCPs that were activated by the VTAM on the host where NTuneMON is running.

NTuneNCP. A feature of NTuneMON that runs in a communication controller and, with VTAM, enables a network administrator to tune NCP interactively.

Open Systems Adapter. See *OSA*.

OSA. Open Systems Adapter (OSA) is the term applied to the family of LAN adapters supported on the mainframe. An OSA feature provides industry-standard direct connectivity to local area networks (LANs).

OSA-2. The OSA-2 features are offered on the S/390 servers (Ethernet, Token Ring, FDDI, and ATM) and supported in the compatibility I/O cage of the z900 (Token Ring, FDDI).

OSA-Express. As technology has improved new generations of OSA features have been introduced with faster microprocessors and PCI buses. The IBM Open Systems Adapter-Express features are a newer generation of LAN adapters. They may or may not be unique to a server family. OSA-Express features are offered on G5, G6, z900, z800,

z990, and z890. Depending upon the server, they include Fast Ethernet, 1000BASE-T Ethernet, Gigabit Ethernet, ATM, and Token Ring.

The OSA-Express features plug into an I/O slot just like a channel card, providing a direct, peer-to-peer network connection.

OSA-Express2. OSA-Express2 is the newest generation of OSA features. OSA-Express2 10 Gigabit Ethernet, and Gigabit Ethernet are offered on z990 and z890. OSA-Express2 10 Gigabit Ethernet, Gigabit Ethernet, and 1000BASE-T Ethernet are offered on z9-109.

OSN. OSA-Express2 OSN (OSA for NCP) is designed to appear to operating systems as an ESCON channel connected to an IBM 3745 Communication Controller device type which exploits existing CDLC protocols. OSN is a new CHPID type, exclusive to z9-109, and is supported by the OSA-Express2 Gigabit Ethernet and 1000BASE-T Ethernet features. A feature is configured on a port-by-port basis to be used for direct, *internal*, connectivity between z9-109 LPARs (running Communications Server for z/OS or TPF) and NCPs running under CCL.

z/OS. Pertaining to the IBM operating system that includes and integrates functions previously provided by many IBM software products (including the MVS operating system).

packet. In data communication, a sequence of binary digits, including data and control signals, that is transmitted and switched as a composite whole. The data, control signals, and, possibly, error control information are arranged in a specific format.

packet switching. The process of routing and transferring data by means of addressed packets so that a channel is occupied only during transmission of a packet. On completion of the transmission, the channel is made available for transfer of other packets.

PAD. A functional unit (packet assembler/disassembler) that enables data terminal equipment (DTEs) not equipped for packet switching to access a packet switched network. For example, an X.3 PAD enables start-stop terminals to access X.25 networks by packetizing the individual characters transmitted by the start-stop device, and by de-packetizing X.25 and converting back to individual async-ascii characters for receipt by the start-stop terminal.

parallel channel. A channel having a System/360™ and System/370 channel-to-control-unit I/O interface that uses bus-and-tag cables as a transmission medium.

Parallel Sysplex. An enhanced z/OS version of *sysplex* (systems complex), originally introduced by IBM in 1990 as a platform for MVS/ESA™ mainframe servers. A sysplex consists of the multiple computers or *systems* that make up a *complex*, and is designed to be a solution for business needs involving any or all of the following:

- ▶ Parallel processing
- ▶ Online transaction processing (OLTP)
- ▶ Very high transaction volumes
- ▶ Very numerous small work units—for example, online transactions or large work units that can be broken up into multiple small work units
- ▶ Applications running simultaneously on separate systems that must be able to update to a single database without compromising data integrity

PCNE. An NPSI function that simulates a logical unit (LU) for the non-SNA device to the host, so the host LU believes that it is communicating with an SNA LU type 1, rather than with a non-SNA X.25 DTE.

For data sent from the host to the X.25 DTE, the PCNE replaces the SNA headers with packet headers. The data is then sent over the network to the X.25 DTE using X.25 protocols.

For data sent from the X.25 DTE to the host, the PCNE replaces the packet headers with SNA headers. The SNA data is then sent to the host.

You can also establish a PCNE-to-PCNE connection for communication between SNA application programs in two subarea nodes as an alternative to using cross-domain SNA networking facilities.

Like NTO, PCNE provides an IBM 3767 image to the SNA application. The IBM 3767 is a line-by-line mode device.

peripheral device. Any device that can communicate with a particular computer, for example input/output units, auxiliary storage.

peripheral node. A node that uses local addresses for routing and therefore is not affected by changes in network addresses. A peripheral node requires boundary-function assistance from an adjacent subarea node. A peripheral node can be a type 1, 2.0, or 2.1 node connected to a subarea boundary node.

permanent virtual circuit. See *PVC*.

persistent. Pertaining to data that is maintained across session boundaries, usually in nonvolatile storage such as a database system or a directory

point-to-point. Pertaining to data transmission between two locations without the use of any intermediate display station or computer.

Point-to-Point Protocol. See *PPP*.

PPP. A protocol that provides a method for encapsulating and transmitting packets over serial point-to-point links.

private network. A network established and operated by an organization or corporation for users within that organization or corporation.

protocol. The meanings of, and the sequencing rules for, requests and responses used for managing a network, transferring data, and synchronizing the states of network components.

PTF. A temporary solution or bypass of a problem diagnosed by IBM in a current unaltered release of the program.

PU 4. SNA architecture physical unit type 4. Also called subarea nodes. In the SNA architecture, PU 4s interact with host nodes (PU 5) and peripheral nodes (PU 2 and PU 2.1) to provide the SNA functions that route and control the flow of data in a subarea network. An IBM 3745 Communication Controller running NCP code is an example of a PU4.

public network. A network established and operated by a telecommunication administration or by a Recognized Private Operating Agency (RPOA) for the specific purpose of providing circuit-switched, packet-switched, and leased-circuit services to the public.

PVC. In X.25 and frame-relay communications, a virtual circuit that has a logical channel permanently assigned to it at each data terminal equipment (DTE). Call-establishment protocols are not required. Contrast with switched virtual circuit. The logical connection between two frame-relay terminating equipment stations, either directly or through one or more frame-relay frame handlers. A PVC consists of one or more PVC segments.

QLLC. An X.25 protocol that allows the transfer of data link control information between two adjoining SNA nodes that are connected through an X.25 packet-switching data network. The QLLC provides the qualifier “Q” bit in X.25 data packets to identify packets that carry logical link protocol information.

redundant. A term describing computer or network system components, such as fans, hard disk drives, servers, operating systems, switches, and telecommunication links that are installed to back up primary resources in case they fail.

router. On the Internet, a router is a device or, in some cases, software in a computer, that determines the next network point to which a packet should be forwarded toward its destination. The router is connected to at least two networks and decides which way to send each information packet based on its current understanding of the state of the networks it is connected to. A router is located at any gateway (where one network meets another), including each Internet point-of-presence. A router is often included as part of a network switch.

A router may create or maintain a table of the available routes and their conditions and use this information along with distance and cost algorithms to determine the best route for a given packet. Typically, a packet may travel through a number of network points with routers before arriving at its destination.

routing. Routing is a function associated with the Network layer (layer 3) in the Open Systems Interconnection (OSI) model (see router above). A layer-3 switch is a switch that can perform routing functions.

RS-232C. A long-established standard (“C” is the current version) that describes the physical interface and protocol for relatively low-speed serial data communication between computers and related devices. It was defined by an industry trade group, the Electronic Industries Association (EIA), originally for teletypewriter devices.

RS-232C is the interface that your computer uses to talk to and exchange data with your modem and other serial devices. Somewhere in your PC, typically on a Universal Asynchronous Receiver/Transmitter chip on your motherboard, the data from your computer is transmitted to an internal or external modem (or other serial device) from its Data Terminal Equipment (Data Terminal Equipment) interface. Since data in your computer flows along parallel circuits and serial devices can handle only one bit at a time, the UART chip converts the groups of bits in parallel to a serial stream of bits. As your PC’s DTE agent, it also communicates with the modem or other serial device, which, in accordance with the RS-232C standard, has a complementary interface called the Data Communications Equipment (DCE) interface.

SDLC. A discipline conforming to subsets of the Advanced Data Communication Control Procedures (ADCCP) of the American National Standards Institute (ANSI) and High-level Data Link Control (HDLC) of the International Organization for Standardization, 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 non-switched links. The configuration of the link connection may be point-to-point, multipoint, or loop.

server. A computer that provides services to one or more clients over a network. Examples include a file server, a print server, and a mail server.

SNA. The IBM architecture that defines logical structure, formats, protocols, and operational sequences for transmitting information units through, and controlling the configuration and operation of, networks. The layered structure of SNA allows the ultimate origins and destinations of information (the users) to be independent of and unaffected by the specific SNA network services and facilities that are used for information exchange.

SNA network. The part of a user-application network that conforms to the formats and protocols of Systems Network Architecture. It enables reliable transfer of data among users and provides protocols for controlling the resources of various network configurations. The SNA network consists of network accessible units (NAUs), boundary function, gateway function, and intermediate session routing function components; and the transport network.

SNA network interconnection. See *SNI*.

SNI. The connection, by gateways, of two or more independent SNA networks to allow communication between logical units in those networks. The individual SNA networks retain their independence.

source route bridging. In LANs, a bridging method that uses the routing information field in the IEEE 802.5 medium access control (MAC) header of a frame to determine which rings or token-ring segments the frame must transit. The routing information field is inserted into the MAC header by the source node. The information in the routing information field is derived from explorer packets generated by the source host.

SSCP. A component within a subarea network for managing the configuration, coordinating network operator and problem determination requests, and providing directory services and other session services for users of the network. Multiple SSCPs, cooperating as peers with one another, can divide the network into domains of control, with each SSCP having a hierarchical control relationship to the physical units and logical units within its own domain. Usually, VTAM (CS/390) provides SSCP services for the network.

SSP. An IBM licensed program, made up of a collection of utilities and small programs, that supports the operation of the NCP as well as certain trace/dump formatting services.

start-stop transmission. Asynchronous transmission such that each group of signals representing a character is preceded by a start signal and is followed by a stop signal. Asynchronous transmission in which a group of bits is (a) preceded by a start bit that prepares the receiving mechanism for the reception and registration of a character, and

(b) followed by at least one stop bit that enables the receiving mechanism to come to an idle condition pending reception of the next character.

subarea. A portion of the SNA network consisting of a subarea node, attached peripheral nodes, and associated resources. Within a subarea node, all network accessible units (NAUs), links, and adjacent link stations (in attached peripheral or subarea nodes) that are addressable within the subarea share a common subarea address and have distinct element addresses.

switched line. A telecommunication line in which the connection is established by dialing.

S.W.I.F.T. Refers to the S.W.I.F.T. network, an industry-owned co-operative supplying secure messaging services and interface software to over 7,000 financial institutions in 193 countries.

Synchronous Data Link Control. See *SDLC*.

system services control point. See *SSCP*.

Systems Network Architecture. See *SNA*.

TCAM. An access method used to transfer data between main storage and remote or local terminals.

TCP. A communications protocol used in the Internet and in any network that follows the Internet Engineering Task Force (IETF) standards for internetwork protocol. TCP provides a protocol for reliable exchanges of information between hosts in packet-switched communications networks and in interconnected systems of such networks. It uses the Internet Protocol (IP) as the underlying transport protocol.

TCP/IP. The Transmission Control Protocol and the Internet Protocol, which together provide reliable end-to-end connections between applications over interconnected networks of different types. The term is often used to refer to the full suite of transport and application protocols that run over the Internet Protocol.

Tivoli NetView for z/OS. A Tivoli product that enables centralized systems and network management from an z/OS environment. Through its MultiSystem Manager component, Tivoli NetView for z/OS enables management of distributed resources, such as Internet Protocol (IP) resources, NetWare resources, asynchronous transfer mode (ATM) resources, and others.

TN3270. A standard protocol for transmitting 3270 data streams over Telnet. TN3270 is an emulation of an 3270-type terminal via TCP/IP.

Telnet requires a client and a server. The Server is normally called the TELNET Daemon or just TELNETD and maps the 3270 data stream to the Telnet Protocol. The Client is called *TN3270* and displays the 3270 data stream.

TN3270E. Enhances the traditional TN3270 as follows:

- ▶ Provides capability to emulate the 328x printers
- ▶ Enables the Telnet client to request a 3270 device name
- ▶ Supports ATTN and SYSREQ keys
- ▶ Adds support for SNA positive/negative responses

TN3270E server. An IBM software component that enables a TCP/IP client workstation to communicate with a host processor. The TN3270E server accepts SNA traffic from the host processor and converts it

into Telnet format for the client workstation. It also accepts Telnet traffic from the client workstation and converts it into SNA format for the host processor. The TN3270E server implements the protocols defined in Requests for Comments (RFCs) 1646, 1647, and/or 2355.

token-ring. According to IEEE 802.5, network technology that controls media access by passing a token (a special frame) between media-attached stations. A FDDI or IEEE 802.5 network with a ring topology that passes tokens from one attaching ring station (node) to another.

Token-Ring Adapter Type 1. A token-ring interface coupler (TIC) for an IBM Communication Controller that operates at 4 Mbps (megabits per second) token-ring speed.

Token-Ring Adapter Type 2. A token-ring interface coupler (TIC) for an IBM Communication Controller. The adapter can be configured to support 4 Mbps (megabits per second) or 16 Mbps token-ring speed and to support subarea and peripheral nodes on the same adapter. When configured for 16 Mbps, the Token-Ring Adapter Type 2 provides the capability for early token release.

Token-Ring Adapter Type 3. A token-ring interface coupler (TIC) that is supported only on an IBM 3746 Model 900 expansion frame. The adapter can be configured to support 4 Mbps (megabits per second) or 16 Mbps token-ring speed and to support subarea and peripheral nodes on the same adapter. When configured for 16 Mbps, the Token-Ring Adapter Type 3 supports early token release.

token-ring interface coupler. An adapter that can connect a 3720, 3725, or 3745 Communication Controller to an IBM Token-Ring Network.

token-ring network. A ring network that allows unidirectional data transmission between data stations, by a token passing procedure, such that the transmitted data return to the transmitting station. A network that uses a ring topology, in which tokens are passed in a circuit from node to node. A node that is ready to send can capture the token and insert data for transmission.

TPF. A high-availability, high-performance system designed to support real-time, transaction-driven applications. The specialized architecture of TPF is intended to optimize system efficiency, reliability, and responsiveness for data communication and database processing. TPF provides real-time inquiry and update to a large, centralized database, where message length is relatively short in both directions, and response time is generally less than three seconds. Formerly known as the Airline Control Program/Transaction Processing Facility (ACP/TPF).

Transaction Processing Facility. See *TPF*.

transmission group. A connection between adjacent nodes that is identified by a transmission group number. In a subarea SNA network, a single link or a group of links between adjacent nodes. In an APPN network, a single link between adjacent nodes.

TWX. Teletypewriter service in which suitably arranged teletypewriter stations are provided with lines to a central office for access to other such stations throughout the U.S. and Canada. Both baudot- and ASCII-coded machines are used. Business machines may also be used, with certain restrictions. TWX devices are considered to be start-stop terminals.

T1. In the United States, a 1.544 Mbps public access line. It is organized into twenty-four 64 Kbps channels. The European version (E1) transmits 2.048 Mbps (30 channels at 64 Kbps). The Japanese version (J1), like the US version, transmits 1.544 Mbps.

UDP. In the Internet suite of protocols, a protocol that provides unreliable, connectionless datagram service. It enables an application program on one machine or process to send a datagram to an application program on another machine or process. UDP uses the Internet Protocol (IP) to deliver datagrams.

value-added network. A network service provider that is hired by a company to provide network services such as facilitating electronic data interchange (EDI). Before the arrival of the World Wide Web, some companies contracted with value-added networks to move data from their company to other companies. With the arrival of the World Wide Web, many companies found it more cost-efficient to move their data over the Internet instead of paying for VAN services. In response, contemporary value-added network providers now focus on offering EDI translation, encryption, secure e-mail, management reporting, and other extra services for their clients.

virtual circuit. In packet switching, the facilities provided by a network that give the appearance to the user of an actual connection. A logical connection established between two DTEs.

VM. A data processing system that appears to be at the exclusive disposal of a particular user, but whose functions are accomplished by sharing the resources of a real data processing system. In VM/ESA, the virtual

processors, virtual storage, virtual devices, and virtual channel subsystem allocated to a single user. A virtual machine also includes any expanded storage dedicated to it.

VM/ESA. An IBM licensed program that manages the resources of a single computer so that multiple computing systems appear to exist. Each virtual machine is the functional equivalent of a real machine.

VTAM. IBM software that controls communication and the flow of data in an SNA network by providing the SNA application programming interfaces and SNA networking functions. Beginning with Release 5 of the z/OS operating system, the VTAM for MVS/ESA function was included, along with host TCP/IP support, in Communications Server for z/OS.

WAN. A network that provides communication services to a geographic area larger than that served by a local area network or a metropolitan area network, and that may use or provide public communication facilities. A data communication network designed to serve an area of hundreds or thousands of miles, for example public and private packet-switching networks, and national telephone networks.

wide area network. See *WAN*.

Workload Manager. An z/OS component whose purpose is to balance the available system resources to meet the demands of S/390 subsystems.

XCF. See “*cross-system coupling facility*.”

XI. Resides in one or more communication controllers in an SNA network. It opens IBM SNA networks to X.25 traffic and provides resource sharing for X.25 and SNA traffic.

X.21. An International Telegraph and Telephone Consultative Committee (CCITT) recommendation for a general-purpose interface between data terminal equipment and data circuit-terminating equipment for synchronous operations on a public data network.

X.25. An International Telegraph and Telephone Consultative Committee (CCITT) recommendation for the interface between data terminal equipment and packet-switched data networks. The X.25 protocol allows computers on different public networks (such as CompuServe, Tymnet, or a TCP/IP network) to communicate through an intermediary computer at the network layer level. X.25's protocols correspond closely to the data-link and physical-layer protocols defined in the Open Systems Interconnection (OSI) communication model.

X.25 SNA Interconnection. See *XI*.

X.25 interface. An interface consisting of a data terminal equipment (DTE) and a data circuit-terminating equipment (DCE) in communication over a link using the procedures described in the CCITT Recommendation X.25.

X.25 NCP Packet Switching Interface. See *NPSI*.

X.3. An International Telegraph and Telephone Consultative Committee (CCITT) recommendation for packet assembly/disassembly (PAD) in a public data network.

Abbreviations and acronyms

ACF	Advanced Communications Function	CEPT	Conference of European Postal Telecommunications Administration
AHDBEE	APPN, HPR, DLUR, BX, EE, and EBN	CIP	Classical IP. Channel Interface Processor (Cisco).
ALCI	Airlines Line Control Interconnection	CLA	Communication line adapter
ALCS	Airline Control System	CLP	Communication line processor
ANSI	American National Standards Institute	CNN	Composite network node
APPC	Advanced program-to-program communication	COS	Class of service
API	Application programming interface	CPU	Central processing unit
APPN	Advanced Peer-to-Peer Networking	CSPDN	Circuit Switched Public Data Network
ASCII	American National Standard Code for Information Interchange	DCE	Data circuit-terminating equipment
ASP	Application service provider	DHCP	Dynamic Host Configuration Protocol
ATM	Asynchronous transfer mode	DiffServ	Differentiated Services
A2CS	Advanced Application Communication System	DLC	Data link control
BAN	Boundary access node	DLSw	Data link switching
BF	Boundary function	DLUR	Dependent LU requester
BN	Border node	DNS	Domain Name System
BNN	Boundary network node	DTE	Data terminal equipment
bps	Bits per second	EBN	Extended Border Node
BSC	Binary synchronous communication	EDI	Electronic Data Interchange
BX	Branch Extender	EE	Enterprise Extender
CCL	Communication Controller for Linux on System z9 and zSeries	EP	Emulation Program
CCU	Central control unit	ESA	Enterprise Systems Architecture
		ESCD	ESCON Director
		ESCON	Enterprise Systems Connection
		FDDI	Fiber Distributed Data Interface
		FP	Focal point

FRAD	Frame relay access device	NDF	NCP/EP definition facility
FRFH	Frame relay frame handler	NNP	Network Node Processor
Gb	Gigabit	NPSI	X.25 NCP Packet Switching Interface
HPR	High-Performance Routing	NRF	Network Routing Facility
HSSI	High-Speed Serial Interface	NSF	X.25 SNA Network Supervisory Function
IBM	International Business Machines Corporation	NTO	Network Terminal Option
IEEE	Institute of Electrical and Electronics Engineers	MERVA	Message Entry and Routing with Interfaces to Various Applications
IP	Internet Protocol	MLTG	Multilink transmission group
ISDN	Integrated Services Digital Network	MVS	Multiple Virtual Storage
ISDN-BRI	ISDN Basic Rate Interface	OSA	Open Systems Adapter
ISDN-PRI	ISDN Primary Rate Interface	OSA-E	OSA-Express
ITSO	International Technical Support Organization	PAD	Packet assembler/disassembler
ITU-T	See <i>ITU-TS</i> .	PCNE	Protocol converter for non-SNA equipment
ITU-TS	International Telecommunication Union - Telecommunication Standardization Sector	PEP	Partitioned emulation programming
JES	Job Entry Subsystem	PPP	Point-to-Point Protocol
Kbps	Kilobits per second	PTF	Program temporary fix
LAN	Local area network	PVC	Permanent virtual circuit
LCS	LAN Channel Station	PVM	Parallel Virtual Machine
LEN	Low-entry networking	QLLC	Qualified logical link control
LLC	Logical link control	RJE	Remote job entry
LLC2	Logical link control - type 2	RSCS	Remote Spooling Communications Subsystem
LPAR	Logical partition	SDLC	Synchronous Data Link Control
LSA	Link Services Architecture	SNA	Systems Network Architecture
LU	Logical unit	SNI	SNA network interconnection
MAC	Medium access control	SP	(1) Service point (2) Service Processor
MAE	3746 Multi-Access Enclosure	SSCP	System services control point
Mbps	Megabits per second	SSP	System Support Programs
MATIP	Mapping of Airline Traffic over Internet Protocol	SHM	Short Hold Mode
MLTG	Multi-link transmission group		
NCP	Network Control Program		

SH/MPS	Short Hold Mode/Multiple Port Sharing
SONET	Synchronous Optical Network
S.W.I.F.T.	Society for Worldwide Interbank Financial Telecommunication
TAP	Trace Analysis Program
TCAM	An access method used to transfer data between main storage and remote or local terminals.
TCP/IP	Transmission Control Protocol/Internet Protocol
TIC	Token-ring interface coupler
TPF	Transaction processing facility
TWX	Teletypewriter exchange service
UDP	User Datagram Protocol
VAN	Value-added network
VM	Virtual machine
VM/ESA	Virtual Machine/Enterprise Systems Architecture
VTAM	Virtual Telecommunications Access Method
WAN	Wide area network
WLM	Workload Manager
XA	Extended architecture
XCF	Cross-system coupling facility
XI	X.25 SNA Interconnection

Related publications

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

IBM Redbooks

For information on ordering these publications, see “How to get IBM Redbooks” on page 456.

- ▶ *3746-9x0 IP Implementation Guide*, SG24-4845
- ▶ *Inside APPN - The Essential Guide to the Next-Generation SNA*, SG24-3669
- ▶ *IBM 3746 Nways Controller Models 900 and 950: APPN Implementation Guide*, SG24-2536
- ▶ *Migrating Subarea Networks to an IP Infrastructure Using Enterprise Extender*, SG24-5957
- ▶ *MQSeries Primer*, REDP0021
- ▶ *New and Improved! IBM Multisegment LAN Design Guidelines*, GG24-3398
- ▶ *SNA in a Parallel Sysplex Environment*, SG24-2113
- ▶ *Subarea to APPN Migration: HPR and DLUR Implementation*, SG24-5204
- ▶ *Subarea to APPN Migration: VTAM and APPN Implementation*, SG24-4656
- ▶ *TCP/IP Tutorial and Technical Overview*, GG24-3376
- ▶ *zSeries Connectivity Handbook*, SG24-5444

Other resources

These publications are also relevant as further information sources:

- ▶ *Communication Controller for Linux on System z9 and zSeries Implementation and User's Guide*, SC31-6872
- ▶ *3745 Communication Controller Models A and 170, 3746 Nways Multiprotocol Controller Models 900 and 950 Overview*, GA33-0180
- ▶ *IBM 3745 Communication Controller Models A, IBM 3746 Expansion Unit Model 900, IBM 3746 Models 900 and 950 Planning Series: Overview, Installation, and Integration*, GA27-4234

- ▶ *Network Control Program, System Support Programs, Emulation Program Resource Definition Guide*, SC31-6223
- ▶ *Network Control Program, System Support Programs, Emulation Program Resource Definition Reference*, SC31-6224
- ▶ *VTAM Network Implementation Guide*, SC31-6494
- ▶ *X.25 NPSI Version 3 Licensed Program Specifications*, GC30-9605
- ▶ *X.25 NPSI Version 3 General Information*, GC30-3469
- ▶ *X.25 NPSI Version 3 Planning and Installation*, SC30-3470
- ▶ *X.25 NPSI Version 3 Host Programming*, SC30-3502
- ▶ *X.25 SNA Interconnection X.25 SNA Supervisory Function Planning*, GH11-3033
- ▶ *Network Terminal Option Planning, Migration, and Resource Definition (Release 11)*, SC30-3347
- ▶ *NRF General Information*, GC27-0594
- ▶ *NRF Licensed Program Specifications*, GC27-0595
- ▶ *NRF Planning*, SC27-0593
- ▶ *NRF Migration, Resource Definition, and Customization*, SC31-6203
- ▶ *MERVA Extended Connectivity for ESA Installation and User's Guide*, SH12-6157, only at:
<http://publibfi.boulder.ibm.com/epubs/pdf/cmvu0m01.pdf>
- ▶ *TPNS Teleprocessing Network Simulator General Information, Version 3 Release 5*, GH20-2487
- ▶ *OS/390 eNetwork Communications Server: SNA Network Implementation Guide*, SC31-8563
- ▶ *OS/390 eNetwork Communications Server: SNA Resource Definition Reference*, SC31-8565
- ▶ *z/OS IBM Communications Server: IP Migration*, SC31-8512

Referenced Web sites

These Web sites are also relevant as further information sources:

- ▶ IBM Networking Hardware
<http://www.networking.ibm.com/>

- ▶ IBM Communication Controllers
<http://www.prodguide/server-access.html>
- ▶ zSeries Networking
<http://www.ibm.com/servers/eserver/zseries/networking/>
- ▶ Networking & Communications Software
<http://www.ibm.com/software/network/>
- ▶ Advanced Technical Support (Flashes, Presentations, White Papers, etc.)
<http://www.ibm.com/support/techdocs/atmastr.nsf>
- ▶ TPF
<http://www.ibm.com/software/ts/tpf/index.html>
- ▶ Advanced Application Communication System (A2CS)
<http://www.ibm.com/software/ts/tpf/euss/products/a2csnt.html>
- ▶ Cisco Systems, Inc.
<http://www.cisco.com>
- ▶ WebSphere MQ Product
<http://www.ibm.com/software/ts/mqseries>
- ▶ INETCO Systems, Ltd.
<http://www.inetco.com>
- ▶ Datalex, Inc.
<http://www.datalex.com>
- ▶ WebSphere MQ Services
<http://www.ibm.com/software/integration/websphere/services/>
- ▶ Overlap Company, France
<http://www.overlap.fr>
- ▶ Products tested with HNAS
<http://evolution.sna.free.fr/hnase.html>
- ▶ Computer Associates International, Inc.
<http://ca.com>
- ▶ Comm-Pro Associates, Inc.
<http://www.comm-pro.com>
- ▶ Hydra Systems, Inc.
<http://www.hydrasystems.com>

- ▶ JBM Electronics, Co.
<http://www.jbmelectronics.com>
- ▶ TPS Systems, Inc.
<http://www.tps.com>
- ▶ CQ Computer Communications, Inc.
<http://www.cq-comm.com>
- ▶ Society for Worldwide Interbank Financial Telecommunication
<http://www.swift.com>
- ▶ MERVA Product
<http://www.ibm.com/software/solutions/finance/merva>
- ▶ Alebra Technologies, Inc.
<http://www.alebra.com>
- ▶ Enterprise Extender white paper
<http://www.ibm.com/servers/eserver/zseries/networking/pdf/ee-snasw.PDF>
- ▶ OSA-Express performance table
http://www.ibm.com/servers/eserver/zseries/networking/pdf/OSA-Express_perf.pdf

How to get IBM Redbooks

Search for additional IBM Redbooks, Drafts or Redpapers, view, download, or order hardcopy from the IBM Redbooks Web site:

ibm.com/redbooks

Also download additional materials (code samples or diskette/CD-ROM images) from this IBM Redbooks Web site.

Drafts are IBM Redbooks in progress. Not all Redpapers become IBM Redbooks and sometimes just a few chapters will be published this way. IBM Redbook collections are also available on CD-ROMs. Click the CD-ROMs button on the IBM Redbooks Web site for information about all the CD-ROMs offered, as well as updates and formats.

Index

A

- ABC Company 12–17, 19–28, 30–44
 - notable savings 43
 - sample logical and functional inventory sheets 22
- adaptive rate-based (ARB) congestion control 332, 337
- Address Resolution Protocol (ARP) 274
- Airlines Line Control (ALC) 97, 104, 143–144
- Airlines Line Control Interconnection (ALCI) 141, 144–145
- A-ITSO 47
- A-ITSO Company 47
- APPN
 - 3746 Network Node Processor 131
 - Border Node 173
 - Composite Network Node 176
 - For SNA COS 161
 - For SNA routing 165
 - Glossary 427
 - Net390 265
 - router APPN as Composite Network Node alternative 178
- APPN network 174, 340, 343, 407, 416, 420
 - DLUR/DLUS function 416
 - logical structures 340
 - significant numbers 420

B

- binary synchronous communication (BSC) 89–90, 105, 157, 194, 212–217, 219–221, 246, 251
- Border Node (BN) 98, 169, 173, 175, 305
 - Optimization Examples 88
 - SNI alternative 173
- boundary access node (BAN) 147
- boundary function (BF) 96, 98, 107, 110, 135–137, 141, 143, 149, 151, 158, 165–166, 182, 191, 276, 342
 - controller replacement 105, 107–108, 110, 113, 115, 117
 - DLUR 266, 341
 - Glossary 428
 - NCP 149, 151, 155, 157–158, 162–163,

165–166

NNP DLUR 131–132, 135–136

NPSI 191–192, 194, 196

OSA 275–276

boundary network node (BNN) 147, 190

BrNN 340, 343

BSC 3270 99, 146, 214–215, 217

connection 90

device 97, 99, 146, 212, 214–215, 217

functional alternatives 214

Hydra 3000 215

NCP support 146

non-SNA applications (EP) 214

protocol conversion solution 146, 217

SNA applications (NCP) 146

support 146, 214

terminal connection 97, 99, 143

traffic 146, 214–215, 265

VM console solution 217

BSC RJE 99, 217, 219–220, 227, 229–233, 235, 348

non-SNA applications (EP) 217

peer-to-peer (NRF and NTO) 245

SNA applications (NTO) 229

BTAM 214–215

BTAM application 99, 212, 216–217

Bus and Tag 26

byte channel

EP 212

NSI 252

C

CCL

defined 284

CCL Engine 284

communications path 284

CCL NCP 96, 118–120, 132, 135, 137, 182, 287, 289, 296

NNP APPN 137

CCL NCPs

OSA ports 119

CCL option 94, 141, 237, 283

CCL product 169, 283, 286, 404

- additional information 404
- CCM 6
- Channel attachment 96, 165
- Channel Interface Processor (CIP) 138
- Channel Port Adapter (CPA) 138
- Channel Request Handler (CRH) 271
- Channel-attached router 96, 111, 113, 116–117, 123–128
- Channel-attached server 96, 111, 113–114, 117, 123–124, 126–128
- Channel-to-Channel (CTC) 96, 129, 257
- Cisco Systems 138–139, 144, 156, 165, 193, 209, 228–229, 245, 253, 278, 334
- class of service (COS) 98, 141, 143, 158–161, 171
 - DLSw 178, 334
 - EE 179, 333, 341
 - Glossary 430
- CMC 264, 266
- CNN 176
- Communication and Transmission Control Program (CTCP) 196–197, 201
- communication controller
 - and NCP 107
 - CCL Ethernet 118
 - channel attachments 128
 - channel-attached routers 116, 127
 - channel-attached servers 113, 126
 - communication line connections 106
 - Ethernet, FDDI, and ATM 124
 - expansion frames 103
 - for Linux on System z9 283
 - for Linux on zSeries 283
 - frame handler function 148, 150
 - frame relay frame handler function 150
 - functional capabilities 103
 - future role 8
 - hardware 94
 - hardware overview 102
 - host CPU cycles 110
 - IBM 3745 family 286
 - ICA replacement 158
 - logical and functional characteristics 392
 - logical and functional inventory 5
 - M Manufacturing Company example 45
 - MERVA Extended Connectivity 204–205
 - migration 104
 - models 102
 - multi-link transmission group capability 174
 - Net290 262
 - network node processor 131
 - NTO 226
 - Open Systems Adapter 125
 - optimizing 4
 - OSA-Express 270
 - physical inventory 47
 - reconciling and optimizing 72
 - replacing token ring 110
 - router replacement 155, 157
 - saving CPU cycles 136
 - separate token-ring segment 112, 114, 116
 - SNA subarea connectivity 265
 - SNI connectivity 266
 - specific application 252
 - specific support 112, 114, 116
 - staff requirements 5–6, 8
 - subarea SNA networking 163
 - token-ring interfaces 108, 110–111
 - TPNS 256
 - versus routers 105
 - virtualized 288
 - XI and NSF 208
 - ZYX Company example 12
- Communication Subsystem For Interconnection (CSFI) 197
- Communications Management Configuration 264
- Communications Server (CS) 119, 132–133, 138, 319, 324, 338, 414
- Communications Server for AIX 338, 407
- Communications Server for Linux 338, 407
- Communications Server for OS/390 338
- Communications Server for Windows 338, 407
- Composite Network Node (CNN) 176
- composite network node (CNN) 98, 131, 158, 177–178
- Configuration Definition File-Extended (CDF-E) 355
- Connection Network 338
- Content Switching Module (CSM) 139, 327
- controller environment 2, 4, 6, 8–9, 90, 391
 - significant costs 2
- Controller strategic planning 8
- Controllor Configuration and Management (CCM) 6
- Cross-System Coupling Facility (XCF) 129
- Customer Engineer (CE) 12, 47, 355

D

data center 6, 45, 104, 108, 130, 132–134, 145, 153, 166, 188, 324, 334, 340–342
 APPN function 133
 fixed-function terminals 108
 host systems 132
 more expensive routers 134
Data Facility Product (DFP) 180
Data Link Switching (DLSw) 334
Data Terminal Equipment (DTE) 152, 210
DATE 189
dependent LU 333, 341, 407
 high-order element addresses 164
 network route 342
 pre-assigned network address 407
Dependent LU Requester (DLUR) 105, 110, 131, 136, 149, 151, 166, 276, 333
Dependent LU Server (DLUS) 266, 341, 407
dependent PU 406–407
Digital Certificate Authentication Server (DCAS) 319
Direct Memory Access (DMA) 274
display LU
 name 321, 328
DLSw 134, 176
 Composite Network Node alternative 178
 Glossary 431
 NNP alternative 132
 Router alternative to NCP X.25 156
 Router alternative to NPSI X.25 194
 SNI adjacent network connectivity 172
 SNI consolidation 176
 to centralized APPN/DLUR 134
 to VTAM boundary function 135
 versus EE 334
DLUR
 Boundary function 166
 Glossary 431
 SNA addressing 164
DLUR-served LU 408–409, 416
DLUR-served LUs
 low-order addresses 424
 LU-LU sessions 408
DLUR-served PUs 407, 417, 420, 424
 high-order network addresses 417, 421
 low-order addresses 424
DLUS 266, 341
dynamic VIPA 310

E

EE 133, 331
 architecture 336
 class of service (COS) 160
 Connection Network 337
 Glossary 432
 within sysplex 339
EE technology 72, 134, 280
electronic data interchange (EDI) 154, 167–168, 190, 218, 230
element address 163–165, 403, 411, 422
 logical extension 403
Emulation Program
 continued support 26
Emulation Program (EP) 26, 89–90, 94
Encapsulating Security Payload (ESP) 343
End of Record (EOR) 316
Enterprise Extender 72–73, 98, 118, 123, 130, 134, 160, 162, 165, 174–175, 177–179, 257, 259, 265–268, 275, 308, 324, 331–339, 343, 420–421, 425
 extremely important aspect 338
Enterprise Systems CONnectivity (ESCON) 129
EP 211
ESCON channel
 adapter line group 393
 attachment interface 270
 connectivity 74
 technology 26, 28
Ethernet LAN
 Emulation 276
 interface 285
Extended Border Node (EBN) 173

F

Fibre Channel (FICON) 129
fixed-function terminals 106
Frame Relay 97, 141, 143, 147–151, 155, 161, 176, 209, 287
 BAN 147
 BNN 147
 Frame handling 148
 Glossary 433
 INN 147
Functional alternative migrations 9

G

GATE 189

generic resource (GR) 328
generic resources 266–267, 328
Gigabit Ethernet 118, 122, 125, 128–130, 273, 276,
280
Global Connection Network (GCN) 174

H

High Performance Routing (HPR) 118
high-order address 165, 404, 408, 410–411, 413,
422–423
high-order network 411–412, 415–418, 420,
422–424
high-order network address 164, 404, 406,
410–411, 414, 416–417, 419–421, 423–424
high-performance routing (HPR) 131, 267, 333
Host Access Transformation Server (HATS) 306
host environment 113, 116, 122–123, 126–128,
137, 180, 197, 199, 218, 222, 230, 235, 262, 266
 LAN gateways 126–127
Host Network Access Support (HNAS) 200, 202
Host On Demand (HOD) 317, 325
Hydra 3000 99–100, 214–215, 219–220, 222, 232,
235, 253

I

IBM 3745 communication controller 141, 199, 237,
239–240, 244, 283, 285
 hardware 164, 166, 171, 176, 183, 404
 transparent, migration path 244
ICA 97–100, 157–158, 190, 194–196, 214–215,
219–220, 222, 232, 235, 253
independent LU 409–410, 413
INETCO System 146, 152, 227, 231
infrastructure will provide (IP) 96–98, 270,
273–275, 277, 279–280
Integrated Facility for Linux (IFL) 137, 166, 172,
183, 199, 241, 244
Integrated Services Digital Network (ISDN) 97,
143, 151
Interchange Node (ICN) 266
Intermediate network node (INN) 143, 147, 190
Internet Protocol (IP) 331
interorganizational coordination 154, 168, 190,
218, 230
Interpret As Command (IAC) 316
IP network 134, 156, 178–179, 228, 257, 312, 324,
327, 332–339, 343
 common occurrence 337

HPR transport 333
infrastructure 324
logical bridge 312
SNA data 335
IP routing
 NCP 180
 NNP 137

L

LAN Channel Station
 SNMP support 278
LAN Channel Station (LCS) 275, 278
line printer
 daemon 305
 requester 305
Link Service Architecture (LSA) 289
LLC type
 0 196, 201–202
 1 334
 2 traffic 334
 3 91, 153, 189–190
 4 196, 201–202
 5 196, 201
load balancing 109, 119, 139
locally administered address (LAA) 112, 114, 116
Logical and functional inventory
 example 21, 60
 task description 5
 worksheets 391
logical unit (LU) 154, 163, 314–316, 319, 321–323,
325–328, 406
low-order address 165
 significant number 415
low-order network 415–416, 418, 420, 424
 address 410, 412, 415–418, 420–424
low-order network address 416–417, 421, 423, 425
low-speed line
 concentration 25, 44
 hardware 26
LU name 321–323, 325, 327–328
 assignment 321–323, 327
 assignment configuration 321, 323
 assignment configuration definition 323
 nailing 322
 prefix 322
 Specific request 327
LU-LU session 168, 405, 407, 409–411, 414,
416–417, 419–420, 422–424

- DLUR-served LUs 408, 416
- DLUR-served LUs use 417
- eligible DLUR-served dependent LUs 419
- establishment 405
- first network address 418
- low-order addresses 424
- network address pair uniqueness 409
- network addresses 405
- session path 407
- traffic 341

M

- MAE 138
- mainframe host 97, 108, 111–113, 125–126, 135, 176, 199, 238, 240, 242, 244, 247–248, 256, 312, 323–324, 334
 - availability characteristics 324
- mainframe SNA application
 - access 195
- Maintenance and Operations SubSystem (MOSS) 12
- Management Information Base (MIB) 277
- Mbps 102–103, 113, 271, 276, 348
- Media Access Control (MAC) 109, 118–120, 274, 277
- MERVA 203
- message queuing 345
 - EP-based BSC RJE alternative 221
 - Glossary 437
 - NRF alternative 243
 - NRF/NTO alternative 246
 - NTO-based BSC RJE alternative 233
 - SNI alternative 176
- message queuing (MQ) 97, 99–100, 145–146, 154, 168, 190, 218, 221–222, 230, 233–234, 239, 243–244, 246–247, 249, 259, 307–308, 345–348
- migration-strategy alternative 95
- MLTG 161
- MOSS 12
- MQ 345
 - EP-based BSC RJE alternative 221
 - Glossary 438
 - NRF alternative 243
 - NRF/NTO alternative 246
 - NTO-based BSC RJE alternative 233
- MQ alternative 145, 222, 235
- MQ application program interface (API) 346
- multi-access enclosure (MAE) 102, 104, 138

- Multi-link transmission group 161
- Multi-link transmission group (MLTG) 98, 158, 174, 287
- Multiple Image Facility (MIF) 273
- MVS system 216, 220, 232–233

N

- native SNA
 - application 332
- NCP Packet Switching Interface (NPSI) 91, 94, 98, 152–153, 157, 187–192, 194, 196–197, 200
- NCP software 104, 142, 151–153, 162
 - license 2
 - license tier level 4
- Net390 261
- Net390 architecture 121, 162, 259, 261–262, 265, 268
 - Key features 265
- NetView Performance Monitor (NPM) 317, 393
- Network address 163–165, 403–407, 409–415, 417–421, 423–424
 - unique pair 405
- Network Addressable Unit (NAU) 163, 405
- Network Control Program (NCP) 21, 60, 90, 94, 103, 213–214, 237–239, 244, 283, 285–286, 296
- Network Device Handler (NDH) 284
- Network Dispatcher 138
- network node
 - branch extender 340
 - composite network node (CNN) 176
 - Glossary 438
 - Net390 266
 - NNP 131
 - router versus server 343
- network node (NN) 131, 303, 340
- Network Node Processor (NNP) 392–393
- Network Routing Facility (NRF) 90, 94, 100, 227, 235, 237–239, 241–247, 249, 283, 285–286, 393
- Network Session Accounting (NSA) 393
- Network Supervisory Function (NSF) 94, 207–208
- Network Terminal Option (NTO) 90, 94, 99, 393
- Network Virtual Terminal (NVT) 313
- NNP APPN
 - capability 131–132
 - environment 132
 - function 133, 135
 - support 137
- non-SNA application 90, 99, 212, 223

- exchange information 212
- non-SNA device 99–100, 152, 188, 196, 216, 226–227, 235, 245
 - peer-to-peer communication 235
 - Peer-to-peer connection 99, 235
- Non-SNA Interconnection (NSI) 94
- NPSI 187
- NRF 237
- NSF 210
- NSI 251
- NTO 225
- NTuneMON 286, 296

O

- Open Systems Adapter (OSA) 72, 96, 98, 111, 117, 124–125, 129, 164, 259, 269–270, 325
- Optimization 7, 11
- OS/390 V2R5 329, 417, 421, 424
 - DLUR-attached dependent LUs 416
 - enhancement 417
- OS/390 V2R7 411, 417–418, 420, 422–424
 - cause VTAM 419
 - change 418
 - TN3270 APPLs 419
- OSA 269
 - Ethernet, FDDI, or ATM 125
 - Glossary 439
 - Net390 122
 - token-ring 111
- OSA-Express 269

P

- Packet Level Protocol (PLP) 196
- packet switched data network (PSDN) 209
- Parallel channel 129
- parallel session 405–406, 415, 418
 - same network address 405
- Parallel Sysplex
 - Net390 265
 - Net390 benefits 267
 - Net390 NN placement 266
 - Network Dispatcher 138
 - OSA 275
- partitioned emulation program (PEP) 213
- PCNE 189
- Peer-to-peer connection 100
- PEP 213
- Peripheral Component Interconnect (PCI) 278

- physical interface 7, 101, 103–105, 107, 142, 147, 176, 240, 287, 339
 - migration alternatives 104
 - SNI traffic 176
 - wide variety 240
- Physical inventory
 - example 12, 47
 - overall goal 5
 - task description 5
 - worksheets 353
- physical inventory 4–5, 7–8, 11–13, 47, 353–354, 356, 384
- Physical Unit (PU) 142, 162–163, 176, 178, 314, 326
- protocol conversion 97, 99–100, 143, 145–146, 152, 209, 217–219, 223, 226–231, 235, 242, 246–249
 - NRF replacement 248
- PSH 189
- PU 2 162
- PU 2.1 162, 165, 176–177
- PU 4 142, 162
- PU 5 162
- public switched data network
 - SNA devices 152
- public switched data network (PSDN) 152
- pure NN 418, 424

Q

- QLLC
 - Glossary 442
 - NCP 152
 - NPSI 190
- qualified logical link control (QLLC) 153–154, 156–157, 189–194, 196–197, 336
- Quality of Service (QOS) 179, 301, 303, 309
- Queued Direct I/O (QDIO) 273–274, 276, 278–279

R

- Redbooks Web site 456
 - Contact us xix
- remote job entry (RJE) 212, 217, 219–222, 251
- remote location 102, 134, 324, 333, 343, 407
 - APPN function 134
 - APPN network 407
 - IP network infrastructure 324
 - utility server 343
- Remote Printing (LPR/LPD) 305

Resource Measurement Facility (RMF) 278
 Responsive Mode ARB 337
 RJE-based user 348
 router-based solution 155, 158, 192, 196
 routers 105, 301, 332

- APPN 132, 178
- communication lines 106, 108
- DLSw 334
- DLUR 342
- Enterprise Extender 174, 179, 332
- Ethernet, FDDI, ATM 127
- frame relay 147
- Glossary 442
- IP routing 137, 180
- ISDN 151
- MATIP 144
- migrating from X.25 156, 193
- SNA addresses 163
- SNA class-of-service (COS) 160, 333
- SNA Switching Services (SNASw) 338, 407
- SNI consolidation 176
- TN3270 138
- token-ring 116
- tunneling 245, 253
- VTAM enhanced addressing 164
- X.25 154, 156, 191, 193
- X.25 over TCP/IP (XOT) 197
- X.25 transport 209

 RTP pipe 421
 RTP PU 421
 Running NCP 98, 100

S

S.W.I.F.T. 203
 secondary LU 314, 320, 405
 secondary LUs

- message flows 239

 Secure Sockets Layer (SSL) 303
 Security Access Facility (SAF) 325
 Self Timed Interface (STI) 271, 275
 Server/Application State Protocol (SASP) 97, 139
 Service Access Point (SAP) 119
 session partner 406

- network address 415, 418–420
- network addresses 415
- parallel sessions 406
- subarea number 417

 session-capable resource 164, 404, 406, 410–411,

414, 420, 422

- various types 415

 Simple Network Management Protocol (SNMP) 278
 single NCP 29
 single network address (SNA) 105, 107–108, 110, 112, 118, 123, 131–132, 134–137, 405, 409, 418
 SNA (Systems Network Architecture) 443
 SNA 3270 146, 214, 216–217, 311–312, 314, 323
 SNA application 90, 97, 99, 119, 143, 176, 187, 189, 209, 212, 214, 217, 223, 225–226, 275, 306–307, 314–315, 320, 322–324, 326, 328, 332, 403, 405

- connection 189
- data 332
- environment 266
- line-mode SNA session 315
- migration 169
- program 405
- subsystem 320
- traffic 118

 SNA Class 335, 421
 SNA class

- detailed discussion 335

 SNA Class of Service 98, 158–159, 171, 176
 SNA data

- stream 319
- unit 335

 SNA device 99–100, 110, 131, 136, 147, 149, 151–152, 154, 188–189, 191, 193, 196–197, 238–239, 241–242, 245, 276, 289, 403–404
 SNA interconnection 88, 91, 94, 207–208, 393
 SNA Network

- Interconnect 26–27, 44, 72–74, 88
- Supervisory Function 208

 SNA network 26, 91, 146, 159, 173, 187, 191, 207, 214, 237, 239, 251, 253, 256, 266, 314, 324, 347, 405

- CMC ownership 266
- connected TN3270 client 314
- connection 253
- control mode 213
- dependency 308
- element address 403
- flow 314
- host applications 251
- infrastructure 314
- interconnection 89–90, 98, 105, 141, 143, 158, 285, 305, 393

- management focal point 266
 - non-SNA host NJE traffic 253
 - packet forwarding 207
 - peer-to-peer communication 239
 - specific resources 405
 - transport 209
 - SNA Network Interconnect (SNI) 167
 - SNA networking
 - environment 159, 296
 - SNA Primary LU Support (SPLS) 342
 - SNA session
 - disruption 334
 - manager application 403
 - setup failure 110, 136, 276, 403
 - traffic 132
 - SNA subarea 98, 110, 136, 158, 257, 276
 - SNA Switching Services (SNASw) 132, 338, 407
 - SNA traffic 112, 118, 134, 147, 159, 161, 163–165, 178, 194, 257, 265, 275, 285, 324, 331–332, 334, 336–337
 - different types 334
 - physical links 161
 - SNASw
 - EE implementations 338, 407
 - Router replacement of APPN functions 132
 - SNI 167
 - adjacent network connectivity 172
 - alternatives 169
 - border node 173
 - controller host-to-host 129
 - controller replacement 105
 - EBN over VPN using EE 343
 - Extended Border Node (EBN) 174
 - Glossary 443
 - Logical Inventory 6
 - MQ 347
 - Net390 266
 - optimization examples 26, 44, 72, 74, 88
 - over frame relay 147
 - SNI partner 26, 73, 88, 169, 171–172, 176
 - source route bridging (SRB) 334
 - SSCP-LU session 407–408, 416, 419
 - DLUR-served LUs 417
 - Start-stop
 - EP 212
 - Glossary 443
 - NCP to TCAM 152
 - NPSI 196
 - NTO 226
 - start-stop terminal 152, 200, 227–229
 - PCNE 3767 function 200
 - storage requirement 29, 110, 136, 149, 151, 166, 276
 - subarea networking and APPN (SNA) 141, 143, 146–147, 149, 151–171, 173, 176–178, 262, 264, 266–268
 - subarea node 162, 165, 173, 405–407, 410, 414
 - support new applications (SNA) 107–108, 110, 112–113, 115, 117–118, 122–124, 129–130, 132–138, 142, 146–147, 149, 151–158, 160–170, 172–174, 176–178, 182, 331–339, 341–342
 - Synchronous Data Link Control (SDLC) 97, 141, 143, 152, 154, 157, 161, 190, 194
 - sysplex
 - generic resources for TN3270 329
 - Sysplex Distributor 97, 138–139, 267, 310, 326–328
 - Systems Network Architecture (SNA) 44, 188, 237–242, 245, 247, 270, 274–276, 280, 305–308, 427
- ## T
- TCP connection 313–314, 317, 328, 335
 - additional burden 335
 - secure sockets 317
 - unmodified 3270 data stream 314
 - technical solution 95, 101, 142, 187, 203, 207, 211, 214, 225, 237, 251, 255
 - Teleprocessing Network Simulator (TPNS) 94, 100, 255–257
 - Telnet 3270 (TN3270) 311
 - Telnet client 315–316, 328
 - terminal controllers 106
 - TN3270 135, 138, 311
 - BSC 3270 conversion 146
 - Enhancements (TN3270E) 316
 - Glossary 444
 - Overview 312
 - SNA addressing 164
 - SNA boundary function 166
 - TN3270 client 275, 307, 314, 317–320, 322
 - data stream 315
 - IP-based communication 324
 - network protocols 324
 - secure TCP connection 317
 - SSL connection 319
 - TCP connection 317

TN3270 emulation 315
TN3270 server 97, 135–136, 138, 263, 275, 310,
312–315, 317–318, 321–328, 423
 Configuration definitions 321
 load balancing 323
 SNA communication 314
 SNA network 324
 strategic platform 326
token ring 110–113, 285, 287–288
 connectivity 111
 defined 445
 Duplicate MAC addressing 108
 Open Systems Adapter 111
Token-ring 102, 108
token-ring gateway 44, 113–114, 116–117, 403
token-ring interface 108, 111–112, 114, 116, 120
TPNS 255
traditional Telnet
 and TN3270 terminal emulation 315
 client 315
 data stream 315
 server 315
 uses line mode 315
Transaction Processing Facility (TPF) 144–146,
175, 273
Transmission Control Protocol (TCP) 332

V

value-added network (VAN) 167, 172
Virtual IP Address (VIPA) 310
virtual private network (VPN) 108, 145, 154–155,
168, 190, 193, 218, 230, 302
Virtual Router Redundancy Protocol (VRRP) 310
virtualized communication controller 288
VTAM 119, 131–132, 134–137, 146, 152, 157,
164–165, 168, 173–176, 178–179, 314, 317, 326,
329, 331, 340, 403–406, 408–411, 413–424
VTAM application
 definition file 201
 program 406
VTAM Enhanced Addressing (ENHADDR) 403
VTAM subarea 110, 112–113, 115, 117, 136–137,
164, 276, 404, 408, 411, 413, 415–417
 NCP subarea 164
 SNA element addresses 110, 136

W

Web Site 154, 191, 204, 210, 231–232, 347

wide-area network
 infrastructure 420
 LINE attachment 420
wide-area network (WAN) 170, 240, 334
Workload Manager 267, 327

X

X.25
 Glossary 447
 NCP and 3746 feature 152
 NPSI 188
 Non-SNA support 196
 SNA support 190
 S.W.I.F.T. connectivity 203
 XI and NSF 207
X.25 over TCP/IP (XOT) 199, 228
X.3 PAD 189
XOT 199, 228
XRF environment 180, 182

Z

z/OS Communications Server 338
z/OS Telnet 316
zero-hop route 329
zSeries server 264, 267–268, 271–273
 OSA-Express feature support 272



Redbooks

IBM Communication Controller Migration Guide:

Updated for Communication Controller for Linux on System z9 and zSeries (V1.2)

(1.0" spine)
0.875" <-> 1.498"
460 <-> 788 pages



IBM Communication Controller Migration Guide: Updated for Communication Controller for Linux on System z9 and zSeries (V1.2)



**Meet changing
business needs with
clear migration
options and
strategies**

**Reduce costs by
simplifying your
networking
environment**

**Enhance
performance by
exploiting new
technologies**

IBM communication controllers have reliably carried the bulk of the world's business traffic for more than 30 years. Over the years, IBM controllers have been enhanced to the point that the functional capabilities of the current products, the 3745 Communication Controller and the 3746 Nways Multiprotocol Controller, surpass the capabilities of any other data networking equipment ever developed. Beyond the SNA architecture PU Type 4, beyond APPN, even beyond IP routing, these controllers support an extraordinary set of functions and protocols. Because of their long history and their functional richness, IBM controllers continue to play a critical role in the networks of most of the largest companies in the world.

Over the past decade, however, focus has shifted from SNA networks and applications to TCP/IP and Internet technologies. In some cases, SNA application traffic now runs over IP-based networks using technologies such as TN3270 and Data Link Switching (DLSw). In other cases, applications have been changed, or business processes reengineered, using TCP/IP rather than SNA. Consequently, for some organizations, the network traffic that traverses IBM communication controllers has declined to the point where it is in the organization's best interest to find functional alternatives for the remaining uses of their controllers so they can consolidate and possibly eliminate controllers from their environments.

This IBM Redbook provides you with a starting point to help in your efforts to optimize your communication controller environment, whether simply consolidating them or migrating from them altogether. We discuss alternative means to provide the communication controller functions that you use or ways to eliminate the need for those functions outright. Where multiple options exist, we discuss the relative advantages and disadvantages of each.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

**For more information:
ibm.com/redbooks**