

IBM SecureWay Communications Server for Windows



Network Administration Guide

Version 6.1

IBM SecureWay Communications Server for Windows



Network Administration Guide

Version 6.1

Note

Before using this information and the product it supports, read the information in Appendix B, "Notices", on page 187.

Third Edition (December 2003)

This edition applies to Version 6.1 of IBM SecureWay Communications Server for Windows and to all subsequent releases and modifications until otherwise indicated in new editions.

© Copyright International Business Machines Corporation 1998, 2003. All rights reserved.

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

Contents

| | |
|--------------------------|------------|
| Figures | vii |
|--------------------------|------------|

| | |
|-------------------------|-----------|
| Tables | ix |
|-------------------------|-----------|

| | |
|-----------------------------------|-----------|
| About This Book. | xi |
|-----------------------------------|-----------|

| | |
|--|------|
| How to Use This Book | xi |
| Conventions Used in This Book | xii |
| Text Conventions | xii |
| Number Conventions | xii |
| Where to Find More Information | xii |
| What's New in This Book | xiii |

Part 1. Overview and Concepts 1

Chapter 1. Introducing Communications

| | |
|------------------------------------|----------|
| Server Functions. | 3 |
|------------------------------------|----------|

| | |
|---|----|
| Supported Functions. | 3 |
| SNA Communications Support | 3 |
| APPC/APPN | 4 |
| Advanced Peer-to-Peer Networking (APPN). | 8 |
| Programming Interfaces | 9 |
| High Performance Routing | 9 |
| SNA Data Compression. | 9 |
| Discovery of Service Providers | 10 |
| Dependent LU Requester (DLUR) | 10 |
| SNA Gateway | 10 |
| AnyNet Support. | 11 |
| TN3270E Server | 13 |
| TN5250 Server | 14 |
| Legacy Data Access. | 15 |
| Communications Server Client Support | 16 |
| Configuration and Administration Support. | 17 |
| Load Balancing | 18 |
| Data Security. | 18 |
| Adapter Support | 18 |
| Data Link Control (DLC) Support | 19 |
| APPN MIB Support | 19 |
| Hot Standby | 19 |
| Entry-Level Emulator Functions | 20 |

Chapter 2. Communications Server and

| | |
|-----------------------|-----------|
| SNA. | 21 |
|-----------------------|-----------|

| | |
|---|----|
| Overview of SNA Functions. | 21 |
| Data Link Control Profiles | 22 |
| Advanced Peer-to-Peer Networking | 22 |
| APPN Node Types | 23 |
| Control Points | 25 |
| Branch Extender. | 27 |
| Supported Functions | 30 |
| Data Link Control (DLC). | 31 |
| Connections | 31 |
| Link Types | 32 |
| Link Definition and Activation Parameters | 33 |
| Activate at Startup | 33 |

| | |
|--|----|
| Activate on Demand | 33 |
| Automatic Link Retry | 34 |
| Maximum Activation Attempts. | 34 |
| User Requested Reactivation. | 35 |
| Link Deactivation Parameters | 35 |
| Inactivity Timeout | 35 |
| Limited Resource | 36 |
| Connection Networks | 36 |
| Other Link Parameters | 37 |
| Adjacent Node Type | 38 |
| Preferred Network Node Server | 38 |
| Solicit SSCP Sessions | 38 |
| Directory Services | 38 |
| Topology and Route-Selection Services | 40 |
| Topology Database | 41 |
| Modes | 42 |
| Class of Service | 42 |
| SNA Transmission Priority | 43 |
| Route Selection | 43 |
| Route Selection for VTAM Users | 44 |
| Intermediate Session Routing | 44 |
| High Performance Routing (HPR) Support | 45 |
| Rapid Transport Protocol (RTP). | 46 |
| Automatic Network Routing (ANR) | 47 |
| LU Support | 47 |
| SDDLU Support. | 48 |
| Dependent Logical Unit Requester Support. | 49 |
| LU-LU Sessions | 50 |
| LU 6.2 | 51 |
| Other LUs. | 52 |
| APPN Network Node and T2.1 Support. | 52 |
| Data Compression | 52 |
| SNA Session-Level Compression Architecture | 53 |
| Communications Server Data Compression. | 54 |
| SNA Session-Level Encryption | 55 |
| Management Services | 56 |
| Focal Points, Service Points, and Entry Points | 56 |
| Levels of SNA Management Services Architecture | 57 |
| Flow Control | 57 |
| Session-Level Pacing | 57 |
| Adaptive BIND Pacing | 58 |
| Segmenting and Reassembly. | 58 |
| High Performance Routing Pacing. | 58 |
| Fixed Pacing | 58 |
| Partitioning LUs among Hosts | 59 |
| SNA Gateway Support | 60 |

Part 2. Planning and Installation 63

Chapter 3. Planning for AnyNet

| | |
|--------------------------|-----------|
| Support | 65 |
|--------------------------|-----------|

| | |
|---|----|
| Configuring AnyNet SNA over TCP/IP | 65 |
| Mapping SNA Resources to IP Addresses | 65 |
| SNA over TCP/IP Gateway Considerations. | 67 |

| | |
|--|----|
| SNA over TCP/IP Access Node Function Considerations | 70 |
| AnyNet SNA over TCP/IP Configuration Examples | 71 |
| Example 1. Running APPC or CPI-C Applications over a TCP/IP Network | 71 |
| Example 2. 3270 Emulation via DLUR over a TCP/IP Network | 72 |
| Example 3. Using an SNA Gateway to Enable 3270 Emulation Between SNA and TCP/IP Networks | 73 |
| Example 4. Using an SNA Gateway for 3270 Emulation over a TCP/IP Network | 74 |
| Example 5. 3270 Emulation from Two Windows NT Workstations on Different IP Networks | 75 |
| Helpful Hints. | 76 |
| Configuring AnyNet Sockets over SNA | 77 |
| How Does Sockets over SNA Work? | 77 |
| Application Program Support Provided by Sockets over SNA | 78 |
| Planning for Sockets over SNA | 79 |
| Routing and Mapping Overview | 79 |

Chapter 4. Planning for Client/Server Communication 85

| | |
|---|----|
| SNA API Clients. | 85 |
| Installation and Configuration | 85 |
| Lightweight Directory Access Protocol | 85 |
| Common API Client for Communications Server for Windows and NetWare for SAA | 85 |
| Assigning a Default Local LU | 86 |
| Assigning a Default Partner LU | 87 |
| Assigning a Default LUA Session Name | 87 |
| Client/Server Data Encryption | 87 |
| Client Server Security | 87 |
| Novell NetWare for SAA Clients | 88 |

Chapter 5. Planning for Legacy Data Access. 89

| | |
|--|----|
| iSeries OLE DB Provider | 89 |
| iSeries Shared Folder Server | 90 |

Chapter 6. Planning for HPR 91

| | |
|--|----|
| HPR Link Characteristics | 91 |
| HPR Considerations | 92 |
| Considerations for Tuning Parameters on LAN DLCs | 93 |
| HPR Path Switch Processing. | 94 |

Chapter 7. Planning for SNA Gateway 95

| | |
|---|-----|
| Supported Connections to Workstations | 95 |
| SDLC Workstation Connections. | 95 |
| X.25 Workstation Connections | 96 |
| Supported Connections to Hosts | 96 |
| Host Considerations | 97 |
| LAN Destination Address Configuration | 98 |
| Planning Gateway LU Definitions | 98 |
| Using Explicit and Implicit Connections | 99 |
| SNA Gateway Performance. | 100 |
| DLUR Considerations | 101 |

Chapter 8. Planning for TN3270E Server 103

| | |
|--|-----|
| Supported Client Workstations under TN3270E Server | 105 |
| Highlights | 105 |
| Changing the Default Port Number | 105 |
| Managing System Traffic | 106 |
| IP Filtering | 106 |
| Secure Sockets Layer (SSL) Support | 107 |
| Configuring SNA Connections. | 107 |
| Pooling | 107 |
| Load Balancing of TN3270E Servers | 108 |

Chapter 9. Planning for TN5250 Server 109

| | |
|--|-----|
| Supported Client Workstations under TN5250 Server | 111 |
| Highlights | 111 |
| Changing the Default Port Number | 111 |
| Managing System Traffic. | 111 |
| Specifying Access to iSeries Hosts with Multiple Ports | 112 |
| IP Filtering | 112 |
| Secure Sockets Layer (SSL) Support | 112 |
| Configuring the SNA Network | 112 |
| Load Balancing of TN5250 Servers | 113 |

Chapter 10. Planning for Secure Sockets Layer-based Security 115

| | |
|---|-----|
| How SSL Security Works | 115 |
| Communications Server SSL Support | 116 |
| Configuring SSL Security | 116 |
| Server Authentication. | 116 |
| Client Authentication | 117 |
| Key Management Utility. | 117 |
| Changing the Password | 118 |
| Configuring SSL Using a Well-Known Trusted CA | 118 |
| Creating a Key and Certificate Request | 118 |
| Submitting a Certificate Request | 119 |
| Storing a Certificate in the Key Database | 119 |
| Configuring SSL Using an Unknown CA | 120 |
| Creating a Key and Certificate Request. | 120 |
| Submitting a Certificate Request | 120 |
| Storing a Certificate in the Key Database | 121 |
| Creating a Self-Signed Certificate. | 121 |
| Storing Certificates for Clients. | 122 |

Chapter 11. Planning for Load Balancing 123

| | |
|---|-----|
| Dependent LU Load Balancing | 123 |
| LU 6.2 Load Balancing | 123 |
| Planning for TCP/IP Scopes | 124 |
| Scopes and Client Server Security | 124 |
| Where Scope is Configured. | 124 |
| How Scope Relates to SLP | 125 |
| Tracking the Current Load | 125 |

Chapter 12. Planning for Backup Host Connections. 127

| | |
|---|-----|
| Planning for Backup Connections between Servers Using Hot Standby | 127 |
| Using a Production or Dedicated Backup Server for Hot Standby | 129 |
| Planning for the Complete Hot Standby Environment | 130 |
| Configuring Hot Standby on the Backup Server Configuration of a Hot Standby Environment | 131 |
| Planning for Backup Connections between Servers without Using Hot Standby | 132 |
| Planning for Backup Connections in a Single Server | 132 |

Chapter 13. Planning for X.25 133

| | |
|-------------------------------------|-----|
| Possible X.25 Connections | 134 |
| X.25 Network Charges | 134 |

Chapter 14. Configuring Host Connectivity 135

| | |
|---|-----|
| NCP Token Ring (NTRI) with VTAM and Communications Server | 135 |
| Token Ring with VTAM Switched Network and Communications Server | 135 |
| 9370 IBM Token-Ring Network Adapter and Communications Server | 136 |
| 3174 Controller and Communications Server | 137 |
| 3174 Controller VTAM PU and Communications Server | 137 |
| Ethernet with VTAM Switched Network and Communications Server | 138 |
| VTAM/NCP and Communications Server (SDLC) VTAM and Communications Server (APPC) | 140 |
| VTAM and Communications Server (MPC) | 141 |
| ES/9000 and Communications Server (SDLC) | 142 |
| Host and Communications Server (DLUR) | 144 |
| NPSI and Communications Server (X.25) | 145 |
| NPSI Connections over X.25 | 145 |
| VTAM/NCP and Communications Server (X.25) APPC over X.25 | 147 |
| iSeries Host Line Definitions | 148 |
| iSeries Host Controller Definitions | 149 |
| iSeries Host Device Definitions | 150 |
| iSeries Host Line Description and Communications Server | 150 |
| iSeries Host Controller Description and Communications Server | 151 |
| iSeries Twinaxial Descriptions | 151 |
| iSeries Host Line Definitions (X.25) | 152 |
| iSeries Host Controller Definitions (X.25 PVC) | 153 |
| iSeries Host Controller Definitions (X.25 SVC) | 154 |

Chapter 15. Implementing Communications Server Configurations 155

| | |
|--|-----|
| Introduction | 155 |
| Step 1. Planning for Hardware | 155 |
| Storage | 156 |
| Memory Capacity | 156 |
| CPU Speed | 156 |
| Step 2. Planning for Software Applications | 157 |

| | |
|--|-----|
| Step 3. Creating Your Naming Conventions | 157 |
| Criteria for Naming | 157 |
| Types of Names and Their Restrictions | 158 |
| Step 4. Defining Network Addresses | 161 |
| LAN Adapter Addresses | 161 |
| SDLC Secondary Station Addresses | 162 |
| X.25 Addresses | 162 |
| Internet Addresses | 162 |
| Step 5. Selecting Configuration and Installation Tools | 163 |
| Node Configuration Application | 163 |
| Remote Configuration | 163 |
| ASCII Configuration File | 164 |
| Web Administration | 164 |
| Configuration with Template and Response Files | 164 |
| Step 6. Choosing Communications Server Features | 167 |
| Step 6.1 Planning for Application Programs | 169 |
| Step 7. Creating Configurations and Installing | 170 |
| Using Response File Configuration and Installation | 170 |
| Step 8. Creating User Materials | 170 |
| Preparing Documentation | 170 |
| Communications Server Materials and Procedures | 170 |
| Preparing Backup Procedures | 171 |
| Step 9. Maintaining Your Network | 171 |

Part 3. Controlling and Maintaining a Network 173

Chapter 16. System Management Facilities 175

| | |
|--|-----|
| Common Capabilities of System Management Facilities | 175 |
| SNA Node Operations | 176 |
| Remote Administration | 176 |
| Command-Line Programs | 177 |
| Web Administration | 177 |
| SNA Node Operations ActiveX Control | 177 |
| Tivoli Plus Module | 177 |
| APPN MIB Support | 178 |
| zSeries Remote Operations Support (ROPS) for the NetView Program | 178 |

Part 4. Appendixes 179

Appendix A. Planning for Communications Server 181

| | |
|---|-----|
| Modems | 181 |
| Session-Level Encryption Adapters | 181 |
| Data Compression Considerations | 181 |
| Planning the Host/VTAM Connection | 182 |
| Planning for iSeries Data Compression | 184 |
| Performance Considerations | 185 |

Appendix B. Notices 187

| | |
|----------------------|-----|
| Trademarks | 188 |
|----------------------|-----|

Index 191

Figures

| | | |
|-----|---|-----|
| 1. | A Session between Two LUs (LU-LU) | 6 |
| 2. | Transaction Program Conversation Occurring over a Session | 7 |
| 3. | Parallel Sessions between LUs. | 8 |
| 4. | Example of SNA Gateway Connections | 11 |
| 5. | SNA over TCP/IP Gateway | 12 |
| 6. | Sockets over SNA Gateway | 13 |
| 7. | TN3270E Server Connections. | 14 |
| 8. | TN5250 Server Connections | 15 |
| 9. | A Portion of a Sample APPN Network | 24 |
| 10. | Conceptual Overview of Branch Uplinks and Branch Downlinks | 27 |
| 11. | The Branch Extender in a Network. | 28 |
| 12. | A Sample Connection Network | 37 |
| 13. | LEN Node Directory | 39 |
| 14. | End Node Directory. | 39 |
| 15. | Network Node Directory | 40 |
| 16. | Local Configuration Database and Network Topology Database in Network Nodes | 42 |
| 17. | DLUR Connection to a Host through a Communications Server Gateway | 50 |
| 18. | Multiple and Parallel Sessions | 51 |
| 19. | Example of SNA Gateway Configuration | 60 |
| 20. | Formats of the Domain Names That SNA over TCP/IP Builds | 66 |
| 21. | Defining a CP Name and a Connection Network Name | 68 |
| 22. | Domain Name Server Definitions for a Single Gateway Connected to an SNA Network with Two Network IDs | 69 |
| 23. | Domain Name Server Definitions for Parallel Gateways Connected to an SNA Network with Two Network IDs | 70 |
| 24. | Structure of a Windows NT Node Running Sockets over SNA | 77 |
| 25. | Example of an IP Routing Panel. | 81 |
| 26. | Example of a Network Using the Sockets over SNA Route Discovery Function | 82 |
| 27. | iSeries OLE DB Provider | 90 |
| 28. | Perspective to Use for Destination Address Entries | 98 |
| 29. | A Simple Scenario Using Pooled and Dedicated LUs | 99 |
| 30. | Communications Server Configured as a TN3270E Server with Wide Area SNA Network | 103 |
| 31. | Communications Server Configured as a TN3270E Server with Wide Area TCP/IP Network | 104 |
| 32. | Communications Server Configured as a TN5250 Server with Wide Area SNA Network | 109 |
| 33. | Communications Server Configured as a TN5250 Server with Wide Area TCP/IP Network | 110 |
| 34. | A Critical Server with One Backup Server | 128 |
| 35. | A Critical Server with Two Backup Servers | 128 |
| 36. | X.25 and SNA Gateway Connectivity | 134 |
| 37. | VTAM Startup Parameter | 183 |
| 38. | VTAM Logmode Table | 184 |
| 39. | VTAM Application Statement | 184 |
| 40. | VTAM PU and LU Definitions | 184 |

Tables

| | | | | | |
|-----|--|-----|-----|--|-----|
| 1. | Supported APPC Connections | 4 | 21. | Host/PC Parameters Cross-Reference: X.25 (VTAM/NCP) | 146 |
| 2. | Local NN Configuration Database | 41 | 22. | Host/PC Parameters Cross-Reference: APPC on an IBM Token-Ring Network Using X.25 to a Host | 147 |
| 3. | Local NN Network Topology Database | 42 | 23. | iSeries Communications Server Parameters Cross-Reference: OS/400 Host Program Line Definition | 148 |
| 4. | SNA Gateway Summary | 61 | 24. | iSeries Communications Server Parameters Cross-Reference: Controller Parameters | 149 |
| 5. | IP Address Masks Supported by Sockets over SNA | 80 | 25. | iSeries Communications Server Parameters Cross-Reference: Device Parameters | 150 |
| 6. | Command Equivalents | 105 | 26. | Host/PC Parameters Cross-Reference: Line Description | 151 |
| 7. | Host/PC Parameters Cross-Reference: IBM Token-Ring Network NTRI Definition | 135 | 27. | Host/PC Parameters Cross-Reference: Controller Description | 151 |
| 8. | Host/PC Parameters Cross-Reference: IBM Token-Ring Network VTAM Switched Net | 135 | 28. | Host/PC Parameters Cross-Reference: Controller Description | 151 |
| 9. | Host/PC Parameters Cross-Reference: IBM Token-Ring Network 9370 VTAM LAN | 136 | 29. | iSeries Communications Server Parameters Cross-Reference: OS/400 Host Program Line Definition (X.25) | 152 |
| 10. | Parameter Matching between a 3174 Token Ring and Communications Server | 137 | 30. | iSeries Communications Server X.25 PVC Parameters Cross-Reference: Controller Parameters | 153 |
| 11. | Host/PC Parameters Cross-Reference: IBM Token-Ring Network 3174 Controller VTAM PU and Communications Server | 137 | 31. | iSeries Communications Server X.25 SVC Parameters Cross-Reference: Controller Parameters | 154 |
| 12. | Host/PC Parameters Cross-Reference: Ethernet Network VTAM Switched Network | 138 | 32. | Memory Capacity | 156 |
| 13. | Host/PC Parameters Cross-Reference: SDLC | 139 | 33. | IP Address Classes Supported by Communications Server | 163 |
| 14. | Host/PC Parameters Cross-Reference: APPC | 141 | 34. | Management Actions for Communications Server Resources | 175 |
| 15. | IOCP/HCD Parameters Cross-Reference: MPC | 141 | | | |
| 16. | VTAM Parameters Cross-Reference: MPC | 142 | | | |
| 17. | Host/PC Parameters Cross-Reference: 9370 Information System via SDLC | 143 | | | |
| 18. | VTAM Communications Server Parameter Cross-Reference: Using DLUR | 144 | | | |
| 19. | VTAM Communications Server DLUR Parameter Matching Connection from Host to Workstation | 144 | | | |
| 20. | Host/PC Parameter Cross-Reference: NPSI Connections over X.25 | 145 | | | |

About This Book

IBM® SecureWay® Communications Server for Windows® (termed *Communications Server* in this book) is a communications services platform. This platform provides a wide range of services for Windows workstations that communicate with host computers and with other workstations. Communications Server users can choose from among a variety of remote connectivity options.

The *Network Administration Guide* is provided as a planning and administration source for the person using or planning to use Communications Server. Administrative personnel, technical and service support personnel, service coordinators, and IBM personnel, as well as anyone who has responsibility for making decisions on information systems will find this book helpful. This book provides an overview of Communications Server functions and the steps required to create the Communications Server configurations for the workstations you support.

How to Use This Book

Network Administration Guide helps you perform the planning required to install and configure a network of workstations. This book helps you:

- Plan your network
- Choose configuration and installation tools
- Create a template configuration file
- Implement your network plan
- Maintain your network

Setting up a network of workstations requires careful and thorough planning. When you create a network, you need to know the following things:

- The workstation names, profile names, configuration file names, and other names in your network
- Where the focal point is for sending alerts
- The workstations you have that you can continue to use, and how many new workstations you need
- The role you want each workstation to have in the network
- The package to be installed on each workstation so that it can fulfill the role you have assigned to it
- The hardware and software resources you already have that you want to preserve or upgrade
- The new hardware and software resources you need before you can implement your network plan
- The method to maintain your network and its resources over time
- The number and type of host connections and which workstations are to be attached to the hosts through the gateway and gateway definition

The steps in Chapter 15, “Implementing Communications Server Configurations”, on page 155 provide one procedure to determine this information for your

network. If you already have a procedure that works, you might want to use it instead, but you should review the procedure in the chapter before you make that decision.

Conventions Used in This Book

The following conventions are used in this book.

Text Conventions

| | |
|----------------------------|---|
| Bold | Bold type indicates the following: <ul style="list-style-type: none">• Verbs, functions, keywords, and parameters that you can use in a program or at a command prompt. These values are case sensitive and should be entered exactly as they appear in the text.• The names of window controls, such as lists, check boxes, entry fields, push buttons, and menu choices. |
| <i>Italics</i> | Italic type indicates the following: <ul style="list-style-type: none">• A variable that you supply a value for.• Book titles.• A letter is being used as a letter or a word is being used as a word. Example: When you see an <i>a</i>, make sure it is not supposed to be an <i>an</i>. |
| <i>Bold italics</i> | Bold italic type is used to emphasize a word. |
| UPPERCASE | Uppercase indicates constants, file names, and options that you can use in a program or at a command prompt. You can enter these values in uppercase or lowercase. |
| Double quotation marks | Double quotation marks indicate messages you see in a window. An example of this would be the messages that appear in the operator information area (OIA) of an emulator session. |
| Example type | Example type indicates information that you are instructed to type at a command prompt or in a window. |

Number Conventions

| | |
|---------------------|---|
| Binary numbers | Represented as B'xxxx xxxx' or B'x' except in certain instances where they are represented with text ("A value of binary xxxx xxxx is..."). |
| Bit positions | Start with 0 at the rightmost position (least significant bit). |
| Decimal numbers | Decimal numbers over 4 digits are represented in metric style. A space is used rather than a comma to separate groups of 3 digits. For example, the number sixteen thousand, one hundred forty-seven is written 16 147. |
| Hexadecimal numbers | Represented in text as hex xxxx or X'xxxx' ("The address of the adjacent node is hex 5D, which is specified as X'5D'.") |

Where to Find More Information

For more information, see *Quick Beginnings* which contains a complete description of both the Communications Server library and related publications.

The Communications Server home page on the Internet has general product information as well as service information about APARs and fixes. You can access the Communications Server home page at the following URL:

<http://www.ibm.com/software/network/commserver/>

What's New in This Book

The following new functions and enhancements were added for Version 6.1:

- Telnet enhancements
 - Configuration of multiple TN3270E server ports
 - Specification of a default pool to be associated with the TN3270E server port
 - Client authentication support for TN3270E and TN5250 servers to verify that a client is authorized to establish a secure connection to the server
 - Specification of the level of security supported for connections on the TN3270E and TN5250 server ports
 - Support for certificate revocation list (CRL) processing for TN3270E and TN5250 servers when client authentication is requested.

Part 1. Overview and Concepts

Chapter 1. Introducing Communications Server Functions

This chapter provides an overview of the communications functions supported by Communications Server and of the methods that you can use to configure your systems to accomplish these functions. Later chapters in this book describe the functions and the configuration methods in more detail.

Supported Functions

This section describes the following Communications Server functions:

- SNA communication support
- Advanced Program-to-Program Communications (APPC), which includes Advanced Peer-to-Peer Networking (APPN)
- 32-bit application programming interfaces (APIs)
- High performance routing (HPR)
- SNA data compression
- Discovery of service providers
- Conventional LU application (LUA)
- Dependent logical unit requester (DLUR)
- SNA gateway
- AnyNet support
 - AnyNet SNA over TCP/IP (access node and gateway)
 - AnyNet Sockets over SNA (access node and gateway)
- TN3270E server
- TN5250 server
- Legacy data access
 - iSeries OLE DB provider
 - iSeries shared folder server
- SNA API client support
- Novell NetWare for SAA client support
- Configuration and administration support
- Load balancing
- Data security
- Adapter support
- Data Link Control (DLC) support
- APPN MIB support
- Hot standby
- Entry-level emulator functions

SNA Communications Support

Communications Server can act as a Systems Network Architecture (SNA) type 2.0 and SNA type 2.1 node. This support lets you write programs to communicate with many other IBM SNA products.

APPC/APPN

Communications Server provides Advanced Peer-to-Peer Networking (APPN) end node and network node support for workstations, enabling them to communicate more flexibly with other systems in the network. Additionally, a branch extender function enables you to isolate branches to avoid unnecessary CP-CP traffic.

Communications Server provides advanced program-to-program communications (APPC) to support communications between distributed processing programs, called transaction programs (TPs). The TPs can be located at any node in the network that provides APPC. APPC uses the LU 6.2 protocol for exchanging data between programs located at different logical units (LUs). In addition, APPC supports multiple concurrent links and parallel sessions. Conversation or session security between the communicating programs is also supported through APPC.

Communications Server provides APPC throughput in performance-critical LAN environments. Communications Server supports the following connectivities:

Table 1. Supported APPC Connections

| Protocols | Connection | Connection Type |
|-------------|---|---|
| Channel | <ul style="list-style-type: none">• CDLC (OEM)• MPC | |
| LAN (802.2) | <ul style="list-style-type: none">• Token-Ring• Ethernet• Frame Relay• ATM Lan Emulation | |
| SDLC | <ul style="list-style-type: none">• Synchronous• Autosynchronous• Modems | <ul style="list-style-type: none">• Autodialed, DTR dialed, Permanent• Autodialed• Autodialed, Permanent• Autodialed |
| X.25 | <ul style="list-style-type: none">• Synchronous• Hayes AutoSync• X.25 network | <ul style="list-style-type: none">• Autodialed, Semipermanent• Autodialed, Permanent• Autodialed |
| IP | SNA over IP | IP-provided connections |
| IP | HPR over IP | IP-provided connections |

Note: APPC/APPN also has data compression capability. Refer to “Data Compression” on page 52 for more information on data compression. Also, refer to “SNA Session-Level Encryption” on page 55 for information on encryption.

SNA LU Type 6.2 Support

LU 6.2 is an architecture for program-to-program communications. Communications Server supports the following optional SNA LU 6.2 features:

- Basic and mapped conversations
- Synchronization level of confirmation
- Security support at session and conversation levels
- Multiple LUs

Introducing Communications Server Functions

- Parallel sessions, including the ability to use a remote system to change the number of sessions
- Transactions in either a one-way or two-way simultaneous mode
- Concurrent multiple links with SSCP-PU sessions, each of which can be connected to a subarea network

Refer to *Communications Server Programming Guide and Reference* for a complete list of features.

APPC Configuration Enhancements

Basic end node configuration requires as few as four parameters: network ID, local node name, link type, and destination address. System definition is reduced by:

- Modes supplied by IBM
- Implicit creation of partner LU and mode definitions when a conversation is allocated
- Implicit initialization of session limits (the CNOS function), which removes the requirement to explicitly initialize the session limits before conversations can be allocated
- Removal of session limits for local and partner LUs, which enables greater freedom in configuring and initializing mode session limits
- Defaults for transaction program names, operation, and type, which remove the requirement to configure transaction program definitions for applicable programs
- Dynamic configuration updates
- SNA gateway implicit workstation definitions
- Ability to edit an .ACG file containing configuration parameters
- Ability to find system definition information with Discovery

Because the SNA configuration is stored as a text file, you can quickly and easily modify the file by using an editor or a program written by a user. You can then verify and dynamically update an active configuration (without stopping Communications Server).

You can now configure connections to multiple hosts, and multiple host connections can be active at the same time. Connections can be set to start on demand, or as a result of a hot standby failure.

Communications Server support of Discovery enables a node to dynamically find the control point name, medium access control (MAC) address and service access point (SAP) address of another Communications Server network node server on a Token-Ring or Ethernet LAN. This means the user does not have to know the control point name, MAC and SAP address of a partner machine before being able to define a connection to it. Currently, iSeries Client Access/400, Host On-Demand 4.0 and Personal Communications users can use this facility.

Fundamental APPC Concepts

This section introduces APPC concepts and terms.

Transaction Programs: A *transaction program* (TP) is a program, or part of an application program, that uses APPC communications functions. Application programs use these functions to communicate with application programs on other systems that support APPC.

Introducing Communications Server Functions

Communications Server provides the APPC API and supports the IBM Systems Application Architecture (SAA) Common Programming Interface for Communications (CPI-C) calls for transaction programs.

Transaction programs issue APPC parameters to invoke APPC functions. A *parameter* is a formatted request that a transaction program issues and APPC executes. A program uses APPC parameter sequences to communicate with another program. Two programs that communicate with each other can be located at different systems or on the same system. The APPC API is the same in both cases.

When a transaction program exchanges data with another transaction program, the other transaction program is called a *partner* transaction program.

Transaction programs can issue CPI-C calls. These calls let application programs take advantage of the consistency that SAA provides.

Logical Units: Every transaction program gains access to an SNA network through a *logical unit* (LU). An LU is SNA software that accepts parameters from your programs and acts on those parameters. A transaction program issues APPC parameters to its LU. These parameters cause commands and data to flow across the network to a partner LU. An LU also acts as an intermediary between the transaction programs and the network to manage the exchange of data between transaction programs. A single LU can provide services for multiple transaction programs. Multiple LUs can be active in the node simultaneously.

LU Types: Communications Server supports LU types 0, 1, 2, 3, and 6.2. LU types 0, 1, 2, and 3 support communication between host application programs and different kinds of devices, such as terminals and printers.

LU 6.2 supports communications between two programs located at type 5 subarea nodes, type 2.1 peripheral nodes, or both, and between programs and devices. APPC is an implementation of the LU 6.2 architecture.

LU Sessions: Before transaction programs can communicate with each other, their LUs must be connected in a mutual relationship called a *session*. A session connects two LUs, so it is called an *LU-LU* session. Figure 1 illustrates this communication relationship.




Figure 1. A Session between Two LUs (LU-LU)

Sessions act as conduits that manage the movement of data between a pair of LUs in an SNA network. Specifically, sessions deal with things such as the quantity of data transmitted, data security, network routing, and traffic congestion.

Sessions are maintained by LUs. Normally, transaction programs do not work with session characteristics. You define session characteristics when you:

- Configure your system

Introducing Communications Server Functions

- Use SNA Node Operations
- Use the management parameters

Transaction Program Conversations: The communication between transaction programs is called a *conversation*. Like a telephone conversation, one transaction program calls the other, and they “converse”, one transaction program talking at a time, until a transaction program ends the conversation. A conversation starts when a transaction program issues an APPC parameter or CPI-C call that allocates a conversation. Conversations occur across LU-LU sessions.

Allocating a conversation to a session establishes a send-receive relationship between the transaction programs connected to the conversation. One transaction program issues parameters to send data. The other transaction program issues parameters to receive data. When the sending transaction program finishes sending data, it can transfer send control of the conversation to the receiving transaction program. Conversations can exchange control information and data.

Figure 2 shows a conversation between two transaction programs occurring over a session.




Figure 2. Transaction Program Conversation Occurring over a Session

A session can support only one conversation at a time, but one session can support many conversations in sequence. Because multiple conversations can reuse sessions, a session is a long-lived connection compared to a conversation. When a program allocates a conversation and all applicable sessions are in use, the LU puts the incoming attach (allocation request) on a queue. It completes the allocation when a session becomes available.

Two LUs can also establish parallel sessions with each other to support multiple concurrent conversations. A parallel session occurs when either transaction program allocates a conversation and a session exists, but is being used by a conversation. The LU can request a new session to satisfy the allocation.

Figure 3 on page 8 shows three parallel sessions between two LUs; each session carries a conversation.

Introducing Communications Server Functions




Figure 3. Parallel Sessions between LUs

Advanced Peer-to-Peer Networking (APPN)

Advanced peer-to-peer networking (APPN) is a set of functions, formats, and protocols that greatly enhances the managing of an SNA network and the usability of APPC applications running in the network. APPN does this through reduced configuration requirements, dynamic directory searches, route calculation capabilities, and intermediate session routing.

With APPN, you can write programs without knowing the details of the underlying network. All you need to know is the name of the partner LU; you do not need to know its location. SNA determines the partner LU location and the best path for routing data. A change to the underlying network, such as a physical address change, the addition of a new adapter, or the relocation of a machine, does not affect APPC programs.

Communications Server provides APPN end node and network node support for workstations, enabling them to communicate more flexibly with other systems in the network. Additionally, a branch extender function enables you to isolate branches to avoid unnecessary CP-CP traffic.

- End node capability provides the following services:
 - Connection of the local node to a serving network node for directory and routing services
 - Automatic reactivation of CP-CP sessions when an end node loses its network node server.
- Network node capability provides the following services:
 - Route selection services that calculate the best route for a session through an APPN network
 - Intermediate session routing that enables two LUs in nonadjacent nodes to be in session with each other, where the session traverses one or more intermediate nodes
 - Directory services that provide the capability to learn partner LU locations dynamically
 - Connection network services for the LAN that provide the capability to learn LAN destination addresses without having to configure them
This capability enables a node to establish a link connection directly to another node with no LAN destination address configured.
 - Branch extender support that enables you to efficiently interconnect a branch office with LANs, end nodes, and low end network nodes with dependent and independent LUs, and PUs such as teller machines, to one or several

WANs. This enables more network nodes to function as gateways to branch offices. A LAN at each branch office can be configured as a connection network.

Programming Interfaces

Communications Server supports a wide range of 32-bit application programming interfaces (APIs) on the server for the application program developer. These APIs provide convenient ways for application programs to access Communications Server functions and allow applications to address the communication needs of connections to both IBM and other computers. In addition, the provided interfaces support SNA protocols so that standardization is ensured.

The APIs supported include:

- Advanced program-to-program communications (APPC)
- Common Programming Interface for Communications (CPI-C)
- Conventional LU Application Interface (LUA) RUI and SLI
- Host Access Class Libraries (HACL)
- Java CPI-C (JCPI-C)
- WinSock (in conjunction with AnyNet Sockets over SNA)
- Network Operator Facility
- Management Services
- Common Services

On the clients, the Enhanced APPC (EHNAPPC) API is also provided.

The Communications Server Software Developers Tool Kit (which can be separately installed from the Communications Server CD-ROM) is also available for application developers to use. This tool kit contains samples, header files, library files, and online manuals for each of the APIs.

For more information on Communications Server programming interfaces, refer to *Client/Server Communications Programming* and *System Management Programming*.

High Performance Routing

High performance routing (HPR) is an enhancement to APPN that increases data routing performance and reliability and establishes a virtual link between rapid transport protocol (RTP) nodes. HPR replaces intermediate session routing, which is the routing technique used in APPN.

HPR provides faster transmission at intermediate nodes, nondisruptively reroutes sessions around failed nodes and links, and regulates traffic flow by predicting and reducing congestion in the network.

Communications Server supports HPR connections over Enterprise Extender (IP), synchronous data link control (SDLC), LAN, WAN, channel, Multi-Path Channel (MPC), and X.25 connections.

SNA Data Compression

Data compression at the session level increases throughput for large amounts of data across communication links, resulting in the following benefits:

- Enhanced data throughput on low-speed lines
- Reduced costs on high-cost lines
- Faster response times, resulting in productivity improvements

Introducing Communications Server Functions

SNA data compression is compatible with the zSeries and iSeries implementations and can be used with all LU types.

Discovery of Service Providers

Discovery is a LAN address resolution protocol that can be used by a node on the LAN to find another node that matches given search criteria. By adjusting the search parameter, a node can search for APPN network nodes, nodes that provide SNA boundary function, iSeries hosts, SNA gateways, or user-defined classes of server. A Communications Server for Windows server can respond to requests from clients as a network node server, a PU 2.0 gateway, or as a user-defined class of server. A Communications Server can also use discovery to find APPN nodes and SNA gateways.

Dependent LU Requester (DLUR)

Communications Server provides dependent LU requester (DLUR) end node and network node support for workstations, enabling them to take advantage of the enhanced system services control point (SSCP) support provided by a dependent LU server (DLUS). DLUS is supported by VTAM V4R2 and later. With this support, traditional SNA dependent LUs, such as emulators and even printers, can gain the many advantages of an APPN network.

Some of these benefits include:

- Enabling dependent LUs to reside on nodes that are nonadjacent to the host
- Using APPN search logic to provide the best path for LU-LU sessions

See “Dependent Logical Unit Requester Support” on page 49 for more information about DLUR.

SNA Gateway

A gateway permits communication between hosts that support PU 2.0 workstations and workstations that use different DLC types. An SNA gateway can do the following things:

- Enable workstations to share the same host link
- Provide multiple host connectivity to workstations that do not have multiple PU capability
- Act as a PU pool for dedicated PU downstream workstations
- Act as a protocol converter between workstations that use DLCs on their links that are different from the DLC that is used on the host link
- Reduce the amount of system definition at the host and workstations
- Enable you to dynamically change network definitions and add workstations
- Reduce host resources and reduce the number of host connections through the use of pooled LUs and by automatic logoff of sessions that are unused for a user-specified length of time
- Improve reliability by providing a host backup link

The SNA gateway enables a zSeries family host to support workstations that implement LU 0, 1, 2, 3, or dependent LU 6.2 (APPC). The SNA gateway also supports LU 0, 1, 2, or 3 to an iSeries host. The iSeries host passes the data through to a zSeries family host.

Each host views the SNA gateway as an SNA PU 2.0 node, supporting one or more LUs per workstation. As far as the host is concerned, all LUs belong to the SNA

gateway PU. The SNA gateway can have multiple host connections simultaneously and direct different workstation sessions to specified hosts. However, only one host (and it must be on a link with a CP PU) can act as the focal point, and the control point name is appended to all network management vector transports (NMVTs) routed through the gateway.

To the supported workstations, the SNA gateway looks like an SNA PU 4 communications controller and forwards such host messages as BIND and UNBIND. The network LUs are not aware of the SNA gateway. The SNA gateway, however, is aware of all LUs at the workstations.

In reality, the SNA gateway is a special type of PU 2.0. As long as a dependent workstation is inactive, the SNA gateway implements the LU functions for the workstation, just as a real PU 2.0 would. However, as soon as a workstation is online to the host, the SNA gateway enables the workstation to implement LU functions and merely passes data between workstations and the host.

An SNA gateway enables supported workstation applications to access remote supported applications on a subarea network without requiring a separate direct connection to each host in each workstation. From a host point-of-view, the host has a single connection to the gateway.

See Chapter 7, “Planning for SNA Gateway”, on page 95 for more information about using an SNA gateway.

Figure 4 shows an example of a connection using an SNA gateway.

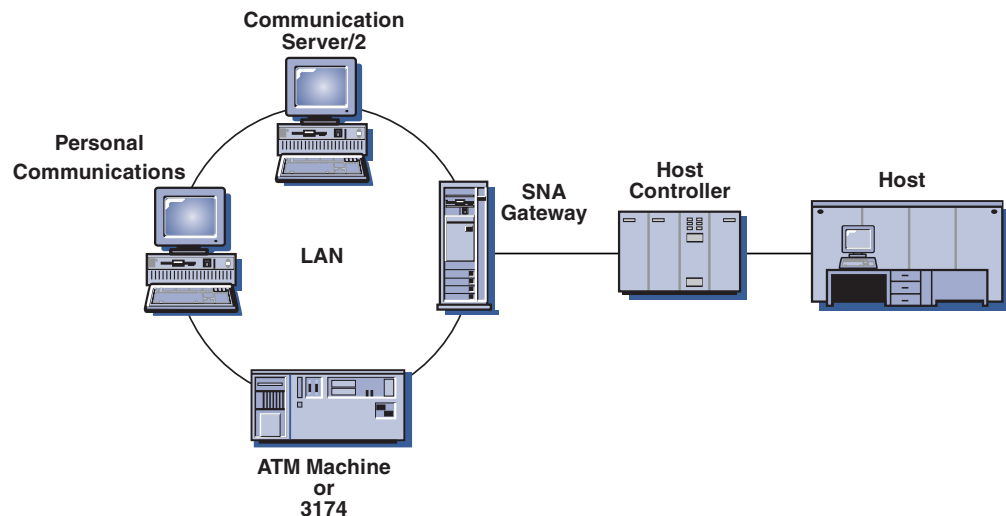


Figure 4. Example of SNA Gateway Connections

AnyNet Support

Communications Server incorporates the SNA over TCP/IP and Sockets over SNA functions from the AnyNet product family. This support enables you to extend and simplify your network by allowing SNA applications to communicate across a TCP/IP network and Sockets applications to communicate across an SNA network without changes to the applications.

SNA over TCP/IP

The SNA over TCP/IP access node function allows SNA applications residing on an IP network to communicate. This function supports independent LU 6.2 and

Introducing Communications Server Functions

dependent LU 0, 1, 2, 3, or 6.2 either with or without dependent LU requester (DLUR). In addition, the SNA over TCP/IP access node can be used in conjunction with SNA gateway to enable SNA gateway sessions over TCP/IP.

The SNA over TCP/IP gateway function extends the reach of SNA applications by allowing SNA applications in an SNA network to communicate with SNA applications in an IP network. The SNA over TCP/IP gateway supports independent LU 6.2 sessions.

Figure 5 shows SNA applications communicating through an SNA over TCP/IP Gateway across IP and SNA networks.

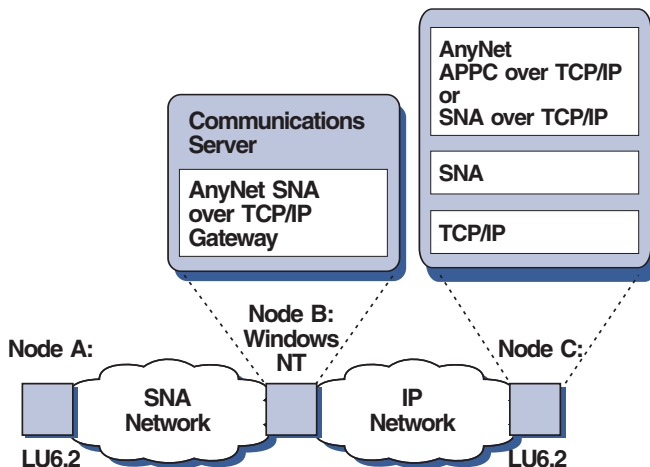


Figure 5. SNA over TCP/IP Gateway

Sockets over SNA

The Sockets over SNA access node function enables TCP/IP application programs using the WinSock 1.1 and WinSock 2.0 socket interface to communicate over an SNA network.

The Sockets over SNA gateway function enables sockets applications in SNA and TCP/IP networks to communicate. Sockets over SNA gateways are often used to connect isolated TCP/IP networks using an SNA backbone network.

Figure 6 on page 13 shows Socket applications communicating through a Sockets over SNA gateway across IP and SNA networks.

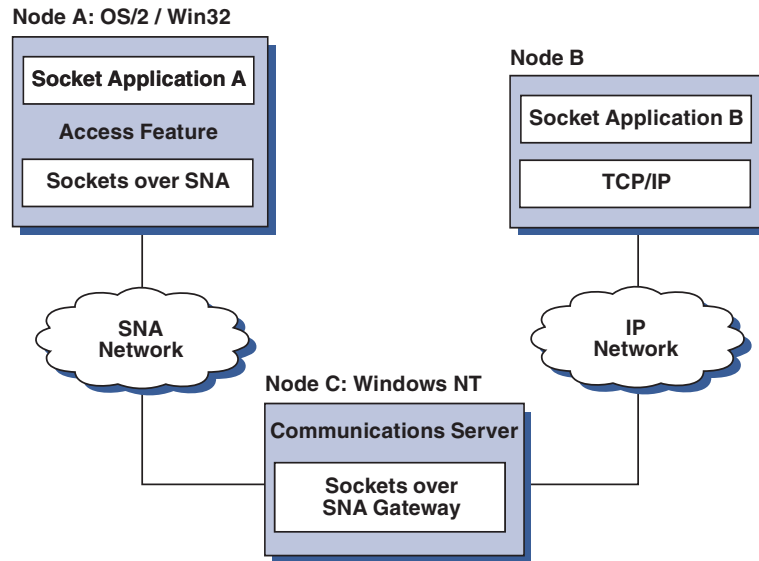


Figure 6. Sockets over SNA Gateway

TN3270E Server

The TN3270E server function enables TCP/IP users to access applications on a host machine in an SNA network. Any industry-standard TN3270 or TN3270E client workstation can connect to the TN3270E server for access to SNA networks. The TN3270E server supports ATTN and SYSREQ key handling and enables users to print from host applications to printers attached to their workstation. These printers may be locally-attached or network-attached.

Communications Server supports load balancing for client connections of a TN3270E server that connects to the same host resources, if the client is enabled for load balancing.

TN3270E server supports IP and hostname filtering that allows controlled access to LUs without modifying client configurations.

TN3270E server also supports Secure Sockets Layer (SSL) authentication and encryption, providing secure access across the TCP/IP network. If you specify security, the server must have an authenticated certificate provided by a certificate authority such as IBM Vault Registry or Verisign. Communications Server provides a utility that generates and manages keys and certificates used by SSL Version 3. For more information on using SSL authentication and encryption, refer to Chapter 10, "Planning for Secure Sockets Layer-based Security", on page 115.

Figure 7 shows an example of TN3270E server connections.

Introducing Communications Server Functions

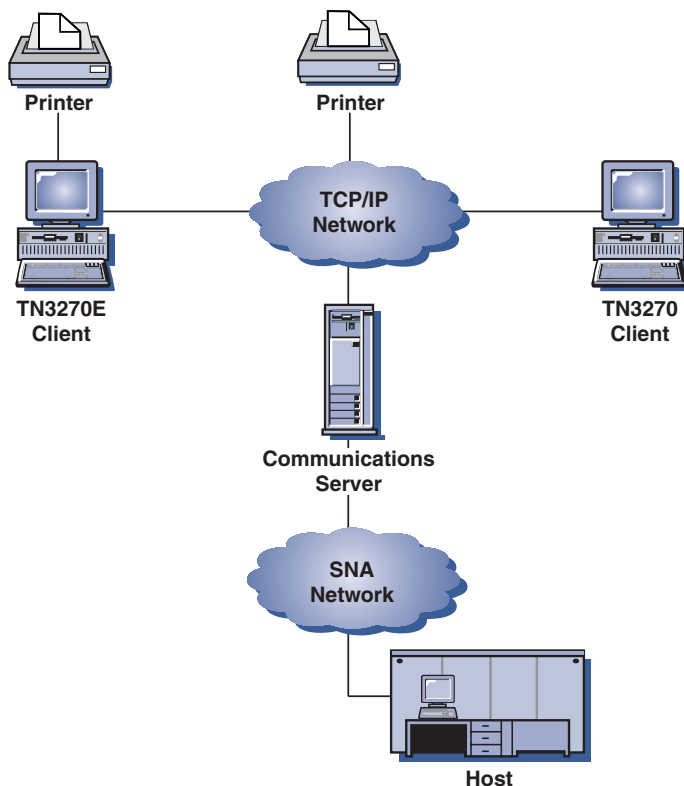


Figure 7. TN3270E Server Connections

TN5250 Server

The TN5250 server function enables TCP/IP users to access applications on an iSeries in an SNA network. Any industry-standard TN5250 client workstation can connect to the TN5250 server for access to SNA networks.

Communications Server supports load balancing for client connections of TN5250 servers that connect to the same iSeries hosts, if the client is enabled for load balancing.

TN5250 server supports IP and hostname filtering that allows central administration of client access to the server, as well as directing clients to specific iSeries hosts.

TN5250 server also supports Secure Sockets Layer (SSL) authentication and encryption, providing secure access across the TCP/IP network. If you specify security, the server must have an authenticated certificate provided by a certificate authority such as IBM Vault Registry or Verisign. Communications Server provides a utility that generates and manages keys and certificates used by SSL Version 3. For more information on using SSL authentication and encryption, refer to Chapter 10, "Planning for Secure Sockets Layer-based Security", on page 115.

Figure 8 shows an example of TN5250 server connections.

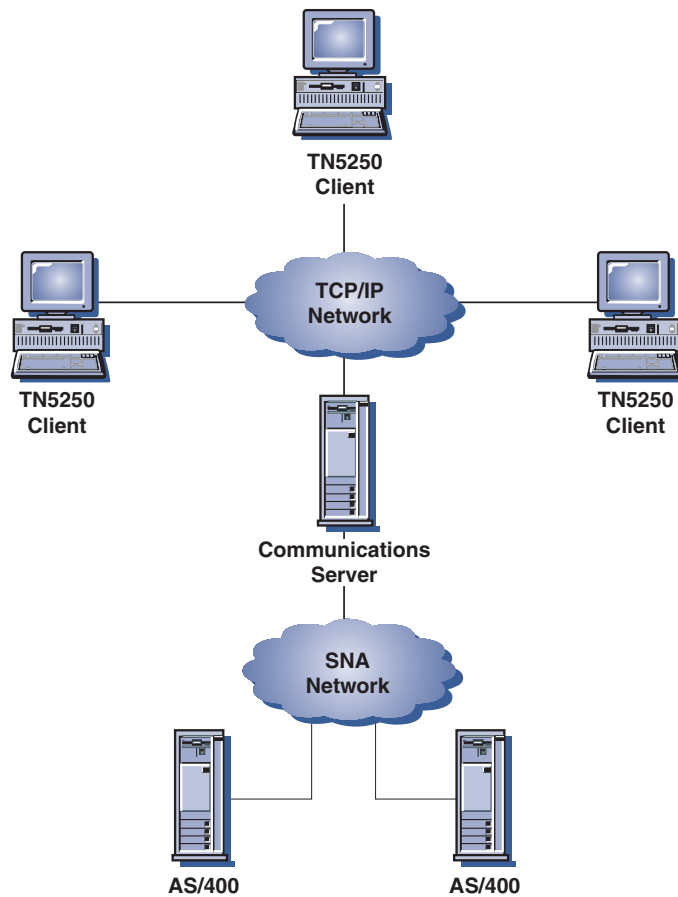


Figure 8. TN5250 Server Connections

Legacy Data Access

Communications Server provides access to data on host machines, iSeries hosts, and workstations on SNA networks through the following functions:

- iSeries OLE DB provider
- iSeries Shared Folder Server

iSeries OLE DB Provider

Applications that use OLE DB or ActiveX can communicate through Communications Server for record-level access to files on an iSeries. Documentation for this function, as well as information about developing these applications using Client Access, is provided in the `ibmcs\sdk\pubs\oledb` directory.

iSeries Shared Folder Server

The iSeries system uses a structure called a folder to store and organize documents, mail, and other related objects. Communications Server enables you to create disk devices on the server that communicate with iSeries folders through the iSeries Integrated File System (IFS). Additionally, if the server shares these disk devices, clients can NET USE to them. Multiple clients can connect to folders on the iSeries system as if they were drives on their workstations.

You can use shared folders to:

Introducing Communications Server Functions

- Use iSeries security to limit access to workstation files
- Share data with multiple users at the same time
- Back up workstation files to an iSeries folder

Communications Server Client Support

Communications Server provides support for SNA API clients (available on the CD-ROM,) and for Novell NetWare for SAA clients.

SNA API Client Support

The Communications Server SNA API client support allows TCP/IP- and IPX-attached clients to access SNA APIs without requiring SNA protocols to flow between the clients and the server. This allows most SNA configuration to take place at the central server.

Communications Server supports SNA API clients on Windows 95, Windows 98, Windows NT, Windows 2000, Windows 3.1, and OS/2.

The SNA clients provide support for CPI-C, APPC, EHNAPPC, LUA RUI, JCPI-C, and HACL API interfaces, while providing the actual SNA processing at the server. These clients are delivered as part of the server but are actually installed and configured at the client.

The 32-bit Windows and OS/2 clients have additional enhancements:

- Encryption between the client and server
- More detailed trace formatting that can help you debug problems faster
- Ability to start and stop tracing while the client is running
- Ability to reinstall the client without uninstalling

The Windows 95, Windows 98, Windows NT, and Windows 2000 clients run from the same executable. This executable can be installed on a shared drive; any fixes apply to all clients. The 32-bit Windows client can communicate with either Communications Servers or Novell NetWare for SAA servers.

For more information about the API clients in Communications Server, see Chapter 4, “Planning for Client/Server Communication”, on page 85.

Novell NetWare for SAA Client Support

Communications Server supports IPX- or TCP/IP-attached clients running emulator software packages that implement Novell’s Queue Element/Message Unit (QEL/MU) architecture for 3270 emulation, enabling the clients to access mainframe host data. This includes support for popular client features, including dedicated, pooled, and public LU categories (sometimes referred to as resource types).

Communications Server supports Novell NetWare for SAA clients on Windows 95, Windows NT, Windows 3.1, and OS/2.

Refer to *Novell NetWare for SAA 3270 Client Interface Guide and Reference P/N 100-002018-001* for more information about developing these clients.

Load Balancing

Communications Server now supports load balancing for all client types. Load balancing enables you to distribute LU 0 to 3 and LU 6.2 sessions across

Communications Server and NetWare for SAA servers. The server advertises services including load factors, which the clients or servers can gather and organize to select a server.

Directory Exploitation

You can configure multiple clients from a central location using the Lightweight Directory Access Protocol (LDAP) to simplify the configuration process.

For information about using directory exploitation, refer to Chapter 4, “Planning for Client/Server Communication”, on page 85.

Configuration and Administration Support

Communications Server provides facilities for the configuration and administration of resources.

Communications Server Configuration

This section provides an overview of the Communications Server configuration components and the methods used to create or change them. The configuration is composed of a single file (ACG) stored in the PRIVATE subdirectory of the directory where you installed the product (for example, C:\IBMCS\PRIVATE). The ACG file can be created using **SNA Node Configuration**. You can modify the ACG file using **SNA Node Configuration**, **Web Administration**, or by using an ASCII editor. A verification program is available to check the validity of the ACG file prior to use.

You can use the following methods to create or change a Communications Server configuration:

- **SNA Node Configuration**
- **Web Administration**
- Response file configuration

Node Configuration: Communications Server provides the **SNA Node Configuration** application (PCSCFG) that enables you to configure the Communications Server functions using a graphical interface and supplies defaults so you can configure them easily using a minimum number of parameters. During configuration, each definition you create is checked to ensure that it is valid. When you save the configuration, the required configuration files are created.

Local configuration is supported at both the client and server level. Remote configuration of the server is supported from Windows 95, Windows 98, Windows NT, and Windows 2000 clients.

Most configurations can be created using **SNA Node Configuration**. However, a few keywords and some keyword parameters are not supported by **SNA Node Configuration**.

Web Administration: **Web Administration** provides the capability to modify your configuration.

Response File Configuration: Response file configuration enables you to customize a template configuration file to meet the needs of specific users. For more information on using response files for configuration, refer to “Configuration with Template and Response Files” on page 164.

SNA Node Operations: **SNA Node Operations** provides the capability to create and modify select resources.

Introducing Communications Server Functions

Communications Server Administration

Communications Server provides the following facilities for administration of resources.

- SNA Node Operations
- Command line utilities
- Web Administration
- Tivoli Plus Module

For more information about the capabilities of these facilities, refer to Chapter 16, “System Management Facilities”, on page 175.

Load Balancing

Load balancing is a function of Communications Server that dynamically balances dependent LU (host-to-workstation) sessions and independent LU 6.2 sessions by distributing them to the communications server with the smallest load.

Communications Server performs load balancing for Communications Server API programs and third-party 3270 emulators that connect over TCP/IP protocols, or third-party TN3270 and TN5250 emulators. The resources across which balancing occurs depend on the session type:

- For dependent LU sessions, the load is distributed across servers within a named LU pool and a named scope.
- For LU 6.2 sessions, the load is distributed across all available servers in a named scope or across selected servers within a named server list.

The load balancing capabilities of Communications Server are built into the SNA client APIs. Load balancing is configured for the clients using **SNA Client Configuration**.

For dependent LU sessions, emulators that use SNA client APIs can participate in load balancing. Otherwise, you must purchase third-party 3270, TN3270, or TN5250 emulator software that supports load balancing.

For LU 6.2 sessions, the initial connection established by a SNA API client determines which server manages all subsequent LU 6.2 sessions.

For more information on load balancing, refer to Chapter 11, “Planning for Load Balancing”, on page 123.

Data Security

Communications Server provides basic and enhanced security support at session and conversation levels. There is security in limiting which Windows users may access SNA resources through the SNA API clients. Conversation security includes support for password substitution. There is also enhanced LU-LU security.

Communications Server provides support for Secure Sockets Layer-based (SSL-based) security on connections between TN clients and the TN3270E server or TN5250 server. This security uses SSL Version 3 to provide data encryption and server authentication using signed certificates.

Adapter Support

Communications Server provides an open interface for adapter manufacturers to build connectivity solutions. A shallow (nonprogrammable) adapter interface is provided for adapter manufacturers to work with Communications Server’s SDLC

and X.25 protocol stack. A deep (programmable) adapter interface is provided for adapter manufacturers to build connectivity solutions using manufacturer-supplied data link controls.

Data Link Control (DLC) Support

Communications Server enables communications over the following DLCs:

- AnyNet (SNA over TCP/IP)
- Twinaxial (upstream only)
- LAN (Any NDIS compliant network adapter)
- X.25 SDLC (synchronous, asynchronous, and AutoSync)
- OEM (third-party adapter support)
- Channel (upstream only)
- Multi-Path Channel (HPR only)
- Enterprise Extender

AnyNet (SNA over TCP/IP) DLC

For more information about the AnyNet SNA over TCP/IP DLC, refer to “SNA over TCP/IP” on page 11.

IBM Multi-Path Channel DLC

The Multi-Path Channel (MPC) DLC provides high-capacity, high-availability fiber connections to one or more zSeries MPC-capable hosts over the ESCON channel adapter card (P/N 9663 001). MPC connections provide high data transmission rates with transparent backup when physical connections break or become temporarily unavailable. This channel-to-channel connection enables you to provide LAN clients ready access to zSeries resources and services.

Enterprise Extender DLC

Communications Server now provides HPR connections on IP networks, using UDP/IP packets. To the HPR network, the IP backbone appears to be a logical link. To the IP network, the SNA traffic appears to be UDP datagrams. These datagrams are routed without changes to the IP backbone. Because there is no protocol transformation and because packaging takes place at the routing layer without the overhead of additional transport layers, this results in efficient use of the intranet infrastructure for IP clients that access SNA-based data (TN3270 clients or Web browsers using IBM Host on Demand, for example), as well as for SNA clients.

APPN MIB Support

Communications Server supports simple network management protocol (SNMP) requests for APPN management information from any SNMP management system.

Hot Standby

In Communications Server, you can configure certain host links to activate automatically if a specified critical server fails. Configured connections to a host can continue to function by activation of alternative connections on a backup server. This function is known as hot standby.

The connections named in a critical server configuration on the backup server are activated when the backup server detects a loss of contact with the critical server and licensing charges for the critical server are managed on the backup server.

Introducing Communications Server Functions

Note: The hot standby feature only provides for activation of host connections on a backup server and depends on the use of emulator software that supports alternate routing to the backup server when a critical server becomes inactive.

For more information about using hot standby for backup connections, refer to Chapter 12, “Planning for Backup Host Connections”, on page 127.

Entry-Level Emulator Functions

Communications Server includes an entry-level version of the popular Personal Communications 3270 and 5250 emulator for administrative purposes. This emulator provides basic 5250 and 3270 support on the server that includes a subset of the features and functions that are in the full-function IBM Personal Communications family of emulators.

The entry-level emulation functions provided include:

- Color mapping
- Command line transfer (3270 only)
- Full font set
- Screen sizes mod 2-5
- Two sessions

Although graphical keyboard remapping is not supported for the entry-level emulator, you can use the remap files generated by the full-function emulator.

Chapter 2. Communications Server and SNA

This chapter discusses the SNA network functions provided by Communications Server and includes the following topics:

- Overview of SNA functions
- Data link control (DLC) definitions
- Advanced Peer-to-Peer Networking (APPN)
- High Performance Routing (HPR)
- LU support
- Data compression
- Encryption
- Management services (MS)
- Flow control
- SNA gateway support

Overview of SNA Functions

This section provides an overview of how Communications Server implements SNA on a workstation. It is not a comprehensive discussion of the SNA functions. For more detailed information on SNA, refer to the following books:

- *Systems Network Architecture Concepts and Products*
- *Systems Network Architecture Technical Overview*
- *Systems Network Architecture Network Product Formats* (available online in the Communications Server documentation list)
- *Systems Network Architecture Format and Protocol Reference Manual: Architecture Logic for LU Type 6.2*

SNA defines the standards, protocols, and functions used by devices in the network (from mainframes to terminals) to communicate with one another. This communication enables these devices to transparently share information and process resources. In other words, a user at a workstation does not have to know what happens in the background to access information at a host or to communicate with another user.

An SNA network is organized as a system of nodes and links. It is important to remember that this organization is logical. SNA classifies the nodes according to their capabilities and the amount of control that they have over other nodes in the network. The node type is not necessarily associated with a specific type of hardware. The capabilities of a node can be performed by different devices. A workstation acting as a gateway can perform the same functions as a communications controller. It is even possible for a single device to contain more than one node.

The SNA network is responsible for moving data between two end users in an efficient, orderly, and reliable manner. For example, when a user at a workstation sends a message to another workstation, SNA does the following things:

- Packages the message in a form that can be used on the link type
- Addresses the package

Communications Server and SNA

- Selects a route for the message
- Monitors the communication to ensure that the message arrives at its destination
- Converts the message back to a form that can be used at the destination

These tasks are defined in SNA as separate functional layers. These layers are not part of the discussion here, but it is important to remember that the SNA layers are all part of a logical link.

As mentioned previously, the SNA nodes are classified according to their communications capabilities and the amount of control the node has over other nodes in the network. The SNA nodes are broadly classified as subarea nodes and peripheral nodes. The subarea nodes are like hubs and can communicate with the peripheral nodes and with other subarea nodes. The subarea nodes activate and control resources at the peripheral nodes. Subarea nodes are also classified as type 4 or type 5 nodes. Type 5 nodes contain a system service control point (SSCP) that provides a central control point for the type 4 nodes attached to it. A type 5 node is sometimes referred to as a host node. A peripheral node can only communicate directly with the subarea node that it is attached to. However, a peripheral node can control devices to which it is attached. For example, a cluster controller acting as a peripheral node can support the terminals that are attached to it. Peripheral nodes are also referred to as type 2 or type 2.1 nodes.

Each node contains network accessible units (NAUs) that perform control and communication functions. One of these NAUs is a physical unit (PU). The PU manages the physical resources of the node. Other NAUs called logical units (LUs) provide logical access points to the network that enable communication between users and applications at each node. The communication between the logical units is referred to as a session. Sessions not only support communication between users and applications, but also support communication between applications in order to share processing resources. Communication between applications is known as advanced program-to-program communication (APPC). APPC is a set of programming conventions and protocols that implement LU 6.2. (APPC is the name given to the LU 6.2 capability in products that implement this LU type.)

Data Link Control Profiles

Data link control (DLC) enables orderly exchanges of data between two nodes through a logical link. The DLC provides the protocols necessary for reliable delivery of basic transmission units (BTUs) between a pair of nodes in the SNA network. You must configure the appropriate Communications Server DLC profiles for a workstation to access an SNA network.

Refer to the online *Tutorial* for more information on configuring the appropriate DLC profiles.

Advanced Peer-to-Peer Networking

Advanced Peer-to-Peer Networking (APPN) is an extension of SNA that adds communications functions beyond those described in the previous sections. Its basic components include:

- APPN node types
- Control points
- Branch extender
- Data link control
- Logical links

- Directory services
- Topology and route-selection services

APPN Node Types

This section discusses the three types of nodes implemented by Communications Server that can participate in an APPN network:

- Network node
- End node
- Low entry networking (LEN) node

In addition, branch extender is an extension to a network node. Though it provides services to the end nodes as a network node, it appears in the network like an end node connected to other network nodes.

Each node is distinguished from other nodes in the network by a unique name consisting of two parts—a network ID and a local node name (also known as a control point [CP] name). The name identifies each node to all other nodes in the network. Also, the node can have multiple PU names for simultaneous access to multiple PU T4/5 hosts.

A node can be configured to be an end node or a network node, but when an end node does not have CP-CP sessions (see “CP-CP Sessions” on page 25) to an APPN network node, it acts as a LEN node. A LEN node does not support APPN functions.

The node types are described in more detail in the following sections. Figure 9 on page 24 illustrates a sample APPN network that includes all of these node types.

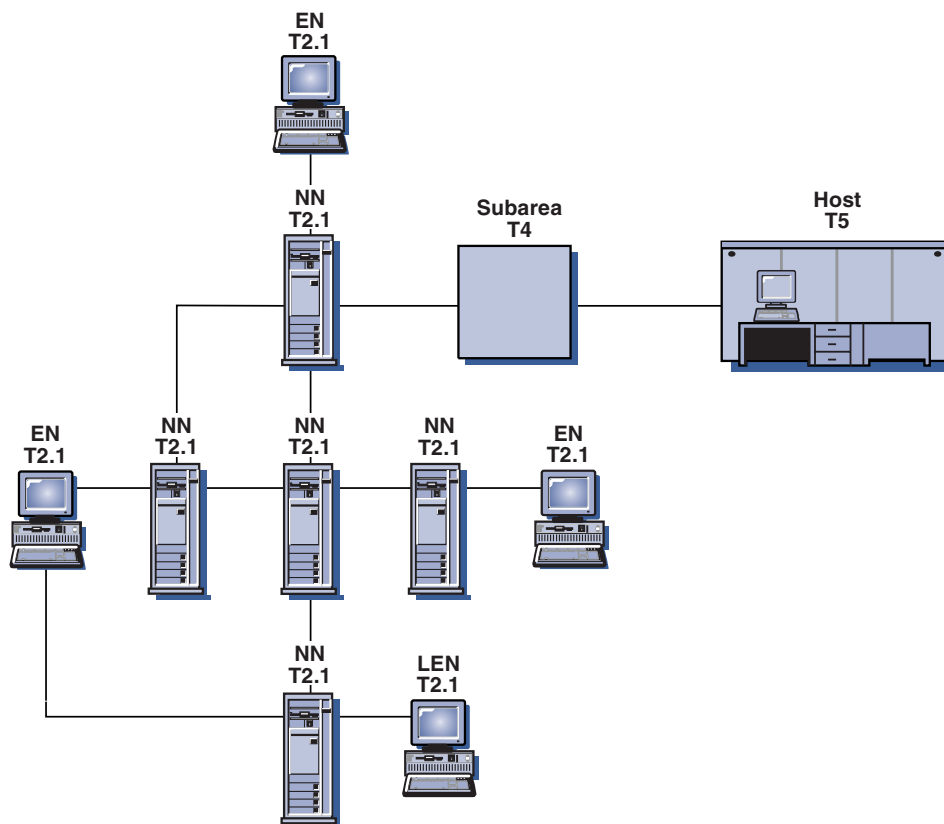


Figure 9. A Portion of a Sample APPN Network. This view of network topology shows five network nodes (NNs). Three end nodes (ENs) are connected, as well as a LEN node and subarea. APPC application programs written for any node in this network can communicate with any other.

Network Node

A network node supports its own end users; provides directory, route selection, and management services to end nodes; and performs intermediate routing of data on sessions that traverse it. The network node performs distributed searches of the network to locate partner LUs and calculates the best route from origin node to destination node based on user-specified criteria.

A network node server refers to the role of a network node in acting as the network entry point for specific end nodes attached to it. These end nodes are defined as being in its domain. For example, all directory requests regarding resources (such as LUs) in these end nodes (as well as its own resources) pass through directory services in the network node server. The network nodes are able to collect and control directory information that passes into the APPN network.

A network node provides the following:

- LU-LU session services for its locally resident LUs
- Intermediate session routing
- Network server functions (to perform network directory searches and route selection) for attached end nodes or LEN nodes (as well as for its own locally resident LUs)

- Management services (MS) routing capability for forwarding management services data (such as alerts) between a served end node and a management services focal point (FP).

End Node

An end node operates in a peer environment for LU-LU sessions (using LU 6.2 protocols) while providing additional APPN functions. An end node provides APPN functions such as directory services and route selection services to end users at its own node. It can participate in the APPN network by using the services of an attached network node server for session requests that involve nodes not directly connected; it does this by exchanging requests and replies for directory services with an adjacent network node (its server) using CP-CP sessions.

APPN end nodes can register their local LUs with their network node server. By registering the end nodes, the network operator at the network node server does not need to predefine the LU names for the LUs in all the attached end nodes for which the network node provides services.

An APPN end node can be attached to multiple network nodes, but it can only have CP-CP sessions active with one network node at a time—its network node server. The other network nodes can be used to provide intermediate session routing for the end node, or as a substitute network node server if the main network node server becomes unavailable. CP-CP sessions are never established between two end nodes.

LEN Node

A LEN node is a node that implements the basic T2.1 protocols without the APPN enhancements. In a LEN node, all potential connections with partner LUs are predefined before initiating sessions to them. A LEN node, connected to an adjacent APPN network node, uses the advanced functions of APPN by predefining potential connections with partner LUs as if they existed at that network node. The network node, in turn, can automatically act as the LEN node's network node server and locate the actual destination of the partner LU and select the best route to it. By going through a network node, the LEN node can participate in an APPN network without requiring direct connections to all nodes.

Control Points

The control point (CP) is responsible for managing the node and its resources. To obtain APPN network services, the control point in an APPN end node must communicate with the control point in an adjacent network node. Also, to manage the network, the control point in an APPN network node must communicate with the control points in adjacent network nodes. The control point directs such functions as adapter activation and deactivation and link activation and deactivation, and assists LUs in session initiation and termination.

When setting up a workstation, you must define the control point name (also known as the local node name). The control point is also an LU, and you can choose to have the control point LU be the only LU defined in your workstation.

CP-CP Sessions

To perform directory services and topology and route-selection services, adjacent nodes throughout the APPN network use a pair of parallel CP-CP sessions to exchange network information. Network nodes use CP-CP sessions to monitor nodes in a network link, as well as to track directory and session services. A network node establishes two parallel sessions with each adjacent network node

Communications Server and SNA

and with each served end node. An APPN end node establishes two parallel sessions with a single adjacent network node acting as its current server. LEN nodes do not support CP-CP sessions.

After a connection has been established, the nodes exchange identification information (XID). Then, CP-CP sessions are started between the control points in the directly attached nodes. The CP-CP sessions use LU 6.2 protocols and both sessions of a given pair must be active for the partner control points to begin and sustain their interactions. All CP-CP sessions are used to conduct directory searches.

After the CP-CP sessions are established, the two nodes exchange control point capability messages that inform each node of the other's capabilities. When both nodes are network nodes, they exchange topology database update (TDU) messages. The TDU messages contain identifying information, node and link characteristics, and resource sequence numbers to identify the most recent updates for each of the resources described in the TDU.

CP-CP Connection Activation

When Communications Server is started, it attempts to activate the preferred NN server connection first. Communications Server tries to activate all other connections defined as *activate at startup*. If an alternate parallel link to the preferred NN server connection exists, rather than waiting for the results of the activation attempt for the preferred NN server connection before attempting an alternate link, Communications Server attempts to activate the CP-CP sessions on the alternate link.

Note: If the connection was deactivated by operator request from the local node, CP-CP sessions are not redriven. If the connection was deactivated by operator request from the remote node, CP-CP sessions are redriven at the local node. For links between NN nodes, only demand activated (links with adjacent CP name specified and it is not defined as *activate at startup*) links are activated.

CP-CP Connection Reactivation

Communications Server provides support for CP-CP connection reactivation. Loss of CP-CP sessions between an end node and its network node server and between adjacent network nodes can interfere with the operation of an APPN network. CP-CP connection reactivation support improves reliability of an APPN network by reestablishing these important sessions when they are terminated due to failure or connection inactivation.

A CP-CP connection reactivation attempt is initiated by a CP-CP link activation, CP-CP session failure, or by a CP-CP retry timer expiring. CP-CP sessions are initiated by Communications Server with the first of the following:

1. The preferred server CP-CP (if it has not previously been attempted).
2. If DLUR is configured, use an adjacent CP that supports DLUR registration.
3. The last (most recent) activation of NNs that has not been attempted, to which an active connection exists.
4. The first of any other uplevel NNs for branch extender.

Note: If the CP-CP connections were terminated due to a link failure, Communications Server does not reactivate the link. You can configure a connection as *auto-reactivate* (infinite retry) to keep important connections active.

Branch Extender

The branch extender is a border node subset that is designed to interconnect a branch office to an APPN WAN backbone network. The interconnected networks can be native (that is, they have the same network ID) or nonnative. A node that supports the branch extender is a branch network node that typically has LAN and WAN interfaces, and can also include DLUR and HPR.

Links at a node that support branch extender are defined *branch uplinks* or *branch downlinks*. Figure 10 provides an example of the way that a branch network node works in a network. In this figure, the node in the center is a branch network node. Usually, the adjacent CP (branch uplink node) will be the network node server (NNS) for the branch network node, which looks like an end node to the branch uplink node.




Figure 10. Conceptual Overview of Branch Uplinks and Branch Downlinks

Branch uplinks are defined at the branch network node as upstream to the backbone network. You can consider a node with an uplink to be peripherally attached to the backbone network.

Branch downlinks are defined from the branch network node as downstream. The node sees downlinks as connections to end nodes (control points) in the domain. Branch downlinks are typically LAN links (but are not required to be). You can consider end nodes attached through branch downlinks to be local resources. The branch network node is the network node server for these end nodes. On branch downlinks, it provides network node services for domain end nodes, LEN end nodes, dependent T2.0 nodes and T2.1 nodes, and local LUs and PUs.

A branch network node works as a network node server for its domain. It maintains topological information about all of its branch downlink nodes, but does not maintain complete information about the entire uplink network. If the information the node has is not sufficient, it passes the LOCATE requests to its uplink network node server, which may be another branch network node or an APPN network node.

The branch extender optimizes the peer-to-peer communication environment for administrators who want to connect LAN-based branches to one large WAN primarily based on a switched network. The branch extender enhances performance in large APPN networks. Specifically, it:

Communications Server and SNA

- Reduces the number of network nodes in large APPN networks, enabling you to add additional branch networks
- Hides branch topology information
- Enables peer-to-peer communication between branches connected to the same APPN connection network
- Enables coexistence with PU gateway servers
- Reduces uplink CP-CP session traffic (WAN traffic)
- Isolates the branch network from backbone WAN traffic overhead

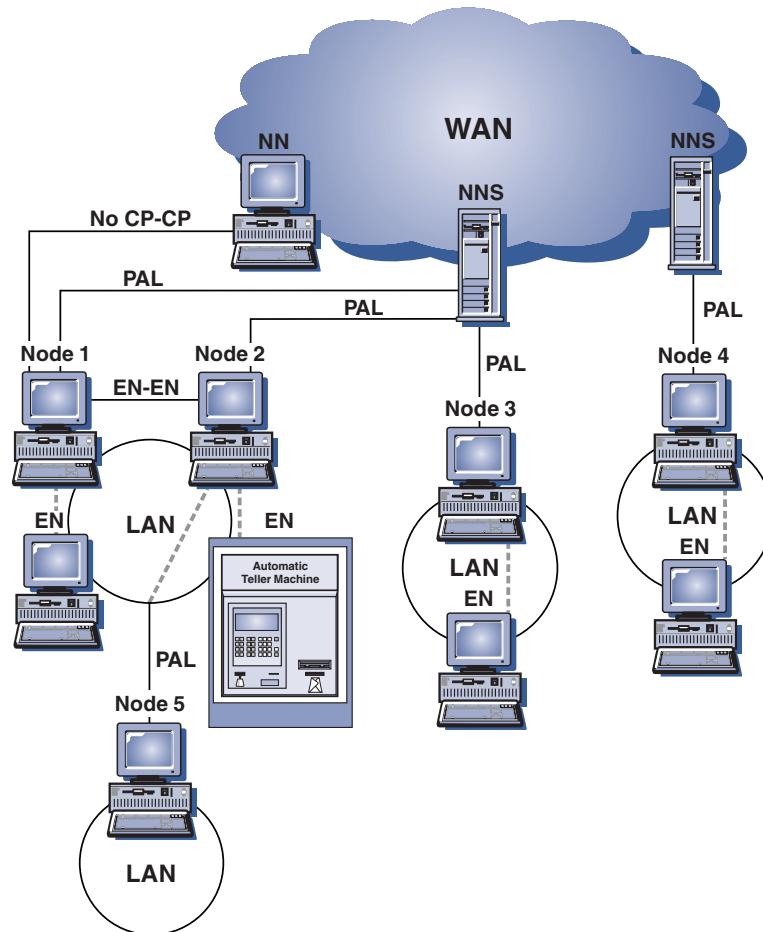


Figure 11. The Branch Extender in a Network

Figure 11 shows how branch network nodes work in a network. Dashed lines represent logical links. In the figure, nodes 1, 2, 3, and 4 are configured to support the branch extender and function as both end nodes and network nodes. They hide their downlink topology from the WAN network (that is, to the upstream network node servers, they appear as end nodes). To the nodes on the LANs downstream, the nodes function as network node servers. Node 5 appears as an end node to Node 2, but is a network node server (NNS) for other end nodes on its LAN. To the NNS, Node 5 appears to be an LU on Node 2.

When an end node served by Node 1 attempts to establish a session to an end node across the WAN, Node 1 can send a Locate (Send) request on its branch extender link to its network node server. If the target CP is found, the network

node server determines a route from Node 1 to the target CP. Node 1 will modify the route before returning it to the source end node. The source end node uses this route for its session.

Branch Extender Restrictions

The following restrictions apply to networks configured to use the branch extender:

- A branch network node can only connect to a network node over a branch uplink.
- To prevent search looping, downlink end nodes must be real end nodes, not nodes presenting an end node image.
- A node cannot have CP-CP sessions with the branch network node through both branch uplinks and branch downlinks simultaneously.
- The branch network node can have CP-CP sessions over only one branch uplink at a time; that is, it can have CP-CP sessions with at most one network node server.
- Branch extender does not provide Twinax support.
- A branch extender node cannot have DLUR sessions downstream.

Branch Extender Configuration

To configure the branch extender, you must first configure a branch network node. Then you must configure either a DLC (for an implicit link) or a branch network node. A link configured to support the branch extender is a branch uplink. On a branch network node, any links that are not configured to support the feature are branch downlinks.

If branch network nodes have links defined between one another, they must be defined to be *peer* connections, which give them a link type of **LEARN** in the ACG file. Alternatively, you can define the link as **ACTIVATE_AT_STARTUP=1** so the link is always active. When branch nodes connect to each other, loops in the topology occur. This is acceptable as long as the links are always active or the links are learned as they are activated. If links are defined between branch network nodes as **END_NODE** or **NETWORK_NODE** links, the topology reported upstream might interpret the links inappropriately and cause allocation failures.

You can also use an ACG file to configure branch extender.

Branch Extender Administration

The following sections describe how to verify a configuration and restrictions on how you configure your network.

Verifying the Configuration: You can use **SNA Node Operations** at a node that supports the branch extender to determine whether a local branch has been configured successfully. At run time, a display of the topology from the node should never have more than two network nodes, itself and the uplink network node server.

Note: Only one uplink is available for CP-CP. Each workstation should be configured as an end node with the branch network node defined as its preferred network node server.

You can use **SNA Node Operations** to verify whether the DLC or link has been configured to support branch extender correctly. You can also use **SNA Node**

Communications Server and SNA

Operations to determine whether an active link is a branch uplink or branch downlink. Downstream end nodes registered using AnyNet will not register their resources.

Supported Functions

Communications Server supports all the APPN Version 2 base functions (both end node and network node). Additionally, the following options are supported:

- End Node

| Function Set | Option |
|---------------------|--|
| 162 | MS_CAPS Be Sphere of Control (SOC) EN |
| 171 | Problem Diagnosis Data in Alert |
| 177 | LAN Alert |
| 178 | SDLC/LAN LLC Alert |
| 181 | X.25 Alert |
| 182 | Held Alert for CPMS |
| 1002 | Adjacent Link Station Name |
| 1007 | Parallel TGs |
| 1011 | Multiple Local Independent LUs |
| 1012 | CP=LU |
| 1018 | Delete EN Resources Before Registering |
| 1067 | Dependent LU Requester |
| 1070 | Session Cryptography |
| 1107 | Central Resource Registration (of LUs) |
| 1116 | DLUS-Served LU Registration |
| 1200 | Tree and TG Caching |
| 1400 | HPR Base (ANR) |
| 1401 | Rapid Transport Protocol |
| 1402 | Control Flows over RTP |
| 1520 | Common Operations Services |

- Network Node

| Function Set | Option |
|---------------------|------------------------------------|
| 164 | MS_CAPS Have a Subarea Focal Point |
| 171 | Problem Diagnosis Data in Alert |
| 177 | LAN Alert |
| 178 | SDLC/LAN LLC Alert |
| 181 | X.25 Alert |
| 182 | Held Alert for CPMS |
| 1002 | Adjacent Link Station Name |
| 1007 | Parallel TGs |

| | |
|------|--|
| 1011 | Multiple Local Independent LUs |
| 1012 | CP=LU |
| 1018 | Delete EN Resources Before Registering |
| 1067 | Dependent LU Requester |
| 1070 | Session Cryptography |
| 1100 | Safe-Store of Directory Cache |
| 1101 | Preload Directory Cache |
| 1107 | Central Resource Registration (of LUs) |
| 1116 | DLUS-Served LU Registration |
| 1118 | EN TG Vector Registration |
| 1121 | Branch Extender Function |
| 1200 | Tree and TG Caching |
| 1203 | Detection and Elimination of TDU Wars |
| 1301 | Nonpaced Intermediate Session Traffic |
| 1400 | HPR Base (ANR) |
| 1401 | Rapid Transport Protocol |
| 1402 | Control Flows over RTP |
| 1520 | Common Operations Services |

Data Link Control (DLC)

The DLC provides the protocols necessary for reliable delivery of basic transmission units (BTUs) between a pair of nodes in the APPN network and maintaining the logical connections between nodes.

Connections

A connection links a pair of adjacent nodes across the underlying DLC.

Parallel Links

Your local node can have multiple links to an adjacent node. This association is referred to as *parallel links*. Parallel links are each assigned a unique number (transmission group number) and can have different link characteristics assigned to them. To have two parallel links between two nodes, the link stations for the links can be on a single adapter in one node but must be on separate adapters in the other node; that is, the combination of adapter number and adjacent (or destination) link station address must be unique for each link.

Link Activation

The message unit that is used to convey node and link characteristics to an adjacent node is referred to as an *Exchange Identification (XID)*. If `USE_PU_NAME_IN_XID=1`, the PU name is used in the CP name field of the XID. Otherwise, the control point name is used in that field. XIDs are exchanged between nodes before and during link activation to establish and negotiate link and node characteristics, and after link activation to communicate changes in these characteristics.

APPN nodes exchange XID format 3 (XID3) with other T2.1 or boundary nodes to perform role negotiation. For PU 2.0 connections, you use the `LINK_STATION`

Communications Server and SNA

keyword to specify a PU name and node ID that are exchanged on XID3. If `USE_PU_NAME_IN_XID=1`, the PU name is used in the name field of the XID. Otherwise, the control point name is used in that field. Information about the sending node's characteristics is contained in the XID3, including link station role (primary, secondary, or negotiable), TG number, node type, logical link number, the maximum basic transmission unit size that can be received, node ID, and PU name. The PU name is normally the control point name, but alternate PU name and node ID can be specified on the `LINK_STATION` keyword to support simultaneous PU 2.0 attachments.

Link Types

The following six types of links are typically defined in Communication Server nodes:

- Primary network access, initially activated (APPN or host)
- Primary network access, host link
- Primary network access, in use full time (use when required)
- Secondary network access, inbound link (predefined link characteristics)
- Secondary network access, limited use (also referred to as limited resource link)

Communications Server configuration provides a way to define and control the use of these link types. This section describes the node, DLC, and link configuration parameters used as well as the relationships and dependencies that exist between these parameters. The following categories are described:

- Link definition and activation parameters
- Link deactivation parameters
- Other link parameters

The following parameters are described:

- **Definition and activation parameters**
 - Activate at startup
 - Activate on demand
 - Automatic link retry
 - Maximum activation attempts
 - User requested reactivation
- **Deactivation parameters**
 - Inactivity timeout
 - Limited resource
- **Other link parameters**
 - Connection network
 - CP-CP session support
 - High performance routing (HPR) support
 - Adjacent node type
 - Preferred network node server
 - Solicit SSCP sessions
 - PU name (when solicit SSCP sessions is 1)

Note: These parameters are in the ACG file and might not be available on the panels (through **SNA Node Configuration**).

Link Definition and Activation Parameters

This section describes definition and activation parameters.

Activate at Startup

An *activate at startup* link is typically used for primary network access links that are initially activated when you start Communications Server. The link is activated when Communications Server is started at your machine and stays active as long as Communications Server is running.

To define a link to activate at startup, specify **ACTIVATE_AT_STARTUP=1** on the **LINK_STATION** keyword of the ACG File. Links are generally configured to activate at startup when they are important for network connectivity. An important link can also be configured for automatic link retry (see “Automatic Link Retry” on page 34).

The link from an APPN end node (EN) to its preferred network node (NN) server is an example of this type of link.

Examples of a non-activate at startup links, coded as **ACTIVATE_AT_STARTUP=0**, could be a host link that is not needed immediately when Communications Server is started, an inbound link that a node would like to control the link characteristics of, or a link that may be too costly to have active all the time. These links are defined as non-activate at startup and are activated when the link resources are requested by an application (see “Activate on Demand”) or the partner.

Activate on Demand

An *activate on demand* (also referred to as “auto-activate”) link is typically used for access to a partner LU which requires dynamic activation of the link. When Communications Server is started, the link remains inactive. However, it is placed into the topology as an available link if an adjacent CP name is specified. The link is activated when a transaction program (TP) requests a connection to a remote LU that requires the link to be active. Communications Server uses the fully qualified partner LU defined to activate the link.

To define a link as activate on demand, **ACTIVATE_AT_STARTUP=0** and **FQ_ADJACENT_CP_NAME=(netid.cpname)** must be configured on the **LINK_STATION** keyword in the ACG file of the originating node. If the partner is not the **FQ_ADJACENT_CP_NAME**, configure the **PARTNER_LU** keyword.

Frequently, an activate on demand link is also configured as either a limited resource (see “Limited Resource” on page 36) or with an inactivity timeout (see “Inactivity Timeout” on page 35) so that the link will be deactivated when it is no longer required.

An example of an activate on demand link is a link defining a connection to a partner which needs to be active for a limited amount of time. The link may cost more than you are willing to pay to keep active at all times. For example, you might have a collection of one or more computers communicating on a regular basis. At the end of each day, one of the machines is required to activate a link to some remote machine, in order to send the daily results, or to make a backup of the data.

Another example might be when you have connections to a data server or a print server. The connection requires resources at the server. To avoid limitations on the

maximum number of link stations and sessions at the server, configure an activate on demand link to free the resources at the data server after the requests over the activate on demand link are complete.

An activate on demand link is not necessarily a limited resource link, but it might be defined as one by including the **LIMITED_RESOURCE=1** parameter on the **LINK_STATION** keyword (see "Limited Resource" on page 36).

Automatic Link Retry

Automatic link retry is an error recovery function that enhances the availability of a failed link. If a link needs to be reactivated (without user intervention) after a failure, automatic link retry can be used. Automatic link retry causes automatic reactivation attempts of the link if the following parameters have been specified on either the **LINK_STATION** or **PORT** keywords:

- **DELAY_APPLICATION_RETRIES**
- **RETRY_LINK_ON_DISCONNECT**
- **RETRY_LINK_ON_FAILED_START**
- **RETRY_LINK_ON_FAILURE.**

If one of these parameters is specified on the **PORT** keyword, the values are used by the **LINK_STATION** keyword if the **INHERIT_PORT_RETRY_PARMS** parameter has been specified.

After a successful activation, the interval timer is reset to 0.

It might be beneficial to use automatic link retry on any of the following kinds of links:

- A link to a NN that is activate at startup. If the link fails, only a network node server (NNs) link would be reactivated.
- Host link
- Primary link that has a backup (also normally defined as an activate at startup link)
- Any link that needs to remain active at all times.

Maximum Activation Attempts

Maximum activation attempts is a link activation parameter that provides a mechanism to prevent dependent LU host traffic (for example, LUA, 3270, an LU 2 gateway and/or DLUR trying to activate a host link for a downstream application) from indefinitely retrying link activation. The **INHERIT_PORT_RETRY_PARMS** maximum activation attempts parameter represents the number of times an activate link request is attempted. After this number of attempts is reached, subsequent requests are rejected until the number of attempts is reset. Dependent LU link activation requests issued after maximum activation attempts is reached are immediately rejected without an actual attempt to activate the link. In this case, a primary return code of X'0003' and a secondary code of X'00000005', DLC retry, with a sense code of X'00000000' is returned. Independent LU link activation requests issued after maximum activation attempts is reached are attempted, but along with the primary X'0003' and secondary X'00000005' return code, a sense code of 081C0001 is returned to indicate the maximum activation attempts limit has been reached.

Note: If a TP, such as an LUA application, is in a loop trying to activate a session to the same host as the dependent LU application, the TP causes the maximum activation attempts number to be exceeded before the dependent application can make its first request.

The maximum activation attempts number is reset as follows:

- When an independent TP activates the link. It is the TP's responsibility to check the sense code and implement a reasonable protocol for link activation that does not infinitely retry the activation.
- When the user activates the link using **SNA Node Operations**.
- When the link is activated from the other direction (for example, by the partner).

If **SNA Node Operations** is attempting a link activation, the maximum activation attempts number is decremented by 1, and the result is ignored. If the limit has been exceeded on a gateway host link, activation of the link will be attempted after 30 minutes if a workstation link has become active. This enables retries to cease, although they are restarted later when the workstations are restarted. When gateway, LUA, or **SNA Node Operations** successfully activates a link, the maximum activation attempts number is set to 0.

To configure maximum activation attempts on a link, the **MAX_ACTIVATION_ATTEMPTS= n** parameter is configured on the **LINK_STATION** keyword, where n is the number of attempts from -1—127. A -1 indicates that the value on the **PORT** keyword should be used, and 0 indicates infinite retry.

User Requested Reactivation

User requested reactivation is when the user requests to reactivate a link from either **SNA Node Operations** or from the command line using, for example, Communications Server.

Link Deactivation Parameters

The following sections describe deactivation parameters.

Inactivity Timeout

Inactivity timeout is an SDLC link deactivation parameter that controls when links are deactivated. The inactivity timeout value specifies the time (in seconds) that a link may be idle before it is deactivated. This is similar to the **LINK_DEACT_TIMER** on the **LINK_STATION** keyword. The difference is the **LINK_DEACT_TIMER** waits for all sessions to end (session count reaches 0) before it triggers the link deactivation. The **INACTIVITY_TIMER** ignores the session count and triggers deactivation after the link has been idle for the specified time.

The inactivity timeout function was implemented to handle the situation where an emulator session, LUA, 3270, or a LEN connection, was accidentally left active for long periods of time. When the node detects no activity over this type of connection for the **INACTIVITY_TIMER** duration, the link is automatically deactivated, regardless of whether or not sessions and conversations exist on the link. By definition, the **LINK_DEACT_TIMER** is considered non-disruptive, but the **INACTIVITY_TIMER** is considered disruptive.

Notes:

1. Inactivity timeout can be used on limited resource or non-limited resource links. If a link is defined as limited resource and conversations remain active, limited resource timeout will not expire, and the link will be deactivated when the inactivity timeout expires. If a link is defined as non-limited resource, inactivity timeout is used to deactivate the link to free resources at the remote end.
2. Currently with HPR, the inactivity timeout is ignored. This is because HPR cannot tell what type of traffic is on the link, and the HPR keep alive protocol generates enough traffic so that the link will never be idle.

To configure inactivity timeout on an SDLC connection, the **INACTIVITY_TIMER= n** parameter is coded on the **LINK_STATION_SDL_C_SPECIFIC_DATA** parameter of the **LINK_STATION** keyword in the ACG file, where n is 40-160. To configure inactivity timeout on a link, the **LINK_DEACT_TIMER= n** parameter is coded on the **LINK_STATION** keyword of the ACG file, where n is 0—1 000. The 0 indicates no timeout (the link will stay active). The SDLC default is 80 and the default for the link is 10.

Limited Resource

A limited resource link is typically used for limited use primary network access links and secondary network access links. A limited resource link is a link that is automatically deactivated when its session count reaches 0. The limited resource link can be defined as **ACTIVATE_AT_STARTUP=1** or **ACTIVATE_AT_STARTUP=0**. If the link is **ACTIVATE_AT_STARTUP=1**, it is started when Communications Server is started. If the link is **ACTIVATE_AT_STARTUP=0**, it is placed in the topology when Communications Server is started, if the adjacent CP name is specified, and activated when services are requested.

Note: Activate at startup links are not placed into the topology unless they are active.

To configure a limited resource link, **LIMITED_RESOURCE=1** is specified on the **LINK_STATION** keyword of the ACG file. The **LINK_DEACT_TIMER= n** is specified on the **LINK_STATION** keyword and the **ADJACENT_NODE_TYPE=LEARN** parameter must be specified.

Note: If **CP_CP_SESSION_SUPPORT=1**, the link is not a limited resource link. Configuration verification will flag this as a warning. Active CP-CP sessions will keep the link from deactivating.

Connection Networks

Connection networks enable APPN nodes in a LAN to have direct links with each other without requiring logical link definitions at each node. This feature greatly reduces system definition without adding the performance burden of routing all sessions through a network node. It also enables new nodes that are added to the LAN to fully participate in APPC conversations without requiring definition changes at every other node.

A network node in the connection network assumes that all the nodes in one connection network can have links directly between one another. When calculating the route for a session, the network node considers the direct link and normally

selects the direct link as the optimal route. Having calculated the direct route, the network node simply sends the end node the address of the partner to use for activating the link.

The connection network route might not be taken when connection network security is less than required. If the connection network DLC is not secure and a mode like #BATCHSC is used on the **MODE_NAME** parameter, the network node attempts to find a secure route, ignoring the connection network.

If LAN bridges are being used, APPN views the entire bridged LAN as a single logical network. Because links can be activated between any two systems on the LAN, only one connection network is needed. The connection network should be defined at all the APPN systems on the LAN.

A network node learns connection network information during EN registration and APPN directory searches. The network node server then has enough information to calculate a direct connection between the session endpoint nodes without routing through intermediate nodes.

Only end nodes and network nodes can take advantage of the connection network; links to LEN nodes must still be explicitly defined.

Figure 12 illustrates a sample connection network. This view of a LAN shows a connection network given a name of LOCALNET.IBMLAN. With this type of definition, any EN can connect directly to any other EN as long as NN1 is the active network node server for all the end nodes.




Figure 12. A Sample Connection Network

Other Link Parameters

Other **LINK_STATION** keyword parameters that are mentioned in the preceding sections, but may not be obvious or implicitly defined are:

- Adjacent node type
- Preferred network node server
- Solicit SSCP sessions

This section contains a brief definition of each.

Adjacent Node Type

The adjacent node type specifies the type of node that is adjacent to the node defining the link. Valid types include:

- DSPU_NOXID
- DSPU_XID
- END_NODE
- HOST_DEP_LU_ONLY
- HOST_XID0
- LEARN
- NETWORK_NODE.
- SUBAREA_LEN

See the *Configuration File Reference* for further details.

Preferred Network Node Server

The preferred network node server specifies whether the adjacent network node is to be used as the network node server over the link being defined.

Solicit SSCP Sessions

Solicit SSCP sessions specifies whether or not SSCP-PU sessions are requested from the host over the link being defined.

Directory Services

A network node provides directory services to the LUs located in the network node and to the LUs in the end nodes that the network node serves. The network node also assists in the directory services provided by the other network nodes in the network by responding positively to received directory search requests when the resource named is found in the local directory. The local directory maps an LU name to the control point name of the node where that LU is located. If the destination control point is a LEN or end node, the directory includes the name of the serving network node.

The directory services component resides in every node; however, its scope and functions vary depending on the level of directory support in the node.

An end node maintains a local directory containing entries for locally resident LUs. In addition, the end node maintains directory entries for LUs in adjacent nodes with which the end node has been in session. For an LU-LU session with an adjacent peer node, a search of the local directory returns the appropriate destination control point associated with the LU searched for, permitting the proper logical link to be selected.

In a LEN node, all partner LUs are entered in the directory, as the example in Figure 13 on page 39 shows. Those not in an adjacent peer end node but out in the APPN network are associated in its directory with its designated network node server. The LEN node sends an LU-LU session activation (BIND) request to its network node server for any LU associated in its directory with its server; the server automatically locates the destination LU for it and forwards the BIND

appropriately. The network node can send a Locate search, wait for a response, and then send the BIND.




Figure 13. LEN Node Directory. The LEN node directory must contain all the LUs with which it communicates. Because the adjacent network node (NN) serves the LEN node even without CP-CP sessions, the LEN node must define the network node control point as the “owning control point” of all the LUs, including LUs located at the end nodes (ENs).

When an LU is not represented in an end node directory, the end node initiates a Locate search to find the desired LU. To activate the search, the end node invokes the services of its network node server. Figure 14 shows an example of an end node directory.




Figure 14. End Node Directory. The end node (EN) uses the services of its network node server to find the location of the LUs. None of the LUs in the APPN network need to be defined in the end node. The adjacent LEN node LU, however, must be defined because it is not connected to the network node and is not part of the APPN network.

A network node provides distributed directory services to its served end nodes in cooperation with all other network nodes in the APPN network. The origin

Communications Server and SNA

network node receives the name of a destination LU in a Locate search request from a served end node, or the name of a secondary LU in a BIND from a LEN node. The network node verifies the current location of the LU if it is represented in the network node's directory (but is not in the network node itself). The verification is done by sending a directed search to the destination network node server.

If the LU is not in the origin network node's directory, the network node initiates a search of the network. The search is initiated by sending a broadcast search to every adjacent network node, each of which in turn propagates the broadcast and returns replies indicating success or failure. For its future needs, a network node caches information obtained from successful broadcast searches.

An APPN end node can also receive (and respond to) Locate search requests from its network node server to search for, or ensure the continued presence of, specific LUs in the end node.

Each end node can register its LUs with its network node server by sending the network node a registration message. If the end node is registered with the network node server, the network node maintains current directory information pertaining to the end nodes in its domain.

Figure 15 shows an example of a network node directory.




Figure 15. Network Node Directory. The network node (NN) directory contains all the LUs it serves. The end nodes (ENs) register their LUs; the LEN node LU must be configured.

Topology and Route-Selection Services

A network node provides route selection services to itself and to the end nodes it serves. It maintains an internal network topology database that has complete and current topology information about the network. This topology information consists of the characteristics of all network nodes in the network and of all links between network nodes. All network nodes contain a copy of the topology database.

A network node uses the network topology database to compute routes for sessions that originate at the LUs in it and at the end nodes that it serves. Each route that a network node computes is the current least-weight route from the node containing the origin LU to the node containing the destination LU. To provide an appropriate path through the network, the algorithm used to select the

route first assigns weights to links and nodes. Based on the relative significance of the characteristics for the requested class of service, the weighting algorithm computes a scalar value for each node and logical link.

Topology Database

The network topology database in a network node contains information about all network nodes and all transmission groups interconnecting them. It is a fully replicated database that is shared among all network nodes in the network and used for route selection. The maintenance of the database requires broadcast updates among all network nodes. The updates are accomplished through topology database update (TDU) messages, which contain node-identifying information, node and link characteristics, and update-sequence numbers to identify the most recent changes for each of the resources described in a TDU.

A local topology database in an end node contains information about itself and directly attached nodes only.

The topology and routing services component uses the CP-CP sessions between network nodes to exchange information to build and maintain a topology database. This topology database in network nodes is kept current using updates that are transmitted among all network nodes whenever a resource (node or link) is activated or deactivated, or the characteristics of an existing resource change.

A local configuration database and a network topology database are maintained at each network node as illustrated in Figure 16 on page 42. The local configuration database is unique to the node, while the network topology database is replicated at all network nodes.

Table 2 shows the information contained in the configuration database at the local network node.

Table 2. Local NN Configuration Database

| Node | Links | Connection |
|------|-------|------------|
| NN5 | e | NN5—EN1 |
| | a | NN5—NN7 |
| | b | NN5—NN6 |
| NN7 | a | NN7—NN5 |
| | d | NN7—NN8 |
| NN6 | b | NN6—NN5 |
| | f | NN6—EN2 |
| | c | NN6—NN8 |
| | g | NN6—EN3 |
| NN8 | c | NN8—NN6 |
| | d | NN8—NN7 |
| | j | NN8—EN3 |
| | h | NN8—EN4 |

Table 3 on page 42 shows the information contained in the network topology database at the local network node.

Communications Server and SNA

Table 3. Local NN Network Topology Database

| Node | Links | Connection |
|--------------------|-------|------------|
| NN5, NN6, NN7, NN8 | a | NN5—NN7 |
| | a | NN7—NN5 |
| | b | NN5—NN6 |
| | b | NN6—NN5 |
| | c | NN6—NN8 |
| | c | NN8—NN6 |
| | d | NN7—NN8 |
| | d | NN8—NN7 |




Figure 16. Local Configuration Database and Network Topology Database in Network Nodes

Modes

The mode determines the values for the session characteristics and number of sessions between session partners. For example, the size of the largest request unit (RU) to be exchanged on a session (that is, the *maximum RU size*) is one of the characteristics of a mode. The mode also specifies a class of service, which is used to select the route for the session.

Class of Service

At session initiation time, the BIND specifies a mode name. This mode name is associated with a class-of-service (COS) definition that is used to determine the most desirable route between the origin and destination nodes of the session. The

COS definitions specify the characteristics that nodes and links must possess to be included in the route selected for the session. This specification enables the route-selection algorithm to determine if a node or link is acceptable. From the set that is acceptable, the algorithm calculates the best route for the session.

Because COS definitions can vary, different sessions can use different routes between the same origin and destination nodes, depending on the specified mode name. Each network node is capable of computing the least-weight (the most desirable) route to any destination.

When a session goes through both an APPN network and a subarea network, it uses two classes of service:

- In the APPN network, the session uses the class of service (COS) defined for the mode used by the transaction program. This definition is in the active configuration of the workstation. This COS is used to route the session in the APPN network.

If you use the default blank mode, the associated COS is #CONNECT.

- In the subarea network, the name of the COS can be written in the logon mode table, in the entry for the mode name used by the session. This COS name is used by the VTAM program to route the session through the subarea network.

If you use the default blank mode or if you do not define any COS in an entry of your logon mode table, the COS defaults to the blank entry in the ISTSDCOS table.

In both cases, each network uses the mode name to find the COS name, but the two COS names are not necessarily the same.

SNA Transmission Priority

The transmission priority is a value specified in the class of service. The transmission priority is sent in the BIND in the Class of Service/Transmission Priority (COS/TPF) control vector. Once the session is established, subsequent session data flows at the transmission priority specified in the COS/TPF control vector.

Data flowing on sessions using a class of service with high priority can pass data on sessions with lower priority. You should give high priority to sessions carrying interactive traffic where response time is important for example, emulator sessions. Sessions carrying high volumes of data, for example file transfers for NetView Distribution Manager, should be given lower priority. Transmission priority support helps to prevent high volume sessions from blocking traffic on the interactive sessions.

The four transmission priorities are network, high, medium, and low. Network priority is used for network control data such as topology and directory services. The other priorities are used for user data.

Communications Server supports transmission priority for LAN, SDLC, and X.25 links. The benefit is most apparent when the network contains congested low speed links.

Route Selection

After the network node server receives a response from its locate search, the topology and routing services component calculates the best route from the origin node to the destination node for the COS requested. Because the topology and

routing services component sends and receives topology database updates as characteristics of any resource change, every route is calculated with the most current information.

Route Selection for VTAM Users

To route APPC traffic through a subarea, the workstations connected to the subarea must be defined as network nodes in Communications Server. In each network node, a link is defined that connects the node to the subarea. From the viewpoint of the network node, partner LUs on the other side of the subarea are defined as being located at the host (a LEN node). From the viewpoint of the host, each network node connected to the subarea must be defined to the VTAM program with a PU macro. All destination LU 6.2 logical units within the APPN network for a particular connection are defined under the PU (network node) as if they are actually located at the PU. However, the LUs can actually be located at other nodes within the APPN network connected to the network node. The host sees only the network node PU. The network node PU can also be a gateway PU. SETN traffic (**CP_CP_SESS_SUPPORT=NO**) is not allowed when the parameter is set to YES on the NCP and there are no PU or control point sessions.

If the PU name in the VTAM definition is the same as the control point name defined in Communications Server, be aware that you will not be allowed to define the control point as an LU in the VTAM definitions. Names must be unique in the VTAM program, whether they are PU or LU names.

The PU macro must contain **XID=YES** to use an XID exchange during activation of the PU. This parameter is coded in the NCP major node. It must not be in the PU statement of a switched major node.

For switched SNA devices, you can use a new parameter in the PU macro: **CPNAME=ccccccc**. It specifies the control point name of the network node connected to the subarea. Either **CPNAME** or **IDBLK** and **IDNUM** must be specified on a switched PU definition statement. Both can be specified. The network node provides its control point name to the VTAM program in the XID exchange during the connection sequence. The VTAM program uses the control point name to locate the corresponding PU macro. If there is no PU macro with the corresponding control point name, the VTAM program uses **IDNUM** and **IDBLK** to locate the PU macro.

To route APPC traffic from an APPN network through the subarea and out to another portion of the APPN network, the network name (**NETID**) of the owning VTAM must agree with the network ID of the APPN network. In Communications Server, the network ID (of the network node connected to the subarea) can be found using the SNA local node characteristics profile.

Intermediate Session Routing

Intermediate session routing is a function performed by a network node. This capability enables a network node to receive and route data destined for another node. The origin and destination of the data can either be an end node, a network node, or a LEN node. The piece of the session between two adjacent nodes is called a session stage.

High Performance Routing (HPR) Support

Communications Server supports high performance routing (HPR) over Enterprise Extender (IP), synchronous data link control (SDLC), LAN, WAN, channel, Multi-Path Channel (MPC), and X.25 connections.

HPR automatic network routing (ANR) minimizes the storage and processing requirements in intermediate nodes, which is a better solution than APPN intermediate session routing (ISR) for high-speed networks with low error rates.

HPR improves SNA routing with these major features:

- Stateless intermediate nodes

Stateless intermediate nodes have no knowledge of the connections that pass through them.

An HPR intermediate node does not need any control blocks or buffer pools to support the connections that pass through it. It also does not need routing tables because the routing information is carried in each packet. Using ANR, an HPR intermediate node simply receives a packet, examines the packet header to find an identifier for the logical link to the next node and sends the packet on that link.

- An HPR intermediate node:
 - Does not recover packets because the HPR connection end points will detect and recover lost packets using the rapid transport protocol (RTP).
If your LAN has high error rates you can also choose link-level error recovery protocol (ERP) support. However, this should not be necessary on links with low error rates.
 - Does not reassemble segmented packets and then segment them again because the HPR connection end points will automatically choose a packet size supported by all intermediate nodes in the route.
 - Does not actively manage buffers and prevent congestion because the HPR connection end points employ flow control algorithms that prevent congestion at intermediate nodes.
 - Does not determine a packet's route by performing a table search because the HPR connection end points perform source routing and build the packet routing information using labels chosen for optimal performance by each intermediate node.

To achieve stateless intermediate nodes:

- The end points of the route must have buffer capacities that reflect the combination of route delay and application throughput, a value that depends on your network configuration.
- The data links must support I-field sizes large enough to contain the HPR headers and still carry application data efficiently.
- The data links must be very reliable because a high error rate will adversely affect the HPR throughput.

For configurations where these conditions cannot be met, APPN without HPR is still a valid implementation choice. HPR technology is an APPN option and APPN without HPR is still available and supported.

- Congestion avoidance

An HPR network can achieve very high link utilizations. Using a technique called automatic rate based (ARB) flow control, HPR connection end points

automatically sense network congestion and reduce transmission load. This avoids lost packets that, in other routing algorithms, may be caused by the lack of available buffers in intermediate nodes.

- Automatic rerouting without disruption

When an HPR connection end point senses that a route has failed, it automatically computes a new route, reestablishes the connection, and recovers any packets lost due to the failure. Because HPR intermediate nodes are stateless (that is, they have no knowledge of the connections that pass through them), they do not participate in this recovery protocol. Because no packets are lost if the route is switched successfully, applications using the route are not disrupted by the failure.

- Software compatible with existing hardware

HPR is a software extension of APPN; you can use HPR on your existing LAN or WAN adapters and frame relay links. Because HPR does not attempt to provide dedicated bandwidth and guaranteed real time transport (leaving this level of function to Broad Band Networking Services), the performance constraints on packet forwarding delays are the same as for APPN without HPR. The control point protocols used by HPR are those of base APPN with minor additions, so HPR nodes can be installed incrementally within an existing APPN network without preplanning or coordination.

Rapid Transport Protocol (RTP)

RTP is a set of message formats and protocols designed to utilize modern data communication media, minimize overhead in intermediate nodes, and automatically switch paths when a link in the path fails.

RTP connections are established within an *HPR subnet* and are used to transport session traffic. An HPR subnet is the portion of an APPN network that is capable of establishing RTP connections and transporting HPR session traffic. RTP connections can be thought of as *transport pipes* over which sessions are carried. These connections can carry data at very high speeds using low-level intermediate routing and minimize traffic over the links for error recovery and flow control. These flows are managed by the RTP connection endpoints.

An RTP connection's physical path can be switched automatically to reroute data around a failed node or link without disrupting the sessions. Data in the network at the time of the failure is recovered automatically.

RTP does error recovery on an end-to-end basis, rather than on a link-level basis. Performance is improved by reducing the number of flows required to do error recovery. Link-level error recovery protocols (ERPs) are also supported for all connections. ERP is a method of detecting a lost packet at one end of a link and recovering by asking the other end of the link to retransmit the packet. If ERP is used, HPR packets are sent as numbered information frames (I-FRAMES). When a frame is lost, the DLC detects the loss and the sender retransmits the frame. If ERP is not used, HPR packets are sent as unnumbered information frames (UI-FRAMES). When a frame is lost, the DLC cannot detect the loss and HPR rapid transport protocol (RTP) must detect and recover lost packets at connection end points.

In either case, RTP always detects and recovers lost packets at connection end points. For any given connection, there are no restrictions on the number of links that use ERP, or do not use ERP.

ERP can be enabled or disabled on a link-by-link basis. Because RTP detects and recovers lost packets at connection end points, you can use either ERP links or non-ERP links when you build the network. This lets you specify link-level ERP on links that have a high rate of packet loss and maximize throughput on other links by specifying that they do not use link-level ERP. In general, the use of ERP is not recommended in LANS.

Note: ERP is always enabled in a wide area network (WAN) environment.

Flow control and congestion control are also done by RTP on an end-to-end basis. RTP uses a technique called automatic rate based (ARB) flow control to fully utilize network bandwidth when possible. RTP increases the rate at which packets are sent when the network supports this increased send rate. Congestion is automatically recognized and the send rate will be decreased accordingly when congestion occurs. The configured effective capacities of links in the connection path are used to determine both the initial send rate and the send rate increment.

Support for control flows (CF) over RTP connections is now available with HPR in Communications Server. Previously, control flows, including CP-CP sessions and route setup messages, used APPN connections while the data flows used HPR connections. Now, both control flows and data flows can use RTP connections. The benefits of this support include automatic path switching for CP-CP sessions.

Control flows automatically flow over RTP if both endpoints of the connection support this function.

Automatic Network Routing (ANR)

Automatic network routing (ANR) is a stateless routing technique enabled by RTP where a message arrives with a label that uniquely identifies the next hop in the path. Because of its simplicity, ANR can be performed at a low level with no knowledge of the connections using the path. ANR minimizes cycles and storage requirements for routing packets through intermediate nodes.

The ANR fast packet switching function improves performance in intermediate nodes by routing at a lower level than APPN and performing error recovery, segmentation, flow control, and congestion control at the end node, rather than on the intermediate node.

Intermediate ANR nodes are not aware of the SNA sessions or the RTP connections. Routing information for each packet is carried in a network header with the packet. Each node strips off the information it has used in the header before forwarding the packet, so the next node can find its routing information at a fixed place in the header. There is no need to keep the routing tables for session connectors as in base APPN, so switching packets through nodes can be done more quickly.

LU Support

SNA defines LU types 0, 1, 2, 3, 4, 6.0, 6.1, 6.2, and 7. LU types 0, 1, 2, 3, 4, and 7 support communications between application programs and different kinds of workstations. LU types 6.0 and 6.1 provide communications between programs located at type 5 subarea nodes. LU type 6.2 supports communications between two programs located at type 5 subarea nodes or type 2.1 peripheral nodes, or both, and between programs and devices.

Communications Server and SNA

Communications Server supports LU types 0, 1, 2, and 3, which support communications with host applications that support devices such as:

LU type 0 3650 and 4700 financial terminals

LU type 1 3270 printers

LU type 2 3270 interactive displays

LU type 3 3270 printers

Communication occurs only between LUs of the same LU type. For example, an LU 2 communicates with another LU 2; it does not communicate with an LU 3. Communications Server also supports LU type 6.2 or APPC.

The Communications Server SNA functions enable applications to use the APPC application programming interface (API) to provide a distributed transaction processing capability in which two or more programs cooperate to carry out a processing function. This capability involves communication between the two programs so they can share local resources such as processor cycles, databases, work queues, and physical interfaces such as keyboards and displays.

Communications Server supports APPC through the APPC APIs. Refer to the following publications for more information:

- *Client/Server Communications Programming*
- *System Management Programming*

The following Communications Server functions support a range of LU types:

- SNA Gateway

The SNA gateway enables IBM System/370 (S/370) hosts, or zSeries architecture host computers, or both, to support LAN-attached workstations that use LU types 0, 1, 2, 3, or 6.2 for communicating through an SNA gateway workstation.

The SNA gateway also supports LU types 1, 2, or 3 to iSeries host computers, which can process the data or pass it through to S/370 or zSeries architecture host computers.

- Conventional LU Application (LUA)

The LUA consists of system software that supplies service routines to support LUs 0, 1, 2, and 3.

SDDLU Support

The self-defining dependent LU (SDDLU) support enables you to dynamically define and activate a dependent LU at the host (VTAM). In VTAM, this is known as dynamic definition of dependent LUs (DDDLU). SDDLU is enabled in Communications Server by coding an LU_MODEL statement on an LU definition.

To enable the DDDLU facility in VTAM, code the LUGROUP operand on the PU definition statement for the PU, and code an LU group major node. To use the IBM-supplied SDDLU exit routine that generates the LU names for you, you should also code the LUSEED operand on the PU statement.

The LUGROUP operand specifies the name of the model LU definition group that VTAM will use when dynamically defining LUs for this PU. The LU group major node contains the model definition statements. Dynamic definitions for LUs are built using the model LU definitions contained in this major node.

The LUSEED operand provides a pattern name that is used with the SDDLUI exit routine to create a name for the dynamically created LU. Once the correct statements have been added to the PU statement and the LU group major node coded, these major nodes need to be active for the SDDLUI function to be enabled.

Dependent Logical Unit Requester Support

Dependent Logical Unit Requester (DLUR) is an architecture intended to provide dependent LU support in an APPN network. Communications Server supports all base DLUR functions and the following optional functions:

- Self-defining dependent LU (SDDLUI) Support
By using SDDLUI support the DLUR-supported LUs can be dynamically defined to VTAM, thus eliminating the need to have these LUs predefined.
- TakeOver/GiveBack Support
This support enables LU-LU sessions to remain active even when the connection between the DLUR and DLUS nodes has failed. The connection can be reestablished with the same or another DLUS, without the LU-LU sessions ever being disrupted.
- MultiSubnet Support
This support enables the DLUR node, DLUS node, and the node containing the application to be in different subnetworks.
- DLUS-served LU registration
An end node DLUR registers its LUs so that the network node can locate these LUs without having to pass the locate requests to the DLUR.
- Backup DLUS Support
When the connection to the primary DLUS fails or does not become active, Communications Server will automatically try to establish a connection with the backup DLUS.

Using DLUR

To use the DLUR function, you configure a **DLUR_DEFAULTS** definition and use the link name from that definition as the host link for your LUA, dependent LU 6.2, or gateway definitions. Communications Server sends the PUNAME, CPNAME, and NODEID to the DLUS. The PUNAME is sent as part of the signalling information (CV X'0E').

- If the DLUS is uplevel (supports checking of the CV X'0E' on REQACTPU) it uses the PUNAME in its search algorithm. This is available in VTAM 4.3 with PTF or higher.
- If the DLUS is downlevel, it ignores the CV X'0E'.
- If there is no match on PUNAME or the DLUS is downlevel, the DLUS then attempts to locate a PU with matching CP name or NODEID (IDBLK/IDNUM) either predefined in a VTAM switch major node or dynamically created using the ISTECCS exit.

Connections to the network using the connectivity of your choice (Token Ring, SDLC, AnyNet, and so on) must be configured and active before the DLUR-to-DLUS connection can be established. Once an APPN connection exists between the DLUR and DLUS, a pair of control sessions are established between the DLUR and DLUS using a special mode, CPSVRMGR. This pair of control sessions is also referred to as the CP-SVR pipe and appears as a link to Communications Server. It can therefore be activated, deactivated, and displayed using **SNA Node Operations**.

Communications Server and SNA

Once the pipe is activated, SSCP-to-PU and SSCP-to-LU support can be provided to PUs and LUs that have defined the pipe as their host link. LU-to-LU sessions do not use the pipe, but will use the best path available through the network.

In the DLUR environment, any number of dedicated PUs can be defined on the LU 6.2 sessions. This enables the gateway to provide network management access through the dedicated PU to downstream workstations without requiring numerous physical links to the hosts.

Figure 17 shows a Communications Server workstation acting as a DLUR gateway for both a workstation and a 4702 controller.




Figure 17. DLUR Connection to a Host through a Communications Server Gateway

LU-LU Sessions

Communications Server LUs can both initiate sessions and respond to session initiation requests. An LU initiates and responds to requests according to the type of LU: independent or dependent.

Independent LU

An independent LU is able to activate an LU-LU session (that is, send a BIND request) without assistance from the SSCP; therefore, it does not have an SSCP-LU session. An independent LU is capable of sending and receiving BINDs. The BIND sender is referred to as the primary LU (PLU); the BIND receiver is referred to as the secondary LU (SLU).

Only an LU 6.2 can be an independent LU. Communications Server supports independent LU protocols to other type 2.1 nodes as well as to low entry networking-level type 5 subarea nodes.

Independent LUs can have parallel sessions between the same pair of LUs and can have multiple sessions between one LU and several other LUs. Their session limits are established on a mode name basis, which can be from 1 to 32 767.

Figure 18 on page 51 shows how multiple and parallel sessions can be established by an independent LU. LU_x supports parallel sessions with LU_y and a single session with LU_z. The direction of the session arrows shows the PLU-SLU relationship. LU_x acts as the PLU for the session with LU_z and for one of the

sessions with LUy. LUx also acts as the SLU for one of the parallel sessions with LUy.




Figure 18. Multiple and Parallel Sessions

Dependent LU

A dependent LU is an LU that is controlled by an SNA host system. To activate an LU-LU session, a dependent LU requires assistance from an SSCP. It requires an SSCP-LU session to send a BIND. Dependent LU protocols are supported by Communications Server, but only to type 5 subarea nodes using type 2.0 protocols, not to other type 2.1 peripheral nodes. Dependent LUs act as SLUs only and have an LU-LU session limit of 1. However, multiple PU support in Communications Server enables you to establish multiple simultaneous SSCP-PU sessions with dependent LU sessions.

The dependent LU requester (DLUR) function enables Communications Server to take advantage of the enhanced SSCP support provided by a dependent LU server (DLUS). Some of the benefits of this function are:

- The dependent LUs can reside on nodes that are not adjacent to the host.
- APPN search logic finds the best path for LU-LU sessions.
- SNA dependent LUs can gain the advantages of an APPN network.

To use DLUR, configure a **DEFINE_DEPENDENT_LU_SERVER** parameter and use the link name from that definition for your LUA, dependent LU 6.2, or gateway definitions.

LU 6.2

Independent LUs are defined to the VTAM program by coding LOCADDR=0. There can be as many LUs defined with LOCADDR=0 as you want. Note, however, that not all LU 6.2s are independent LUs.

When you define the LUs of one part of the APPN network to the VTAM program, you must define them as being in the network node that connects this part of the APPN network to the subarea network. Following the PU definition of this network node, define each LU that you want to reach from the other part of the APPN network. Do not forget that control points are LUs.

Communications Server and SNA

An LU must be defined in VTAM to establish a session with another LU if this session goes through the subarea network. There is no effective way to get around it (such as the wildcards of APPN); VTAM must know the name of each destination LU.

Because an APPN network is intended to change easily, you should define the LUs of the APPN network in a special major node whenever it is possible. You can also define, in VTAM, LUs that do not exist yet.

Other LUs

If the network node uses its connection to the subarea network for 3270 emulation, the LU type 2 LUs of the 3270 emulation are defined in the same PU macro as the LU type 6.2 LUs of the APPN network. The link will also be used for the connection between the 3270 emulation and the host.

APPN Network Node and T2.1 Support

APPN is an enhancement to IBM's SNA and type 2.1 (T2.1) node architecture. APPN enables interconnection of systems of widely differing sizes into networks of a dynamic topology. An APPN network is easier to use, is more reliable, and provides more flexibility than traditional SNA networks.

Refer to *3174 APPN Implementation Guide* for additional information on APPN network nodes.

Data Compression

Data compression is the process of compressing repeated bytes or repeated data strings to shorten the length of records or blocks. This reduces the transfer time needed for communications. By reducing the amount of data transferred between host and workstation sessions, you can increase the throughput on slow speed lines and lower the cost per bit on expensive lines.

The performance gain, as measured by the number of transferred bytes, that you can expect by using data compression is often a ratio of about 2:1. This means that, with data compression active, you save about every second byte in the buffers needed for lower level protocol conversion.

Data compression is beneficial to those who need:

- Improved response times, especially on low-speed lines
- Reduced costs on lines where tariffs are based on volume of data
- Fewer leased lines

However, data compression should not be applied to every session you are running as there are disadvantages you should consider:

- Compression and decompression need additional CPU cycles
- Increased data storage (32 bytes to 9.0 KB) will be needed
- Data compression requires a secure link as defined in your APPN class-of-service (COS) table.
- Tuning of the RU size might be required to obtain the best performance. Each RU is compressed and then sent. If packets are being sent, a larger RU size on the mode will enable more full packets to be sent.

Note: The amount of increased storage depends on the compression algorithm you use. This increase is in addition to the storage required for the message.

For detailed and technical descriptions of different compression algorithms see the following publications:

- *SNA Formats*
- *Better OPM/L Text Compression IEEE Transactions on Communications, vol COM-34, no 12, pgs 1176-1182,1986*
- *IBM ITSC VTAM V3R4 and V3R4.1 Planning Guide*
- *A Technical Guide to ESA/390 Compression*

The following sections describe SNA session-level compression and Communications Server implementation.

SNA Session-Level Compression Architecture

SNA session level compression implements data compression in the LU-LU half session. With Communications Server it is available to all supported LU types, that is: LU types 0, 1, 2, 3, and 6.2. Data compression at the session level provides these advantages:

- It is more efficient to compress data before you encrypt it
- Different algorithms are more effective for different types of data
- Application programs do not have to provide their own compression logic

Two algorithms are generally defined for SNA session level compression, Run length encoding (RLE) and a form of Lempel-Ziv (LZ). Communications Server supports SNA session level compression using the following algorithms:

- NO compression.
- RLE compression; RLE is the simplest and most widely known algorithm, replacing strings or identical bytes with shorter encoded strings.
- LZ9 compression; LZ9 is a dynamic compression algorithm that compresses previously seen strings (in the current or preceding RUs) to a 9 bit code which represents the zero origin index of an entry in the compression/decompression table. The table entries store previously seen strings.
- LZ10 compression; LZ10 is a dynamic compression algorithm that compresses previously seen strings (in the current or preceding RUs) to a 10 bit code which represents the zero origin index of an entry in the compression/decompression table. The table entries store previously seen strings.

Typically, LZ compresses data better than RLE, but at a greater cost of memory and CPU function.

SNA session level compression views the session in two directions, PLU-SLU and SLU-PLU. The primary logical unit (PLU) is the LU responsible for activating the session. The secondary logical unit (SLU) is the responding LU. The PLU activates a session by sending a Bind Session request (BIND) to the SLU, which responds with a BIND. This means that different compression algorithms can be used in the PLU-SLU and SLU-PLU directions. This is performed through BIND negotiation of the compression levels. LU 6.2 can use any combination of compression levels for a session (for example, PLU could use RLE and SLU-PLU could use LZ9). All other LU types have compression enabled or disabled. When enabled, the PLU-SLU compression level is LZ9 and SLU-PLU is RLE.

Communications Server Data Compression

Communications Server supports SNA session level data compression with RLE, LZ9, and LZ10 compression algorithms. With Communications Server, you can specify the use of data compression for communications over CPI-C sessions (through APPC session), APPC (LU 6.2) sessions, and LUA (LU 0, LU 1, LU 2, and LU 3) sessions.

A two-part configuration is used to enable data compression. The Communications Server node must be enabled for data compression and the LU (APPC and LUA) must first be enabled. The two node compression fields (level and tokens) are on the local node characteristics window (**NODE** keyword in the .ACG file).

The compression level field sets the maximum level that any session can be started with: NONE, RLE, LZ9, or LZ10. This field takes precedence over all compression levels configured or attempted (the only exception is stand-alone DFT, which does not require the node definition). If you are configuring a session using LUA (LU 0, LU 1, LU 2, and LU 3) to support 3270 emulation or printers, LZ9 is required for data compression. The other compression levels do not allow data compression for these LU types.

The default request unit size for compressed modes is two times the connection's basic transmission unit (BTU) size. If you are using packet switching, you might not want to use this default size; instead, use a larger size and segment the packets.

LU 6.2 compression can be enabled by:

- Using the modes supplied with Communications Server
 - #INTERC
 - #BATCHC
 - #BATCHCS
 - #INTERCS
- Changing or adding modes with compression enabled

Three mode compression fields, compression need, PLU->SLU compression level, and SLU->PLU compression level, are displayed on the Mode Definition panel (**MODE** keyword in the .ACG file).

Compression need can have two values.

Prohibited

No compression.

Requested

Use this to request data compression with the values defined in PLU->SLU compression level and SLU->PLU compression level. The requested level might not be obtained for the following reasons:

- Node settings have limited the allowable level
- The SLU has negotiated the levels down

The SLU honors the compression levels requested by the PLU, unless limited by its node compression settings.

For more information on Communications Server compression, refer to the online help text for the product or the *Configuration File Reference*.

SNA Session-Level Encryption

SNA session-level encryption enables you to encrypt either all of the data or selected data that is transferred between the workstation and the host. If you want to protect any workstation data by using encryption, the host must also be configured to use encryption.

An IBM SecureWay 4758 PCI Cryptographic Coprocessor (referred to as the IBM 4758) adapter must be installed on the server to enable data confidentiality. This adapter must be initialized by following the instructions provided with the adapter.

In Communications Server, LU 6.2 session-level encryption is configured based on the mode description used for a given transaction program. There are two levels of encryption:

- APPC session level
Strictly between two LUs, no control points get involved in translating keys
- APPN session key translation
The control point of the originating LU (EN or NN) and possibly its NN server translate the keys in addition to the originating LU.

To configure a mode for encryption, bring up the SNA Features window and select **MODES**. Then inside the Mode Definition window, select **Setup...** The Compression and Session-level Encryption Support window appears. The parameters for the encryption configuration are in two parts:

- For session level encryption, you can specify whether encryption is optional (the LUs negotiate) or mandatory (encryption must be used).

Note: Encryption from Communications Server to a host is always set to be mandatory.

- For APPN session key translation, you can specify that only the LU translates the key, that only the end node translates the session key, or that the end node and its network node translate the key.

Communications Server requires other products for key storage and translation. A common cryptographic architecture (CCA) product is required for key storage, managed by the utilities supplied with the IBM 4758 adapter. Communications Server calls a CCA product, which interacts with the IBM 4758 adapter to get the keys and encrypt the data.

VTAM Users:

Communications Server does not encrypt the SNASVCMG session. You must specify ENCR=OPT in the APPL statement of your VTAM application definition. In working with VTAM, you must have encryption specified on the MODEENT statement. For example:

```
ENCR=B'0011' FOR MANDATORY ENCRYPTION
```

To use the VTAM encryption facility, the IBM Programmed Cryptographic Facility (PCF) must be initiated before starting VTAM.

Beginning with VTAM V3R4.1, VTAM uses a new interface to Integrated Cryptographic Service Facility/MVS (ICSF/MVS) for

cryptographic services, such as providing session-level cryptography. This interface complies with the Common Cryptographic Architecture (CCA) as implemented by ICSF/MVS. With this support, you can start and stop the cryptographic service after VTAM has been started, and you can change the master key without disrupting VTAM or active LU-LU sessions.

For information on how to define data encryption, refer to *OS/390 eNetwork Communications Server: SNA Network Implementation*.

Management Services

Communications Server's management services (MS) are functions distributed among network components to operate, manage, and control a network. This capability is based on the SNA management services architecture documented in *Systems Network Architecture Management Services Reference*.

Focal Points, Service Points, and Entry Points

Communications Server provides programming support that enables installation of management services focal point (FP), service point (SP), and entry point (EP) applications. Management services SP applications are just a variation of management services EP applications and differ only in the kinds of function they provide; otherwise, they interact with a management services focal point as management services EP applications.

Focal Point

A management services focal point is a central point of control for managing a network. From a management services SP or management services EP application, the management services focal point can request certain data relating to the operation of a network such as problem and performance data or product identification.

The management services focal point can also accept certain unsolicited management services data from the nodes it manages based on the category of the management services data. An example of a management services category is MS alerts. A management services focal point can manage one or more categories of management services data, and there can be one or more management services focal points in a network. IBM Communications Server, the IBM NetView program, and the IBM OS/400 operating system are examples of products that provide management services focal point capability.

Service Point

A management services SP is the function in a node that can request and capture data from devices that, by themselves, cannot serve as management services EPs, such as devices connected by LAN protocols (but not higher-level SNA protocols) to the management services SP node. Aside from collecting nonlocal data, a management services SP functions like a management services EP in its relation to the management services focal point. The IBM NetView/PC and IBM LAN Network Manager program products are examples of management services EP applications that provide management services SP functions and management services EP functions.

Entry Point

A management services EP is the function in a node that captures local management services data and sends it to a management services focal

point for processing, either upon request or unsolicited. Communications Server provides the management services EP function for sending alerts to the alert management services focal point. These alerts can originate within Communications Server or the DLCs it uses. Communications Server also provides programming support for applications, such as the IBM NetView/PC and IBM LAN Network Manager program products, by supplying alerts to be sent to the alert management services focal point.

Levels of SNA Management Services Architecture

An SNA product implements a particular level (or generation) of the SNA Management Services architecture, and some products support several levels of the architecture. Communications Server can send management services data to, and receive management services data from, SNA products that implement any of three levels of the management services architecture. These levels are:

Multiple Domain Support (MDS) level

An SNA product that implements the MDS level of the management services architecture, such as Communications Server and IBM NetView Version 2 Release 2 (or later). It can send and receive MDS message units (MDS-MUs). IBM NetView Version 2 Release 2 provides MDS level as a subarea LU, not a control point (CP), and uses SNASVCMG-mode sessions for transporting MDS-MUs. As a focal point, it supports explicit, implicit (primary) and implicit (backup) FP-EP relationships. NetView Version 2 Release 2 also continues to support the host FP-EP relationship to EP products that do not have MDS-level support.

Migration level

An SNA product that implements the previous level of the management services architecture, such as IBM OS/400 Version 1 Release 3 Modification Level 0 (or earlier). A migration-level product can support explicit, default, and domain FP-EP relationships. The domain FP-EP relationship is inferred when the CP-CP sessions are activated to a migration-level node. A migration-level serving network node (NN) does not send MS Capabilities for FP Notification to its served end nodes (ENs), and a migration-level served EN does not accept MS Capabilities for FP Notification from its serving NN. It can send and receive CP-MSUs but not MDS-MUs. As a focal point, it supports only the alert MS category.

Network Management Vector Transport (NMVT) level

An SNA product that implements the NMVT level of SNA management services architecture, such as IBM NetView Version 2 Release 1 (or later). NMVT is a management services request unit (RU) that flows over an active session between PU management services and control point management services. If an NMVT is routed from a workstation through a gateway, then the gateway adds its control point name to the NMVT.

Flow Control

To manage the flow of data over a network, Communications Server uses adaptive session-level pacing. The pacing occurs between each pair of adjacent nodes participating in the session route. The pacing between two adjacent nodes is independent of the pacing used between other adjacent nodes in the route.

Session-Level Pacing

Adaptive session-level pacing uses a window-based scheme, where a sender can send only a limited number, or *window*, of request units per explicit grant of

permission to proceed. The window size can be changed based on conditions at the receiver. This function permits a node to control the amount of data that is sent and received during normal session operation. The window control enables the receiving node to manage its rate for receiving data into its session buffers. Adaptive session-level pacing provides a node supporting many sessions a dynamic means to allocate resources to a session that has a burst of activity and to reclaim unused resources from sessions that have no activity. Adaptive session-level pacing enables the receiving node to use its available buffer resources efficiently.

Because each session stage between the endpoints is independently paced, both endpoint nodes and intermediate nodes can adapt the pacing for the sessions they handle in accordance with their own local congestion conditions. This action is the basis for global flow control and congestion management in APPN networks.

If, however, an interactive session and a session transferring a large file share a link, the interactive session data should be transmitted as quickly as possible. There are two ways to do this:

1. Assign a lower priority to the file transfer session. #BATCH uses low priority.
2. Use fixed pacing with a small window size for the file transfer session to enable interactive session data to use the link when the file transfer session is waiting for the pacing response. If connected directly to an NCP host, two-way fixed window pacing can be used to set pacing in both directions to the receive window on the defined mode.

Adaptive BIND Pacing

BIND traffic can occur in bursts, particularly at node or network startup. Therefore, adaptive BIND pacing exists to control the flow of BINDs between two adjacent nodes. The same window algorithm used for session-level pacing is employed.

Segmenting and Reassembly

To transmit RUs longer than the maximum-size basic transmission unit allowed by a particular link, Communications Server supports data segmentation and reassembly. These segments are reassembled into whole RUs at the partner node. This action enables the RU size defined for a session to be independent of the link that is used for the route.

High Performance Routing Pacing

High performance routing (HPR) provides a new method of flow control called adaptive rate-based congestion control (ARB). ARB regulates traffic flow by predicting congestion in the network and reducing a node's sending rate into the network, preventing congestion rather than reacting to it.

Fixed Pacing

Fixed pacing enables you to share a physical connection between two sessions. Without fixed pacing, the data to be transmitted is placed on a common data link control (DLC) queue and interactive data follows previously queued data. Fixed pacing also reduces the amount of storage that can be used to place data on the DLC queue. Two-way fixed pacing can be used with an NCP to avoid defining host fixed pacing. However, in general, adaptive pacing is the most efficient method of data transfer between nodes.

Transmission priority, like fixed pacing, enables sharing of a physical link between sessions. However, it will lock storage as the data is placed on the DLC queues, but it does not require the additional pacing responses that are required for fixed pacing.

Partitioning LUs among Hosts

When you define multiple subarea host connections, it is required that traffic from the domain of a given host enter only on one logical link. Note that manual dial connections appear to be a single link. You must define a different PU to support each different host. Only the host links defined on the control point can have CP-CP sessions and participate in the APPN network. Links that have **USE_PU_NAME_IN_XID=1** cannot have CP-CP sessions. Otherwise, the host links can have CP-CP sessions and can also participate in APPN communication. The **LINK_STATION** keyword specifies the PU name and logical link to be used for the PU. If parallel links are required (in situations where there are more than 254 dependent LUs), one of the links must have CP-CP session support set to No.

Each dependent LU can be defined as associated with only one PU. The **LOCAL_LU** keyword specifies the host link name used for a dependent LU type 6.2. The **LU_0_TO_3** keyword specifies the host link name for LUA, and the 3270 profile specifies the host link for each 3270 emulation session.

The control point automatically defines a PU with the same name as the control point. Dependent LUs that are defined at one subarea SSCP can be associated with the control point PU. (In fact, the automatically defined LU for the control point can also be specified as dependent.) Each additional PU 2.0 for a different subarea host requires a separate link and PU definition. A PU is defined by specifying PU on a **LINK_STATION** keyword. This is unnecessary if all LUs are independent.

If the SSCP-PU session is to send alerts to the host, define the focal point link with the **NODE** keyword in the .ACG file. If a host is to receive the alerts, specify a logical link for the control point to the host. If the link to that host is not available, Communications Server logs the alerts. The only network management requests that are accepted from a PU 2.0 host are the ones for the control point PU. Replies are sent to the same host using the control point PU.

Each dependent logical unit must have a configured local address that is the same as the one configured at the host. However, use of the SNA gateway allows for a gateway address translation. A dependent LU 6.2 can act as an independent LU to a peer node; that is, it is dependent only to the subarea. Such an LU should not be a part of an APPN network with any other connection to the same subarea, that is, the subarea containing the controlling SSCP. Only one subarea connection is permitted for an APPN network, unless the subareas are independent. Although a gateway can have parallel links to the same subarea, only one can have CP-CP sessions and APPN traffic.

An SSCP in the subarea network activates the dependent LUs it controls after the link to it is established. Until an LU is activated, the LU cannot start a session. When APPC is taken down, disconnect is requested from each active host in session with a PU in the node. Each host frees the link after deactivating first the LUs and then the PU on the SSCP-PU session.

SNA Gateway Support

Communications Server provides a full-function Systems Network Architecture (SNA) gateway. The gateway allows multiple LAN-attached workstations to access System/370 or zSeries hosts through one or more physical connections to one or more hosts. This helps reduce the cost per workstation of host connections.

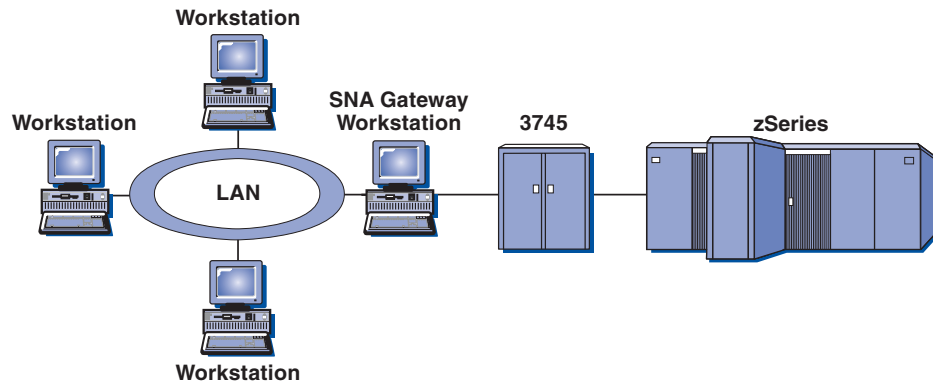


Figure 19. Example of SNA Gateway Configuration

The Communications Server gateway supports the SNA protocols LU 0, 1, 2, 3, and dependent LU 6.2 (APPC). With the AnyNet SNA over TCP/IP function, downstream workstations can now communicate with the SNA gateway over an IP network. The gateway also supports LU 0, 1, 2, or 3 to an iSeries host using SNA pass-through. The iSeries host passes the data through to a zSeries host.

A gateway can also act as a protocol converter between workstations attached to a LAN and a WAN host line.

The LUs defined in the gateway can be dedicated to a particular workstation or pooled among multiple workstations. Pooling allows workstations to share common LUs, which increases the efficiency of the LUs and reduces the configuration and startup requirements at the host. You can also define multiple LU pools, each pool associated with a specific application. And you can define common pools that are associated with multiple hosts. When a client connects to the gateway, the gateway retrieves an LU from the pool to establish a session. The LU is returned to the pool for access by other workstations when the session is ended.

In addition, an SNA gateway can support the forwarding of network management vector transports (NMVTs) between the workstations and the host.

Each host views the SNA gateway as an SNA PU 2.0 node, supporting one or more LUs per workstation. As far as the host is concerned, all LUs belong to the SNA gateway PU. The SNA gateway can have multiple host connections simultaneously and can direct different workstation sessions to specific hosts.

To the supported workstations, the SNA gateway looks like an SNA PU 4 communications controller and forwards such host requests as BIND and UNBIND. The workstation LUs are not aware of the SNA gateway. The SNA gateway, however, is aware of all LUs at the workstations.

Downstream applications using standard SNA connectivity protocols for LU 0, 1, 2, and 3 and dependent 6.2, and communicating through an SNA gateway to a host,

are supported by Communications Server. Table 4 summarizes the SNA gateway features.

Table 4. SNA Gateway Summary

| Feature | Description |
|-------------------------------|---|
| Active workstations | 254 (LAN) per adapter 128 (X.25) |
| DLCs | AnyNet (SNA over TCP/IP) Twinaxial (upstream only) LAN (Any NDIS compliant network adapter) X.25 SDLC (synchronous, asynchronous, and AutoSync) OEM Channel (upstream only) MPC Channel (upstream only, requires DLUR) Enterprise Extender |
| Downstream workstations | Any product that supports standard SNA connectivity protocols for LU 0, 1, 2, 3, and 6.2. |
| Dynamic additions and changes | Yes |
| Implicit workstation support | Yes |
| LU pooling | Yes |
| Maximum number of LUs | 254 per PU; no limit on the number of PUs |
| Mode of operation | Multiple downstream PUs (not apparent to the host) PUs not visible to the host (except when through DLUR) |
| Multiple PU support | Yes |
| Segmenting support | Yes |
| Supported LU types | LU 0, 1, 2, 3, and dependent 6.2 |

Part 2. Planning and Installation

Chapter 3. Planning for AnyNet Support

Communications Server allows for two types of AnyNet support:

- AnyNet SNA over TCP/IP Access Node and Gateway
- AnyNet Sockets over SNA Access Node and Gateway

The AnyNet SNA over TCP/IP function in Communications Server enables SNA applications to communicate over interconnected IP and SNA networks.

The SNA over TCP/IP access node function enables SNA applications residing on an IP network to communicate. This function supports independent LU 6.2 and dependent LU 0, 1, 2, 3, or 6.2 either with or without dependent LU requester (DLUR). In addition, the SNA over TCP/IP access node can be used in conjunction with SNA gateway to enable SNA gateway sessions over TCP/IP.

The SNA over TCP/IP gateway function extends the reach of SNA applications by enabling SNA applications in an SNA network to communicate with SNA applications in an IP network. The SNA over TCP/IP gateway supports independent LU 6.2 sessions.

For more information on how to configure AnyNet SNA over TCP/IP, see “Configuring AnyNet SNA over TCP/IP”.

The Sockets over SNA access node function enables TCP/IP application programs using the WinSock 1.1 and WinSock 2.0 socket interface to communicate over an SNA network.

The Sockets over SNA gateway function enables sockets applications in SNA and TCP/IP networks to communicate. Sockets over SNA gateways are often used to connect isolated TCP/IP networks using an SNA backbone network.

Refer to “Configuring AnyNet Sockets over SNA” on page 77 for more information about configuring Sockets over SNA.

Configuring AnyNet SNA over TCP/IP

This section contains detailed information about configuring AnyNet SNA over TCP/IP.

Mapping SNA Resources to IP Addresses

One of the most crucial steps to enable SNA over TCP/IP communication is not done through Communication Server panels. Before SNA sessions or connections can be established, SNA over TCP/IP must determine the IP address of the partner. This is achieved by mapping the SNA identifier of the partner to an IP address using the following steps:

1. SNA over TCP/IP receives the SNA identifier from Communications Server in one of the following formats:
 - For LU names, `netid.luname`
 - For CP names, `netid.cpname`
 - For IDBLK and IDNUM values, configured as adjacent node ID, with two fields: block ID and physical unit ID.

AnyNet SNA over TCP/IP

2. SNA over TCP/IP takes the identifier and generates a domain name:
 - For LU names, `luname.netid.snasuffix`
 - For CP names, `cpname.netid.snasuffix`
 - For IDBLK and IDNUM values, `bbbnnnn.snasuffix`

Note: The default value for `snasuffix` is **SNA.IBM.COM**. For additional information on the SNA domain name suffix, refer to the online help.

Figure 20 shows examples of domain names generated by SNA over TCP/IP.

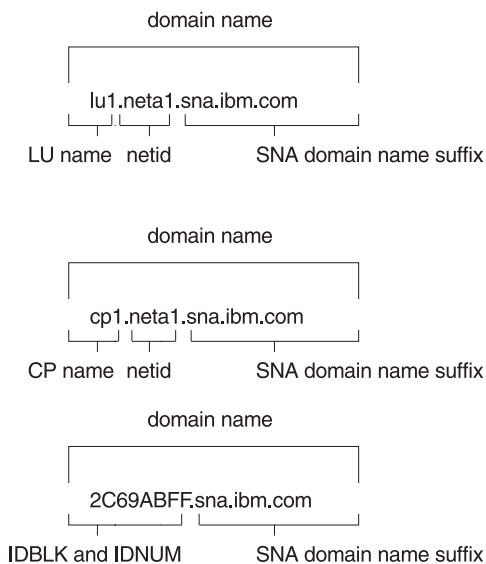


Figure 20. Formats of the Domain Names That SNA over TCP/IP Builds

3. SNA over TCP/IP requests that the domain name be translated to an IP address.
4. TCP/IP uses the HOSTS file or domain name server to translate the domain name into an IP address (for example, `9.67.192.28`).

When the IP network includes SNA over TCP/IP gateways, consider the following additional address mapping issues:

- If the partner LU can be reached through a gateway, the domain name of the partner must be mapped to the IP address of the gateway. If there are parallel gateways, the domain name should be mapped to each gateway IP address.
- To minimize the address mapping information needed for AnyNet, consider using a naming convention in which your LU name is the same as the TCP/IP host name. For example, suppose your computer's host name is `PATR.ANYNET.OURCORP.COM`. If you use the LU name `PATR` and the netid `ANYNET`, and then set the SNA domain name suffix to `OURCORP.COM`, AnyNet will ask TCP/IP to resolve the name `PATR.ANYNET.OURCORP.COM`. This name should already be in your domain name server.

Defining Domain Names and IP Addresses

This section describes the TCP/IP name resolution function, used by AnyNet to map SNA resources to IP addresses. This function queries both the local HOSTS file and any domain name servers to convert a domain name (for example, `lu1.neta1.sna.ibm.com`) into an IP address (for example, `10.1.1.1`).

HOSTS File

You can use TCP/IP HOSTS files to map domain names to IP addresses for your network. However, as your network becomes larger and maintaining the HOSTS file on each end-user workstation becomes too time-consuming, it is recommended that you use a domain name server.

The HOSTS file (in the `drivers\etc` subdirectory of your Windows NT system directory) lists:

- IP address
- Domain name
- Other aliases for the domain name

For example, if your IP address is 10.1.1.1, network ID is NETA1, SNA resource name is LUA1, and SNA domain name suffix is the default (`sna.ibm.com`), enter the following in your HOSTS file:

```
10.1.1.1    lua1.neta1.sna.ibm.com
```

Domain Name Server

Domain names and IP addresses can also be defined in the domain name server database.

Each SNA identifier is mapped to a corresponding IP address by a domain name server. The location of these servers is configured in the Network section of the Control Panel.

For more information on HOSTS files and domain name servers, refer to your TCP/IP documentation. If your workstation is using the TCP/IP support in Windows NT, refer to the online TCP/IP documentation that is included with the Windows NT product.

SNA over TCP/IP Gateway Considerations

The following information pertains to gateways but not to access node functions.

Defining Unique CP Names and Connection Network Names

For configurations that have two or more SNA over TCP/IP gateways connecting an SNA network with two or more IP networks, you need to define a unique SNA control point (CP) name and a unique SNA connection network name for each IP network.

All LUs that reside on access nodes in the IP network appear to reside on a node with this CP name.

Use the reverse data file of the domain name server or the HOSTS file to define the CP name and the connection network name for a given IP network. Map the IP address 127.0.0.3 to the CP name and map the IP address 127.0.0.4 to the connection network name.

The following example shows entries in the reverse data file. For an IP network with SNA network ID NETA, CP name MYCPNAME, and connection network name MYCNET, you would define the following entries:

```
127.0.0.3    NETA.MYCPNAME.
127.0.0.4    NETA.MYCNET.
```

Notes:

1. A period is required at the end of the name only if the definition is in the DNS reverse data file. No period is used in HOSTS file definitions.

AnyNet SNA over TCP/IP

2. Do not include the SNA domain name suffix.

The AnyNet SNA over TCP/IP function of Communications Server provides a default CP name (\$ANYNET.\$GWCP) and a default connection network name (\$ANYNET.\$GWCNET). In configurations with one IP network, you can use the default by not defining a CP name or a connection network name. In configurations with multiple gateways connecting multiple IP networks, one IP network can use the default. You must, however, define a unique CP name and connection network name for all other IP networks.

Figure 21 shows how to define the CP name and the connection network name for a configuration with two IP networks.

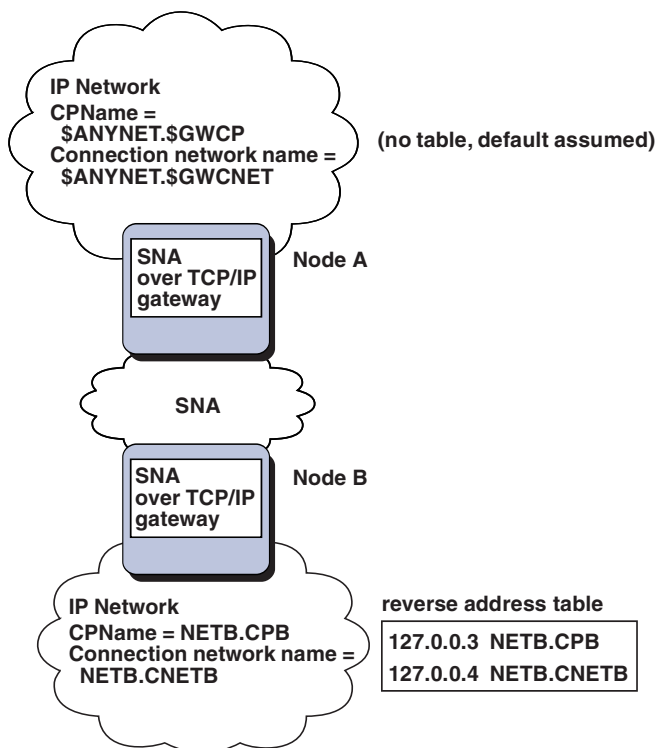


Figure 21. Defining a CP Name and a Connection Network Name

Using the Wildcard Entry to Reduce Domain Name Server Definitions

If you are using the SNA over TCP/IP gateway and your configuration meets the following naming restrictions, you can reduce the number of domain name server entries by defining a domain name entry for each SNA network ID that can be accessed through one or more SNA over TCP/IP gateways.

- Each SNA network must have an SNA network ID that is unique among all SNA networks.
- Each IP network must have an SNA network ID that is unique among all IP networks.
- An SNA network and an IP network can have the same SNA network ID.

By coding a single domain name entry for each SNA network ID, you do not have to define a domain name entry for every LU in the SNA network that you want to communicate with over the IP network. You can use a wildcard entry (*) to specify the LU name of all LUs that have the same SNA network ID. By substituting a

wildcard entry for the *luname* in the entry, you define a single domain name server entry that represents all LUs in that particular network.

Note: If you use the wildcard entry, you must use the full wildcard. Partial wildcards, such as `LUA*`, are not valid.

The wildcard entry is mapped to the IP address of the first SNA over TCP/IP gateway used to reach the network with that SNA network ID. As shown in Figure 22, logical units SNAAPPL1, APPC1, APPC2, and LU5 reside in network NETB and can only be reached from the IP network through an SNA over TCP/IP gateway with IP address IPgwg. If the SNA domain name suffix is SNA.IBM.COM, you define the following entry in the domain name server:

```
*.NETB.SNA.IBM.COM      IPgwg
```

This entry is used for all four logical units.

Note: You have the option of defining each logical unit individually.




Figure 22. Domain Name Server Definitions for a Single Gateway Connected to an SNA Network with Two Network IDs

Each SNA network must have a unique entry. As shown in Figure 22, if you also have SNAAPPLX, APPCX, and LUY in network NETC, which can only be reached through the SNA over TCP/IP gateway with IP address IPgwg, the domain name server entries are as follows:

```
*.NETB.SNA.IBM.COM      IPgwg
*.NETC.SNA.IBM.COM      IPgwg
```

In addition, each gateway must have a unique entry. If you add a parallel SNA over TCP/IP gateway, as shown in Figure 23 on page 70, with IP address IPgwh to the preceding example, the domain name server entries are as follows:

```
*.NETB.SNA.IBM.COM      IPgwg
*.NETC.SNA.IBM.COM      IPgwg
*.NETB.SNA.IBM.COM      IPgwh
*.NETC.SNA.IBM.COM      IPgwh
```




Figure 23. Domain Name Server Definitions for Parallel Gateways Connected to an SNA Network with Two Network IDs

SNA over TCP/IP Access Node Function Considerations

The following information pertains to access nodes only, and not to gateways.

How to Route SNA Sessions over AnyNet SNA over TCP/IP

When an SNA application initiates a session, Communications Server must first determine which transport to use - either SNA, IP, or a combination of the two.

You configure which transport you prefer by setting the routing preference. The routing preference can be set for the whole node via the default routing preference on the AnyNet over TCP/IP device or on a per LU basis when defining partner LUs.

The routing preference table is used only for new sessions. Previously existing sessions use the same transport; they are not brought down and rerouted if the routing preference table is changed.

Note: The routing preference for a node only governs sessions that are initiated from the node (access node sessions). Sessions that go through a node are not affected by the routing preference.

You can set or modify the default routing preference to one of the following:

Native first

Requests are routed over SNA. If no SNA route is available, requests are routed over TCP/IP.

Non-native first

Requests are routed over TCP/IP. If no TCP/IP route is available, requests are routed over SNA.

Native only

Requests are routed only over SNA. If no SNA route is available, the connection will fail.

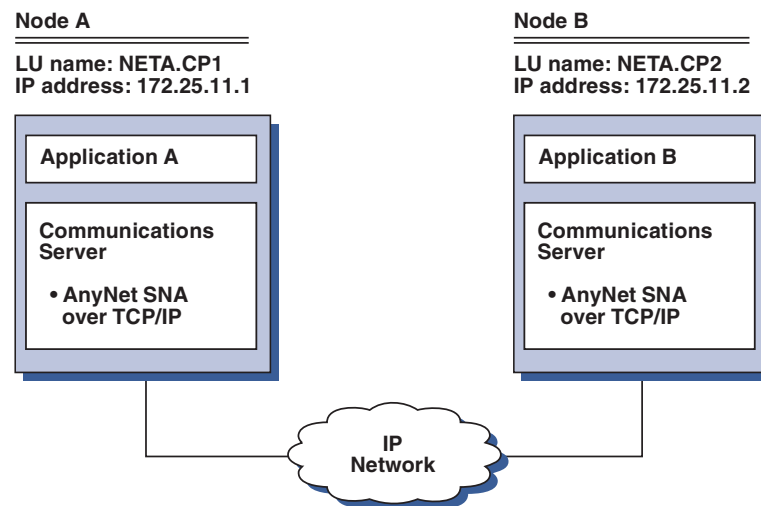
Non-native only

Requests are routed only over TCP/IP. If no TCP/IP route is available, the connection will fail.

AnyNet SNA over TCP/IP Configuration Examples

This section includes examples of AnyNet enabling SNA over IP communication. The following configuration instructions are complete only for the Windows NT operating system. In all examples, the SNA domain name suffix is SNA.IBM.COM.

For more information about configuring AnyNet for any other platforms mentioned in this section (such as VTAM or iSeries), refer to the appropriate product documentation.

Example 1. Running APPC or CPI-C Applications over a TCP/IP Network**Steps**

Follow these steps to establish communications between the two Windows NT nodes. Note that in this example, the CP names are used as LU names.

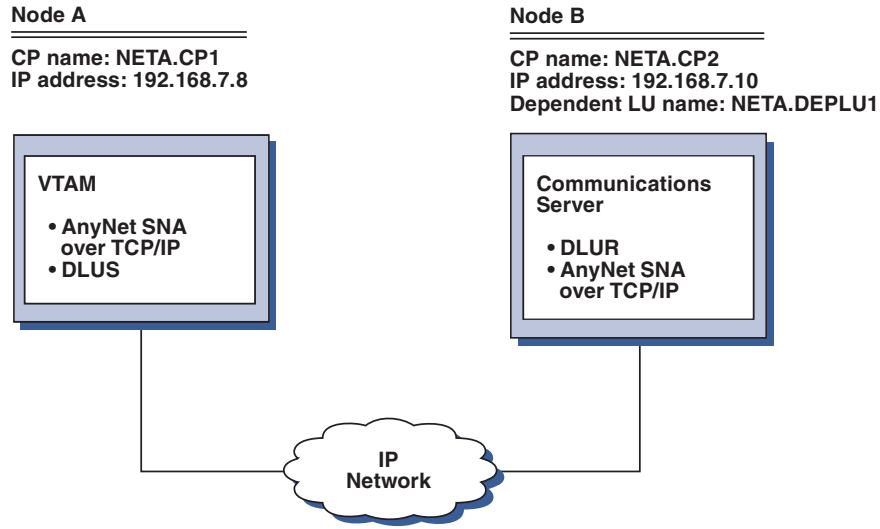
For Node A, do as follows:

1. Add the following entry to the local HOSTS file:
172.25.11.2 CP2.NETA.SNA.IBM.COM
2. Use **NETA.CP1** as the control point name during node setup. Ensure that the routing preference is set to route sessions over TCP/IP. Refer to the **SNA Node Configuration** help panels for more information.

For Node B, do as follows:

1. Add the following entry to the local HOSTS file:
172.25.11.1 CP1.NETA.SNA.IBM.COM
2. Use **NETA.CP2** as the control point name during node setup. Ensure that the routing preference is set to route sessions over TCP/IP. Refer to the **SNA Node Configuration** help panels for more information.

Example 2. 3270 Emulation via DLUR over a TCP/IP Network



Steps

Follow these steps to establish communications between Node A and Node B.

For Node B, do as follows:

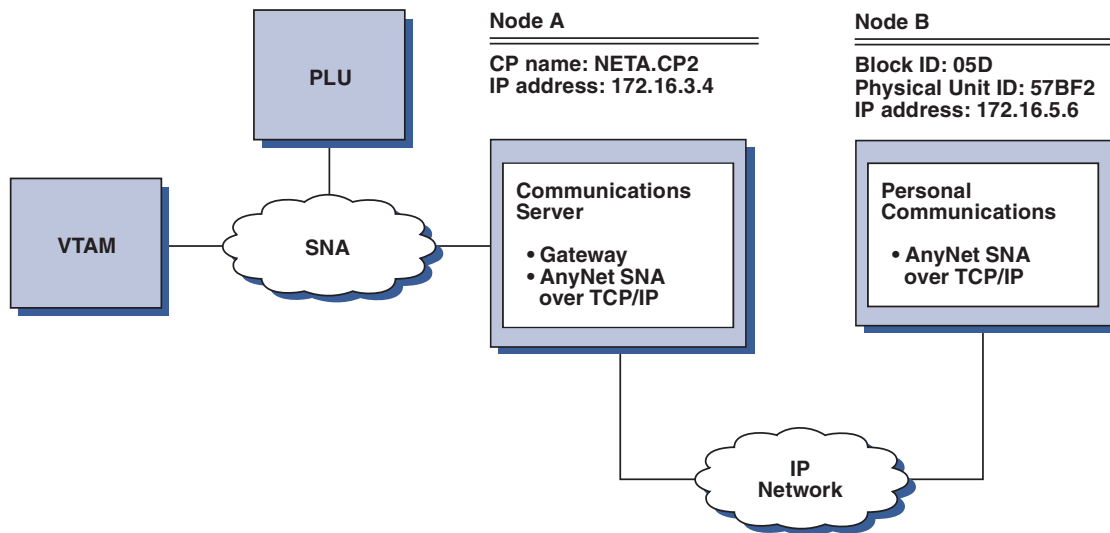
1. Add the following entry to the local HOSTS file:
`192.168.7.8 CP1.NETA.SNA.IBM.COM`
2. Use **NETA.CP2** as the control point name during node setup and **NETA.CP1** as the DLUS name when configuring DLUR PUs. Ensure that the routing preference is set to route sessions over TCP/IP. Refer to the **SNA Node Configuration** help panels for more information.

For Node A, add the following entries to the HOSTS file:

```
192.168.7.10 CP2.NETA.SNA.IBM.COM
192.168.7.10 DEPLU1.NETA.SNA.IBM.COM
```

Note that MVS AnyNet SNA over TCP/IP currently requires DLUS/DLUR for dependent LU communication.

Example 3. Using an SNA Gateway to Enable 3270 Emulation Between SNA and TCP/IP Networks



Steps

Follow these steps to establish communications between Node B and the VTAM host.

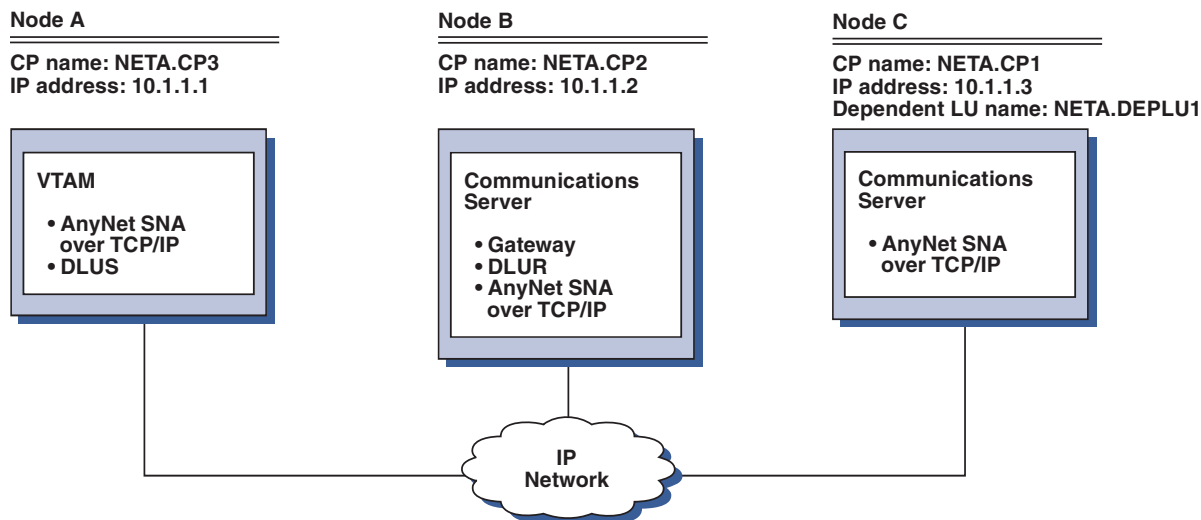
For Node A, do as follows:

1. Add the following entry to the local HOSTS file:
172.16.5.6 05D57BF2.SNA.IBM.COM
2. Use **NETA.CP2** as the control point name during node setup and use ANYNET device to assign implicit templates when defining clients. Refer to the **SNA Node Configuration** help panels for more information.

For Node B, add the following to the local HOSTS file:

```
172.16.3.4 CP2.NETA.SNA.IBM.COM
```

Example 4. Using an SNA Gateway for 3270 Emulation over a TCP/IP Network



Steps

Follow these steps to establish communications between Node C and Node A.

For Node B, do as follows:

1. Add the following entry to the local HOSTS file:
10.1.1.3 CP1.NETA.SNA.IBM.COM
2. Use **NETA.CP2** as the control point name during node setup, **NETA.CP1** as the adjacent CP name when defining the AnyNet SNA over TCP/IP connection definition, and **NETA.CP3** as the DLUS name when assigning the DLUS to a client template. Ensure that routing preference is set to non-native for NETA.CP3. Refer to the **SNA Node Configuration** help panels for more information.

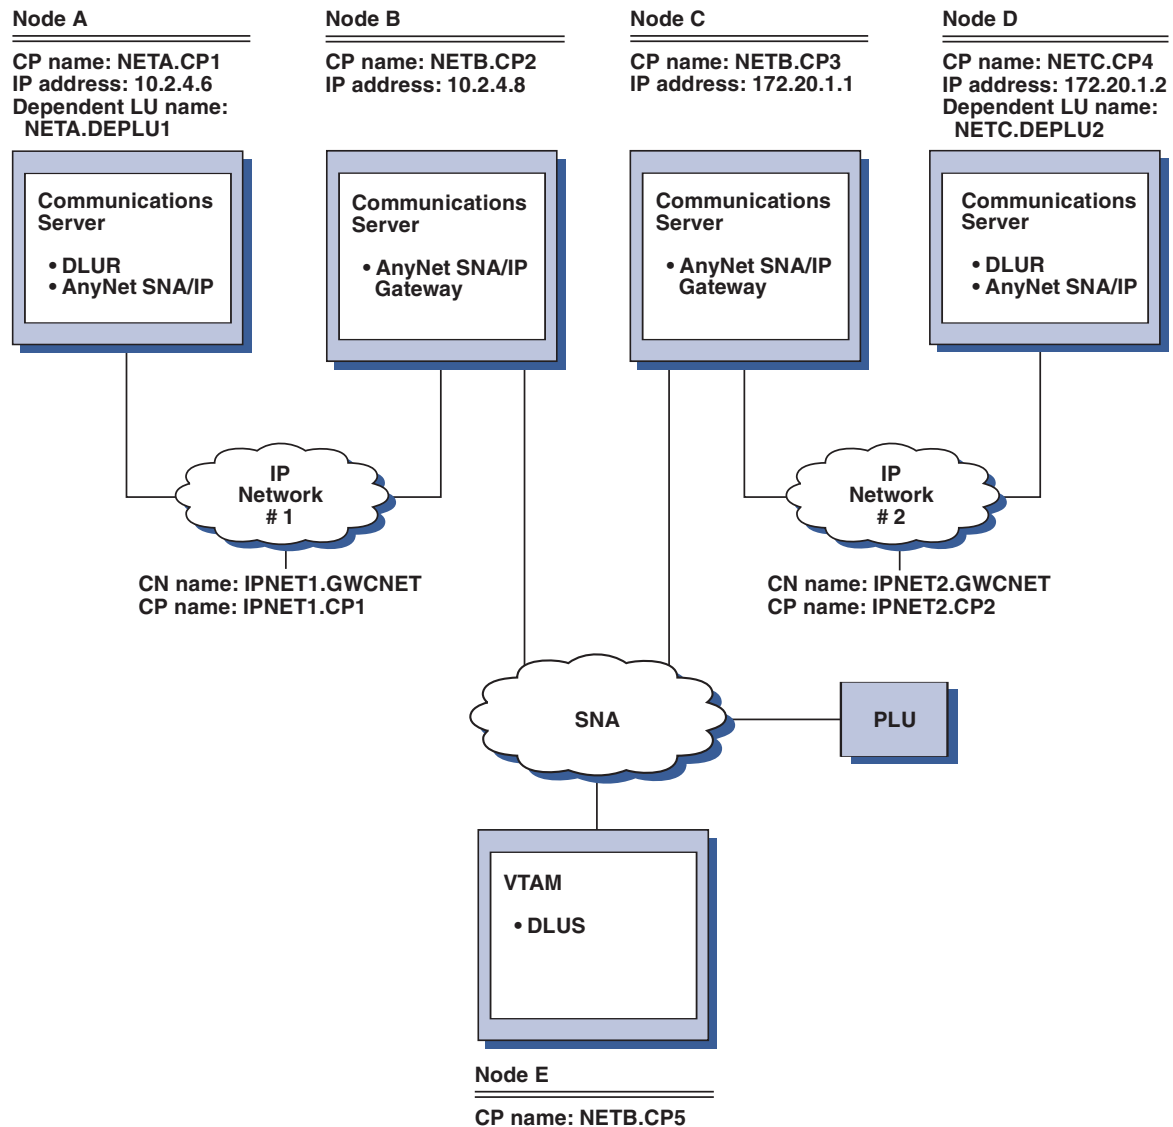
For Node C, do as follows:

1. Add the following entry to the HOSTS file:
10.1.1.2 CP2.NETA.SNA.IBM.COM
2. Use **NETA.CP1** as the control point name during node setup, and **NETA.CP2** as the adjacent CP name when defining the AnyNet SNA over TCP/IP connection definition. Refer to the **SNA Node Configuration** help panels for more information.

For Node A, add the following entry to the HOSTS file:

```
10.1.1.2 CP2.NETA.SNA.IBM.COM
10.1.1.2 DEPLU1.NETA.SNA.IBM.COM
```

Example 5. 3270 Emulation from Two Windows NT Workstations on Different IP Networks



Steps

Follow these steps to establish communications from the Nodes A and D to Node E.

For Node A, do as follows:

1. Add the following entries to the local HOSTS file:

```
10.2.4.8    CP5.NETB.SNA.IBM.COM
127.0.0.4  IPNET1.GWCNET
127.0.0.3  IPNET1.CP1
```

2. Use **NETA.CP1** as the control point name during node setup, and **NETB.CP5** as the DLUS name when configuring DLUR PUs. Ensure that the routing preference for **NETB.CP5** is set to non-native. Refer to the **SNA Node Configuration** help panels for more information.

For Node B, add the following entries to the HOSTS file:

AnyNet SNA over TCP/IP

```
10.2.4.6    CP1.NETA.SNA.IBM.COM
127.0.0.2  DEPLU1.NETA.SNA.IBM.COM
10.2.4.6    DEPLU1.NETA.SNA.IBM.COM
127.0.0.4  IPNET1.GWCNET
127.0.0.3  IPNET1.CP1
```

For Node C, add the following entries to the HOSTS file:

```
172.20.1.2 CP4.NETC.SNA.IBM.COM
127.0.0.2  DEPLU2.NETC.SNA.IBM.COM
172.20.1.2 DEPLU2.NETC.SNA.IBM.COM
127.0.0.4  IPNET2.GWCNET
127.0.0.3  IPNET2.CP2
```

For Node D, do as follows:

1. Add the following entry to the local HOSTS file:

```
172.20.1.1 CP5.NETB.SNA.IBM.COM
127.0.0.4  IPNET2.GWCNET
127.0.0.3  IPNET2.CP2
```

2. Use **NETC.CP4** as the control point name during node setup, and **NETB.CP5** as the DLUS name when configuring DLUR PUs. Ensure that the routing preference for NETB.CP5 is set to non-native. Refer to the **SNA Node Configuration** help panels for more information.

Helpful Hints

This section contains helpful hints on tuning, TCP/IP connectivity via SLIP or PPP, and dynamic IP addresses.

Tuning

If you can access an LU through multiple SNA over TCP/IP gateways, and you have mapped that LU name to multiple IP addresses, increase the **CONN_RETRY_SECS** and **CONNWAIT_SECS** parameters of the **ANYNET_COMMOM_PARAMETERS** keyword when configuring the AnyNet base parameters. This ensures that TCP connections will be attempted to all possible adapters and gateways. It can take as long as 90 seconds for a TCP connection to fail to an inactive IP address.

Dynamic IP Addresses

Generally, AnyNet SNA over TCP/IP depends on SNA resources (for example, LU names, CP names, or idblk/num) being statically mapped to IP addresses. However, depending on your configuration and how your connections are initiated, you might be able to use AnyNet SNA over TCP/IP in environments where IP addresses are dynamically assigned (for example, DHCP).

APPC or CPIC Applications: An SNA over TCP/IP access node with a dynamically assigned IP address may always initiate sessions to another SNA over TCP/IP access node or gateway with a static IP address.

The only way an SNA over TCP/IP access node or gateway with a static IP address (node A) can initiate a session to a partner with a dynamic IP address (node B) is:

1. Node B initiated a session to or through node A first.
2. The session initiated in Step 1 is still active.

Dependent LU Applications:

Note: The information in this section applies to LU 0, 1, 2, 3, or dependent 6.2 applications.

SNA over TCP/IP access nodes with dynamically assigned IP addresses can support dependent LU communications if the following criteria are met:

- Dependent LU communication is occurring through an SNA gateway either through:
 - Communications Server/2 4.1 with APAR JR10461 applied
 - Communications Server/2 V5
 - Communications Server for Windows NT 5.01 or higher
- The SNA over TCP/IP access node with a dynamically assigned IP address initiates connections to an SNA gateway with a static IP address.

Dependent LU communication through DLUS/DLUR over AnyNet SNA over TCP/IP is not supported if the DLUR node has a dynamically assigned IP address.

Configuring AnyNet Sockets over SNA

This section contains information about how to configure Sockets over SNA.

The Sockets over SNA access node function of Communications Server enables WinSock compliant applications to communicate over SNA networks. The Sockets over SNA gateway function enables socket applications in SNA and IP networks to communicate.

How Does Sockets over SNA Work?

Figure 24 shows the structure of a Windows NT node that is running Sockets over SNA and illustrates how socket application programs and Sockets over SNA operate on a Windows NT node.

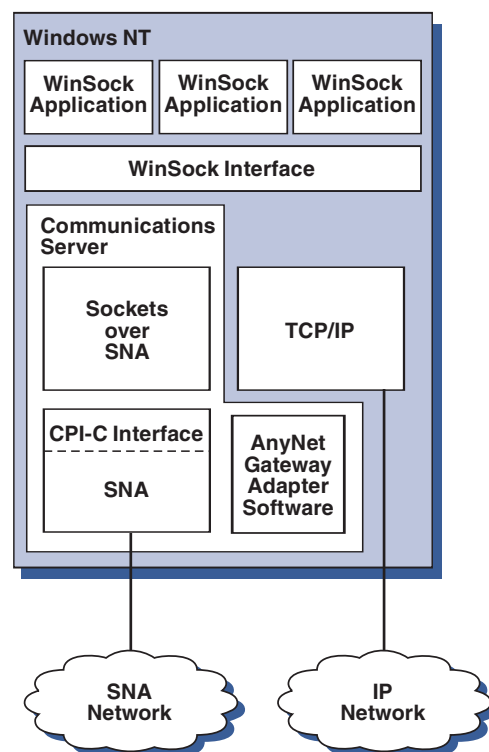


Figure 24. Structure of a Windows NT Node Running Sockets over SNA

AnyNet Sockets over SNA

WinSock is an API that lets socket applications run in a Windows environment.

Sockets over SNA does not provide a WinSock interface and does not process socket calls. Instead, WinSock applications use the WinSock interface of the native TCP/IP stack. The Sockets over SNA gateway code enables these applications to communicate across the SNA network.

Sockets over SNA gateway will allow socket applications running in an IP network to communicate with socket applications running on Sockets over SNA nodes. This is accomplished by routing packets between the SNA and IP networks and transforming them between SNA and IP protocols. The AnyNet gateway device driver assists in routing packets between TCP/IP and SNA networks, and the Sockets over SNA code converts between the two protocols.

Generating an LU 6.2 Call from a Socket Call

To enable TCP/IP-formatted information to route over SNA, Sockets over SNA maps IP addresses to SNA network-qualified LU names. When an application program invokes Sockets over SNA to establish a stream connection with another application program, Sockets over SNA establishes two half-duplex LU 6.2 conversations for the stream connection.

Sockets over SNA establishes one LU 6.2 conversation for all datagrams sent to a single destination. Conversations dedicated to datagram traffic are deallocated if they are unused for some specified period of time.

Mapping an IP Address to an SNA Network-Qualified Name

When an application program invokes Sockets over SNA to communicate with another application program, it supplies the IP address of the destination node. Sockets over SNA must map the IP address to an SNA address to issue an appropriate LU 6.2 call. For every IP address that identifies a node, there will be a corresponding SNA network-qualified name.

“Routing and Mapping Overview” on page 79 explains how address mapping works and provides guidelines and requirements for setting up IP-LU address mapping.

Routing and Mapping Data over SNA and IP Networks

Sockets over SNA gateways enable communication between socket application programs in IP and SNA networks by combining the routing function of TCP/IP with the protocol translation and address mapping capabilities of Sockets over SNA.

Protocol translation and address mapping are required when data is routed between nodes that use different transport protocols. The Sockets over SNA gateway automatically performs protocol translation after determining the type of transport associated with the destination IP address. For a summary of the routing and mapping process, see “How the Sockets over SNA Gateway Routes and Maps Data” on page 81.

Application Program Support Provided by Sockets over SNA

Sockets over SNA supports WinSock 1.1 and WinSock 2.0 (Windows NT 4.0 only) applications that use AF_INET sockets.

Sockets over SNA does not support applications that use broadcasting.

If you intend to use Sockets over SNA gateway to route information to and from an MVS/ESA node configured with the VTAM V3R4.2 Sockets over SNA function, you must first install the route function on the MVS/ESA node. To install the route function on MVS/ESA, install the program temporary fix (PTF) UW03567. You can obtain PTFs through any of the following sources:

- Information Access
- SoftwareXcel Extended
- IBMLink (ServiceLink)

If you do not have access to these sources, contact the IBM Support Center.

Planning for Sockets over SNA

This section describes what the network planner should consider before configuring a network with Sockets over SNA.

Routing and Mapping Overview

This section explains basic concepts of Internet addressing and how these concepts relate to routing and mapping. It includes the following information:

- “Internet Addressing”
- “IP Routing Table” on page 80
- “SNA Network ID Used by Sockets over SNA” on page 80
- “How an IP Address Is Mapped to an LU Name” on page 80
- “How the Sockets over SNA Gateway Routes and Maps Data” on page 81

Internet Addressing

Every host is assigned at least one unique Internet Protocol (IP) address, which is used to route data through the network.

Note: In the IP suite of protocols, *host* refers to an end system and can be any workstation; it does not have to be a mainframe.

The IP address assigned to the host does not define a host on the network; it defines a network interface on that host to a network. For example, the address of the SNA network interface identifies a node’s connection to the SNA network.

A gateway host has a unique IP address for each network interface. Because Sockets over SNA gateway routes SNA and TCP/IP data, you must set up unique IP addresses for the TCP/IP and SNA interfaces.

The following section describes the IP address format, address classes, and network masks. For more detailed information, see your TCP/IP documentation.

IP Address Format and Classes: An IP address consists of a 2-part, 32-bit address field:

- The first part of the address field contains the network address; the second part contains the host address.
- The number of bits used for the network and host portions of an IP address are variable, depending on the address class of the IP address.
- A network mask enables you to use part of the host portion of the IP address as a subnetwork address.

AnyNet Sockets over SNA

Default network masks are shown in Table 5.

Table 5. IP Address Masks Supported by Sockets over SNA

| For a dotted decimal IP address in the form <i>a.b.c.d</i> , the range of values for <i>a</i> is: | Default Network Mask |
|---|----------------------|
| 1–127 | 255.0.0.0 |
| 128–191 | 255.255.0.0 |
| 192–223 | 255.255.255.0 |

Masks Used by Sockets over SNA: Sockets over SNA uses two types of masks:

- **Subnet mask**

The subnet mask is used in routing and is specified during the configuration of the local node and routes. You can accept the default subnet mask or specify a value other than the default to define subnetwork addresses.

- **Address mask**

The address mask is used for generated IP-LU address mapping and is specified during configuration.

IP Routing Table

Each host has an IP routing table that stores information about possible destinations and how to reach them. Route entries are added when:

- The local IP address is defined. For more information, see the online help.
- Route statements are defined from the **SNA Node Configuration** window. Click **Configure AnyNet Sockets** and then click **Routes**. For more information, see the online help.
- A route with fewer hops is found. This occurs when Sockets over SNA gateway receives an ICMP redirect or RIP message. For more information, see “Route Discovery Function” on page 81.

For a routing table example, see Figure 25 on page 81.

SNA Network ID Used by Sockets over SNA

For each route you define through the SNA interface (sna0), there must be a corresponding SNA network ID to which the IP network address is mapped. The number of SNA network IDs you define depends on how you want to map the IP network to the SNA network.

For example, if the socket applications using SNA are configured to use IP subnetworks 9.67.0.0 and 9.77.0.0, you can define an SNA network ID that corresponds to each IP subnetwork, or you can define one SNA network ID that corresponds to both subnetworks. Sockets over SNA does not require a unique one-to-one mapping between an IP network address and an SNA network ID.

How an IP Address Is Mapped to an LU Name

You can use either explicit or generated mapping to map IP addresses to SNA LU names:

- In explicit mapping, all LU names are defined, and you assign each IP host address to a specific SNA LU name during configuration.
- In generated mapping, LU names are automatically generated for you. During configuration, you specify a network IP address, address mask, network ID, and LU template value.

Sockets over SNA uses the address mask to map the network portion of the IP address to the SNA network ID and the host portion to the SNA LU name. The LU template value is used to determine the characters and the positions of characters used in the LU name.

You can display the generated LU name for a given IP address using the **sxmap** command line utility. The syntax for this utility is:

```
sxmap convert <IP address> <address mask> <LU template>
```

How the Sockets over SNA Gateway Routes and Maps Data

The following steps briefly describe how Sockets over SNA Gateway determines whether to route data over SNA or TCP/IP, and how address mapping is handled:

1. Sockets over SNA searches its routing table to find a route that enables data to reach the destination IP address. If Sockets over SNA does not find any matching routes, the connection request is forwarded to the native TCP/IP stack.
2. If Sockets over SNA finds a matching route, the route entry indicates how the destination can be reached:
 - a. If the router address is the address of a local network interface, such as `sna0`, the destination network, subnetwork, or host address can be reached directly.
 - b. If the router address is the address of a gateway or router, the destination can only be reached through that intermediate gateway or router.

Figure 25 shows an example of an IP routing table.

| Destination IP Address | Destination Mask | Gateway IP Address | Use Count |
|------------------------|------------------|--------------------|-----------|
| 1.2.3.4 | 255.255.255.255 | 199.245.253.1 | 10 |
| 10.0.0.0 | 255.0.0.0 | 199.245.253.2 | 0 |
| 10.11.0.0 | 255.255.0.0 | 199.245.253.113 | 37 |
| 127.0.0.1 | 255.255.255.255 | 127.0.0.1 | 8 |
| 128.1.0.0 | 255.255.0.0 | 199.245.253.3 | 0 |
| 199.245.253.0 | 255.255.255.0 | 199.245.253.113 | 368 |

Figure 25. Example of an IP Routing Panel

3. If no route is found in the Sockets over SNA routing table, then Sockets over SNA assumes that the TCP/IP destination can be reached through a native IP network. Refer to your TCP/IP documentation for more information on how TCP/IP routes data.
4. If the chosen route indicates that data should go over the SNA interface (`sna0`), Sockets over SNA looks up the next-hop address in the IP-LU mapping table:
 - a. If Sockets over SNA finds a matching entry, an LU 6.2 connection is established.
 - b. If Sockets over SNA does not find a matching entry, the connection attempt fails.
 - c. Sockets over SNA passes the destination address and data to Communications Server.
5. All routes defined to Sockets over SNA are mirrored to the native TCP/IP stack so that packets from IP can be routed over SNA.

Route Discovery Function: The route discovery function provided by Sockets over SNA gateway can help you route TCP/IP traffic more efficiently and reduce the number of explicitly defined route statements in your network. You do not have to select or configure this function.

AnyNet Sockets over SNA

One of the problems with large networks is how to discover the addition of new networks or subnetworks and which router to use to get to the new network or subnetwork. Sockets over SNA solves this problem by having all nodes initially use a default router that notifies other nodes when a more direct route is discovered. This is more efficient than using the typical TCP/IP solution of broadcasting routing information.

Note: To effectively use this function, algorithmic mapping of IP addresses to LU names and an APPN backbone network should be used. Otherwise, nodes must explicitly define LU names and IP addresses for all remote nodes with which they communicate.

Figure 26 shows a configuration example.

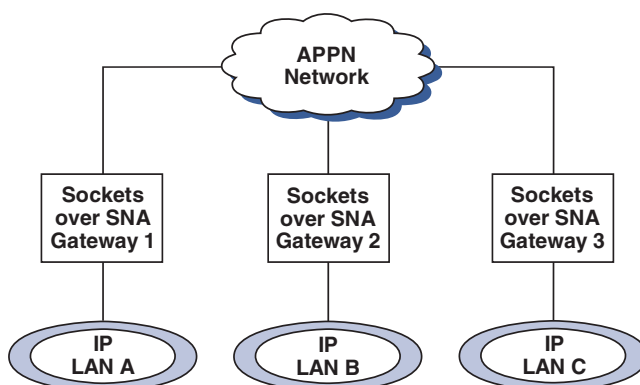


Figure 26. Example of a Network Using the Sockets over SNA Route Discovery Function

In this scenario:

- Gateways 1, 2, and 3 are Sockets over SNA gateways connected to IP LANs. These gateways connect the IP LANs to an APPN backbone network.
- Gateway 1 is the only gateway where you maintain a complete, permanent routing table.

Gateways 2 and 3 define Gateway 1 as their default router. If a remote network or subnetwork is known to Gateway 1, Gateways 2 and 3 do not have to explicitly define these routes.

- When Gateways 2 and 3 route data to undefined networks or subnetworks, these requests are sent to their default router, Gateway 1.
If the network or subnetwork is known to Gateway 1 and a more direct path is available, Gateway 1 sends an ICMP redirect message back to the requester indicating the path to take in the future. This ICMP redirect message updates the requester's routing table. Therefore, Gateways 2 and 3 dynamically build their routing tables for remote networks and subnetworks as needed.
- If new gateways are added, update the permanent routing table of Gateway 1. No explicit route statements are needed for Gateways 2 and 3.

Defining Sockets over SNA Modes

Sockets over SNA uses LU 6.2 conversations to enable communication between socket application programs. When an LU 6.2 conversation is established, Sockets over SNA defines the mode and associated session characteristics of the connection. Communications Server uses the mode name to identify characteristics of the connection between the two Sockets over SNA nodes.

The default Sockets over SNA mode is BLANK. You can use the default mode for Sockets over SNA or define your own. To change the default Sockets over SNA mode, from the **SNA Node Configuration** window, click **Configure AnyNet Sockets Over SNA** and then click **Modes**. You can define another default mode for all TCP/IP traffic and assign a specific mode to a specific TCP/IP port.

If you specify an alternative mode that is not defined by Communications Server, you must define the session characteristics associated with that mode to Communications Server.

Changing the Idle-Timeout Interval

The idle timeout start option enables you to adjust the number of idle seconds before Sockets over SNA deallocates a datagram conversation. This interval enables you to balance using system resources to maintain an existing datagram conversation and taking longer to reestablish a new datagram conversation. For example, if you set this value low, unused datagram conversations end faster, but it takes longer to send the next datagram. The default idle-timeout interval is 90 seconds.

To modify the start option, from the **SNA Node Configuration** window, click **Configure AnyNet Sockets over SNA**, then click **View/Change/Add**, click the **Advanced** tab and select a new value for this option.

Chapter 4. Planning for Client/Server Communication

This chapter contains information about planning for SNA API clients for Communications Server and Novell NetWare for SAA.

SNA API Clients

This section contains information specific to SNA API clients.

Installation and Configuration

For information about installing and configuring Communications Server SNA API clients, refer to *Quick Beginnings*.

Lightweight Directory Access Protocol

Lightweight Directory Access Protocol (LDAP) provides a standard way to access directory services. LDAP enables you to read, search, add, and remove information from centralized or distributed databases. LDAP is similar to the older X.500 Directory Access Protocol (DAP), but without some of the overhead (such as session and presentation OSI layers). You can use **LDAP Client Configuration** to configure SNA API client information and store it on an LDAP server.

LDAP Client Configuration shows the logical organization and relationships between the clients and their configurations.

For more information on using these major areas, refer to the help panels for **LDAP Client Configuration**.

Directory Information Tree

The Directory Information Tree (DIT) connects all Directory Service Agents (DSAs) in a predefined model. The model is an extensible hierarchical structure consisting of a number of base objects. A typical LDAP DIT would have a root, under which *countries* (c) are defined. Below the countries are usually *organizations* (o), and below an organization are *individuals*, ((cn) or (uid)), or *organizational units* (ou). For example, an fully-distinguished LDAP entry would be specified as:

```
c=US, o=company.com, ou=Sales, cn=temp
```

Client Hierarchy

The client hierarchy is relative to the LDAP DIT you logged in to. You can add users or organizational units to the client hierarchy or delete users or organization units from the client hierarchy.

Configuration Hierarchy

The configuration hierarchy logically shows, through a graphical tree view, the client configuration and all the relationships between the configuration definitions. The configuration hierarchy enables you to create, modify, or remove definitions.

Common API Client for Communications Server for Windows and NetWare for SAA

The client shipped with Communications Server permits clients to locate and access IBM Communications Servers and NetWare for SAA (NWSAA) servers in the network.

Planning for Client Server Communication

Install Options for Configuration Location

During installation of the client, the user can choose the client configuration location. There are three choices:

- INI file configuration, which is the method used with previous versions of Communications Server and is the default method when installing the SNA API client that is included with Communications Server
- Netware Directory Services (NDS), which is the method used with previous versions of NWSAA and is the default method when installing the SNA API client that is included with NetWare for SAA
- Lightweight Directory Access Protocol (LDAP)

The configuration location contains server lists, user names, and all other client configuration information.

Configuration Locations

Because configuration information is independent of the server type (Communication Server for Windows or NetWare for SAA), the information can be stored in either NDS, INI, or LDAP. For example, Communication Servers names and addresses can be stored in NDS, INI, or LDAP for the API client. Since NDS administration programs are only shipped with NetWare for SAA, configuration of SNA API client information in NDS requires that the NetWare for SAA product exist on the network. The common SNA API client enables you to easily integrate different server types on the same network.

Configuration Utility

Use the **Set Configuration Options** utility to change the client configuration location (INI, NDS, or LDAP) after installation.

Locating Servers

The client enables the user to locate and access IBM Communications Servers and NWSAA servers using both TCP/IP and SPX transport protocols.

TCP/IP: The method used to locate servers using the TCP/IP transport protocol is identical. The network administrator can control access to the servers through server configuration. See “Planning for TCP/IP Scopes” on page 124 for more information.

SPX: The methods used to locate servers using the SPX transport protocol are different. Locating both IBM Communications Servers and NWSAA servers takes more time than locating only IBM Communications Servers or NWSAA servers.

The Communications Server SNA API client tries to locate IBM Communication Servers first, followed by NWSAA servers. If no NWSAA servers exist in the network or should not be accessed, the user can set an environmental variable to prevent locating NWSAA servers. Setting the variable improves performance in an environment where only IBM Communication Servers are accessed by the SPX transport protocol. The environmental variable is **SNA_API_CLIENT_NO_CC**. Setting the variable does not affect server location using the TCP/IP transport protocol.

Assigning a Default Local LU

You can assign a default local LU alias for each user through either the INI file configuration or LDAP configuration for 32-bit Windows SNA API clients.

APPC programs can use a default local LU alias rather than specify one directly. When an APPC program issues a TP_START verb with the local LU alias field set to binary zeroes or all ASCII blanks, the APPC API uses the configured default local LU alias.

EHNAPPC programs can use a default local LU alias rather than specify one directly. When the local LU field on the iSeries configuration panel is left blank, the configured default local LU alias is used.

CPI-C programs can use a default local LU alias rather than specify one directly. When the local LU alias field in the CPI-C side record is left blank, and no value has been defined using the APPCLLU environment variable, the CPI-C API uses the configured default local LU alias when initiating a conversation.

Attach Manager started programs can use a default local LU alias rather than specify one directly. When the local LU alias field in the attach manager record is left blank, the Attach Manager uses the configured default local LU alias when processing incoming conversation requests.

Assigning a Default Partner LU

You can assign a default partner LU alias for each user through either the INI file configuration or LDAP configuration for 32-bit Windows SNA API clients.

APPC programs can use a default partner LU alias rather than specify one directly. When an APPC program issues an ALLOCATE verb with the partner LU alias field and the fully qualified partner LU field set to binary zeroes or all ASCII blanks, the APPC API uses the configured default partner LU alias.

CPI-C programs can use a default partner LU alias rather than specify one directly. When the partner LU alias field in the CPI-C side record is left blank, the CPI-C API uses the configured default partner LU alias when initiating a conversation.

Assigning a Default LUA Session Name

You can assign a default partner LU alias for each user through either the INI file configuration or LDAP configuration for 32-bit Windows SNA API clients.

LUA programs, such as 3270 emulators, can use a default LUA session name rather than specify one directly. When a LUA program issues an RUI_INIT or SLI_OPEN verb with the LUA LU name field set to binary zeroes or all ASCII blanks, the RUI/SLI API uses the configured default LUA session name.

Client/Server Data Encryption

Client/server data encryption enables the client to request encryption of the application data in communications between the SNA API client and Communications Server.

Data encryption can be on, off, or optional. Optional data encryption means that encryption is used if it is supported by the server.

Client Server Security

Communications Server enables Windows 95 and Windows NT SNA API clients to use Windows NT domain security to authenticate the client connection to the server without reentering the userid and password. The client must be part of a

Planning for Client Server Communication

Windows NT domain, either by participation in a Communications Server domain or logging in locally with a synchronized userid and password.

SNA API client users outside of the Windows NT domain are required to provide the userid and password, either through a prompt or by storing these values in the client configuration file.

The authorized users for client server are maintained in the IBMCSAPI local group, which is located either directly on the Communications Server or on the domain controller where Communications Server participates. This user group is created during installation and can be administered using the Windows NT User Manager application.

You can determine whether the userid and password are required for connection to the server by changing the **LogonControl** value entry at the server to zero (0), which removes the requirement of the userid and password specification by the client.

Audit trail records of client connections to the server are tracked in the Windows NT Event Log. You can view the records using the Windows NT Event Viewer under the **Application** log. If you do not want to track client connections to the server, you can change the **AuditTrail** value entry at the server to zero (0).

You can access the **LogonControl** and **AuditTrail** value entries through the Windows NT Registry Editor under the key:
HKEY_LOCAL_MACHINE/SYSTEM/CurrentControlSet/Services/IBM SNA Client Services.

Novell NetWare for SAA Clients

Communications Server supports IPX- or TCP/IP-attached clients running emulator software packages that implement Novell's Queue Element/Message Unit (QEL/MU) architecture for 3270 emulation, enabling the clients to access mainframe host data. This includes support for popular client features, including dedicated, pooled, and public LU categories (sometimes referred to as resource types).

Communications Server supports Novell NetWare for SAA clients on Windows 95, Windows NT, Windows 3.1, and OS/2.

TCP/IP, IPX/SPX, or both are required for Novell NetWare for SAA clients to communicate with Communications Server. If you are running IPX, you also need the following software enabled on the server:

- Gateway Service for Netware
- SAP agent
- NWLink IPX/SPX or a compatible transport

If Communications Server is running on a primary or backup controller, users must be given user rights in the IBMCSAPI group to log on locally to the server.

Chapter 5. Planning for Legacy Data Access

This chapter contains information about functions within Communications Server that enable you to access host information.

iSeries OLE DB Provider

Applications that use OLE DB or ActiveX can communicate through Communications Server for record-level access to files on iSeries hosts. Access to iSeries files requires you to define the remote iSeries, where the files are stored, as a partner LU.

The iSeries OLE DB Provider can be installed as a component of Communications Server, or as a component of the SNA API client for Windows 95 or Windows 98. It can also be installed as a component of the SNA API client for Windows NT or Windows 2000; however, it requires Windows NT Workstation or Server 4.0 or later or Windows 2000. If you do not have the Microsoft Data Access Components (MDAC) Version 2.0 installed, Communications Server will install it for you when you install this component.

If you install the iSeries OLE DB Provider on a machine that already has Client Access Version 3 Release 2 M0 (or earlier) installed, some Client Access function will be disabled, including support for ActiveX and OLE DB programs that use Client Access for access to iSeries data queues, remote commands, distributed program calls, stored procedures, and SQL statements.

If you want to continue using these functions through Client Access Version 3 Release 2 M0 (or earlier), do not install the iSeries OLE DB Provider component during the installation of Communications Server. If you install the iSeries OLE DB Provider component during the installation of Communications Server, and you want to recover the Client Access functions, you can re-install Client Access. However, if you re-install Client Access, you can not use iSeries OLE DB Provider over Communications Server.

When Client Access Version 3 Release 2 M0 (or earlier) is installed after Communications Server, registry information is overwritten. If you want to use iSeries OLE DB Provider over Communications Server after Client Access is installed, you can run the REGSVR32.EXE tool against the cwazzodb.dll and cwazzidx.dll files to restore the iSeries OLE DB Provider registry information.

You can use the iSeries OLE DB Provider to connect Windows clients through Communications Server to one or more iSeries hosts, as shown in Figure 27 on page 90. This configuration gives clients that have the iSeries OLE DB Provider installed access to specific records within files on the iSeries.

Planning for Legacy Data Access




Figure 27. iSeries OLE DB Provider

In this figure, all clients have the iSeries OLE DB Provider installed. Alternatively, you can install and use the function from the server.

Documentation for this function, as well as information about developing OLE DB and ActiveX applications using Client Access, is provided in the `ibmcs\sdk\pubs\oledb` directory.

iSeries Shared Folder Server

You can create disk devices on the server that communicate with iSeries folders through the iSeries Integrated File System (IFS). If the server shares these disk devices, clients can connect to and use these devices to get access to iSeries drives without additional configuration or code installation on the client machine. Multiple clients can connect to folders on the iSeries system as if they were drives on their workstations.

You enable shared folder support through the **SNA Node Configuration** application. When a disk device is available to the user, it acts as a system supported disk drive. Use the interfaces provided by Windows NT to control the device. Users can share the drive across the network, but maintain individual access rights.

The iSeries shared folder server requires OS/400 Version 3.1 or higher.

Chapter 6. Planning for HPR

Communications Server supports high performance routing (HPR) over Enterprise Extender (IP), synchronous data link control (SDLC), LAN, WAN, channel, Multi-Path Channel (MPC), and X.25 connections.

HPR automatic network routing (ANR) minimizes the storage and processing requirements in intermediate nodes, which is an improvement over APPN intermediate session routing (ISR) for high-speed networks with low error rates.

Communications Server HPR support includes support for rapid transport protocol (RTP) and automatic network routing (ANR).

HPR Link Characteristics

To provide the benefits listed above, HPR assumes that data links provide low rates of lost frames. HPR recovers from conditions that cause lost frames, but if the conditions persist for long periods, HPR assumes that network congestion is too high and HPR throughput is lowered. Achieving good HPR throughput on a given link involves proper configuration of the sending node, the facilities that constitute the link between the two nodes, and the receiving node.

Note: You might want to set ERP values in cases where LAN line facilities have high frame loss.

Sending Node

Typical adapters used with HPR will send frames onto the line with low rates of failure. For Ethernets, HPR behavior will be degraded if the LAN is heavily loaded with many collisions (for example, if the load is more than about 30% to 40% of the nominal Ethernet speed). For frame relay, facilities should be configured so that individual connections can burst to the access line rate without causing congestion conditions in the frame-relay network.

Line Facilities

LAN facilities can have substantial hidden complexity, ranging from large numbers of stations sharing a single LAN segment to bridges between LAN segments, perhaps with hidden lower speed links, to simulated LAN segments that operate over complex subnets. The low frame loss assumption must be met. For example, WAN links between split bridges must use facilities with low error rates or must compensate by using link level error recovery between the bridges. As another example, conditions that cause bridge buffers to overrun (broadcast storms or large bursts of broadcast frames) must be prevented.

Receiving Node

Typical adapters used with HPR will receive frames from the line with low rates of failure. Some adapters, however, will fail to receive frames when consecutive frames arrive too quickly. The best HPR performance is realized when the receiving adapter can handle its traffic without missing inbound frames.

Planning for HPR

If your network has a high error rate, systems administrators may notice lower than expected throughput on HPR connections.

HPR Considerations

The HPR features are achieved by including additional information in frame headers. If you do not need the HPR features, the additional overhead may reduce your network performance. The HPR features are an advantage whenever there are intermediate nodes or alternate routes. However, if you require high performance between two nodes connected by a single link without non-disruptive recovery during a link failure, you should disable HPR when you configure that link.

HPR implements priority queuing. In general, interactive traffic passes through the network faster than lower priority batch traffic. Adapters that queue frames do so without regard to priority and may interfere with HPR's priority queuing. You may need to reduce the adapter's queue depth to prevent low priority traffic from delaying high priority traffic. (See the adapter documentation for instructions on setting the adapter queue depth.)

Many PCs now have advanced power management features that will suspend power to devices (like disk drives) and reduce the clock rate on a CPU significantly to save power if the keyboard or mouse is not being used. On a workstation, this is a good feature, but on a server platform, this can occur while it is routing data from one adapter to another (and no disk activity is occurring). For Communications Server, disable advanced power management and hardware power management by entering the BIOS configuration during boot up (usually pressing F1 while initializing hardware) and disable these options.

HPR will use more of the LAN adapter memory due to increased performance and queuing. Many shared RAM adapters default to using 8 kilobytes of shared RAM. It is required that the shared RAM size is 16 kilobytes. To adjust this configuration, refer to the adapter configuration manual. This setting is located in either your machine BIOS, or on the adapter switch settings.

Note: If you do not increase the shared RAM to 16 kilobytes, network performance is degraded.

Do not adjust the shared RAM value above 16 kilobytes, the optimal setting for the adapter memory.

If HPR links are failing under heavy traffic, you can set your receive timers to a higher value. This allows more time for the links to receive an acknowledgement for sent frames.

For SDLC links, this is the Primary Receive Timer. The receive timer can be set on either the DLC or a logical link statement.

To calculate the minimum receive timer value in tenths of seconds, use the following formula:

$$[(\text{send_window_count} * \text{max_i_field_size} * 8) / \text{actual line speed}] * 10$$

Considerations for Tuning Parameters on LAN DLCs

You can specify eight parameters to affect how quickly Communications Server handles recovery and failure of LAN links. You might need to adjust these parameters to achieve the best results if you are using slow links, HPR, or backup host connections.

You should consider the following parameters when defining a LAN device in **SNA Node Configuration**:

Under the **Advanced** tab:

Test retry interval

The test retry interval specifies the time between attempts to find the adjacent link station on the local area network (LAN). The number of attempts made are based on the value specified for the test retry limit.

The default value is 8.

Test retry limit

The test retry limit specifies the maximum number of attempts to find the adjacent link station on the local area network (LAN) without receiving an acknowledgment in the time set by the value for the test retry interval.

The default value is 5.

Maximum PIU size

The maximum path information unit (PIU) size is the maximum number of bytes that are contained in the data buffer used by SNA sessions for this link. If the value of the maximum PIU size exceeds the frame size supported by your local device driver, the value will be reduced to match the frame size.

The default value is 65 535.

Under the **Performance** tab:

Idle timeout

The idle timeout specifies the time that the LAN device driver waits for a frame to be received before declaring the link to be inoperative.

The default value is 30 seconds.

Acknowledgment delay

The acknowledgment delay specifies the time that the LAN device withholds a response to a received frame in order to allow more frames to be received and acknowledged with the same Request Ready (RR).

The default value is 100 milliseconds.

Poll response timeout

The poll response timeout specifies the time that the LAN device waits for a response to a frame sent with the POLL bit set.

The default value is 8 000 milliseconds.

Anticipated outstanding transmits

The anticipated outstanding transmits specifies the maximum number of frames the LAN device queues to a link station before sending a Receive Not Ready (RNR) to the adjacent link station.

The default value is 16 frames.

Planning for HPR

Receive buffer count

The receive buffer count specifies the number of buffers that are reserved in memory to hold data received from the host until it can be processed. Each buffer is the size of the PIU.

The default value is 32 buffers.

HPR Path Switch Processing

HPR provides a non-disruptive path switch to a different connection if one connection fails. A retry of the failed connection is attempted before a path switch occurs. An explanation of the parameters involved in determining whether a connection has failed and the time to initiate a path switch follows.

When the value specified for **Idle timeout** is reached, a keepalive frame is sent. The default for **Idle timeout** is 30 seconds.

The value specified for **Poll response timeout** determines the time that the device waits for a response to the keepalive frame. The default for **Poll response timeout** is eight seconds (8 000 milliseconds). A keepalive frame is retried 10 times. This retry count can not be changed.

If there is no response to the keepalive frame, TEST commands are sent to the partner system. The value specified for **Test retry interval** determines the time between TEST command attempts, and the value specified for **Test retry limit** determines the maximum number of TEST command attempts. The default for **Test retry interval** is 8 seconds, and the default for **Test retry limit** is five attempts.

The following calculation indicates the time required to determine whether a connection has failed and initiate a path switch:

$$(\text{Idle timeout}) + (\text{Poll response timeout} * 10 \text{ retries}) + (\text{Test retry interval} * \text{Test retry limit}) + 1$$

Using the defaults for the parameters, the time required to determine a connection failure and initiate a path switch can be up to 151 seconds.

$$(30 \text{ seconds}) + (8 \text{ seconds} * 10 \text{ retries}) + (8 \text{ seconds} * 5 \text{ attempts}) + 1 = 151 \text{ seconds}$$

If you are experiencing long path switch times (or failures) with HPR when a connection fails, you can reduce the time by decreasing the values for the parameters in the calculation.

Chapter 7. Planning for SNA Gateway

The SNA gateway supports:

- Various SNA protocols
- Various downstream workstations
- Various link types between the gateway and the host
- Pooling of LUs (sessions) at the gateway
- Definition of implicit workstations
- Dynamically changing or adding definitions for nonactive workstations
- Multiple hosts simultaneously
- Various link types (which can be combined) between the workstations and the gateway

Refer to *Quick Beginnings* for instructions on how to configure the SNA gateway using **SNA Node Configuration**.

Supported Connections to Workstations

The SNA gateway supports workstations connected by:

- LAN connectivities
 - Token ring
 - Ethernet
 - ATM (LAN Emulation)
- Frame Relay
- SDLC
 - Asynchronous
 - Synchronous
 - Hayes Autosync
 - Multipoint Primary
- X.25 network
- AnyNet TCP/IP (for Microsoft TCP/IP supported connectivities)
- ISDN connections

The SNA gateway can support as many as 254 LUs for each defined PU. You can define a different PU for each of your host connections. The LUs are used by the downstream workstations to connect to the host. The number of downstream workstations supported is dependent on several factors, including the type of connectivity and the number of adapters on the gateway. For instance, if you have one LAN adapter on the gateway, one host link with 254 LUs can be used by 253 downstream workstations concurrently. With two adapters, you can double the number of downstream workstations.

SDLC Workstation Connections

You can choose to connect using SDLC in point-to-point and Multipoint configurations. After configuration has been completed, no special consideration is necessary to run SDLC between the workstation and the gateway.

Planning for SNA Gateway

When setting up multidrop secondary workstations, consider the various factors that control polling turnaround time. There are physical limitations that affect how quickly the primary can poll the secondary that is farthest away. Defining all of the secondary workstations using the same parameters will simplify the following calculation.

To calculate the minimum inactivity time in seconds, use this formula:

$$NS * (SW + RW) * (IS + 2) / (LS / 8)$$

where:

NS = Number of stations
SW = Send window size
RW = Receive window size
IS = I-field size
LS = Line speed (bps)

If these values are not the same for all secondary workstations, a separate computation must be made for each group or workstation and added together to find the correct value for the inactivity timer.

For example, for a multidrop link with 16 secondary workstations at 14.4 Kbps, with each workstation set for a send window size of 7, a receive window size of 7, and I-field size of 521, the calculation for the inactivity timer would be:

$$16 * (7 + 7) * (521 + 2) / (14400 / 8) = 65.1 \text{ seconds}$$

Note: Always add a few extra seconds for a buffer.

X.25 Workstation Connections

If you configure the SNA gateway to use only X.25 permanent virtual circuit (PVC) connections between the gateway and the workstations, it is recommended that you configure all the workstation and gateway PVCs with negotiable link station roles. If you do not configure the gateway and workstation PVCs with negotiable link station roles, the workstations might not be able to establish the X.25 PVC links to the gateway.

See Chapter 13, "Planning for X.25" for more information about PVC connections and the configuration.

Supported Connections to Hosts

The SNA gateway supports the following connections to a host:

- LAN connectivities
 - Token ring
 - Ethernet
 - ATM (LAN Emulation)
- Frame Relay
- SDLC
 - Synchronous
 - Hayes Autosync
- X.25 network
- Twinaxial (to iSeries)
- AnyNet TCP/IP (for Microsoft TCP/IP supported connectivities)
- Enterprise Extender (for Microsoft IP supported connectivities)
- Channel (CDLC)

- Channel (MPC) with a DLUR connection
- ISDN connections

If there are connections to multiple hosts from the gateway being used by dependent LU traffic, only the host link with the control point PU name can do the following:

- Act as a host focal point
- Have CP-CP sessions with the host
- Be used for APPN functions

The links that are defined with a PU name that is not the same as the control point name can be used only for additional dependent LU connections to a host.

If a gateway host link is defined as a limited resource link, it will send a request disconnect to the host after the last LU-LU session is unbound. The gateway then passes the DACTLUs to the workstation and a DACTPU to workstations that have sessions only with that host link. When the workstation connected to the gateway is defined as a limited resource, the workstation link is disconnected if there are no other sessions. In Communications Server you can designate a workstation as a limited resource when defining the explicit client connection.

However, if the workstation has an application (or emulator) that automatically reactivates its host link, the gateway will reactivate the gateway link to its host when the workstation activates its link to the gateway. This means that incompatible workstation definitions would effectively inhibit the limited resource function at the gateway.

Host Considerations

Consider the following things when planning for your host connection:

- If your environment uses a network control point (NCP), the 37xx NCPGEN, **LINE** macro, **MAXLU** value must be at least as large as the number of LUs on the Communications Server.
- VTAM users can control the time interval during which VTAM would buffer data before sending it out to the controller (coattailing). When the VTAM coattailing parameter **DELAY=0** is specified on any PU definition, this might provide faster host send response time at the expense of more work for the host processor.

When a nonzero value for DELAY is specified, this saves the host processor instructions at the expense of increased receive response time on the Communications Server workstation.

For coattailing, begin with default value 0.2. A 0.2 second delay has a moderate effect on response time, but if the rate is approximately one transaction per second or greater, coattailing occurs.

- Although each downstream SNA workstation on the LAN implements a PU with its own set of LUs, the SNA gateway is implemented so that the host sees only the gateway PU. Therefore, only one PU definition is coded at the host. This PU definition contains the LUs for *all* the stations on the LAN that will connect to the host through this gateway.

If your host VTAM supports self-defining dependent LUs (SDDL), you might want to take advantage of the function. When configuring host LUs, specify an

Planning for SNA Gateway

LU model type or provide an LU model name that matches the LUSEED operand value defined in the VTAM switched major node and used by the VTAM SDDL U exit routine.

LAN Destination Address Configuration

You need to provide a destination address if you use one of the supported LAN connections between the gateway and the host or between the gateway and an explicitly defined workstation. When determining the correct destination address to enter in each profile, remember that the proper perspective for both addresses is to view the destination from the SNA gateway. Figure 28 shows this view.

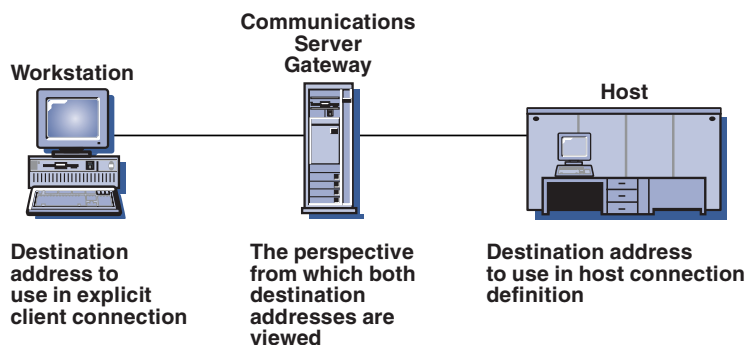


Figure 28. Perspective to Use for Destination Address Entries

Planning Gateway LU Definitions

The SNA gateway supports both *pooled* host LUs and *dedicated* host LUs. When LUs are configured for each host connection at the SNA gateway, they can be grouped into pools. Creating pools is often beneficial for the following reasons:

- If all sessions on all workstations do not need to communicate with the host at the same time, pooling can conserve system resources. The workstation sessions on the LAN can share the same LUs, but not at the same time. This function allows the number of LUs configured at the SNA gateway for workstation sessions to be greater than the number of LUs configured between the host and the gateway, reducing host resource requirements, reducing host workstation definition, and speeding restarts.
- Pools are easier to administer and configure. You can adjust the size of the pool by adding more LUs to it or deleting LUs from it without affecting other definitions.

Pooled LUs are not dedicated to any particular workstation, nor do they have to be dedicated for use by only downstream gateway users. A single pool can be shared by downstream TN3270 users and SNA gateway workstations, as well as SNA API clients and local emulator sessions. If you want to configure one pool (for example, PUBLIC) for all dependent LU uses, you do not need to know how the users will be distributed across those types.

- Pools can provide better availability for the downstream workstations if you define LUs from multiple connections in the same pool. If one of the connections is active but the other is not, users can access the LUs on the active link.

Consider the following scenarios:

- The Communications Server has two host connections with 100 LUs defined on each connection, all belonging to pool A. The first link is configured as *activate at*

startup and is operational. All workstations are using that link with its LUs. The second link is configured as *auto-activate*. This combination of pools with link configuration can provide better availability in two ways. If the first link fails, the recovering dependent sessions automatically activate the second link. If the first link depletes its defined LUs, the backup link will be automatically activated when the next LU (the 101st) is required.

- You can also use the combination of pools and link configuration to provide load balancing for downstream workstations across available host connections. An example of this is for the Communications Server to have two host connections with 100 LUs defined on each connection, all belonging to pool A. Both links are configured as *activate at startup* and are operational. The workstations that are mapped to the pool of LUs will be distributed across both connections to share the load across them.

The SNA gateway assigns pooled LUs to workstation sessions when the downstream workstation connects to the gateway. The workstation sessions can be defined to use LUs on different hosts with either dedicated or pooled LUs.

Dedicated LUs do not belong to a pool. A dedicated LU can be configured for use by an explicitly defined client.

Figure 29 shows a simple configuration with workstations using dedicated LUs, pooled LUs, or both. (The connecting lines represent the LUs.)




Figure 29. A Simple Scenario Using Pooled and Dedicated LUs

Using Explicit and Implicit Connections

There are two types of gateway-supported downstream workstations: explicit and implicit. Explicit workstations are workstations that have defined destination addresses over a particular DLC type (for example, token-ring network and SDLC). To configure explicit workstations, you must know the destination address or fully qualified adjacent control point name or adjacent node ID of each workstation and also define a logical link to the gateway for each workstation. LUs that are defined for explicit workstations can be pooled or dedicated.

Implicit workstations are easier to configure, but they can use only pooled LUs. Instead of defining a link to each workstation using the gateway, define a host LU pool (or pools) and configure the devices (DLCs) used by the workstation

Planning for SNA Gateway

connections. Configure an implicit client template to be used as a model for LU definition for each workstation that connects to the gateway and does not match an explicit definition. For example, if each workstation in an Ethernet LAN has two 3270 sessions configured with NAU addresses 2 and 3, you would configure a client template with two LUs (one for address 2 and one for address 3). If both addresses are used for sessions to a single host, map both addresses to the same host pool. If however, address 2 is used to go to HOST_A and address 3 is used to go to HOST_B, map each to the appropriate host pool. In this example, each time a workstation connects to the gateway over Ethernet, which does not match an explicit definition, a link is dynamically created and the two LUs for NAU 2 and 3 are allocated from the host LU pool(s).

For implicit workstations, users connecting to the gateway need to know only the adapter address of the gateway DLC that is configured for the implicit workstations and what NAU values have been defined on the gateway. They must use these NAU values when defining their 3270 sessions and logical printers.

A downstream Personal Communications workstation connecting to Communications Server can use LAN discovery, searching for the group name IG02HOST, to find the adapter address.

SNA Gateway Performance

Performance through any gateway is dependent upon many factors, including:

- Configuration of the gateway
- Connection to the host or hosts
- Overall processing demand placed on each host
- Connections to the supported workstations
- Configurations of the supported workstations
- Number of supported workstations
- Number of sessions with each workstation
- SNA transmission priority
- Flow control
- Pacing window for all sessions
- Traffic on non-gateway sessions

Using a workstation that implements a software gateway for other functions can also impair the performance of the gateway.

If none of the previously mentioned factors is causing a negative impact on performance, an individual workstation using supported SNA protocols should not experience any noticeable performance difference between a direct connection to the host and an indirect connection by way of a LAN through an SNA gateway connected to the host. In fact, if the host links are active at the gateway, the workstation activation might improve by removing host overhead delay. However, because of the many variables involved, you might want to conduct performance tests in your operational environment so that you can attain the desired balance between function and performance.

DLUR Considerations

By using DLUR for the connection to the host, you gain flexibility in where the SNA gateway can be placed. The connection to the host can traverse any APPN network, and is not restricted to being adjacent to a HOST/NCP. A DLUR to DLUS pipe is created to the DLUS VTAM that is used for dependent session control flows.

The recommended configuration for having a DLUR connection to the host is to define the gateway to be a network node and configure the DLUS information. During configuration of the explicit and implicit clients, map them to the DLUS. This configuration is easiest because it requires no knowledge of the downstream LUs at the gateway and it provides the highest visibility of the downstream devices to VTAM, since VTAM is aware of the PU.

If you do not want VTAM to be aware of the downstream PUs, you can configure an internal PU at the gateway (rather than a host connection), and map the downstream workstations to that internal PU.

If a downstream workstation, such as Personal Communications, is DLUR-capable, you might want that workstation to route through the Communications Server using the network node capability rather than the SNA gateway function.

Chapter 8. Planning for TN3270E Server

The TN3270E server function enables you to configure your network as shown in Figure 30 or Figure 31 on page 104.




Figure 30. Communications Server Configured as a TN3270E Server with Wide Area SNA Network




Figure 31. Communications Server Configured as a TN3270E Server with Wide Area TCP/IP Network

The TN3270E server function supports:

- Multiple ports
TN3270E server allows the configuration of multiple ports with specifications for pool name, security, client authentication, the level of security, and Certificate Revocation List (CRL) support for each port configured.
- LU emulation
TN3270E server supports LU 2, which enables users to run interactive 3270 display application programs.
- Host print
The Telnet 3270 standard extensions (TN3270E) enable users to print from host applications to printers attached to their workstation. These printers may be locally attached or network-attached. LU 1 and LU 3 sessions are supported.
TN3270E server implements the protocols outlined in RFC 1646 and RFC 1647, enabling the server to pass LU 1 and LU 3 session data to TN3270E-enabled clients, wait for client confirmation of the print request, and respond to the host.
- IP filtering
Communications Server supports the specification of TCP/IP client filters, using IP addresses, subnetworks, hostnames, or domain names. This function enables

central administration of the clients allowed to connect into the TN3270E server and specification of the LU names and pools clients can access.

- Response handling
TN3270E-enabled clients can send both positive and negative responses, which TN3270E server sends on to the host. TN3270E server generates request responses for standard TN3270 clients.
- ATTN and SYSREQ key handling
TN3270E server can convert and forward information to the host when the client sends an ATTN or SYSREQ key. Although clients that are not TN3270E-enabled have no explicit definition for ATTN and SYSREQ, the TN3270E server uses the following Telnet commands to implement these functions:

Table 6. Command Equivalents

| Telnet | TN3270E | Standard TN3270 |
|--------|---------|-----------------|
| IP | ATTN | SYSREQ |
| AO | SYSREQ | SYSREQ |
| BREAK | N/A | ATTN |

- LU classes
Communications Server categorizes user connection with LU classes. Classes consist of LUs configured with common characteristics; those that require a specific host connection, for example. This simplifies user access, groups users by application needs, and maximizes host resources.
- Secure Sockets Layer-based (SSL-based) security
You can specify that a TN3270E server port is secure, indicating that SSL Version 3 be used for connections on the port to provide data encryption and server authentication using digital certificates. Additional configuration parameters are supported to specify the level of security for connections on the port, whether client authentication should be processed, and whether a certificate revocation list (CRL) should be checked before accepting the client connection.

Refer to *Quick Beginnings* for instructions on how to configure a TN3270E server.

Supported Client Workstations under TN3270E Server

TN3270E server supports any TN3270E or TN3270 client that is fully compliant with RFC 1576, 1646, 1647 or 2355.

Highlights

This section provides more information about some of the features you can configure for TN3270E server.

Changing the Default Port Number

You can configure the ports your server uses for new connections from the TN3270E Ports configuration panel. The default port number is 23, but other Telnet applications, such as the TN5250 server support, may also use this port. If other Telnet applications are using port 23, you must use another port.

If you change the port number, avoid numbers that you know are used by other applications. If two applications use the same port number, one of the applications will fail.

Planning for TN3270E Server

If you change the port number, use a number greater than 1 024. Numbers less than 1 024 are reserved. For more information about reserved port numbers, see the following location on the Internet:

<http://www.ietf.org/>

At the time of publication of this document, the most recent RFC for assigned numbers is RFC 1700.

Notify TN3270E client users when you change the port number, because they will have to configure their emulator applications to match.

Managing System Traffic

There are three ways to control how often unused connections are disconnected: keepalive processing, automatic logoff, and LU takeover processing.

By default, TN3270E server does not use keepalive processing. If you use keepalive processing, you can choose either NOP or timing mark.

- NOP processing sends a Telnet NOP command after a specified keepalive frequency. This causes data to be transmitted on the connection, which causes TCP/IP to detect that the connection has broken. The server does not expect a response from the client. It can take an unpredictable amount of time for TCP/IP to detect connection outage.
- Timing mark processing sends a Telnet timing mark command to the client. If the client does not respond within the specified period, the connection is closed.

Timing mark processing causes more traffic on the system than NOP processing, but frees unused connections more quickly.

If you choose automatic logoff, the server disconnects any session that has no traffic for the specified period. Traffic from keepalive processing does not keep the connection open; data must be sent to or from the host. Printer sessions are not automatically logged off.

If your client emulators are configured to do keepalive processing, you might want to turn it off at the server, and if keepalive processing is done at the server, you might want to turn it off at the client to reduce network traffic.

If you choose LU takeover processing, the server will disconnect the session if certain conditions are satisfied. For example, when a new TN3270E session request specifies an LU that is in use at the server, the TN3270E server sends a timing mark to the client of that connection. If the client does not respond to the timing mark within the number of seconds that you specify, the server disconnects the session and assigns the LU to the new connection request. LU takeover processing enables you to disconnect unused sessions when they are needed, without the network traffic generated by keepalive detection. Keepalive detection enables you to free unused connections more frequently, but generates more network traffic.

IP Filtering

Communications Server enables you to configure TN3270E filters to specify which TCP/IP clients can connect into the server. Clients can be specified using individual IP addresses, IP subnetworks, or TCP/IP host names or domain names.

Secure Sockets Layer (SSL) Support

When you configure a new TN3270E server port, you can specify that the port be used for secure connections. More than one port can be specified as secure.

To enable security, Communications Server provides a **Key Management** utility to generate the certificate and keys required by SSL. Refer to Chapter 10, “Planning for Secure Sockets Layer-based Security”, on page 115 for more information.

When you have configured and enabled security, TN3270E clients supporting SSL Version 3 can connect into a secure port number and establish secure connections.

Configuring SNA Connections

You need to configure your host connection and host LUs before you use the TN3270E server function. For a table that describes and explains host parameters, refer to *Quick Beginnings*.

Pooling

There are four classes of LU definitions specific to TN3270E server: implicit workstation, explicit workstation, implicit printer, and explicit printer. These classes correspond to the terminal-generic, terminal-specific, printer-generic, and printer-specific classes specified in RFC 1647.

Implicit workstation definitions are defined in a pool that the TN3270E server uses to satisfy requests for connections. The default workstation pool can be configured to satisfy requests where the client does not specify a workstation LU or workstation pool name. All other workstation pools must be requested by name.

You can also define a set of LU definitions used to satisfy requests for a specific LU name. These *explicit workstation* definitions ensure that a terminal device needed by a host application is not assigned to a client that does not specifically request it.

Similarly, you can define a pool of printer definitions that will be used to satisfy request for connections that do not require specific LU names (*implicit printer*) and a set used to satisfy requests for a specific LU name (*explicit printer*). The default printer pool can be configured to satisfy requests where the client does not specify a printer LU or printer pool name. All other printer pools must be requested by name.

Implicit and explicit workstation definitions can have printers that are associated with them. Each terminal definition can have a printer that is assigned to it and each printer can have an associated terminal definition. These printers are not included in the explicit or implicit printer definitions.

Associated printer definitions can only be accessed by referencing the terminal LU name. They reduce the amount of information the client user needs, because he or she only needs to know the LU name of the terminal to connect to both the terminal and the printer sessions.

Load Balancing of TN3270E Servers

Communications Server also supports load balancing of TN3270 clients across multiple TN3270E servers. Refer to Chapter 11, “Planning for Load Balancing”, on page 123 for more information.

Chapter 9. Planning for TN5250 Server

The TN5250 server function enables you to configure your network as shown in Figure 32 or Figure 33 on page 110.




Figure 32. Communications Server Configured as a TN5250 Server with Wide Area SNA Network

Planning for TN5250 Server




Figure 33. Communications Server Configured as a TN5250 Server with Wide Area TCP/IP Network

The TN5250 server function supports:

- Multiple ports
TN5250 server allows the configuration of multiple ports with specifications for server name, security, client authentication, the level of security, and Certificate Revocation List (CRL) support for each port configured.
- Connecting to TN5250 clients with 5250 workstation data stream
In implementing the protocols outlined in RFC 1205, the server passes 5250 workstation data to and from a TCP/IP client emulating an IBM 5250 workstation. This client is commonly known as a TN5250 client.
- Access to one or more iSeries hosts on a SNA network
The server connects to iSeries hosts using SNA LU 6.2 protocol. Support is provided to access one or more iSeries hosts from the same or different client workstations.
- IP filtering
Communications Server enables the specification of TCP/IP client filters, using IP addresses, subnetworks, hostnames, or domain names. This function enables central administration of the clients allowed to connect into the server, as well as support for directing clients to specific iSeries hosts.
- Secure Sockets Layer-based (SSL-based) security

You can specify that a TN5250 server port is secure, indicating that SSL Version 3 be used for connections on the port to provide data encryption and server authentication using digital certificates. Additional configuration parameters are supported to specify the level of security for connections on the port, whether client authentication should be processed, and whether a certificate revocation list (CRL) should be checked before accepting the client connection.

Refer to *Quick Beginnings* for instructions on how to configure a TN5250 server.

Supported Client Workstations under TN5250 Server

TN5250 server supports any TN5250 client that is fully compliant with RFC 1205.

Highlights

This section provides more information about some of the features you can configure for TN5250 server.

Changing the Default Port Number

When you configure the TN5250 server support, a default port is defined with port number 23, but other Telnet applications, such as the TN3270E server support, may also use this port. If other Telnet applications are running using port 23, you must use another port.

If you change the port number, avoid numbers that you know are used by other applications. If two applications use the same port number, one of the applications will fail.

If you change the port number, use a number greater than 1 024. Numbers less than 1 024 are reserved. For more information about reserved port numbers, see the following location on the Internet:

<http://www.ietf.org/>

At the time of publication of this document, the most recent RFC for assigned numbers is RFC 1700.

Notify TN5250 client users when you change the port number, because they will have to configure their emulator applications to match.

Managing System Traffic

There are two ways to control how often unused connections are disconnected: keepalive processing and automatic logoff.

By default, TN5250 server does not use keepalive processing. If you use keepalive processing, you can choose either NOP or timing mark.

- NOP processing sends a Telnet NOP command after a specified keepalive frequency. This causes data to be transmitted on the connection, which causes TCP/IP to detect that the connection has broken. The server does not expect a response from the client. It can take an unpredictable amount of time for TCP/IP to detect connection outage.
- Timing mark processing sends a Telnet timing mark command to the client. If the client does not respond within the specified period, the connection is closed.

Planning for TN5250 Server

Timing mark processing causes more traffic on the system than NOP processing, but frees unused connections more quickly.

If you choose automatic logoff, the server disconnects any session that has no traffic for the specified period. Traffic from keepalive processing does not keep the connection open; data must be sent to or from the host.

If your client emulators are configured to do keepalive processing, you might want to turn it off at the server, and if keepalive processing is done at the server, you might want to turn it off at the client to reduce network traffic.

Specifying Access to iSeries Hosts with Multiple Ports

Communications Server enables you to configure more than one iSeries, and enables you to configure more than one port for the TN5250 server to listen for incoming connections. When you configure a new TN5250 server port, you can specify the iSeries to be associated with the port. Incoming TCP/IP connections received on that port by the TN5250 server will access the iSeries associated with that port.

Notify TN5250 client users which port numbers to configure for their emulator applications to connect to specific iSeries hosts.

IP Filtering

Communications Server enables you to configure TN5250 filters to specify which TCP/IP clients can connect into the server. Clients can be specified using individual IP addresses, IP subnetworks, or TCP/IP host names or domain names.

You can also configure an iSeries with a filter to specify a different iSeries than the one associated with the port. This can be used to direct the clients associated with a filter to a specific iSeries.

More than one iSeries can be specified in a filter. If a client connection attempt to the first iSeries specified is unsuccessful, connection to other iSeries hosts specified in the filter are attempted.

Secure Sockets Layer (SSL) Support

When you configure a new TN5250 server port, you can specify that the port be used for secure connections. More than one port can be specified as secure.

To enable security, Communications Server provides a **Key Management** utility to generate the certificate and keys required by SSL. Refer to Chapter 10, "Planning for Secure Sockets Layer-based Security", on page 115 for more information.

When you have configured and enabled security, TN5250 clients supporting SSL Version 3 can connect into a secure port number and establish secure connections.

Configuring the SNA Network

You need to configure your SNA network to allow the TN5250 server to connect to the iSeries hosts. Refer to *Quick Beginnings* for information on configuring your SNA network for the TN5250 server.

Load Balancing of TN5250 Servers

Communications Server also supports load balancing of TN5250 clients across multiple TN5250 servers. Refer to Chapter 11, “Planning for Load Balancing”, on page 123 for more information.

Chapter 10. Planning for Secure Sockets Layer-based Security

You can use Secure Sockets Layer-based (SSL-based) security on connections between TN clients and the TN3270E server or TN5250 server. This security uses SSL Version 3 to provide data encryption and authentication using signed certificates.

You can configure specific ports to support SSL-based security. If you specify security, the server must have an authenticated certificate provided by a certificate authority (CA). Communications Server provides a utility that generates and manages keys and certificates used by SSL Version 3.

You can also configure ports to request client authentication to verify that the client is authorized to establish a secure connection to the server. In addition, if client authentication is configured, you can request certificate revocation list (CRL) processing of client certificates. You specify the location of the LDAP server that maintains the CRL to determine if the client certificate has been revoked.

Note: Support for CRL processing of client certificates requires the use of the IBM Vault Registry product, which publishes Certificate Revocation Lists to an LDAP server.

You can also configure the level of security required by the port. Options include HIGHEST, MEDIUM, and AUTHENTICATION ONLY. The definition of HIGHEST is dependent on the version of Communications Server that is installed. For Communications Server supporting strong encryption, HIGHEST specifies that the port will only accept connections from clients supporting strong encryption. For Communications Server supporting export encryption only, HIGHEST is the same as MEDIUM. MEDIUM indicates that the port can establish connections with any supported level of encryption. AUTHENTICATION ONLY indicates that certificates to authenticate one or both ends of the connection are exchanged, but data is not encrypted.

This chapter gives an overview of planning for Secure Sockets Layer-based (SSL-based) security. SSL-based security is provided as an option when configuring TN3270E server and TN5250 server.

How SSL Security Works

Secure sockets layer (SSL) is an industry-standard protocol that uses symmetric-key and public-key cryptographic technology. Symmetric-key cryptography uses the same key to encrypt and decrypt messages. Public-key cryptography uses a pair of keys, a public key and a private key. Each server's public key is published, and the private key is kept secret. To send a secure message to the server, the client encrypts the message using the server's public key. When the server receives the message, it decrypts the message using its private key.

SSL provides three basic security services:

- Message privacy

SSL-based Security

Message privacy is achieved through a combination of public-key and symmetric key encryption. All traffic between an SSL client and an SSL server is encrypted using a key and an encryption algorithm negotiated during session setup.

- Message integrity

The message integrity service ensures that SSL session traffic does not change in route to its final destination. SSL uses a combination of public/private keys and hash functions to ensure message integrity.

- Authentication

Authentication is the process whereby the client and the server convince each other of their identities. The client and server identities are encoded in public-key certificates. A public-key certificate contains the following components:

- Subject's distinguished name
- Issuer's distinguished name
- Subject's public key
- Issuer's signature
- Validity period
- Serial number

Communications Server SSL Support

Communications Server supports the following components of SSL security:

- Secure SSL sessions between TN3270E server and any TN3270 and TN3270E clients supporting SSL Version 3
- Secure SSL sessions between TN5250 server and any TN5250 clients supporting SSL Version 3
- SSL Version 3
- Message privacy and integrity
- Server-side authentication
- Optional client-side authentication
- Optional certificate revocation list (CRL) processing, when client authentication is requested.

Configuring SSL Security

During TN3270E server configuration and TN5250 server configuration, you specify which ports will use SSL security for their connections. See Chapter 8, "Planning for TN3270E Server", on page 103 and Chapter 9, "Planning for TN5250 Server", on page 109 for more information on configuring secure ports.

Server Authentication

To enable SSL security with server authentication, a set of public/private keys and an associated authenticated public-key certificate is required. Communications Server provides a **Key Management** utility to create, manage, and store public/private keys and certificates required by SSL.

Using the **Key Management** utility, you must do one of the following:

- Create a public/private key pair and obtain and store a certificate from one of the predefined (well-known) certificate authorities (CAs). See "Configuring SSL Using a Well-Known Trusted CA" on page 118.

This procedure requires less setup because the key database file is pre-configured with the well-known CA root certificates required to identify the CAs from whom the certificate is issued.

- Create a public/private key pair, define an unknown CA by obtaining and storing the CA's root certificate in the key database file, and obtain and store a certificate from the unknown CA. See "Configuring SSL Using an Unknown CA" on page 120.
- Create a self-signed certificate and store the certificate in the key database file. See "Creating a Self-Signed Certificate" on page 121.

Note: To ensure adequate security for your site, a self-signed certificate should be used for controlled testing purposes only.

Client Authentication

If a secure port is configured for client authentication, SSL requests the client certificate. The **Key Database File** on the server is checked to verify the signer of the client certificate.

If the client certificate is :

- Issued by a well-known CA, no additional configuration is needed
- Issued by an unknown CA, the root certificate of the unknown CA must be obtained and stored in the key database as a signer certificate. See "Storing Certificates for Clients" on page 122.
- A self-signed client certificate, the self-signed certificate must be obtained and stored in the key database as a signer certificate. See "Storing Certificates for Clients" on page 122.

If the certificate revocation list (CRL) support is configured, the LDAP server that maintains the CRL is contacted to determine whether the client certificate has been revoked.

Note: Support for CRL processing of client certificates requires the use of the IBM Vault Registry product, which publishes Certificate Revocation Lists to an LDAP server.

Key Management Utility

The Communications Server **Key Management** utility allows you to create, manage and store public/private keys and certificates required for SSL communications between Communications Server and the client. Before you begin configuring SSL communications, you must open the Communications Server key database file. Use the following steps to open the file:

1. Select **Key Database File** from the main menu, and select **Open** to open an existing key database.
2. From the **Open** window, choose the **private** subdirectory of the directory where you installed the product (for example, C:\IBMCS\PRIVATE).
3. Choose **ibmcs.kdb** as the key database.
4. At the password prompt, enter your password.

Note: The password is initially set to **ibmcs**. For adequate security, you *must* change the password before performing any of the **Key Management** tasks.

Changing the Password

Use the following steps to change the password of the current database:

1. Select **Key Database File** from the menu, and select **Change Password** to change the key database password for the current database. The **Change Password** file dialog appears.
2. Enter the new password you want to use.
3. Enter the same password again to verify.
4. If you want the password to have an expiration date, click **Set expire time**. To ensure adequate security, the password should have an expiration date.
5. If you choose to allow the password to expire, specify the number of days until the password expires.
6. Click **Stash the password** to stash the password of the current database into the **ibmcs.sth** file. Communications Server requires that the password be stashed to access the key database file.

Configuring SSL Using a Well-Known Trusted CA

Follow the procedures in this section to set up SSL security using a certificate issued by a well-known CA. The following CA signer certificates are already stored in the key database and marked as trusted certificates:

- Thawte Personal Premium CA
- Thawte Personal Freemail CA
- Thawte Personal Basic CA
- Thawte Premium Server CA
- Thawte Server CA
- RSA secure server CA (also obtained from VeriSign)
- VeriSign class 4 public primary CA
- VeriSign class 3 public primary CA
- VeriSign class 2 public primary CA
- VeriSign class 1 public primary CA

To set up SSL security using a well-known CA, the following procedures are necessary:

- Create a key and certificate request
- Submit a certificate request to the CA
- Obtain and store a certificate in the server key database

Creating a Key and Certificate Request

Use the following steps in the **Key Management** utility to create the public/private keys and certificate request:

1. Select **Personal Certificate Requests** from the drop-down list of the main menu and click **New** to create a new key pair and certificate request. The **Create New Key and Certificate Request** dialog appears.
2. Enter the name (label) that is used to identify the key and certificate within the database.
3. Enter the number of the key size you want to use. Choosing a larger key size results in stronger security, but requires more processing on the client and the server to establish a connection.

4. Enter the TCP/IP host name of the communications server as the common name (for example, wtr05306.raleigh.ibm.com).
5. Enter an organization name.
6. Enter an organization unit (optional).
7. Enter a city or locality (optional).
8. Enter a state or province (optional).
9. Enter a ZIP code (optional).
10. Enter a country code. You must specify at least 2 characters (for example, US).
11. Enter a certificate request file name, or use the default file name.

When you click **OK**, the information you supplied is processed. Two files are produced:

| | |
|--------------------|--|
| ibmcs.rdb | Private key file |
| certreq.arm | This is the default name of the certificate request file. If you assigned a file name, a file with that name is created. The certificate request file is a PKCS 10-type file in armored 64 format. |

Do not attempt to edit or move these files. If the **ibmcs.rdb** file can not be found or has been corrupted when you attempt to enter the certificate into the key database, you will have to resubmit your certificate request to the CA.

Submitting a Certificate Request

Start a Web browser and access the Web page for the CA. Follow the instructions provided to submit the certificate request. The following are the URLs of some well-known CAs:

- VeriSign: <http://www.verisign.com/>
- Thawte: <http://www.thawte.com/>

Depending on the CA you choose, either e-mail the certificate request generated by the **Key Management** utility or incorporate the certificate request into the form or file provided by the CA.

After submitting a certificate request to a CA, you can enable SSL security by creating and storing a self-signed certificate. A self-signed certificate should be used for controlled testing purposes only. Refer to “Creating a Self-Signed Certificate” on page 121 for more information.

Storing a Certificate in the Key Database

When you receive a certificate from a CA, use the **Key Management** utility to put the certificate into the key database file, **ibmcs.kdb**, located on the server.

1. Select **Personal Certificates** from the drop-down list of the **Key Management** utility and click **Receive** to receive the key pair and certificate request. The **Receive Certificate from a File** dialog appears.
2. Ensure that the data type is **BASE64-encoded ASCII data** (encoded 64 format).
3. Enter the certificate file name.
4. Enter the location (path name) of the certificate. Click **OK**. The stored certificate displays as the first item.
5. Highlight the stored certificate and click **View/Edit**. The **Key information** dialog appears.
6. Click **Set the certificate as the default**. The selected key becomes the default.

Configuring SSL Using an Unknown CA

To set up SSL security using an unknown CA (not already defined in the database), the following procedures are necessary:

- Create a key and certificate request
- Submit a certificate request to the CA
- Obtain the CA's root certificate and your certificate and store them in the server key database

Creating a Key and Certificate Request

Use the following steps in the **Key Management** utility to create the public/private keys and certificate request:

1. Select **Personal Certificate Requests** from the drop-down list of the main menu and click **New** to create a new key pair and certificate request. The **Create New Key and Certificate Request** dialog appears.
2. Enter the name (label) that is used to identify the key and certificate within the database.
3. Enter the number of the key size you want to use. Choosing a larger key size results in stronger security, but requires more processing on the client and the server to establish a connection.
4. Enter the TCP/IP host name of the communications server as the common name (for example, wtr05306.raleigh.ibm.com).
5. Enter an organization name.
6. Enter an organization unit (optional).
7. Enter a city or locality (optional).
8. Enter a state or province (optional).
9. Enter a ZIP code (optional).
10. Enter a country code. You must specify at least 2 characters (for example, US).
11. Enter a certificate request file name, or use the default file name.

When you click **OK**, the information you supplied is processed. Two files are produced:

| | |
|--------------------|--|
| ibmcs.rdb | Private key file |
| certreq.arm | This is the default name of the certificate request file. If you assigned a file name, a file with that name is created. The certificate request file is a PKCS 10-type file in armored 64 format. |

Do not attempt to edit or move these files. If the **ibmcs.rdb** file can not be found or has been corrupted when you attempt to enter the certificate into the key database, you will have to resubmit your certificate request to the CA.

Submitting a Certificate Request

Follow the procedures of the unknown CA to submit the certificate request.

Depending on the CA you choose, either e-mail the certificate request generated by the **Key Management** utility or incorporate the certificate request into the form or file provided by the CA.

After submitting a certificate request to a CA, you can enable SSL security by creating and storing a self-signed certificate. A self-signed certificate should be

used for controlled testing purposes only. Refer to “Creating a Self-Signed Certificate” for more information.

Storing a Certificate in the Key Database

When you receive a certificate from a CA, contact the CA to obtain the root certificate of the CA. You must store the CA root certificate in the key database before you store the certificate that you applied for. The CA root certificate validates the certificate you applied for. Use the **Key Management** utility to store the CA root certificate.

1. Select **Signer Certificates** from the drop-down list and click **Add** to receive the the CA root certificate. The **Add CA Certificate from a File** dialog appears.
2. Ensure that the data type is **BASE64-encoded ASCII data** (encoded 64 format).
3. Enter the certificate file name.
4. Enter the location (path name) of the certificate. Click **OK**. The file is marked as *trusted* and is stored.

To store the certificate you applied for, use the **Key Management** utility to put the certificate into the key database file, **ibmcs.kdb**, located on the server.

1. Select **Personal Certificates** from the drop-down list of the main menu and click **Receive** to receive the key pair and certificate request. The **Receive Certificate from a File** dialog appears.
2. Ensure that the data type is **BASE64-encoded ASCII data** (encoded 64 format).
3. Enter the certificate file name.
4. Enter the location (path name) of the certificate. Click **OK**. The stored certificate displays as the first item.
5. Highlight the stored certificate and click **View/Edit**. The **Key information** dialog appears.
6. Click **Set the certificate as the default**. The selected key becomes the default.

Creating a Self-Signed Certificate

Receipt of a certificate from a well-known trusted CA can take up to three weeks. Until you receive the public server certificate, you can create a self-signed certificate to enable SSL sessions between clients and the server. A self-signed certificate should be used for controlled testing purposes only. To ensure adequate security for your site, a self-signed certificate should not be used in a production environment. To set up your site to use a self-signed certificate using the **Key Management** utility, follow these steps:

1. Select **New Self-signed Certificate** from the **Create** menu option to create a new self-signed certificate. The **Create New Self-Signed Certificate** dialog appears.
2. Enter the name (label) that is used to identify the key and certificate within the database. Select **X509 V3** as the certificate version.
3. Enter the number of the key size you want to use. Choosing a larger key size results in stronger security, but requires more processing on the client and the server to establish a connection.
4. Enter the TCP/IP host name of the communications server as the common name (for example, wtr05306.raleigh.ibm.com).
5. Enter an organization name.
6. Enter an organization unit (optional).
7. Enter a city or locality (optional).

SSL-based Security

8. Enter a state or province (optional).
9. Enter a ZIP code (optional).
10. Enter a country code. You must specify at least 2 characters (for example, US).
11. Enter the number of days the self-signed certificate is to be valid.
12. Click **OK**.
13. Click **Yes** to set the key as the default key in the key database.

If you are using IBM SecureWay Host On-Demand Version 4 or some other SSL client that requires the self-signed certificate to be designated, using the **Key Management** utility, follow these steps:

1. Select **Personal Certificates** from the drop-down list, highlight your self-signed certificate and click **Extract Certificate**. The **Extract Certificate to a File** dialog appears.
2. Select **Binary DER data** as the data type.
3. Enter the certificate file name. The file should have a file type of **der** (for example, ibmcs.der).
4. Enter the location (path name) of the certificate.
5. Click **OK**.

The certificate file must be supplied to each client authorized to connect to the server. Consult the client documentation for instructions on storing the self-signed certificate as a CA root certificate in the client's class file or key database.

Storing Certificates for Clients

If a client certificate is issued by an unknown CA or is a self-signed certificate, the root certificate of the unknown CA or the self-signed certificate must be stored in the key database as a signer certificate. Use the **Key Management** utility to store the certificate.

1. Select **Signer Certificates** from the drop-down list and click **Add** to receive the the CA root certificate. The **Add CA Certificate from a File** dialog appears.
2. Ensure that the data type is **BASE64-encoded ASCII data** (encoded 64 format).
3. Enter the certificate file name.
4. Enter the location (path name) of the certificate. Click **OK**. The file is marked as *trusted* and is stored.

Chapter 11. Planning for Load Balancing

This chapter gives an overview of planning for load balancing. It describes planning for dependent LU and LU 6.2 load balancing.

Dependent LU Load Balancing

Load balancing allows you to balance dependent host sessions across servers. Load balancing is supported for LUx API applications through SNA API clients, 3270 emulators, and TN3270 emulators.

When Communications Server responds to client requests, it sorts servers supporting the requested pool by load.

The load for dependent LUs represents a percentage of available resources from a particular server. The load percentage is calculated by dividing the number of active application connections by the total number of LUs (0 to 3) available.

You can influence the calculated load by specifying an LU0- to-3 load factor (host session load factor) to compensate for differences between two servers, such as available memory, processor speed, and CPU utilization. You can also use the LU0-to-3 load factor to balance the loads between Communications Servers and Novell NetWare for SAA servers.

Notes:

1. The servers participating in load balancing are determined by LU pool membership and by scope membership. All Communications Servers with LUs in pools of the same name are subject to load balancing.
2. Clients using TCP/IP protocol can participate in load balancing. However, TCP/IP clients can load balance only between IBM Communications Servers. Load balancing is determined by configured scopes.

In Communications Server, if LU pools on different servers have the same name, they function as one pool for load balancing. When assigning LUs to an LU pool, select from the host links on a server, and assign LUs from the selected host link to the LU pool. These LUs do not need to be of the same LU type. Also, the resulting pool can contain LUs from multiple links.

LU 6.2 Load Balancing

Load balancing allows you to balance independent LU 6.2 sessions by distributing them across multiple servers. Load balancing is supported for APPC API applications through SNA API clients and TN5250 emulators.

The load for LU6.2 represents a percentage of available resources from a particular server. The load percentage is calculated by dividing the total number of conversations over all local LUs on a particular server by the cumulative maximum session limit for all local LUs. The maximum session limit is the LU 6.2 session limit specified during configuration. If the maximum session limit is specified as zero (0), indicating there is no session limit, the default maximum local LU session limit of 512 per local LU is used when the load is calculated. The default maximum local LU session limit can also be specified during configuration.

Planning for Load Balancing

You can influence the calculated load by specifying an LU6.2 load factor (APPC session load factor) to compensate for differences between two servers, such as available memory, processor speed, and CPU utilization. You can also use the LU6.2 load factor to balance the loads between Communications Servers and Novell NetWare for SAA servers.

Note: Clients using TCP/IP protocol can participate in LU 6.2 load balancing.

Planning for TCP/IP Scopes

Scope is a parameter used to control and manage access by TCP/IP clients to servers in a network. It is the same as the service location protocol (SLP) scope as referenced in RFC 2165. The SLP is used by Communications Server to implement TCP/IP-based load balancing.

The control scope provided is necessary for two reasons:

- As your network, the number of clients, and the number of servers grow, it becomes necessary to partition access to those servers by the growing number of clients in order to reduce overall traffic on the network.
- The control scope allows administrators to organize users and servers in to administrative groups.

The meaning of the scope values are defined by the administrator of the network. These values can represent any entity. Commonly, they fall along either departmental, geographical, or organizational lines.

Once configured, clients are able to reach the SNA network through servers that are configured with the same scope or that have no associated scope (*unscoped* services or servers).

Note: If a server is configured as unscoped, it replies to SLP scoped **and** unscoped requests. If the SNA API client is configured to connect to unscoped servers, only unscoped servers will reply.

Scopes and Client Server Security

Communications Server enables Novell NetWare for SAA clients to use Windows NT domain security to authenticate the client connection to the server without reentering the userid and password. The client must be part of a Windows NT domain, either by participation in a Communications Server domain or logging in locally with a synchronized userid and password.

The authorized users for client server are maintained in the IBMCSAPI local group, which is located either directly on the Communications Server or on the domain controller where Communications Server participates. This user group is created during installation and can be administered using the Windows NT User Manager application.

Novell NetWare for SAA client users outside of the Windows NT domain must replicate the userid and password across all servers configured with the same scope.

Where Scope is Configured

Each server is assigned to a scope or scopes through **SNA Node Configuration**. Clients using these servers must be configured to connect to servers within a single

specific scope or to unscoped servers. For more information on configuring clients, refer to *Client/Server Communications Programming*.

How Scope Relates to SLP

Communications Server scope relates directly to service location protocol (SLP) scope. Therefore, SLP directory agents may reside in the network that support the Communications Server configured scopes. If you plan to allow clients to locate Communications Server services based on scopes, consider how scope relates to the network as a whole. If there are unscoped services in a network where scopes are also used, the unscoped services are eligible to satisfy any scoped requests, which can potentially put a burden on those service agents and directory agents that support the unscoped services.

Note: If the SNA API client is configured to connect to unscoped servers, only unscoped servers will reply.

If directory agents are to be used in the site network (for upward scaling), they should be configured to handle the same scopes as are configured for the Communications Server. In addition, if unscoped services are to be used in networks with directory agents, at least one unscoped directory agent should be configured.

Communications Server includes an SLP service agent that is installed on every server where Communications Server is installed. No additional SLP support needs to be installed in the network in order for TCP/IP load balancing and Communications Server discovery to work properly.

Tracking the Current Load

You can track the current load for both LU 6.2 and dependent LUs using the Windows NT Performance Monitor. You can start the performance monitor from the **Administrative Tools (Common)** selection of the **Programs** menu.

From the performance monitor screen, select **Add to Chart** from the **Edit** pulldown. In the **Object** field, select **IBM SNA Load Balancing**. The two counters listed are:

LU0/LU3 Load

Use this counter to track the load for dependent LUs.

LU6.2 Load

Use this counter to track the load for LU 6.2 LUs.

The results of the load tracking can be saved for future reference.

Chapter 12. Planning for Backup Host Connections

For environments where it may be important to provide backup host connections to the primary connections in use by workstation emulator sessions, Communications Server provides you with options that enable configuration of host connections that are activated in response to failure conditions. The options you can use for this depend on the level of backup you want to provide:

- To enable full system backup between server nodes supporting LAN based IP or IPX workstation emulator sessions, include hot standby in your environment.
- To enable full system backup between server nodes supporting other than LAN based IP or IPX workstation emulator sessions, do not include hot standby in your environment. Configure a duplicate Communications Server installation on the backup system.
- To enable backup within a server node, do not use hot standby. Use normal options on connection definitions in your backup environment.

These options are explained in the following sections.

Planning for Backup Connections between Servers Using Hot Standby

In Communications Server, you can configure host links to activate automatically if a critical server fails. Configured connections to a host can continue to function by activating alternate connections on a backup server. This function is known as hot standby.

For the user of a workstation emulator session with a host machine whose initial connection fails due to a Communications Server failure, hot standby enables a subsequent attempt to reestablish the connection with the host machine, automatically routing through an alternate Communications Server. The initial server providing the connection is known as the *critical server*. The alternate server providing a connection to the host machine after failure of the critical server is known as the *backup server*. The backup server provides connections that are automatically started in response to detection of a critical server failure and licensing charges related to critical server host connectivity are automatically managed on the backup server.

A server can be both a critical server and a backup server at the same time. Servers can back each other up.

Hot standby enables backup server support for LAN based IP or IPX emulator sessions. Hot standby does not provide backup server support for the SNA environment.

In Figure 34 on page 128, the backup server is configured with a complete backup connection for HOSTCON1, which provides the primary host connectivity on the critical server. The dotted lines between the backup server and the workstation emulator clients are not active while the critical server is in operation.

Planning for Backup Host Connections




Figure 34. A Critical Server with One Backup Server

In Figure 35, each of the backup servers are configured with a partial backup connection for HOSTCON1, which provides the primary host connectivity on the critical server. The combination of the two backup servers providing a complete backup for HOSTCON1. The dotted lines between the backup servers and the workstation emulator clients are not active while the critical server is in operation.




Figure 35. A Critical Server with Two Backup Servers

If the critical server fails, recovery occurs as follows:

1. The backup server, the host, and the workstation emulator software on the workstation detect the failure at approximately the same time.
2. The connections named in the backup server's critical server configuration are automatically activated and their connectivity to the host machine is established.

3. The workstation establishes a new session with the backup server that has just activated its critical server connections.
4. The workstation contacts the host machine and displays a logon prompt for the user.

Note: When the critical server is active again, the backup server connections activated for it are not automatically deactivated. You must deactivate the backup server connections manually.

Using a Production or Dedicated Backup Server for Hot Standby

You can not install the same Communications Server licenses on more than one server running on the same network. Install Communications Server on multiple servers by purchasing an additional Communications Server package for each additional server.

Before configuring your Communications Server for hot standby, decide whether to run Communications Server as both a production server and a backup server, or as a dedicated backup server.

Using a Production Server

A *production server* is a Communications Server that takes care of your day-to-day communication requirements. This environment configures Communications Server with host connections that are activated outside of a critical server failure scenario. Activation and use of such connections requires purchasing of adequate licenses for them. You can not install the same Communications Server licenses on more than one server. To install Communications Server for use as a production server on multiple systems, you must purchase an additional Communications Server package for each additional server installation.

When you use a Communications Server as both a production server and a backup server, there are no special license considerations. You install Communications Server and the licenses you purchased for use of Communications Server as a production server.

The role of Communications Server as a backup server requires no additional licensing beyond the base licensing purchased for the critical servers it is configured to monitor. You do not need any special license considerations on the backup server. When the backup server detects a failure of a critical server, the backup server automatically activates the connections configured for the critical server and adjusts the backup server's licenses based on those in effect for the critical server. Communications Server manages hot standby licensing without your intervention.

However, if the connections configured for the critical server have already been activated outside of a critical server failure scenario, the backup server's adjustment for the critical server's licenses is not performed and separate licensing on the backup server is required.

Using a Dedicated Server

A *dedicated* backup server does not function as a production server and does not activate connections outside of a critical server failure scenario. No connections can be activated on it at startup or activated manually. Its function is to monitor one or more critical servers.

Planning for Backup Host Connections

A dedicated backup server requires no additional licensing beyond the base licensing purchased for the critical servers it is configured to monitor. Communications Server manages hot standby licensing without your intervention.

Planning for the Complete Hot Standby Environment

Scenarios for using hot standby involve normal Communications Server functional scenarios for support of LAN based emulator clients over IP or IPX, with some specific hot standby operations.

To configure a hot standby environment, there are two basic dependencies:

- Two or more communications servers must be configured to provide compatible connections to a host machine. The server used as the backup must be configured with a critical server configuration for the server that supplies the primary connectivity to the host machine.
- The workstation emulator client software used to connect through the server to provide users with emulator sessions on the host machine must support some means of effecting alternate routing to the backup server in the event of a failure of the critical server.

Some emulator clients may not provide for alternate routing. Planning for a hot standby environment requires careful evaluation of emulator software and the type of connectivity used between the client and the server. The following are some choices for alternate routing:

- **Hot rollover** - This refers to support in an emulator client for configuration of multiple named servers to be contacted, in a specific order, to establish a connection through the server to the host machine for an emulator session. Hot rollover enables you to configure the backup server to be contacted when attempts to contact the critical server are unsuccessful. When the critical server fails and resources on the backup server are activated, the emulator client contacts the backup server and establishes an alternate route through it for connection to the host machine.
- **LU pools** - This refers to support in an emulator client to dynamically establish the connection to the host machine for an emulator session through the server currently responding to queries for available LUs from a named pool. LU pools support enables the backup server and critical servers to provide LUs to the same pool. The backup server is initially defined but not active. When the critical server fails and the resources on the backup server are activated, the backup server responds to requests from the emulator client for an LU from the pool and establish an alternate route through it for connection to the host machine. LU pools includes the following types of client support for LU pools:
 - Pre-SLP support: Before Communications Server support of SLP, client access of LU pools was random or in a specified server order.
 - SLP support: Communications Server SLP support provides client access of LU pools, balancing the session load across the servers contributing to the pool.
- **Multiple emulator sessions** - This refers to support in an emulator client for configuration of multiple emulator sessions, enabling the user to configure separate sessions using the backup server and critical servers to connect to the host machine. Using multiple emulator sessions is not as easy as hot rollover or LU pool support, because it requires manual intervention by the emulator user to assist with the alternate routing when a failure occurs. When the critical server connection fails, the user manually initiates the session configured to use the backup server to reestablish connection to the host machine.

IBM Personal Communications support for TN3270E, TN5250, API, or QEL/MU emulator clients provide options for alternate routing through one or more of the hot rollover, LU pools, or multiple emulator sessions options. Some vendor emulator clients provide options for alternate routing.

Configuring Hot Standby on the Backup Server

From the **SNA Node Configuration** advanced panel, use the Configure Hot Standby Critical Servers configuration option to configure hot standby on the backup server.

You do not need to perform any hot standby configuration on the critical server, but the backup server requires hot standby configuration. Using the **SNA Node Configuration** advanced option, configure the backup server with critical server definitions. A critical server definition specifies the name of the critical server monitored by the backup server and the names of the connections activated when the critical server fails. Multiple critical server definitions can be configured, enabling a backup server to monitor and provide backup for multiple critical servers. After configuring critical server definitions on the backup server, restart the backup server node for the changes to be put into effect. For detailed information on configuring hot standby, see the online help for **SNA Node Configuration**.

Configuration of a Hot Standby Environment

Follow these steps to configure a hot standby environment:

1. **Configure the Critical Server:** The critical server is configured without any hot standby related configuration. Configure the host connectivity to match the alternate routing option used with the emulator clients.
2. **Configure the Backup Server:** Backup server configuration involves two key concepts:
 - a. Configuration of host connectivity on the backup server is similar to the critical server, except that host connections are configured for manual initiation.
 - b. Critical server configuration on the backup server to define the host connections started when a critical server failure is detected.
3. **Configure the emulator clients:** The workstation emulator clients are configured to match an alternate routing option to the connectivity being supplied by the servers.
4. **Start the Critical Server and Backup Server nodes:** Start the nodes normally. **SNA Node Operations** can verify the hot standby status for the critical and backup servers.
5. **Critical Server failure:** When the critical server fails, recovery occurs as follows:
 - a. The backup server, the host, and the workstation emulator software on the client detect the failure at approximately the same time.
 - b. The connections named in the critical server configuration on the backup server are activated and connectivity to the host machine is established.
 - c. The workstation emulator client establishes a new session with the backup server. The backup server routes sessions to the host machine through the connection activated on the backup server when the critical server failed.
 - d. The workstation contacts the host machine and displays a logon prompt for the user.

Planning for Backup Host Connections

6. **Critical Server back online:** When the critical server is active again, new workstation emulator sessions use the critical server host connectivity. Existing workstation emulator sessions routed through the backup server continue operation until the sessions are closed.

You must deactivate of the backup server host connections manually. Manual deactivation of the backup server connections should be performed during off hours. You should notify users of workstation emulator sessions with the backup server prior to the deactivation of the connections.

Planning for Backup Connections between Servers without Using Hot Standby

Because hot standby is limited to environments supporting LAN based IP or IPX workstation emulator sessions, some environments can not use hot standby for backup connectivity. Many considerations for planning a backup environment without hot standby are similar to backup planning using hot standby. Follow these steps to configure a backup environment without using hot standby:

1. On the backup server, install a separately purchased Communications Server package with adequate licenses to provide for backup host connections to the primary server.
2. On the backup server, configure host connections used as backups to activate on demand.
3. Configure the workstation emulator clients to provide automatic routing to the resources on the backup server when the primary server fails. However, routing should be restricted to access the backup server only if a failure of the primary server is detected.
4. Both nodes are started, but only the primary node with its active upstream connections initially supply connectivity to the host.
5. When the primary server fails, the workstation emulator clients attempt subsequent connections through the backup server. Host connections on the backup server are activated.
6. The duplicate licensing installed on the backup server provides availability of the backup server to clients previously routed through the primary server.
7. When the primary server is active again, connections on the backup server are deactivated manually.

Planning for Backup Connections in a Single Server

You can not use hot standby for backup connectivity in a single server, because hot standby only applies to environments with backup connections between servers. However, some considerations for planning a backup environment within a single server are similar to backup planning using hot standby. The differences are summarized by the following steps for backup connections in a single server:

1. On the server acting as the gateway to the host, configure primary and backup connections that supply LUs to the same LU pool, but with the backup connection configured to activate on demand.
2. Configure the workstation emulator clients to connect through the LU pool.
3. When the primary connection fails, or the LUs in the pool are exhausted, an attempt to connect to the pool initiates activation of the backup connection.
4. When the primary server is active again, connections on the backup server are deactivated manually.

Chapter 13. Planning for X.25

You need to plan carefully for X.25 configuration. You need to consider the following:

- To use a public PSDN, you need to apply for a network subscription for each line (link) required. It might take some time for a line to be supplied, so check well in advance with the X.25 network provider.
- What functions are available on the X.25 network and what functions are required. The network provider usually supplies a network technical guide, which you should obtain and consult.
- Collect information on logical channels you require for PVCs and SVCs.
 - Up to 16 groups of 256 logical channels can be defined.
 - The group number ranges from 0 through 15.
 - The logical channel number within the group ranges from 0 through 255.

Multiplying the group number by 256 and adding the logical channel number gives a number in the range from 0 through 4095, which uniquely identifies a logical channel. This combination is sometimes referred to as the logical channel number.

Notes:

1. If you specify the CCITT as 1984 or 1988, the maximum number of logical channels is 4095.
 2. If you specify the CCITT as 1980, the maximum number of logical channels is 1024.
- Groups of logical channels can be incoming-only (reserved for calls from other DTEs), outgoing-only (reserved for calls to other DTEs), or two-way.
 - How many virtual circuits your use of SNA requires
 - What PVCs you require
 - The mix of incoming-only, outgoing-only, and two-way logical channels
 - The maximum packet size you require and what is supported by the X.25 network.
 - Details of which optional X.25 facilities your X.25 application programs will use. The CCITT X.25 recommendations define a number of optional facilities that your network might support.

Some facilities have parameters that can be specified or negotiated by an application program during call setup and clearing using the facilities fields in the call request, call accepted, and clear request packets.

The X.25 connection can be configured to support the following facilities during call setup:

- Closed user group
- Network user identification
- Packet and window size negotiation
- Reverse charging

Define the facilities during X.25 DLC configuration.

Planning for X.25

- Details of any special requirements for frame and packet parameters. The default values for these configuration parameters have been chosen carefully to suit most situations. However, your network provider might require or recommend some values.

To avoid problems, your configuration parameters must match the details of your network subscription. For instance, if you configured 20 SVC logical channels but only subscribed to 10, the DCE would detect an error condition when you attempted to use the 11th logical channel.

Possible X.25 Connections

Figure 36 shows how an X.25 network can provide a connection from an SNA gateway to a host.




Figure 36. X.25 and SNA Gateway Connectivity

X.25 Network Charges

Actual charges depend on the billing practices of the X.25 network provider, but charges usually consist of some or all of the following elements:

- Installation charge
A one-time installation charge when you first subscribe to the X.25 network.
- Periodic rental charge
This charge depends on the facilities supplied by your network provider. For instance, you can expect to pay a charge for each optional network facility that you subscribe to. You can also expect to pay a charge for each available logical channel. PVCs generally cost more than SVCs because they require dedicated network resources.
- Usage charge
This charge depends on the length of time for which you are logically connected to the network, possibly subject to a minimum charge.
- Volume charge
This charge is related to the amount of data transmitted over the network during a billing period.

Chapter 14. Configuring Host Connectivity

You can connect workstations running Communications Server to zSeries (including System/370) or iSeries hosts. Many of the values you specify for Communications Server configuration are also used when you configure the host. This chapter describes the mapping between the configuration parameters you specify during Communications Server configuration and the parameters you specify on the hosts. This chapter also contains tables that list the parameters that can be used to eliminate the common parameters on each macro to be issued between Communications Server and the host.

Note: Many NCP parameters issued from the LINE, PU, and LU macros can be specified in the GROUP macro for readability and ease of coding. This matching of parameters eliminates the need to include common parameters on each subsequent macro and still allows the capability to override when appropriate.

NCP Token Ring (NTRI) with VTAM and Communications Server

Table 7 compares host parameters to Communications Server parameters for IBM Token-Ring Network connections.

Table 7. Host/PC Parameters Cross-Reference: IBM Token-Ring Network NTRI Definition

| NTRI/NCP | Communications Server Profile/Parameters | Remarks |
|---------------------|---|--|
| GROUP ECLTYPE = PHY | DLC Type: LAN | |
| LINE LOCADDR | Configure Connections LAN Destination address | The value set for this parameter under Communications Server must be the address of the IBM Token-Ring Network adapter in the host controller. |

Token Ring with VTAM Switched Network and Communications Server

Table 8 compares host parameters to Communications Server parameters for IBM Token-Ring Network connections using a VTAM switched network.

Table 8. Host/PC Parameters Cross-Reference: IBM Token-Ring Network VTAM Switched Net

| VTAM | Communications Server Profile/Parameters | Remarks |
|--------|---|---|
| IDBLK= | Configure Node Local node ID Block ID | The host IDBLK must match the first 3 hexadecimal characters of the Local node ID , which defaults to X'05D' for Communications Server. Specify this parameter when using switched lines, or omit this parameter if you use CPNAME=. |

Configuring Host Connectivity

Table 8. Host/PC Parameters Cross-Reference: IBM Token-Ring Network VTAM Switched Net (continued)

| VTAM | Communications Server Profile/Parameters | Remarks |
|---------|---|---|
| IDNUM= | Configure Node Local Node ID Physical Unit ID | The host IDNUM must match the last 5 hexadecimal characters of the Local node ID . Specify this parameter when using switched lines, or omit this parameter if you use CPNAME=. |
| CPNAME= | Node Definition Configure Node Fully qualified CP name (second field) | Specify this parameter when using switched lines, or omit when using IDBLK= and IDNUM=. |
| LOCADDR | Configure Host LUs NAU address | These parameters must have the same value. For independent LUs, VTAM recommends creating a CDRSC definition or using the DYNLU operand instead of coding LOCADDR=0. For more information on coding CDRSC definitions for independent LUs, refer to the <i>VTAM Network Implementation Guide</i> . No host definitions are required for dependent LUs if SDDL or dynamic definition is used for independent LUs on VTAM 3.4 or later. |
| MAXDATA | Configure LAN Device Maximum PIU size (265–65535) | These values do not have to match. |
| MAXOUT | Configure LAN Device Receive buffer count (2–64) | The Receive buffer count in Communications Server should be equal to the MAXOUT parameter on the host. |

9370 IBM Token-Ring Network Adapter and Communications Server

Table 9 compares 9370 host parameters to Communications Server parameters for IBM Token-Ring Network connections.

Table 9. Host/PC Parameters Cross-Reference: IBM Token-Ring Network 9370 VTAM LAN

| VTAM | Communications Server Profile/Parameters | Remarks |
|-------------------------|---|--|
| VBUILD TYPE = LAN | DLC Type: LAN | |
| PORT MACADDR= | Configure Connection LAN Destination address | The destination address specified in the Communications Server customization must be the address of the 9370 IBM Token-Ring adapter. |
| PORT SAPADDR=04 | Configure Connection LAN Remote SAP | These values must match. |
| ADDR=04 (outgoing only) | Configure Connection LAN Local SAP | These values must match. |

3174 Controller and Communications Server

Table 10 compares 3174 controller parameters to Communications Server for IBM Token-Ring connections.

Table 10. Parameter Matching between a 3174 Token Ring and Communications Server

| 3174 | Communications Server Profile/Parameters | Remarks |
|--|---|--|
| Upper limit address greater than 104: CUA | DLC Type: LAN | |
| Token-Ring address for gateway | Configure Connections LAN Destination address | These values must match. |
| Ring transmission definition. W= window max out, F=0, W = 1-7, F=1, W = 1-7, F=2, W = 1-4, F=3, W = 1-2 | Configure Devices LAN Receive buffer count (1-8) | The Communications Server receive buffer count should be equal to 3174 W. |

3174 Controller VTAM PU and Communications Server

Table 11 compares 3174 controller VTAM PU parameters to Communications Server for IBM Token-Ring connections.

Table 11. Host/PC Parameters Cross-Reference: IBM Token-Ring Network 3174 Controller VTAM PU and Communications Server

| VTAM | Communications Server Profile/Parameters | Remarks |
|---------------------------------|--|---|
| CUADDR (local) or ADDR (remote) | (none) | Mapped by 3174 controller to IBM Token-Ring address |
| MAXDATA (remote only) | Configure Devices LAN Maximum PIU size (99-65535) | The Maximum PIU size value set on the Communications Server workstation must match the value set on the host. |
| MAXOUT (remote only) | Configure Devices LAN Receive buffer count (2-64) | The Communications Server Receive buffer count should be equal to MAXOUT. |
| LOCADDR | Configure Host LUs NAU address | These parameters must have the same value. For independent LUs, VTAM recommends creating a CDRSC definition or using the DYNLU operand instead of coding LOCADDR=0. For more information on coding CDRSC definitions for independent LUs, refer to the <i>VTAM Network Implementation Guide</i> No host definitions are required for dependent LUs if SDDL or dynamic definition is used for independent LUs on VTAM 3.4 or later. |

Ethernet with VTAM Switched Network and Communications Server

Table 12 compares host parameters to Communications Server parameters for Ethernet Network connections using a VTAM switched network.

Table 12. Host/PC Parameters Cross-Reference: Ethernet Network VTAM Switched Network

| VTAM | Communications Server Profile/Parameters | Remarks |
|---------|---|--|
| IDBLK= | Configure Node Local node ID Block ID | The host IDBLK must match the first 3 hexadecimal characters of the Local node ID , which defaults to X'05D' for Communications Server. Specify this parameter when using switched lines, or omit this parameter if you use CPNAME=. |
| IDNUM= | Configure Node Local Node ID Physical Unit ID | The host IDNUM must match the last 5 hexadecimal characters of the Local node ID . Specify this parameter when using switched lines, or omit this parameter if you use CPNAME=. |
| CPNAME= | Node Definition Configure Node Fully qualified CP name (second field) | Specify this parameter when using switched lines, or omit when using IDBLK= and IDNUM=. |
| LOCADDR | Configure Host LUs NAU address | <p>These parameters must have the same value. For independent LUs, VTAM recommends creating a CDRSC definition or using the DYNLU operand instead of coding LOCADDR=0. For more information on coding CDRSC definitions for independent LUs, refer to the <i>VTAM Network Implementation Guide</i></p> <p>No host definitions are required for dependent LUs if SDDL or dynamic definition is used for independent LUs on VTAM 3.4 or later.</p> |
| MAXDATA | Configure LAN Device Maximum PIU size (265–65535) | These values do not have to match. |
| MAXOUT | Configure LAN Device Receive buffer count (2–64) | The Receive buffer count in Communications Server should be equal to the MAXOUT parameter on the host. |

VTAM/NCP and Communications Server (SDLC)

Table 13 on page 139 compares VTAM/NCP host parameters to Communications Server parameters for SDLC connections.

Table 13. Host/PC Parameters Cross-Reference: SDLC

| VTAM/NCP | Communications Server Profile/Parameters | Remarks |
|------------|---|--|
| NETID= | Configure Node Fully qualified CP name (first field) | This parameter can be used for both leased and switched lines. This parameter identifies the SNA network you connect to using Communications Server. |
| NAME PU | | Always use a NAME PU. This PU name should be a name that is not the same as the local CP name so you can define an LU for the control point. |
| IDBLK= | Configure Node Local node ID Block ID | The host IDBLK must match the first 3 hexadecimal characters of the Local node ID , which defaults to X'05D' for Communications Server. Specify this parameter when using switched lines, or omit this parameter if you use CPNAME=. |
| IDNUM= | Configure Node Local Node ID Physical Unit ID | The host IDNUM must match the last 5 hexadecimal characters of the Local node ID . Specify this parameter when using switched lines, or omit this parameter if you use CPNAME=. |
| CPNAME= | Node Definition Configure Node Fully qualified CP name (second field) | Specify this parameter when using switched lines, or omit when using IDBLK= and IDNUM=. |
| LNCTL=SDLC | DLC Type: SDLC-WAC, SDLC-MPA, COM Port (depends on hardware used) | SDLC must be selected here. This parameter is used for both leased and switched lines. |
| LOCADDR | Configure Host LUs NAU address | <p>These parameters must have the same value. For independent LUs, VTAM recommends creating a CDRSC definition or using the DYNLU operand instead of coding LOCADDR=0. For more information on coding CDRSC definitions for independent LUs, refer to the <i>VTAM Network Implementation Guide</i></p> <p>No host definitions are required for dependent LUs if SDDL or dynamic definition is used for independent LUs on VTAM 3.4 or later.</p> |

Configuring Host Connectivity

Table 13. Host/PC Parameters Cross-Reference: SDLC (continued)

| VTAM/NCP | Communications Server Profile/Parameters | Remarks |
|-----------------------|--|---|
| NRZI=YES/NO | Configure Connections SDLC-WAC, SDLC-MPA, COM Port NRZI encoding | <p>Communications Server uses USE_NRZI_ENCODING=0 and USE_NRZI_ENCODING=1.</p> <p>This parameter is found only on PCMs for Modem Type: Synchronous Switched, Synchronous NonSwitched, and AutoSync.</p> <p>The encoding scheme, NRZI or NRZ, must be configured the same way at both ends of a connection. If the encoding schemes do not match, the modems will connect but data transfer will not succeed and the connection will be disconnected by Communications Server after a link-level timeout (typically 40 seconds).</p> <p>We recommend the NRZ encoding scheme when you connect to an X.25 network using an IBM Wide Area Connector.</p> |
| DUPLEX= FULL/HALF | Configure Connections SDLC-WAC and COM Port only Full-duplex communication | The values set on each side must be consistent with the modem type and configuration. This parameter is used for both leased and switched lines. |
| (TRANSFER * BFRS) -47 | Configure Devices SDLC-WAC, SDLC-MPA, COM Port Maximum PIU size (265-4105) | The Maximum I-field size value must be less than or equal to the (TRANSFER * BFRS) - 47 indicated on the host. This parameter is used for both leased and switched lines. |
| MAXDATA= | Configure Devices SDLC-WAC, SDLC-MPA, COM Port Maximum PIU size (265-4105) | These values do not have to match. This parameter is used for both leased and switched lines. The lowest value will be used. |
| MAXOUT= | Configure Devices SDLC-WAC, SDLC-MPA, COM Port Receive window count (1-30) | The Receive window count in Communications Server should be equal to the MAXOUT parameter on the host. This parameter is used for both leased and switched lines. |
| ADDR= | Configure Devices SDLC-WAC, SDLC-MPA, COM Port Link station address (01-FE) | These parameters must have the same value. This is coded in hexadecimal for both Communications Server and VTAM/NCP. This parameter is used for both leased and switched lines. |

VTAM and Communications Server (APPC)

Table 14 on page 141 compares VTAM host parameters to Communications Server parameters for APPC connections.

Table 14. Host/PC Parameters Cross-Reference: APPC

| VTAM | Communications Server Profile/Parameters | Remarks |
|------------------------|--|---|
| LU NAME | Configure Host LU Local LU name or: Configure Node Fully qualified CP name (second field) | |
| LOCADDR | Configure Host LUs NAU address | These parameters must have the same value. For independent LUs, VTAM recommends creating a CDRSC definition or using the DYNLU operand instead of coding LOCADDR=0. For more information on coding CDRSC definitions for independent LUs, refer to the <i>VTAM Network Implementation Guide</i> No host definitions are required for dependent LUs if SDDL or dynamic definition is used for independent LUs on VTAM 3.4 or later. |
| NETID APPL | Configure Partner LU 6.2 Partner LU name | |
| VTAM Interpret Table | Configure Partner LU 6.2 Partner LU name | This is for dependent LUs only. |
| LOGMODE Table LOGMODE= | Configure Mode Mode name | These parameters must match. |
| LOGMODE Table RUSIZES= | Configure Mode Maximum RU size (256–32767) | Final values are negotiated. |

VTAM and Communications Server (MPC)

Table 15 and Table 16 on page 142 compare IOCP/HCD and VTAM host parameters to Communications Server parameters for Multi-Path Channel (MPC) connections.

Table 15. IOCP/HCD Parameters Cross-Reference: MPC

| IOCP/HCD | Communications Server Profile/Parameters | Remarks |
|-------------------------------|--|---------|
| IOCP IODEVICE UNIT=SCTC | DLC Type: IBM-MPC | |
| HCD Device Type SCTC | DLC Type: IBM-MPC | |

Configuring Host Connectivity

Table 16. VTAM Parameters Cross-Reference: MPC

| VTAM Definition Statement | Operand | Communications Server Profile/Parameters | Remarks |
|---------------------------|---------------|--|--|
| VBUILD,TYPE=TRL | LNCTL=MPC | DLC Type: IBM-MPC | |
| | MPCLEVEL=HPDT | | Communications Server IBM-MPC connections require HPR/HPDT |
| | MAXBFRU | Configure MPC DLC Maximum PIU size | May affect performance and/or VTAM storage use |
| | MAXREADS | Configure MPC DLC Receive Buffer Count | May affect performance and/or VTAM storage use |
| | READ | Configure MPC DLC | The list of VTAM READ IODEVICE addresses must correspond to the Communications Server MPC DLC list of READ subchannels |
| | WRITE | Configure MPC DLC | The list of VTAM WRITE IODEVICE addresses must correspond to the Communications Server MPC DLC list of WRITE subchannels |
| VBUILD,TYPE=LOCAL | CONNTYPE=APPN | Configure Connections IBM-MPC APPN Support | Communications Server IBM-MPC connection requires APPN. |
| | XID=YES | Configure Connections IBM-MPC APPN Support | Communications Server IBM-MPC connection requires PU2.1. |
| | CPCP=YES | Configure Connections IBM-MPC APPN Support | Communications Server IBM-MPC connection requires APPN. |
| | HPR=YES | Configure Connections IBM-MPC HPR Support | Communications Server IBM-MPC connection requires HPR. |
| | NN | Configure Node Node type | VTAM PU node type must match Communications Server definition. |
| | DELAY | | May affect performance and/or VTAM CPU use |

ES/9000 and Communications Server (SDLC)

Table 17 on page 143 compares 9370 Information System parameters to Communications Server parameters for SDLC connections.

Table 17. Host/PC Parameters Cross-Reference: 9370 Information System via SDLC

| 9370/VTAM | Communications Server Profile/Parameters | Remarks |
|--|---|---|
| QFI: Non-return -to-zero inverted (microcode configuration) | Configure Connections SDLC-WAC, SDLC-MPA, COM Port NRZI encoding | These parameters must match. This parameter is used for both leased and switched lines. Communications Server uses USE_NRZI_ENCODING=0 and USE_NRZI_ENCODING=1. This is the synchronous data stream. |
| QFI: Protocol (microcode configuration) VTAM: GROUP LNCTL=SDLC | DLC Type SDLC-WAC, SDLC-MPA COM Port | The microcode and VTAM parameters must match. This parameter is used for both leased and switched lines. |
| MAXDATA | Configure Devices SDLC-WAC, SDLC-MPA, COM Port Maximum PIU Size (265–4105) | These values do not have to match. This parameter is used for both leased and switched lines. |
| MAXBFRU * IOBUF size | Configure Devices SDLC-WAC, SDLC-MPA, COM Port Maximum PIU Size (265–4105) | The Maximum I-field size value must be equal to or less than the MAXBFRU times IOBUF size indicated on the host. This parameter is used for both leased and switched lines. |
| MAXOUT= | Configure Devices Receive window count (1–30) | The Receive window count in Communications Server should be equal to the MAXOUT parameter on the host. This parameter is used for both leased and switched lines. |
| ADDR= | Configure Devices Link station address (01–FE) | The values set for these parameters must match. This parameter is used for both leased and switched lines. |
| IDBLK= | Configure Node Local Node ID Block ID | The host IDBLK must match the first 3 hexadecimal characters of the Local node ID , which defaults to X'05D' for Communications Server. Specify this parameter when using switched lines, or omit this parameter if you use CPNAME=. |
| IDNUM= | Configure Node Local node ID Physical unit ID | The host IDNUM must match the last 5 hexadecimal characters of the Local node ID . Specify this parameter when using switched lines, or omit this parameter if you use CPNAME=. |
| CPNAME= | Configure Node Fully qualified CP name (second field) | Specify this parameter when using switched lines, or omit when using IDBLK= or IDNUM=. |

Configuring Host Connectivity

Table 17. Host/PC Parameters Cross-Reference: 9370 Information System via SDLC (continued)

| 9370/VTAM | Communications Server Profile/Parameters | Remarks |
|-----------|--|--|
| LOCADDR | Configure Host LUs NAU address | LOCADDR and NAU parameters must have the same value. For independent LUs, VTAM recommends creating a CDRSC definition or using the DYNLU operand instead of coding LOCADDR=0. For more information on coding CDRSC definitions for independent LUs, refer to the <i>VTAM Network Implementation Guide</i> . This parameter is used for both leased and switched lines. |

Host and Communications Server (DLUR)

Table 18 compares host and Communications Server parameters for VTAM connections over DLUR.

Table 18. VTAM Communications Server Parameter Cross-Reference: Using DLUR

| VTAM | Communications Server | Remarks |
|---------|--|--|
| IDBLK= | Configure DLUR PUs Block ID | The IDBLK must match the first 3 hexadecimal characters of the Node ID, that defaults to X'05D' for Communications Server Specify this parameter when NOT using CPNAME=. |
| IDNUM= | Configure DLUR PUs Physical unit ID | The IDBLK must match the last 5 hexadecimal characters of the Node ID. Specify this parameter when NOT using CPNAME=. |
| CPNAME= | Configure DLUR PUs PU name | The CPNAME must match the local PU name. Specify this parameter when NOT using IDNUM= and IDBLK=. |

Table 19 lists the following parameters on the PATH statement that are used when establishing a connection from the host to the workstation.

Table 19. VTAM Communications Server DLUR Parameter Matching Connection from Host to Workstation

| VTAM | Communications Server | Remarks |
|-------------------------|---|--|
| DLURNAME= | Configure Node Fully qualified CP name (second field) | The DLURNAME provides the name of the DLUR node that owns the DLUR PU. For Communications Server, this is the fully qualified CP name. |
| DLCADDR= (1,C,INTPU) | | Required. This parameter tells the DLUR that the connection is for an internal PU. |

Table 19. VTAM Communications Server DLUR Parameter Matching Connection from Host to Workstation (continued)

| VTAM | Communications Server | Remarks |
|--------------------------|--|---|
| DLCADDR=(2,X,zzzzzzzz) | Configure DLUR PUs Block ID and Physical unit ID | The zzzzzzzz in DLCADDR=(2,X,zzzzzzzz) must match the Node ID. Specify this parameter when NOT using DLCADDR=(3,C,puname). |
| DLCADDR= (3,C,puname) | Configure DLUR PUs PU name | The PU name in DLCADDR=(3,C,puname) must match Local PU name. Specify this parameter when NOT using DLCADDR=(2,X,zzzzzzzz). |

NPSI and Communications Server (X.25)

Communications Server X.25 support enables the transmission of SNA frames across an X.25 network. X.25 treats SNA frames as raw data and transmits them as data-within-data packets. The two network nodes are not directly connected, but both are connected to an intermediary called the packet-switching data network (PSDN). The X.25 network can comprise any of the X.25 connectivities.

The software that interacts with the PSDN and packages the SNA data into packets and extracts it intact is an integral part of Communications Server. On the S/370 host side, the X.25 functions are handled by a separate software product called X.25 Network Control Program Packet Switching Interface (NPSI). The NPSI definitions are generated along with the NCP and the NPSI software runs in harmony with the NCP in the same 3705, 3725, 3720, or 3745 controller.

The majority of the X.25 parameters specified in both Communications Server and NPSI take values that are agreed upon by you and the X.25 network provider rather than between the two nodes. For example, with an SDLC connection on a nonswitched line, the host controller (NCP) is configured to poll a particular secondary address. Communications Server is configured to respond to that address on a physical link between the Communications Server workstation and the host.

NPSI Connections over X.25

Table 20 compares host and Communications Server parameters for NPSI connections over X.25.

Table 20. Host/PC Parameter Cross-Reference: NPSI Connections over X.25

| NPSI | Communications Server Profile/Parameters | Remarks |
|---------------------------|--|--|
| X25.LINE TYPE= | DLC Type X.25 COM Port, X.25-WAC | Specify TYPE=P for PVCs and TYPE=S for SVCs for NPSI X.25 line definition. |
| CALL= (only for type = S) | Configure Devices X.25 COM Port, X.25-WAC Virtual circuit ranges | Virtual circuit ranges must match your X.25 network subscription. |

Configuring Host Connectivity

Table 20. Host/PC Parameter Cross-Reference: NPSI Connections over X.25 (continued)

| NPSI | Communications Server Profile/Parameters | Remarks |
|----------|--|---|
| MAXDATA= | Configure Devices X.25 COM Port, X.25-WAC Maximum PIU size (256–4105) | These values do not have to match. |
| VWINDOW | Configure Devices X.25 COM Port, X.25-WAC Window size (1–7) | The SVC/PVC window size must match the X.25 subscriptions. |
| MWINDOW | Configure Devices X.25 COM Port, X.25-WAC Frame values Window size (1–7) | |
| FRMLGTH | Configure Devices X.25 COM Port, X.25-WAC Frame sequence Modulo | The host FRMLGTH value must be at least MAXPKT +3 when running Modulo 8 or MAXPKT +4 when running Modulo 128. The host and Communications Server Modulo size and the packet sizes must match your X.25 subscription. |
| MAXPKTL | Configure Devices X.25 COM Port, X.25-WAC Packet size | The host SVC/PVC packet sizes must match the X.25 subscription value. |

VTAM/NCP and Communications Server (X.25)

Table 21 compares host and Communications Server parameters for X.25 (VTAM/NCP).

Table 21. Host/PC Parameters Cross-Reference: X.25 (VTAM/NCP)

| VTAM/NCP | Communications Server Profile/Parameters | Remarks |
|----------|---|---|
| NETID= | Configure Node Fully qualified CP name (first field) | This identifies the SNA network you connect to using Communications Server. |
| PU NAME | | Always use a PU name. This PU name should not be the same as the local CP name so you can define an LU for the control point. |
| IDBLK= | Configure Node Local node ID Block ID | First 3 hex digits of the Local node ID , which defaults to X'05D' for Communications Server. Specify this parameter when using switched lines, or omit this parameter if you use CPNAME=. |

Table 21. Host/PC Parameters Cross-Reference: X.25 (VTAM/NCP) (continued)

| VTAM/NCP | Communications Server Profile/Parameters | Remarks |
|------------------------|--|--|
| IDNUM= | Configure Node Local node ID Physical unit ID | The host IDNUM must match the last 5 hexadecimal characters of the Local node ID . Specify this parameter when using switched lines, or omit this parameter if you use CPNAME=. |
| CPNAME= | Configure Node Fully qualified CP name (second field) | Specify this parameter when using switched lines, or omit when using IDBLK= and IDNUM=. |
| (TRANSFER * BFRS) - 47 | Configure Devices X.25 COM Port, X.25-WAC Maximum PIU size (265–4105) | The Maximum PIU size must be equal to the value set on the host. |
| MAXDATA= | Configure Devices X.25 COM Port, X.25-WAC Maximum PIU size (265–4105) | These values do not have to match. |
| LOCADDR | Configure Host LUs NAU address | <p>These parameters must have the same value. For independent LUs, VTAM recommends creating a CDRSC definition or using the DYNLU operand instead of coding LOCADDR=0. For more information on coding CDRSC definitions for independent LUs, refer to the <i>VTAM Network Implementation Guide</i></p> <p>No host definitions are required for dependent LUs if SDDL or dynamic definition is used for independent LUs on VTAM 3.4 or later.</p> |

APPC over X.25

Table 22 matches the parameters for a Communications Server workstation running advanced program-to-program communications (APPC) connecting to a host by way of X.25.

Table 22. Host/PC Parameters Cross-Reference: APPC on an IBM Token-Ring Network Using X.25 to a Host

| NCP | Communications Server Profile/Parameters | Remarks |
|---------|---|---------|
| LU NAME | Configure Local LU 6.2 Local LU name or: Configure Node Fully qualified CP name (second field) | |

Configuring Host Connectivity

Table 22. Host/PC Parameters Cross-Reference: APPC on an IBM Token-Ring Network Using X.25 to a Host (continued)

| NCP | Communications Server Profile/Parameters | Remarks |
|------------------------|---|--|
| LOCADDR | Configure Host LUs NAU address | These parameters <i>must</i> have the same value. For independent LUs, VTAM recommends creating a CDRSC definition or using the DYNLU operand instead of coding LOCADDR=0. For more information on coding CDRSC definitions for independent LUs, refer to the <i>VTAM Network Implementation Guide</i> No host definitions are required for dependent LUs if SDDL or dynamic definition is used for independent LUs on VTAM 3.4 or later. |
| NETID APPL | Configure Partner LU 6.2 Partner LU name | |
| VTAM Interpret Table | Configure Partner LU 6.2 Partner LU name | This is for dependent LUs only. |
| LOGMODE Table LOGMODE= | Configure Mode Mode name | These parameters <i>must</i> match. |
| LOGMODE Table RUSIZES= | Configure Mode Maximum RU size (256–32767) | These values must match. |

iSeries Host Line Definitions

iSeries host line definitions include matching OS/400 host program parameters with Communications Server parameters. Table 23 contains the SDLC line parameters.

Table 23. iSeries Communications Server Parameters Cross-Reference: OS/400 Host Program Line Definition

| OS/400 | Communications Server Profile/Parameters | Remarks |
|-------------------------------|---|--|
| NRZI= Data Encoding | Configure Connections SDLC-WAC, SDLC-MPA, COM Port NRZI encoding | Communications Server uses USE_NRZI_ENCODING=0 and USE_NRZI_ENCODING=1. This parameter is found only on PCMs for Modem Type: Synchronous Switched, Synchronous NonSwitched, and AutoSync. |
| Modem data rate select | Configure Devices COM Port only Configure modem | The parameters on both systems for modem connection types, asynchronous switched, must match. |
| Maximum frame size (MAXFRAME) | Configure Devices SDLC-WAC, SDLC-MPA, COM Port Receive window count (1–30) | Asynchronous nonswitched values do not have to match. |

Table 23. iSeries Communications Server Parameters Cross-Reference: OS/400 Host Program Line Definition (continued)

| OS/400 | Communications Server Profile/Parameters | Remarks |
|-------------------------------------|---|---|
| Maximum outstanding frames (MAXOUT) | Configure Devices SDLC-WAC, SDLC-MPA, COM Port Receive window count (1-30) | The Receive window count in Communications Server should be equal to the MAXOUT parameter on the iSeries host. |

iSeries Host Controller Definitions

Defining the iSeries host controller includes matching OS/400 host program parameters with Communications Server parameters.

Table 24 shows the OS/400 host program parameters.

Table 24. iSeries Communications Server Parameters Cross-Reference: Controller Parameters

| OS/400 | Communications Server Profile/Parameters | Remarks |
|---------------------------|---|--|
| Linktype: *SDLC | DLC Type SDLC-WAC, SDLC-MPA COM Port | Link type on both systems must match. This parameter is on the line description, but include for completeness. |
| Attached Nonswitched Line | Configure Devices SDLC-WAC, SDLC-MPA, COM Port Switched line | This parameter is for nonswitched connections. This value must match a previously defined host line definition. |
| Switched Line List | (none) | This parameter is for switched connections. This value must match a previously defined host line definition. |
| Maximum frame size | Configure Devices SDLC-WAC, SDLC-MPA, COM Port Maximum PIU size (265-4105) | These values do not have to match. |
| Remote network identifier | Configure Node Fully qualified CP name (first field) | These values must match. |
| Remote control point name | Configure Node Fully qualified CP name (second field) | These values must match. |
| Data link role: | Configure Connections SDLC-WAC, SDLC-MPA, COM Port Link station role | The parameters on both systems are complementary. If the host role is primary the Communications Server workstation must be configured as secondary or negotiable. Unless multipoint (*MP) is used, it is best to set the data link role to negotiable (*NEG). |

Configuring Host Connectivity

Table 24. *iSeries Communications Server Parameters Cross-Reference: Controller Parameters (continued)*

| OS/400 | Communications Server Profile/Parameters | Remarks |
|-----------------------------------|--|---|
| Station address | Configure Devices SDLC-WAC, SDLC-MPA, COM Port Link station address (01-FE) | These values must match. |
| APPN CP session support: (Yes/No) | Configure Connections SDLC-WAC, SDLC-MPA, COM Port APPN Support | These must match if CP-CP sessions are desired. |
| APPN node type: | Configure Connections SDLC-WAC, SDLC-MPA, COM Port Adjacent CP type | These values must match. |

iSeries Host Device Definitions

The iSeries host can create a device definition when Communications Server workstation connects without a previously defined device definition.

Table 25 contains the OS/400 host program device parameters.

Table 25. *iSeries Communications Server Parameters Cross-Reference: Device Parameters*

| OS/400 | Communications Server Profile/Parameters | Remarks |
|---------------------------|---|---|
| Remote location name | Configure Local LU 6.2 Local LU name or: Configure Node Fully qualified CP name (second field) | These values must match. |
| Attached controller | (none) | This parameter is for nonswitched connections. This value must match a previously defined host line definition. |
| Local location name | | Use the default provided. |
| Remote network identifier | Configure Node Fully qualified CP name (first field) | Use the default provided. |
| Mode | Configure Mode Mode name | Use the default provided. |

iSeries Host Line Description and Communications Server

Table 26 on page 151 contains the iSeries host program line description.

Table 26. Host/PC Parameters Cross-Reference: Line Description

| OS/400 | Communications Server Profile/Parameters | Remarks |
|------------------------|---|--|
| Local adapter address: | Configure Connections LAN destination address | The two addresses must match to communicate. |

iSeries Host Controller Description and Communications Server

Table 27 contains the OS/400 host program controller description.

Table 27. Host/PC Parameters Cross-Reference: Controller Description

| OS/400 | Communications Server Profile/Parameters | Remarks |
|--|--|---|
| Active switched line | (none) | This parameter is for nonswitched connections. This value must match a previously defined host line definition. |
| Remote control point | Configure Node Fully qualified CP name (second field) | The names must match. |
| Remote network identifier | Configure Node Fully qualified CP name (first field) | Use the default or match the parameters if different from the iSeries local network ID. |
| Model controller description | | This parameter must be YES if the iSeries uses a connection network. |
| Connection network ID and Connection network CP name | Configure Connection Networks Connection network name | These values must match for all nodes participating in a connection network. |

iSeries Twinaxial Descriptions

Table 28 contains the iSeries twinaxial description.

Table 28. Host/PC Parameters Cross-Reference: Controller Description

| iSeries | Communications Server Profile/Parameters | Remarks |
|---------------------------|--|---|
| Attached nonswitched line | (none) | This parameter is for nonswitched connections. This value must match a previously defined host line definition. |
| Remote control point | Configure Node Fully qualified CP name (second field) | The names must match. |
| Remote network identifier | Configure Node Fully qualified CP name (first field) | Use the default or match the parameters if different from the iSeries local network ID. |

Configuring Host Connectivity

Table 28. Host/PC Parameters Cross-Reference: Controller Description (continued)

| iSeries | Communications Server Profile/Parameters | Remarks |
|----------------------|---|---|
| Data link role=*SEC: | Configure Connections SDLC-WAC, SDLC-MPA COM Port Link station role | The parameters on both systems are complementary. If the host role is primary, the Communications Server workstation must be configured as secondary or negotiable. Unless multipoint (*MP) is used, it is best to set the data link role to negotiable (*NEG). |

iSeries Host Line Definitions (X.25)

iSeries host line definitions include matching OS/400 host program parameters with Communications Server parameters. Table 29 contains the X.25 line parameters.

Table 29. iSeries Communications Server Parameters Cross-Reference: OS/400 Host Program Line Definition (X.25)

| OS/400 | Communications Server Profile/Parameters | Remarks |
|-----------------------|---|---|
| Local Network Address | Configure Devices Local Address | Both addresses are usually determined by your network provider. |
| Connection Type | Configure Devices X.25-COM Port X.25-WAC Switched Line | Connection Type on both systems must match. |
| X.25 DCE Support | N/A | Must be set to YES, if the connection between the two systems is point-to-point, without going through an X.25 network. |
| Maximum Frame Size | Configure Devices X.25-COM Port X.25-WAC Maximum PIU Size | Maximum Frame Size on both systems must match. |
| Default Packet Size | Configure Devices X.25-COM Port X.25-WAC Packet Size | This value is determined by your network provider, and is the same for Transmit and Receive. |
| Maximum Packet Size | N/A | This value is determined by your network provider. |
| Modulus | Configure Devices X.25-COM Port X.25-WAC Sequencing | This value is determined by your network provider. |

Table 29. iSeries Communications Server Parameters Cross-Reference: OS/400 Host Program Line Definition (X.25) (continued)

| OS/400 | Communications Server Profile/Parameters | Remarks |
|-------------------------|---|--|
| Default Window Size | Configure Devices X.25-COM Port X.25-WAC Window Size | This value is determined by your network provider, and is the same for Transmit and Receive. |
| Insert Network Address | Configure Devices X.25-COM Port X.25-WAC | This value is determined by your network provider, and should be the same for both systems. |
| Logical Channel Entries | Configure Devices X.25-COM Port X.25-WAC Virtual Circuit Ranges | These values are determined by your network provider and your X.25 virtual circuit type. |

iSeries Host Controller Definitions (X.25 PVC)

Defining the iSeries host controller includes matching OS/400 host program parameters with Communications Server parameters.

Table 30 shows the OS/400 host program parameters for X.25 PVC.

Table 30. iSeries Communications Server X.25 PVC Parameters Cross-Reference: Controller Parameters

| OS/400 | Communications Server Profile/Parameters | Remarks |
|---------------------------------|--|--|
| X.25 logical channel ID | Configure Connections X.25 COM Port, X.25-WAC Logical Channel Number | These values are determined by your network provider. |
| Linktype: *X25 | Configure Connections | Link types on both systems must match. |
| Switched Connection: *NO | | All PVCs are non-switched. |
| Maximum Frame Size | Configure Devices X.25-COM Port X.25-WAC Maximum PIU Size | The Maximum Frame Size values do not need to match. |
| Data Link Role | Configure Devices X.25-COM Port X.25-WAC Link Station Role | The parameters on both systems are complementary. If the host role is primary the Communications Server workstation must be configured as secondary or negotiable. |
| X.25 Network Level | Configure Devices X.25-COM Port X.25-WAC CCITT Compliance | These values are determined by your network provider. |
| X.25 Link Level Protocol: *QLLC | | This value must be QLLC. |

iSeries Host Controller Definitions (X.25 SVC)

Defining the iSeries host controller includes matching OS/400 host program parameters with Communications Server parameters.

Table 31 shows the OS/400 host program parameters for X.25 SVC.

Table 31. iSeries Communications Server X.25 SVC Parameters Cross-Reference: Controller Parameters

| OS/400 | Communications Server Profile/Parameters | Remarks |
|---------------------------------|--|--|
| Linktype: *X25 | Configure Connections | Link types on both systems must match. |
| Switched Connection: *YES | | All SVCs are switched. |
| Maximum Frame Size | Configure Devices X.25-COM Port X.25-WAC Maximum PIU Size | The Maximum Frame Size values do not need to match. |
| Connection Number | Configure Devices X.25-COM Port X.25-WAC Local DTE address | This value is determined by your network provider. |
| Data Link Role | Configure Devices X.25-COM Port X.25-WAC Link Station Role | The parameters on both systems are complimentary. If the host role is primary the Communications Server workstation must be configured as secondary or negotiable. |
| X.25 Network Level | Configure Devices X.25-COM Port X.25-WAC CCITT Compliance | These values are determined by your network provider. |
| X.25 Link Level Protocol: *QLLC | | This value must be QLLC. |

Chapter 15. Implementing Communications Server Configurations

A suggested procedure for installing and configuring the software for your network follows:

1. Plan for hardware.
2. Plan for software applications.
3. Create your naming conventions.
4. Define network addresses.
5. Select configuration and installation tools.
6. Choosing Communications Server features
7. Create configurations and install.
8. Create user materials.
9. Maintain your network.

This chapter provides the information necessary for completing the suggested procedure.

Introduction

A template configuration defines the configuration information common to a group of servers. You may want to create one or more template configurations. Each server configuration file starts with the template configuration, and can contain the few parameters required to customize the configuration for that server. This results in simplified server configuration.

The template configuration can also be used to specify configuration options that cannot be specified using response files.

For information on how to use response files and template files for configuration and installation, refer to “Configuration with Template and Response Files” on page 164.

Remember to accumulate totals for RAM and fixed-disk space requirements for each workstation as you gather information in the following steps. That way, you can ensure that you have adequate hardware for your users when they start using the software you have chosen for them.

Step 1. Planning for Hardware

If you have already determined that you can use existing servers, you need to make sure that the existing hardware has adequate memory and fixed-disk space. The existing hardware must also have the correct system units, displays, printers, keyboards, adapters, modems, and cables for the software that is going to be installed.

If you are going to acquire new servers, record the hardware on the worksheets you prepare for those users.

Storage

Refer to *Quick Beginnings* for storage requirements for Communications Server.

Memory Capacity

When planning the hardware for running Communications Server, it is important to assess how the server will be used in terms of capacity. Determine what type of sessions Communications Server will support. These include, but are not limited to, TN3270E sessions, SNA API client sessions, and traditional SNA sessions to a host. Based on this number, you can calculate the average load on the server and determine the correct amount of memory for the server.

Note: These estimates are in addition to the requirements of the base operating system and any other applications running simultaneously on the same system.

Use the following table to determine what the average load on the server will be:

Table 32. Memory Capacity

| Client Session Type | Memory Usage per Session |
|---------------------------|--------------------------|
| SNA Gateway (traditional) | 29.9 KB |
| SNA API client | 25.2 KB |
| TN3270E | 15.8 KB |

The estimated memory consumption of an active Communications Server without any active links or sessions is 21 MB. To minimize performance loss, it is best to reduce the amount of paging (swapping of memory segments out to disk) that occurs. Try to have as much real memory available in the system as required by the running software, with an additional 5 - 10 MB left as a buffer. For example, a Communications Server installation supporting 1 000 TN3270E client sessions would require the following amounts of memory (based on the values in Table 32).

| | |
|--|---------------------|
| Microsoft Windows NT Server operating system | 19.0 MB (estimated) |
| Communications Server for Windows NT | 21.0 MB |
| TN3270E sessions (1 000 X 15.8 KB) | 15.8 MB |
| Contingency | 10.0 MB |
| Total Memory Recommended | 65.8 MB |

Note: The 65.8 MB is not a required amount of memory for 1 000 TN3270E sessions. It is a recommendation for best performance.

CPU Speed

When determining the minimum processor speed required, you need to understand the average number of transactions that will have to be processed by the server each minute. A transaction is any exchange of information between the client and the host through the server, whether it is a screen refresh or a requested database entry.

Deploying Communications Server Configurations

When calculating the average CPU load, multiply the average number of sessions with the average number of transactions per minute per session and the scale factor from the following table (based on session type).

| Client Session Type | CPU Load Factor (100 Mhz) | CPU Load Factor (166 Mhz) |
|---------------------------|---------------------------|---------------------------|
| SNA Gateway (traditional) | .004 | .002 |
| SNA API client | .021 | .011 |
| TN3270E | .011 | .006 |

For example, to support 1 000 concurrent TN3270E sessions, with an average of 6.7 transactions per minute, the average load on the 166 Mhz CPU would be $1\ 000 \times 6.7 \times .006$, or 40.2%. Performance degrades quickly when the CPU utilization exceeds 80%. A faster CPU would improve Communications Server and system performance. More users and transactions would be able to take advantage of the server simultaneously. While there is no maximum CPU speed supported by Communications Server, 100 Mhz is the recommended minimum.

Step 2. Planning for Software Applications

Software applications have requirements in addition to the requirements for Communications Server. For more information, refer to the software application documentation.

Step 3. Creating Your Naming Conventions

Networks created with Communications Server require that you create and use numerous names for multiple objects on the network. Some of these names can be the same from server to server, but other names must be unique on the network to avoid conflicts among the servers trying to access network resources.

For example, two LANs might have some of the same domain names defined. As long as these LANs were not connected, no conflicts would occur. However, if they should be connected later to a backbone LAN, their names would then be in conflict on the LAN.

This means that you must create naming conventions. Naming conventions are rules and standards you use to assign names to the various network resources.

The following sections list the naming conventions for which you need to plan. If you already know what naming conventions you are going to use in your planning, record them while you plan for your network.

Criteria for Naming

When you are creating naming conventions, you should determine:

- What criteria to use to create names and addresses for your network
- How you are going to keep these names and addresses consistent and, where necessary, unique within the network

If you are attaching a server to a host computer network, for example, you might find that most of the names you must use in your network are chosen for you by host personnel. In that case, for the physical units and logical units required by the host, you must record those given names for your server configuration files.

Deploying Communications Server Configurations

As another example, you might choose to name the servers of your network after the names of the people who use them. For a small network, this could be workable as long as the number of names is small and you are able to keep them unique. However, this would not work for a larger network because people's names are not generally unique. You have to create another convention for workstation names for a large network.

Whatever criteria you choose for the names, be sure to record them. That way, you can refer to the appropriate information any time you add new resources to your network.

Types of Names and Their Restrictions

Most names have to be unique within the network in which you use them. The following sections list the types of names you might encounter when you are planning for, installing, and configuring a network. You should preview these names before you begin planning your network so that you are familiar with them.

The following information is provided for each name:

- A definition
- Whether the name must be unique within a network
- Naming restrictions

Names Used by More than One Component

Names used by more than one component are:

- Network IDs
- Passwords
- User IDs
- Network addressable unit (NAU) names

Network IDs: Network IDs are names given to networks and are used by all the servers and workstations (nodes) within the specific network to maintain a unique identity throughout all the connected networks. The network ID is also used in error logs and network management alerts associated with network system errors.

There are two ways of looking at networks. One way is as a physical network that consists of a "ring" in a token-ring environment or a "string" in an Ethernet or PC network environment. The other way is as a logical network that might not be the same as the physical network. Two or more physical networks (for example, two token rings and an Ethernet string) could be connected with the intent of keeping them in the same logical network.

The network IDs are unique among logical networks; otherwise the networks would be logically the same network. Within a logical network, LU names must be unique to avoid naming conflicts. Between logical networks, the network ID guarantees unique names. An LU name could be the same on two logical networks; however, the network ID for each logical network makes the fully qualified LU names unique. Even if the networks are not currently connected, the network IDs should be unique if you plan to bridge the networks in the future.

You should register your network IDs with IBM. This ensures that SNA networks can be interconnected at a later time without addressing conflicts. Contact an IBM Branch Office for more information about registering your network IDs.

The restrictions for network IDs are:

Deploying Communications Server Configurations

- **Length:** From 1 to 8 characters
- **Allowable characters:** A to Z, 0 to 9, \$, @, #

Passwords: Passwords are security functions required by appropriately configured applications and services to protect data and to restrict access to resources.

Passwords do not have to be unique within the network. Passwords are user-specific.

The restrictions for LU-LU passwords are:

- **Length:** From 1 to 8 characters
- **Allowable characters:** Any hexadecimal string

The restrictions for other passwords, such as those in CPI-C and iSeries connections, are:

- **Length:** From 1 to 10 characters
- **Allowable characters:** Any ASCII string

User IDs: User IDs are unique identifying names you give to the users of your network resources so that they can access database, LAN, or host resources with emulation.

User IDs must be unique in the network.

The restrictions for user IDs are:

- **Length:** From 1 to 8 characters
- **Allowable characters:** A to Z, a to z, 0 to 9, \$, @, #

Network Addressable Unit (NAU) Names

The types of names that you might have to specify in Communications Server are:

- Control point (CP) names (local node names)
- Logical unit (LU) names

Control Point (CP) Names (Local Node Names): The control point (CP) is responsible for managing the node and its resources. In an APPN end node, the control point must communicate with the control point in an adjacent network node to obtain APPN network services. In an APPN network node, the control point must communicate with the control points in adjacent network nodes to manage the network. The control point directs such functions as adapter activation and deactivation, and link activation and deactivation, and assists LUs in session initiation and termination.

The control point name is the second half of the fully qualified CP name in the NODE definition of Communications Server.

Control point names must be unique within a network. However, a node can have multiple PU names that are defined in the connecting (**LINK_STATION**) definition and exchanged on XID3 to different hosts. These multiple PU names must be unique within the node and in the host being connected.

The PU and control point are not the same to subarea VTAM. A PU name for each peripheral node in a VTAM domain is defined at that VTAM, and represents the view VTAM has of the peripheral nodes. The PU names at VTAM are not known to the peripheral nodes; that is, VTAM does not send the PU names to the

Deploying Communications Server Configurations

peripheral nodes. If you want the PU name at a peripheral node to be the same as the one defined at VTAM, you must coordinate this. It is recommended that you do this, but for SNA it is not required.

To VTAM, the control point is an LU, used for activating LU 6.2 sessions between the control point and a VTAM LU (for example, CICS). The control point name defined at the peripheral node must match an LU definition at VTAM if VTAM initiates LU 6.2 sessions to the control point. Otherwise, VTAM learns the control point name when the peripheral node initiates a session to a VTAM LU.

For Communications Server, the control point name (not including the network ID) is treated as both the local node control point name and its PU name. The only flow that contains the peripheral node PU name is an alert. However, when Communications Server (APPC/APPN) sends an alert, it includes the control point name (resource type is CP) in the alert, not a PU name. If emulators send a PU name in the alerts, the name is the same as the control point name (again, because Communications Server uses the control point name as the node PU name). Host focal point can be defined only on a link where the PU name is the same as the control point name. Further, host links with a PU name other than the control point name cannot have CP-CP sessions with the host or route APPN traffic over the link. All alerts include the control point name even if the alert is caused by a condition on a link using another PU.

The restrictions for local node names are:

- **Length:** From 1 to 8 characters
- **Allowable characters:** A to Z, 0 to 9, \$, @, #
- **First character:** A to Z, \$, @, #

Logical Unit (LU) Names: Logical unit (LU) names are names given to SNA logical entities within a node that provide support functions for transaction processing. This enables them to communicate to other LUs on the network, including host applications.

The restrictions for LU names are:

- **Length:** From 1 to 8 characters
- **Allowable characters:** A to Z, 0 to 9, \$, @, #

Sockets over SNA LU names: The Sockets over SNA Gateway must have an LU name configured for the Gateway to successfully initialize. Sockets over SNA Gateway will dynamically define the configured LU name to Communications Server upon initialization.

Using a predefined naming convention for Sockets over SNA LU, names can help you:

- Control which names are Sockets over SNA Gateway LUs and which names represent other LUs in your network.
- Analyze and resolve any mapping and routing errors should they occur.
- Enable algorithmic mapping to be used when mapping IP addresses to Sockets over SNA LU names.

For more information on mapping IP addresses to LU names, see “Configuring AnyNet Sockets over SNA” on page 77.

Step 4. Defining Network Addresses

Ensure the consistency and uniqueness of the addresses within your network. Each address must be unique. The addresses that you define will depend on how you configure your network. The following sections describe the addresses for:

- LAN
- SDLC
- X.25

Record the addresses that you use to make sure, when necessary, that none of the addresses conflict and that they are consistent with the naming conventions you have chosen.

LAN Adapter Addresses

LAN adapter addresses are 12-character hexadecimal numbers encoded in the adapter card by the manufacturer (universally administered adapter address) or assigned by the network administrator (locally administered address). Each network adapter card in the workstation you are configuring for LAN communications must have a unique address.

You can use the universally administered addresses, also referred to as “burned-in addresses”, for your network adapter cards, or you can assign locally administered addresses. LAN adapter addresses must be unique within the network. If you use locally administered addresses, ensure that the addresses are unique within the network.

Locally administered addresses offer a significant advantage in the event of an adapter failure that requires replacement of the adapter. You can transfer the existing address to the replacement adapter and avoid changing any of the configurations referring to this address. If you use universally administered addresses, you have to change the network adapter card address for all of the workstations that access the defective adapter card.

The limitations for configuring LAN adapter addresses are:

- **Length:** 12 characters
- **Allowable characters:** A, B, C, D, E, F, 0-9
- **Range:** X'400000000000' through X'7FFF FFFF FFFF'

On the LAN, you might choose to use locally administered LAN adapter addresses. Thus, you are assigning hexadecimal numbers as the LAN adapter addresses to each of the LAN adapters instead of using the universal LAN adapter addresses built into the LAN adapters. You might decide to use a convention with the following criteria:

- First 4 digits are always X'4000'
- Next 3 digits are always X'268' to represent your department
- Final digits are in the range from X'01000' to X'FFFFFF' to represent the individual workstations of your network

Universally administered addresses on Ethernet are in Ethernet format. You can specify the format type (Ethernet or Token-ring) when using locally administered addresses. When configuring the destination address in SNA connections, make sure that the address format is the same as that specified at the remote. With

Deploying Communications Server Configurations

bridging, it is possible to be on a token-ring locally and have the remote station on an Ethernet and using an Ethernet format address (byte-swapped).

SDLC Secondary Station Addresses

Station addresses are used to identify a secondary station to the network.

Secondary station addresses must be unique within a network. The primary station will communicate with a secondary station by using the secondary station address. The secondary will communicate with the primary by using its own address.

For point-to-point connections, if the secondary station supports the broadcast address X'FF', the primary will learn the remote secondary address. The secondary can specify any value between X'01' and X'FE'.

For primaries that do not support the broadcast address, the secondary station must be set to the same value that is defined at the primary. The value must be between X'01' and X'FE'.

Note: Most stations will support the broadcast address, so the secondary station address at the primary should use X'FF'.

Negotiable stations have local secondary station addresses configured between X'01' and X'FE'. The secondary address of the station negotiated to secondary will be used.

For secondaries on a multipoint connection, the address must match the value specified at the node providing the multipoint primary server function. The address will be in the X'01' to X'FE' range.

X.25 Addresses

X.25 addresses are used to identify resources communicating on X.25 networks. X.25 networks implement the CCITT recommendation defining the interface between data terminal equipment and packet-switching networks. X.25 addresses must be unique within a network. You can get these addresses from your X.25 network provider.

Internet Addresses

Internet Protocol (IP) addresses are used to route data through the network. Every TCP/IP host is assigned at least one unique IP address. The IP address assigned to the host does not define a host on the network; it defines a network interface on that host to a network.

A Communications Server node must have a unique IP address for each network interface routing TCP/IP data through the node. For example, a Communications Server node that is routing TCP/IP traffic over an SNA network (using the Sockets over SNA gateway function) needs unique IP addresses for both the SNA network and the TCP/IP network. The IP address of the TCP/IP interface identifies the Sockets over SNA Gateway connection to the IP network, while the IP address identifies the connection to the SNA network, which looks like a "virtual" TCP/IP network to the system.

An IP address consists of a 2-part, 32-bit address field:

- The first part of the address field contains the network address; the second part contains the host address.

- The number of bits used for the network and host portions of an IP address are variable, depending on the address class of the IP address.
- A network mask enables you to use part of the host portion of the IP address as a subnetwork address.

The class of the IP address is determined by reading the first 3 (upper) bits of the address. As shown in Table 33, Communications Server supports address classes A, B, and C. For more information, refer to *Guide to Sockets over SNA*.

Table 33. IP Address Classes Supported by Communications Server

| Address Class | Default Network Mask | Distribution of Bits for Network and Host Addresses | For a dotted decimal IP address in the form <i>a.b.c.d</i> , the range of values for <i>a</i> is: |
|---------------|----------------------|---|---|
| A | 255.0.0.0 | 8-bit network address; 24-bit host address | 1–127 For example, 9.0.0.0 is a Class A IP network address. |
| B | 255.255.0.0 | 16-bit network address; 16-bit host address | 128–191 For example, 132.11.0.0 is a Class B IP network address. |
| C | 255.255.255.0 | 24-bit network address; 8-bit host address | 192–223 For example, 220.11.44.0 is a Class C IP network address. |

Step 5. Selecting Configuration and Installation Tools

Communications Server uses the configuration tools explained in the following sections. Refer to *Quick Beginnings* for more information about these tools.

Node Configuration Application

The **SNA Node Configuration** application is a graphic window application that enables you to manage SNA configuration information. The application uses a tree view to organize the SNA configuration data to show the relationships between definitions. The user is given task assistance when building a configuration through an integrated task list, the online *Tutorial*, and context sensitive help. The application is responsible for building the configuration files for the user and verifying the data provided.

Remote Configuration

SNA Node Configuration can also be used to connect to a remote Communications Server and directly configure its resources. The user is able to remotely manage the configuration for a Communications Server anywhere in the network.

A remote administration client installs only the administrative applications on a client, including **SNA Node Configuration**. From this client, a user is able to fully administer and configure any Communications Server in the network.

Windows 95, Windows 98, Windows NT, and Windows 2000 remote administration clients use Windows NT or Windows 2000 domain security to authenticate the

Deploying Communications Server Configurations

client connection to the server without reentering the userid and password. The client must be part of a Windows NT or Windows 2000 domain, either by participation in a Communications Server domain or logging in locally with a synchronized userid and password.

Remote administration client users outside of the Windows NT domain are required to provide the userid and password, either through a prompt or by storing these values in the client configuration file.

The authorized users for remote administration clients are maintained in the IBMCSADMIN local group, which is located either directly on the Communications Server or on the domain controller where Communications Server participates. This user group is created during installation and can be administered using the Windows NT or Windows 2000 User Manager application. Remote administration client users must be given user rights in the IBMCSADMIN group to log on locally to the server.

ASCII Configuration File

Communications Server **SNA Node Configuration** stores its configuration data into a human-readable ASCII configuration file. This enables the user to modify configuration files without using **SNA Node Configuration**. (Refer to the *Configuration File Reference* for more information about this file and its syntax.) Using this file, a network administrator can quickly make changes to the configuration using automated tasks such as scripting or software distribution services such as Tivoli TME10 or Microsoft System Management Server.

When creating configurations for a large number of servers to implement, the network administrator can create a template configuration file that represents the common configuration elements for all servers. Using a response file with only those changes necessary for each server, the administrator can distribute the template and response file and merge the two to create the target configuration. For detailed information on how to use template files and response files for configuration and installation, refer to “Configuration with Template and Response Files”.

Web Administration

Through Web administration, a user is able to modify a Communications Server configuration file by loading the file into an edit window. The changes are sent down to the server, verified, and saved for immediate use. The user can either stop and restart the server using the changes made, or apply the configuration changes to a running system.

Configuration with Template and Response Files

The Communications Server template and response files enable you to create or modify a configuration using an editor. You can configure all of the Communications Server configuration keywords and parameters with response files. Both the template and response files have the same format as Communication Server configuration (.ACG) files.

Note: The format of the .ACG files is documented in the OCDNTS50.DAT file contained in the Communications Server installation directory (for example, C:\IBMCS). Refer to the *Configuration File Reference* for more information on the keywords and parameters used in the .ACG files.

Deploying Communications Server Configurations

Template files can ease the mass distribution of configurations to remote servers. A template file can specify the keywords which are common to several servers. For example, if you have multiple servers to configure as SNA gateways with implicit client support, many of the keywords will be identical. You can create a template configuration file that reflects those common keywords.

You can use response files to add, modify, or delete keywords in a template file. The original template configuration file is left unchanged. A response file is merged into a template file by specifying the `INCLUDE` keyword at the end of the template file. For example, if a response file is named *myconfig.rsp*, the last line of the template file that will use the response file is `INCLUDE = myconfig.rsp`. When the template file and the response file are merged, you can give the resulting configuration file a name with the `.ACG` extension that distinguishes it from other `.ACG` files.

When you create configurations using template and response files, the verification utility searches directories in the following order:

1. The local directory where the verification utility is being invoked
2. The `PRIVATE` subdirectory in the Communications Server installation directory.

To ensure that the verification utility can locate the template and response files, you should store them in the `PRIVATE` subdirectory. The `PRIVATE` subdirectory is also where the configuration (`.ACG`) files are stored.

Key Fields

The key field is the parameter in a keyword that names the keyword and uniquely identifies it from other keywords of the same type. The `@KEY_NAME` parameter specifies the key field for the keyword.

The key field is always the first parameter in a keyword that has a key field (for example, `MODE_NAME` in the `MODE` keyword).

Some keywords do not have key fields because they can only be specified once in a configuration file. An example of a keyword that can only be specified once is the `TN3270E_DEF` keyword.

Adding Keywords to a Template File

When using the response file to add a new keyword definition, the entire keyword must be provided. The key field must be provided along with a unique value. If any subfields are omitted from the keyword, the defaults for those fields are used. For example, to add a `MODE` keyword to the configuration, the response file might contain the following keyword:

```
MODE=(  
  MODE_NAME=MYMODE  
  COS_NAME=#INTER  
  CRYPTOGRAPHY=NONE  
  DEFAULT_RU_SIZE=1  
  MAX_NEGOTIABLE_SESSION_LIMIT=8192  
  MAX_RU_SIZE_UPPER_BOUND=4096  
  MIN_CONWINNERS_SOURCE=4096  
)
```

The content of the response file assumes that a `MODE` keyword with the parameter of `MODE_NAME=MYMODE` does not exist in the template. If it does, the parameters would have been updated with the values provided in the response file.

Deploying Communications Server Configurations

If the `MODE_NAME` parameter was omitted from the response file, an error would occur during the configuration verification because the `MODE_NAME` parameter could not be uniquely identified. Not all parameters available for the `MODE` keyword were specified in the response file. The remaining parameters use the defaults assigned by the `OCDSNT50.DAT` file. The resulting addition to the configuration would look like this:

```
MODE=(
  MODE_NAME=MYMODE
  AUTO_ACT=0
  COMPRESSION=PROHIBITED
  COS_NAME=#INTER
  CRYPTOGRAPHY=NONE
  DEFAULT_RU_SIZE=1
  MAX_NEGOTIABLE_SESSION_LIMIT=8192
  MAX_RU_SIZE_UPPER_BOUND=4096
  MIN_CONWINNERS_SOURCE=4096
  PLU_MODE_SESSION_LIMIT=8192
  RECEIVE_PACING_WINDOW=20
)
```

Modifying a Keyword in a Template File

When using the response file to modify an existing keyword definition, the original keyword should exist in the template file. If it does not exist in the template file, the response file adds an entry to the new configuration. The key parameter must be specified in the response file to identify the target keyword. Only those parameters specified in the response file keyword are updated in the template file's keyword. Parameters not specified in the response file are left unchanged. For example, if the following `MODE` keyword is in the template file:

```
MODE=(
  MODE_NAME=#INTER
  AUTO_ACT=0
  COMPRESSION=PROHIBITED
  COS_NAME=#INTER
  CRYPTOGRAPHY=NONE
  DEFAULT_RU_SIZE=1
  MAX_NEGOTIABLE_SESSION_LIMIT=8192
  MAX_RU_SIZE_UPPER_BOUND=4096
  MIN_CONWINNERS_SOURCE=4096
  PLU_MODE_SESSION_LIMIT=8192
  RECEIVE_PACING_WINDOW=20 )
```

and the following keyword is specified in the response file:

```
MODE=(
  MODE_NAME=#INTER
  AUTO_ACT=10
)
```

the resulting configuration would have the following `MODE` keyword definition:

```
MODE=(
  MODE_NAME=#INTER
  AUTO_ACT=10
  COMPRESSION=PROHIBITED
  COS_NAME=#INTER
  CRYPTOGRAPHY=NONE
  DEFAULT_RU_SIZE=1
  MAX_NEGOTIABLE_SESSION_LIMIT=8192
  MAX_RU_SIZE_UPPER_BOUND=4096
  MIN_CONWINNERS_SOURCE=4096
  PLU_MODE_SESSION_LIMIT=8192
  RECEIVE_PACING_WINDOW=20
)
```

Deleting a Keyword from a Template File

When using the response file to delete a keyword from the template, the key parameter and value that identify the keyword must be specified, along with the keyword DELETE. For example, if the template file specifies the following keyword:

```
MODE=(
  MODE_NAME=#INTER
  AUTO_ACT=0
  COMPRESSION=PROHIBITED
  COS_NAME=#INTER
  CRYPTOGRAPHY=NONE
  DEFAULT_RU_SIZE=1
  MAX_NEGOTIABLE_SESSION_LIMIT=8192
  MAX_RU_SIZE_UPPER_BOUND=4096
  MIN_CONWINNERS_SOURCE=4096
  PLU_MODE_SESSION_LIMIT=8192
  RECEIVE_PACING_WINDOW=20
)
```

and the response file contains the following keyword:

```
MODE=(
  MODE_NAME=#INTER
  DELETE
)
```

the resulting configuration does not contain the #INTER mode definition.

The **DELETE** keyword can appear after a *parameter=value* specification or on a line by itself, either preceding or following the parameter. For example, the following uses of the **DELETE** keyword are valid:

```
MODE=(
  MODE_NAME=#INTER
  DELETE
)
MODE=(
  DELETE
  MODE_NAME=#INTER
)
MODE=(
  MODE_NAME=#INTER DELETE
)
```

The **DELETE** keyword can **not** appear in front of a *parameter=value* specification on the same line. For example, the following use of the **DELETE** keyword is not valid:

```
MODE=(
  DELETE MODE_NAME=#INTER
)
```

To delete all keywords of a particular type, or to delete one keyword that does not have a key field, only the keyword and the **DELETE** keyword are necessary. For example,

```
TN3270E_DEF=(
  DELETE
)
```

Step 6. Choosing Communications Server Features

The supported features for Communications Server are described below. Some or all of these functions might be supported for your connection type:

Deploying Communications Server Configurations

SNA Gateway Configuration

By configuring an SNA gateway, you can exchange data between downstream client workstations and a Systems Network Architecture (SNA) mainframe host computer. The gateway performs the function of a concentrator, making LU resources from one or more hosts and host PUs available to downstream SNA client workstations, over any communications media supported by Communications Server.

TN3270E Server Configuration

By configuring a TN3270E server, you can exchange data between downstream TCP/IP client workstations running TN3270E (or TN3270) applications and a Systems Network Architecture (SNA) mainframe host computer. The server performs the function of a concentrator, making LU resources from one or more hosts and host PUs available to downstream TN3270E (or TN3270) sessions, over any communications media supported by Communications Server.

TN5250 Server Configuration

By configuring a TN5250 server, you can exchange data between downstream client workstations on a TCP/IP network running TN5250 applications and an iSeries on an SNA network. The server performs the function of a gateway, making iSeries display sessions available to downstream TN5250 clients, over any communications media supported by Communications Server.

APPN Network Node Configuration

Network nodes are responsible for maintaining and routing information through an APPN network. They maintain a directory of all resources (LUs) available on directly attached end nodes, as well as maintain a topology of all network nodes in the APPN network. When an independent LU 6.2 session is requested by a node in the network, the network node is responsible for locating the remote resource requested by the session and establishing and routing the session between the two nodes.

DLUR/DLUS Configuration

Communications Server supports dependent LU sessions over APPN networks. This requires Communications Server to function as a dependent LU requester (DLUR) that requests routing services from a dependent LU server (DLUS), which resides on a host.

Communications Server supports DLUR for both local sessions and devices, as well as downstream sessions and devices. In either case, the local node must be configured to connect into an APPN network.

AnyNet SNA over TCP/IP Gateway Configuration

The SNA over TCP/IP gateway routes SNA traffic (independent LU 6.2) between TCP/IP and SNA networks. The SNA over TCP/IP gateway, for example, enables APPC or CPI-C applications running on another AnyNet SNA over TCP/IP access node to communicate with SNA peer nodes on an SNA network.

The AnyNet SNA over TCP/IP Gateway must be configured as an APPN network node. This will enable APPC sessions to be routed to appropriate SNA peer nodes.

AnyNet Sockets over SNA Configuration

There are two types of AnyNet Sockets over SNA nodes: access nodes and gateway nodes. Access nodes enable TCP/IP sockets applications to run locally and communicate with other access nodes or gateways somewhere

Deploying Communications Server Configurations

in the SNA network. The AnyNet Sockets over SNA gateway serves as a protocol converter, enabling applications on a TCP/IP network to communicate with socket applications running on an access node in the SNA network. Using two gateways, separate TCP/IP networks can be connected across an SNA network.

SNA API Client Configuration

SNA API clients are separate machines from the Communications Server which have the ability to run SNA applications while using the SNA code on the Communications Server node to actually manage the session. Communications Server simply establishes and manages a session as if the origin were the local node, but forwards the application data to the SNA API client instead of to its own application APIs.

SNA API clients give you the ability to run SNA applications without having to install an SNA communications stack, such as the Communications Server for Windows, on the same machine. Smaller, less powerful machines can thus be used to run the SNA applications while a centralized, more powerful machine can be dedicated as the SNA server for these SNA API clients.

SNA API clients support two types of applications: APPC (independent LU 6.2) applications and LUA API applications, such as 3270 emulators.

CPI-C or APPC Configuration

Communications Server supports CPI-C or APPC applications (for example, APPC3270 or 5250 emulation to an iSeries system). APPC networking is very flexible and versatile and, therefore, it can be complicated. The complexity of this configuration depends on how you will use Communications Server and how much work your application will do for you.

Dependent LU 6.2 Sessions to a Host

Communications Server supports dependent LU 6.2 sessions with a host PU.

3270 Configuration

Communications Server supports 3270 sessions with a host.

Focal Point Configuration

Communications Server supports remote focal point locations for routing alert information for specific management services applications. The local node will locate remote focal points through an APPN network.

iSeries Shared Folders Configuration

Communications Server enables you to create disk devices on the server that communicate with iSeries folders through the iSeries Integrated File System (IFS). If the server shares these disk devices, clients can NET USE to them. Multiple clients can connect to folders on the iSeries system as if they were drives on their workstations.

Step 6.1 Planning for Application Programs

Review your applications to be sure you comply with any requirements they have. Specifically, check requirements for:

- Modes
- Transaction programs (TPs)
- LU definitions (host, LU 6.2)
- Partner LUs

Step 7. Creating Configurations and Installing

You need to determine how Communications Server will be configured and installed on your user's servers. You can:

- Give your users instructions for installing and configuring Communications Server. You can either prepare your own instructions or print copies of the *Quick Beginnings*. You should also provide the specific information for each server, such as the network ID, control point name, and the network address.
- You should give your users instructions for installing Communications Server with an existing configuration or template file. You must also provide a response file tailored for each server, as well as Communications Server CDs or access to CD images.
- Configure and install Communications Server on user servers yourself. Unless you have a very small group of users, you might want to consider using a product that aids configuration, installation, and distribution of software for multiple workstations. Examples of this type of product include IBM NetView Distribution Manager/2 (NVDM/2), Tivoli TME-10, and Microsoft System Management Server (SMS).

Using Response File Configuration and Installation

If you choose to use response file configuration and installation, refer to "Configuration with Template and Response Files" on page 164.

Step 8. Creating User Materials

After you have planned for your network and determined how you are going to implement your plan, you need to create user materials. This means that you should prepare documentation for installation, configuration, and daily use, and prepare backup procedures.

Preparing Documentation

Prepare a set of customized documentation to assist your users in installing, configuring, and using Communications Server and local applications for their particular needs. The following sections contain suggestions for the kind of information you should include.

Configuration and Installation Materials and Procedures

Quick Beginnings and online installation helps are available to assist your users in installing Communications Server. Your instructions should tell the user which of the following steps to perform when installing software:

- When to back up configuration files
- Whether to accept the defaults to create a Communications Server configuration file

You might need to provide network information, such as LAN adapter addresses, network names, and so forth.

- When to install with a response file that you have provided
- When to install with a customized configuration file that you have provided

If necessary, provide your users with the appropriate documentation.

Communications Server Materials and Procedures

The following materials are recommended for using Communications Server functions and APIs:

Deploying Communications Server Configurations

- *Quick Beginnings*
- User-written programs for the chosen Communications Server APIs

Starting and Stopping Communications Server

For the procedure to start or stop Communications Server, refer to *Quick Beginnings*.

If you have different instructions for starting or stopping, supply these instructions to your users.

- Host help telephone number for problem determination (contact your host personnel)
- Logon procedures for host systems or applications
- Procedure for stopping the host application
Contact your host personnel.
- Logoff procedures
Your instructions should include any special requirements for logging off systems or applications. This information can be obtained from your host personnel.

Materials and Procedures for Application Programs

You should provide your users with procedures and other information for the application programs you might use on your network. Generally, you should provide:

- Procedures for starting application programs
- Procedures for running application programs
- Procedures for responding to messages generated by the application program
- Procedures for problem determination
- If necessary, procedures for stopping the application program

Contact the application programmer for the preceding information.

Problem Determination Materials and Procedures

For the procedures on problem determination and reporting, refer to *Quick Beginnings*.

Preparing Backup Procedures

Occasionally, your users might erase or change the configuration files, registry, file system, application programs, and other locally created programs or files. Also, your servers might experience erasures or unacceptable changes as well, especially when many users are accessing your server workstations throughout the workday.

Because of this, you need to prepare and document the backup procedures for your network. You might also want to change the attributes of selected files on your servers to Read-Only so your users cannot change them.

Step 9. Maintaining Your Network

When you have completed planning for, installing, and configuring your network, and it is running on a daily basis, your remaining task is maintenance. You must plan for adding, changing, or deleting resources and users in your network, and plan for problem-solving.

Deploying Communications Server Configurations

You need to perform the same level of planning and implementation steps for the changes to your network that you performed for initial setup. You should follow the same steps for these changes to your network that you used for the initial planning, installation, and configuration.

Communications Server provides these tools to help you monitor the day-to-day performance of your network:

- **SNA Node Operations**
- Command line utilities
- Web Administrator
- Remote Administration Client
- System Management API
- SNMP subagent
- NT Performance Monitor

The following products can also assist you in daily management of your network:

- IBM NetView
- IBM LAN Network Manager
- Tivoli TME-10

Part 3. Controlling and Maintaining a Network

Chapter 16. System Management Facilities

This chapter provides an overview of the Communications Server system management facilities and describes the capabilities that are common to each facility. Each facility also has some functions that are unique, and those functions are described for each facility.

The system management facilities enable you to monitor and control the communications resources of Communications Server. They also enable you to adjust these resources to improve the efficiency of SNA communication services or to monitor and test these services during problem determination.

Refer to *Quick Beginnings* for a list of resources that you can manage.

You can manage Communications Server with the following facilities:

- **SNA Node Operations**
- Remote administration
- Command-line utilities
- Web Administration
- SNA Node Operations ActiveX control
- Tivoli Plus Module
- APPN MIB support
- zSeries Remote Operations Support (ROPS)

You can use any combination of these facilities for Communications Server management.

Common Capabilities of System Management Facilities

Each system management facility can perform the following actions for resources:

- Display resource information
- Start resources
- Stop resources
- Delete resources
- Initiate path switches

You can perform a subset of these actions for each resource, as shown in Table 34.

Table 34. Management Actions for Communications Server Resources

| Resource | Display | Start | Stop | Delete | Path switch |
|------------------------|---------|-------|------|--------|-------------|
| Connections | X | X | X | X | |
| CPI-C side information | X | | | X | |
| Devices | X | X | X | X | |
| DLUR PUs | X | X | X | | |
| Local LU 0 to 3 | X | | | X | |
| Local LU 6.2 | X | | | X | |

System Management Facilities

Table 34. Management Actions for Communications Server Resources (continued)

| Resource | Display | Start | Stop | Delete | Path switch |
|---------------------|---------|-------|------|--------|-------------|
| LU 6.2 Sessions | X | | X | | |
| PU 2.0/DLUR gateway | X | | | X | |
| RTP connections | X | | | | X |
| TN3270E sessions | X | | X | | |
| TN5250 sessions | X | | X | | |

SNA Node Operations

SNA Node Operations is an online facility for monitoring and controlling communications resources maintained by Communications Server. **SNA Node Operations** allows you to modify, delete, and monitor resources in your network.

You can use **SNA Node Operations** to help manage the resources in your network. The systems management programming capabilities of Communications Server enable you to configure and manage the nodes within your SNA network. **SNA Node Operations** provides a tree-view diagram of your configuration that enables you to select and modify resources in a hierarchical graphic interface.

You can manage multiple windows within the **SNA Node Operations** main window.

You can define more than one Communications Server configuration and switch between configurations as your needs require. When you start Communications Server, the default configuration is offered as the first choice for node configuration or node operation unless you specify another configuration. Once you have completed a new configuration, you can replace the default configuration with your new one. Use **SNA Node Operations** to apply your new configuration and ensure that the relevant part of the product runs as you have configured it.

SNA Node Operations enables you to launch other Communications Server programs, such as **Log Viewer**, **Trace Facility**, or **SNA Node Configuration**.

Remote Administration

SNA Node Operations can also be used to connect to a remote Communications Server and directly administer its resources. The user is able to remotely manage the resources for a Communications Server anywhere in the network.

A remote administration client installs only the administrative applications on a client, including **SNA Node Operations**. From this client, a user is able to fully administer any Communications Server in the network.

Windows 95, Windows 98, Windows NT, and Windows 2000 remote administration clients use Windows NT or Windows 2000 domain security to authenticate the client connection to the server without reentering the userid and password. The client must be part of a Windows NT domain, either by participation in a Communications Server domain or logging in locally with a synchronized userid and password.

Remote administration client users outside of the Windows NT domain are required to provide the userid and password, either through a prompt or by storing these values in the client configuration file.

The authorized users for remote administration clients are maintained in the IBMCSADMIN local group, which is located either directly on the IBM Communications Server or on the domain controller where Communications Server participates. This user group is created during installation and can be administered using the Windows NT or Windows 2000 User Manager application. Remote administration client users must be given user rights in the IBMCSADMIN group to log on locally to the server.

Command-Line Programs

Command line programs enable you to issue commands from a Windows prompt window. The output of commands is displayed on the screen, but the output can be redirected to a file.

Refer to *Quick Beginnings* for a list of command line programs you can use.

Web Administration

The Web Administration function of Communications Server is used to manage your server using a Web browser. This requires a Web server running on the same machine as Communications Server. The Web server, if installed prior to installing Communications Server, will be configured with the necessary virtual directories by the installation program. If you plan to use Web Administration, you should install the function when you install Communications Server by selecting "Web Administration" during the installation process. Refer to *Quick Beginnings* for information about security requirements for specific Web servers.

Most of the **SNA Node Operations** functions are available through a Web browser. You can query node status, start or stop the node, obtain information about resources, modify the resource state, display and edit configuration files, display the message log, and perform other administrative tasks.

SNA Node Operations ActiveX Control

There is an ActiveX control version of **SNA Node Operations**. This control allows you to embed **SNA Node Operations** into any application that allows embedded controls such as the Microsoft Management Console (MMC). The name of the control is CSNCTX and it is located in the product installation directory. This control has all of the **SNA Node Operations** functions and provides the benefit of embedding **SNA Node Operations** into other applications, enabling cross product management.

Tivoli Plus Module

If you use the Tivoli Management Environment (TME) to centrally manage the devices and applications in your network, the Communications Server Tivoli Plus module enables you to manage Communications Server from TME also. The Communications Server Tivoli Plus module can be installed on the TME server. A TME installable image is in the TME directory of the Communications Server CD-ROM. To access installed TME Plus functions, double click on the TivoliPlus

System Management Facilities

icon that was added to the TME **Desktop for Administrator** panel. When you double click on the IBMCSNTPlus icon, you can use the icons in the panel to perform the following actions:

- Distribute, install, and uninstall Communications Server
- Distribute, install, and uninstall the SNA API client
- Start, stop, and query the server
- Display and modify server resources
- List configuration files
- Distribute configuration files
- Route error messages from Communications Server to a Tivoli Enterprise Console
- Establish monitors and thresholds for key Communications Server attributes

When you double click on any icon in the IBMCSNTPlus panel, enter information in the fields of the dialog box presented. Help is available for the fields in the dialog box.

APPN MIB Support

Communications Server supports simple network management protocol (SNMP) requests for APPN management information from any SNMP management system.

The SNMP Service is installed from the **Services** tab under the **Network** pulldown of the Windows NT control panel. After the SNMP Service software is installed on your computer, you must configure it with valid information for SNMP to operate.

You must be logged on as a member of the Administrator group for the local computer to configure SNMP.

The SNMP configuration information identifies communities and trap destinations.

zSeries Remote Operations Support (ROPS) for the NetView Program

When RUNCMD is issued from NetView, the value for the APPL parameter must be specified as APPL=CSRCMDS.

When the RUNCMD command is issued, all information, including the ID and password, is transmitted in the clear and is not secured in any way, unless the NetView suppression character is used when the RUNCMD command is issued. This information is displayed on the screen and entered in the NetView log. If passwords are transmitted to the ROPS service, the user is responsible for securing the NetView terminal and its NetView logs so that sensitive data is not compromised.

Part 4. Appendixes

Appendix A. Planning for Communications Server

This appendix provides the information you need to know for planning your Communications Server environment.

For a list of compatible hardware and software, refer to the following HTTP Web page: www.ibm.com/software/network/commserver/about/comp_products/csnt.html

Modems

Communications Server supports the following types of modems:

- Asynchronous modems:
 - Modems 100% compatible with the Hayes AT command set
 - IBM modems that conform to the ITU-T (CCITT) V.24/V.28 (EIA RS-232-D) and V.35 interface standards
 - Non-IBM asynchronous modems are also supported, as well as the base function of ITU-T (CCITT) V.24 (EIA RS-232-D) standard data terminal equipment (DTE) to data circuit-terminating equipment (DCE) interface for asynchronous communications
 - Modems that conform to the RS-232-D standard DTE to DCE interface
- Synchronous modems and CSU/DSU units that support the V.24bis command set
- Modems 100% compatible with the Hayes AutoSync protocol

X.25 users will require either a synchronous modem and CSU/DSU units that support the V.24bis command set or a modem compatible with the Hayes AutoSync protocol. Refer to your network supplier for the recommended connection type and modem.

Session-Level Encryption Adapters

To use session-level encryption, you must have the Programmed Cryptographic Facility program and one of the following adapters:

- The IBM SecureWay 4758 PCI Cryptographic Coprocessor Adapter
- Other adapters compliant with the IBM Common Cryptographic Architecture

Note: If you use an adapter other than the IBM 4758, you will need to use cryptographic API verbs instead. Refer to the *Communications Server Programming Guide and Reference* for more information.

Data Compression Considerations

Prior to using SNA data compression several configuration steps have to be worked out on the involved systems. SNA data compression is supported on:

- VTAM V3R4.1 or higher
- OS/400 V2R3 or higher

When compression is done, each RU is compressed independently and a compression header is added to the RU. The TH and RH are not compressed. No compression is done when the RU is smaller than 20 bytes, or greater than 64K

Planning for Communications Server

bytes, or, in the case of RLE only, if the size of the compressed data, including the compression header, is greater than the original data.

Before you start to configure the appropriate features of Communications Server for data compression, ensure that your partner system is able and prepared to support data compression.

Planning the Host/VTAM Connection

Before you can request data compression on a session to an S/370 or zSeries® host with 3270 emulation or LU 6.2 type communications, VTAM Version 3 Release 4.1 (VTAM V3R4.1) must be installed and configured to negotiate such requests.

Your host VTAM system programmer should supply you with the proper VTAM definitions required for data compression. Here is a summary list of the definitions required. Examples of these are given following the list.

- The VTAM startup options must contain the **CMPVTAM=x** parameter
- You need a LOGMODE entry with the parameter **COMPRES=**
- This LOGMODE entry must be in the Mode table that you are using
- There must be a VTAM APPL statement with the parameters **CMPAPPLI=** and **CMPAPPLO=** for the application that you will be using
- The VTAM LU definition for your session must specify the LOGMODE and Mode table that were defined for compression.

CMPVTAM

This is the VTAM startup options parameter indicating the maximum compression level allowed for any session connected to this host. The syntax of this parameter is:

```
CMPVTAM=n
```

where n is the level of compression allowed for a request, ranging from 0 to 4. The recommended value is 4, which allows all levels of data compression to be used.

COMPRES

This is the VTAM MODEENT table entry in the LOGMODE definition that will be used by your LU. The syntax is:

```
COMPRES=value
```

where value stands for SYSTEM, REQUESTED or PROHIBITED. SYSTEM should be coded for all supported LU types.

CMPAPPLO

This is the VTAM application (APPL) definition parameter that sets the requested compression level for outbound Request/Response Unit (RU) data. The syntax is:

```
CMPAPPLO=n
```

where "n" is the requested data compression level for the PLU -> SLU direction of data flow, also called outbound data. For all LU type sessions supported by Communications Server, the value of "n" can be equal to 0, 1, or 2, representing the data compression level 0, level 1, and level 2. However, for LU types 0, 1, 2, and 3, Communications Server always reserves compression resources at level 2.

CMPAPPLI

This is the VTAM application (APPL) definition parameter that sets the requested compression level for inbound Request/Response Unit (RU) data. The syntax is:

CMPAPPLI=m

where "m" is the requested data compression level for the SLU -> PLU data flow, also called inbound data. The value for m depends on the LU type used for the session:

- For 3270 LU type sessions, "m" can have the value equal to 0 or 1.
- On LU 6.2 type sessions, "m" can have the values 0, 1 or 2.

Note: VTAM always represents the primary LU (PLU) and Communications Server represents the secondary LU (SLU) for LU types 0, 1, 2 and 3.

Following are coded examples, which were extracted from the original VTAM listings, pointing out the usage and placement of the parameters discussed above. Pay special attention to the highlighted lines.

```
*****
SSCPID=20,HOSTSA=20,XNETALS=YES,                                X
GWSSCP=YES,                                                    X
NODELST=NODES1,                                              X
CMPVTAM=4,                                                  X
CONFIG=K0,SUPP=NOSUP,                                         X
NETID=USIBMRA,HOSTPU=ISTPUS20,HOSTSA=20,SSCPNAME=RAK,        X
SSCPDYN=YES,SSCPORD=PRIORITY,                                X
ASYDE=TERM,                                                  X
NOTRACE,TYPE=VTAM,IOINT=0,                                    X
NOTRACE,TYPE=SMS,ID=VTAMBUF,                                  X
PPOLOG=YES,                                                  X
NODETYPE=NN,                                                 X
CPCP=YES,                                                    X
CSALIMIT=0,                                                  X
NOTNSTAT,DYNLU=YES,                                         X
IOBUF=(3500,256,3,,1,58),                                    X
LPBUF=(1100,,2,,1,4),                                        X
LFBUF=(100,,,1,1),                                          X
CRPLBUF=(2400,,,1,4),                                       X
SFBUF=(60,,,1,1)                                           X
*****
```

Figure 37. VTAM Startup Parameter. Compression level up to 4 is allowed in this VTAM.

Planning for Communications Server

```
*****
* LOGMODE FOR COMPRESSION TEST - BASED ON D4C32XX3 *
*****
*
*          3274 MODEL 1C (REMOTE SNA)                @0Y02946*
*          PRIMARY SCREEN 24 X 80 (1920)             *
*          ALTERNATE SCREEN TO BE DETERMINED BY APPLICATION *
*
*****
D4C3COMP MODEENT LOGMODE=D4C3COMP, *
                FMPROF=X'03', *
                TSPROF=X'03', *
                PRIPROT=X'B1', *
                SECPROT=X'90', *
                COMPROT=X'3080', *
                RUSIZES=X'87F8', *
                PSERVIC=X'0280000000000000000000300', *
                APPNCOS=#CONNECT, *
                COMPRES=REQD *
```

Figure 38. VTAM Logmode Table. The MODEENT macro configured with data compression requested.

```
VBUILD TYPE=APPL
*****
* VTAM APPL STATEMENTS FOR CICS/MVS* 3.3 *
*****
RAKAC001 APPL ACBNAME=RAKAC001, *
          MODETAB=MTAPPC, *
          EAS=20, *
          SONSCIP=YES, *
          AUTH=(ACQ,VPACE,PASS), *
          PARSESS=YES, *
          VPACING=5, *
          CMPAPPLI=1 *
          CMPAPPL0=2
```

Figure 39. VTAM Application Statement. It controls the maximum supported levels for this session.

```
WTCC1102 PU ADDR=13,
          IDBLK=05D,
          IDNUM=32289,
          etc.
          DLOGMOD=D4C3COMP
          MODETAB=AMODETAB
          etc.

RACC1102 LU LOCADDR=2,DLOGMOD=D4C3COMP,MODETAB=AMODETAB
          LU
          etc.
```

Figure 40. VTAM PU and LU Definitions. The LU RACC1102 uses the logmode D4C3COMP, which is in the AMODETAB modetable.

Planning for iSeries Data Compression

When connected to an iSeries system, you need OS/400 V2R3 to be able to use data compression. All that is really required is to ensure that you have a MODE definition that specifies compression support. Three parameters are required on the "Create Mode Description" profile obtained through either the CRTMODD command or by going through the menus. The three parameters are:

DTACPR

This parameter controls the use of data compression. Valid parameters are shown below:

NETATR

Use the system-wide network value for this parameter.

NONE

No data compression allowed.

ALLOW

If the remote system requests data compression, the local system will allow the session to be established.

REQUEST

The local system requests data compression.

REQUIRE

The use of data compression is required. If one of the systems is not able to perform data compression at the requested level, the session will not be established.

Whenever data compression is allowed and requested, both systems negotiate the level of data compression to be used on that session according to settings of the **INDTACPR** and **OUTDTACPR** parameters.

INDTACPR

This value represents the maximum level of compression used for inbound data.

OUTDTACPR

This value represents the maximum level of compression used for outbound data.

For LU 6.2 sessions, data compression levels 0, 1, and 2 are supported in both directions, that is, for both outbound and inbound data flow.

For information about enabling Communications Server for data compression, refer to *Quick Beginnings*.

Performance Considerations

The performance of computer communications in general, and of Communications Server in particular, is subject to a large number of variables. The intent of the following information is to introduce some of the factors that influence the performance of communications in a Communications Server environment.

Communications performance is affected by the speed of the communications line. This speed is usually expressed as the number of bits per second (bps) that can be transmitted over the line. Typically, the higher the bps, the greater the communications performance that can be achieved on the line.

In addition to the bit rate, several other line-related factors can affect communications performance. For instance, some lines can only carry data in one direction at a time. To change the direction of communications (who can send), the station raises a *request to send* (RTS) and waits for *clear to send* (CTS) to be issued. This requires additional time for this *line turnaround* (which is modem dependent). Another factor is the transmission quality of the communications line.

Planning for Communications Server

When a line error occurs, most protocols will detect this and resend the data. A line that has fewer transmission errors will give greater throughput.

Greater throughput can also be achieved if the protocols operate in a *duplex* data transmission mode, where the protocol can receive and transmit simultaneously. LAN and X.25 are duplex protocols. SDLC can operate in a duplex or half-duplex mode.

To use duplex for SDLC, you should use an appropriate adapter (the MPA adapter cannot support duplex). The remote station that can also support duplex.

As line speed increases, the bit-rate capacity of the line becomes less of a constraint in performance. Instead, the constraint becomes related to delays inside the devices that are using the communication line (modems, control units, or computer internal constraints). At very high speeds, these delays can impact communication performance more than the bit rate capacity of the line.

Analyzing these situations takes some effort. For example, longer than expected communication delays can be caused by any of the following:

- An overloaded communication line
- An overloaded host system
- An overloaded 3X74 control unit
- An overloaded 37xx network controller
- Less than optimal VTAM or NCP definition
- Inadequate network design or network management
- Operating system memory overcommitment and swapping

In a high-speed communication environment, such as a LAN, the bit rate of the communication line is less significant because the communication capacity is largely determined by the speed of the communication devices. Better communication performance can often be achieved by increasing the processing power of the workstation, host, or other network components. This does not mean that the bit rate of a LAN network is never a constraint; rather, it might require several computers running at their communication capacity before the performance of the LAN significantly degrades.

To summarize, in many cases the performance of Communications Server is largely determined by factors that are related to the communication network being used. Providing communications performance at a reasonable cost involves both an understanding of the network as a whole and the analysis of its individual components.

Appendix B. 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
500 Columbus Avenue
Thornwood, NY 10594
U.S.A.

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

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

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

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

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

Any references in this publication 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.

Licenses of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created

programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Department TL3B/062
P.O. Box 12195
Research Triangle Park, NC 27709-2195
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, International Programming License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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.

Trademarks

The following terms are trademarks of the IBM Corporation in the United States or other countries or both:

| | |
|------------------------------------|----------------------|
| Advanced Peer-to-Peer Networking | iSeries |
| AIX | Micro Channel |
| AnyNet | NetView |
| APPN | OS/2 |
| AT | OS/400 |
| BookManager | Personal System/2 |
| CICS | Portmaster |
| DB2/2 | Presentation Manager |
| eNetwork | PS/2 |
| Enterprise System/9000 | SecureWay |
| ESCON | System/370 |
| ES/9000 | SystemView |
| FFST/2 | S/370 |
| First Failure Support Technology/2 | TalkLink |
| Global Network | VTAM |
| IBM | WebExplorer |
| IBMLink | zSeries |
| IMS | |
| iSeries | |

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

C-bus is a trademark of Corollary, Inc.

Java and HotJava are trademarks of Sun Microsystems, Inc.

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

PC Direct is a trademark of Ziff Communications Company and is used by IBM Corporation under license.

ActionMedia, LANDesk, MMX, Pentium, and ProShare are trademarks or registered trademarks of Intel Corporation in the U.S. and other countries.

For a complete list of Intel trademarks see www.intel.com/tradmarx.htm.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Index

Numerics

3174 Controller
 Communications Server host
 parameters 137
 VTAM PU host parameters 137
9370
 Communications Server host
 parameters 136
 SDLC host parameters 142

A

access node
 Sockets over SNA 65
activate at startup 33
activate on demand 33
activation parameters 33
adapter addresses, network 161
adapters
 session-level encryption 181
adaptive BIND pacing 58
adaptive session-level pacing 57
addresses
 defining network 161
 network adapter 161
 SDLC station 162
 X.25 162
administration support 18
administration, branch extender 29
ANR 47
AnyNet
 AnyNet examples 71
 defining unique CP names and
 connection network names 67
 mapping SNA resources to IP
 address 65
 routing SNA sessions over AnyNet
 SNA over TCP/IP 70
AnyNet configurations 65
APPC
 connectivities 4
 fundamental concepts 5
 VTAM host parameters 140
 X.25 host parameters 147
application programs, planning for 169
APPN
 connection networks 36
 control point (CP)
 CP-CP sessions 25
 data link control 31
 general discussion 25
 logical links 31
 defined 8
 directory services 38
 general discussion 3, 21
 logical units (LUs) 6
 node types
 end 25
 general discussion 23
 LEN 25

APPN (*continued*)
 node types (*continued*)
 network 24
 topology and route-selection services
 class of service 42
 general discussion 40
 intermediate session routing 44
 modes 42
 route selection 43
 topology database 41
automatic link retry 34
Automatic Network Routing (ANR) 47

B

backup procedures, preparing 171
branch extender 27

C

charges, X.25 network 134
client support
 Novell NetWare for SAA 16
 SNA API 16
CMSETUP 17
Communications Server
 application programming interfaces
 (APIs)
 conversations 7
 configuration overview 3
 features, choosing 167
 NAU names 159
 supported functions 3
Communications Server and SNA
gateway
 general discussion 95
 LAN destination address
 configuration 98
 SNA gateway performance 100
 supported links to hosts 96
 supported links to workstations 95
compression header 182
configuration
 creating 170
 for a template file 163
 host parameters 135
 methods 17
 process 17
 windows, using CMSETUP 17
configuration, local and remote
support 17
configuring a template file
 create configurations and install
 general discussion 170
 response file configuration and
 installation 170
 create naming conventions
 criteria 157
 general discussion 157
 NAU names 159

configuring a template file (*continued*)
 create naming conventions (*continued*)
 types and restrictions 158
 types used by multiple
 components 158
 create user materials
 Communications Server materials
 and procedures 170
 prepare backup procedures 171
 prepare documentation 170
define network addresses
 general discussion 161
 network adapter 161
 SDLC station 162
 X.25 162
define the template configuration
 choose features 167
 plan for application programs 169
maintain your network 171
plan for hardware 155
plan for software applications 157
select configuration and installation
tools
 configuration 163
configuring branch extender 29
connections, X.25 134
considerations, HPR 92
control point (CP)
 CP-CP sessions 25
 data link control 31
 general discussion 25
 logical links
 link activation 31
 parallel links 31
 names 159
conversations, transaction program 7
CP (control point)
 CP-CP sessions 25
 data link control 31
 general discussion 25
 logical links
 link activation 31
 parallel links 31
 names 159
criteria for naming 157

D

data compression 52
data compression, SNA 9
data link control
 definitions 22
 profiles 22
data security 18
DDDLU 48
dedicated LUs
 general discussion 98
definition parameters 33
Dependent Logical Unit Requester 49
dependent LU 51
Dependent LU Requester (DLUR) 51

- Dependent LU Requester (DLUR)
 - (continued)
 - overview 10
 - support 49
- directory services 38
- discovery of service providers 10
- DLUR 49
 - host parameters 144
- documentation, preparing 170
- dynamic definition of dependent LU 48

E

- emulator
 - entry-level functions 20
- encryption 55
- end node
 - capability through APPN 8
 - definition and description 25
- entry point 56
- entry-level emulator 20
- ES/9000 host connection parameters 142
- Ethernet (VTAM) host parameters 138
- explicit terminal and printer
 - definitions 107
- explicit workstations 99

F

- fixed pacing 58
- flow control 57
- focal point 56

G

- gateway
 - Sockets over SNA 65
- gateway support
 - applications on a subarea network 60
 - as a protocol converter 60
 - description 60
 - LAN Network Manager 60
 - NetView commands 60
 - summary of features 61
 - table of features 61
 - transmission of NMVTs 60
- gateway, SNA
 - definition 48
 - overview of 100
 - performance 100
 - support
 - pooled and dedicated LUs 98

H

- hardware, planning for 155
- high performance routing (HPR) 91
 - overview 9, 45
 - pacing 58
 - planning 91
 - support 45
- host connection parameters 135
- HPR considerations 92

I

- implicit terminal and printer
 - definitions 107
- implicit workstations 99
- independent LU 50
- installation charge 134
- installation, response file 170
- installing a template file
 - create configurations and install
 - general discussion 170
 - response file configuration and installation 170
 - create naming conventions
 - criteria 157
 - general discussion 157
 - NAU names 159
 - types and restrictions 158
 - types used by multiple components 158
 - create user materials
 - Communications Server materials and procedures 170
 - prepare backup procedures 171
 - prepare documentation 170
 - define network addresses
 - general discussion 161
 - network adapter 161
 - SDLC station 162
 - X.25 162
 - define the template configuration
 - choose features 167
 - plan for application programs 169
 - maintain your network 171
 - plan for hardware 155
 - plan for software applications 157
 - select configuration and installation tools
 - configuration 163
- installing branch extender 29
- intermediate session routing 44
- internet addresses 162
- IP addresses 162
- iSeries
 - host controller (Communications Server) host parameters 151
 - host controller host parameters 149
 - host controller host parameters (X.25 PVC) 153
 - host controller host parameters (X.25 SVC) 154
 - host device host parameters 150
 - host line (Communications Server) host parameters 150
 - host line host parameters 148
 - host line host parameters (X.25) 152
 - twinaxial host parameters 151
- iSeries data compression 184
- iSeries OLE DB provider
 - description 15
- iSeries shared folder server
 - configuration 90
 - description 15

K

- keep-alive processing
 - TN3270E server 106
 - TN5250 server 111

L

- LAN destination address
 - configuration 98
- LAN Network Manager 60
- LDAP 85
- legacy data access
 - description 15
 - planning for 89
- LEN node 25
- Lightweight Directory Access Protocol (LDAP) 85
- limitations
 - zSeries ROPS support for NetView 178
- link characteristics, HPR 91
- link types 32
- links to hosts 96
- links to workstations 95
- Load balancing 123
- local configuration support 17
- logical links
 - link activation 31
 - parallel 31
- logical units (LUs)
 - general discussion 6
 - names 160
 - support 47
- LU 6.2 51
- LU names
 - sockets over SNA 160
- LU pools 60
- LU-LU sessions 50

M

- maintaining your network 171
- management services
 - focal points, service points, and entry points 56
 - levels of SNA management services architecture 57
 - session flow control 57
- materials and procedures
 - applications programs 171
 - Communications Server 170
 - creating user 170
 - problem determination 171
- maximum activation attempts 34
- migration level for SNA management services 57
- Mode Description 184
- modems, supported 181
- MPC
 - DLC support 19
 - IOCP/HCD parameters 141
 - VTAM parameters 141
- Multi-Path Channel
 - DLC support 19
- multidrop secondary workstations 95

- multiple domain support (MDS)
 - level 57
- multiple PUs
 - partitioning LUs among hosts 59

N

- names
 - NAU 159
 - types and restrictions 158
 - used by multiple components 158
- naming conventions 157
- NCP (NTRI)
 - Communications Server host parameters 135
- NetView
 - RUNCMD support 178
- network addresses, defining 161
- network charges, X.25 134
- network management vector transport (NMVT) level 57
- network management vector transports (NMVTs) 60
- network node
 - capability through APPN 8
 - definition and description 24
- network, maintaining 171
- new for this release xiii
- NMVTs 60
- node types
 - end node 25
 - general discussion 23
 - LEN node 25
 - network node 24
- Notices 187
- Novell NetWare for SAA client support 16
- NPSI host parameters 145

P

- pacing 57
- parameters
 - host connection 135
- performance considerations 185
- periodic rental charge 134
- planning
 - application programs 169
 - for SNA 95
 - hardware 155
 - software applications 157
 - X.25 configuration 133
- PLU 182
- pooled LUs
 - general discussion 98
- port number
 - changing for TN3270E server 105
 - changing for TN5250 server 111
- prepare backup procedures 171
- priority queuing, HPR 92
- programming support 9
- protocol converter 60

R

- Rapid Transport Protocol (RTP) 46
- reassembly, segmenting and 58
- remote configuration support 17
- rental charge 134
- Request Unit (RU) 182
- restrictions, branch extender 29
- RFC 1205 111
- RFC 1576 105
- RFC 1646 105
- RFC 1647 105
- route discovery function 81
- route selection
 - class of service 42
 - general discussion 40
 - intermediate session routing 44
 - modes 42
 - topology database 41
 - VTAM users 44
- RTP 46
- RU compression 181
- RUNCMD support 178

S

- SDDLU 48
- SDDLU support host parameters 144
- SDLC (Synchronous Data Link Control)
 - link station addresses 162
 - switched workstation connections 95
- SDLC host connection
 - 9370 142
 - VTAM/NCP 138
- secure sockets layer based security
 - TN3270E 13
 - TN5250 14
- security
 - conversation 4
 - overview 18
 - session 4
- segmenting and reassembly 58
- self-defining dependent LU 48
- service point 56
- service provider, discovery of 10
- session flow control
 - adaptive BIND pacing 58
 - adaptive session-level pacing 57
 - segmenting and reassembly 58
- session-level encryption
 - adapters 181
- session-level pacing 57
- SLU 182
- SNA
 - API client support 16
 - Communications Server and X.25
 - planning for X.25 configuration 133
 - possible X.25 connections 134
 - configuration process
 - configuration windows (CMSETUP) 17
 - general discussion 17
 - management services 56
 - session flow control 57
 - data compression 9

SNA (*continued*)

- functions
 - APPC API 48
 - data link control definitions 22
 - LU support 47
 - LU-LU sessions 50
 - overview 21
 - SSCP-dependent LU 51
- gateways 60
 - general discussion 95
 - LAN destination address
 - configuration 98
 - supported links to hosts 96
 - supported links to workstations 95
- management services architecture, levels of 57
- SDLC switched workstation connections 95
- session level data confidentiality 55
- transmission priority 43
 - with Communications Server 21
- SNA (Systems Network Architecture)
 - connection networks 36
 - control point (CP)
 - CP-CP sessions 25
 - data link control 31
 - general discussion 25
 - logical links 31
 - defined 8
 - directory services 38
 - general discussion 3, 21
 - logical units (LUs) 6
 - node types
 - end 25
 - general discussion 23
 - LEN 25
 - network 24
 - topology and route-selection services
 - class of service 42
 - general discussion 40
 - intermediate session routing 44
 - modes 42
 - route selection 43
 - topology database 41
- SNA gateway
 - definition 48
 - overview of 100
 - performance 100
 - support
 - pooled and dedicated LUs 98
- SNA gateway and Communications Server
 - general discussion 95
 - LAN destination address
 - configuration 98
 - SNA gateway performance 100
 - supported links to hosts 96
 - supported links to workstations 95
- SNA gateway support
 - description 60
- SNA over TCP/IP configurations 65
- SNA session-level encryption 55
- sockets over SNA
 - LU names 160
 - route discovery function 81

- sockets over SNA (*continued*)
 - using sockets over SNA gateway 12, 65
- Sockets over SNA
 - description 65
- software applications, planning for 157
- SSCP-dependent LU 51
- SSCP-independent LU 50
- subarea network applications 60
- supported links to hosts 96
- supported links to workstations 95
- supported LUs
 - general discussion 47
 - SSCP-dependent 51
 - SSCP-independent 50

T

- template configuration, defining 155
- template file, install and configure
 - create configurations and install
 - general discussion 170
 - response file configuration and installation 170
 - create naming conventions
 - criteria 157
 - general discussion 157
 - NAU names 159
 - types and restrictions 158
 - types used by multiple components 158
 - create user materials
 - Communications Server materials and procedures 170
 - prepare backup procedures 171
 - prepare documentation 170
 - define network addresses
 - general discussion 161
 - network adapter 161
 - SDLC station 162
 - X.25 162
 - define the template configuration
 - choose features 167
 - plan for application programs 169
 - maintain your network 171
 - plan for hardware 155
 - plan for software applications 157
 - select configuration and installation tools
 - configuration 163
- TN3270E Server
 - changing the port number 105, 111
 - description 13
 - IP filtering 106
 - keep-alive processing 106
 - planning for 103, 109
- TN5250 Server
 - description 14
 - IP filtering 112
 - iSeries multiple port access 112
 - keep-alive processing 111
- Token-Ring (VTAM) host
 - parameters 135
- Token-Ring host connection
 - 9370 network adapter 136
 - NCP with VTAM 135
 - VTAM switched network 135

- topology
 - class of service 42
 - general discussion 40
 - intermediate session routing 44
 - modes 42
 - route selection 43
 - topology database 41
- transaction program (TP) 5
- transmission priority 43
- twinaxial host parameters 151

U

- usage charge 134
- user materials, creating 170

V

- volume charge 134
- VTAM (APPC) host parameters 140
- VTAM (MPC) parameters 141
- VTAM switched network host connection
 - Ethernet 138
 - Token-Ring 135
- VTAM/NCP
 - NTRI host parameters 135
 - SDLC host parameters 138
 - X.25 host parameters 146

W

- what's new in this release xiii
- windows, configuration, using
 - CMSETUP 17
 - WinSock 77, 78

X

- X.25
 - APPC host connection
 - parameters 147
 - NPSI host connection parameters 145
 - VTAM/NCP host connection
 - parameters 146
- X.25 networks
 - addresses 162
 - network charges 134
 - planning 133
 - possible connections 134
- X.25 workstation connections 96

Readers' Comments — We'd Like to Hear from You

IBM SecureWay Communications Server for Windows
Network Administration Guide
Version 6.1

Publication No. SC31-8656-02

Overall, how satisfied are you with the information in this book?

| | Very Satisfied | Satisfied | Neutral | Dissatisfied | Very Dissatisfied |
|----------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|
| Overall satisfaction | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |

How satisfied are you that the information in this book is:

| | Very Satisfied | Satisfied | Neutral | Dissatisfied | Very Dissatisfied |
|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|
| Accurate | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| Complete | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| Easy to find | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| Easy to understand | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| Well organized | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| Applicable to your tasks | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |

Please tell us how we can improve this book:

Thank you for your responses. May we contact you? Yes No

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

Name

Address

Company or Organization

Phone No.



Fold and Tape

Please do not staple

Fold and Tape



NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

IBM Corporation
Information Development
Department CGMD / Bldg 500
P.O. Box 12195
Research Triangle Park, NC 27709-9990



Fold and Tape

Please do not staple

Fold and Tape



Printed in U.S.A.

SC31-8656-02

