

z/OS Communications Server



IPv6 Network and Application Design Guide

Version 1 Release 8

z/OS Communications Server



IPv6 Network and Application Design Guide

Version 1 Release 8

Note:

Before using this information and the product it supports, be sure to read the general information under "Notices" on page 161.

Fifth Edition (September 2006)

This edition applies to Version 1 Release 8 of z/OS (5694-A01) and Version 1 Release 8 of z/OS.e (5655-G52) and to all subsequent releases and modifications until otherwise indicated in new editions.

IBM welcomes your comments. You may send your comments to the following address.

International Business Machines Corporation
Attn: z/OS Communications Server Information Development
Department AKCA, Building 501
P.O. Box 12195, 3039 Cornwallis Road
Research Triangle Park, North Carolina 27709-2195

You can send us comments electronically by using one of the following methods:

Fax (USA and Canada):

1+919-254-4028

Send the fax to "Attn: z/OS Communications Server Information Development"

Internet e-mail:

comsvrcf@us.ibm.com

World Wide Web:

<http://www.ibm.com/servers/eserver/zseries/zos/webqs.html>

If you would like a reply, be sure to include your name, address, telephone number, or FAX number. Make sure to include the following in your comment or note:

- Title and order number of this document
- Page number or topic related to your comment

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

© Copyright International Business Machines Corporation 2002, 2006. All rights reserved.

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

Contents

Figures	ix
Tables	xi
About this document	xiii
Who should read this document	xiii
How this document is organized	xiii
How to use this document	xiv
Determining whether a publication is current	xiv
How to contact IBM service	xiv
Conventions and terminology used in this document	xv
Clarification of notes	xv
How to read a syntax diagram	xv
Symbols and punctuation	xvi
Commands	xvi
Parameters	xvi
Syntax examples	xvi
Prerequisite and related information	xix
Required information	xix
Related information	xix
How to send your comments	xxiii
Summary of changes	xxv
Part 1. IPv6 overview	1
Chapter 1. Introduction	3
Expanded routing and addressing	4
Hierarchical addressing and routing infrastructure	4
Simplified IP header format	4
Improved support for options	4
Address autoconfiguration	5
New protocol for neighbor node interaction	5
Comparison of IPv6 and IPv4 characteristics	6
Dual-mode stack support	7
Chapter 2. IPv6 addressing	9
Textual representation of IPv6 addresses ¹	9
Textual representation of IPv6 prefixes ¹	10
IPv6 address space	11
IPv6 addressing model	11
Scope zones	11
Categories of IPv6 addresses	12
Unicast IPv6 addresses	13
Multicast IPv6 Addresses	17
Anycast IPv6 Addresses	19
Typical IPv6 addresses assigned to a node	19
IPv6 address states	19
Tentative	19
Deprecated	20
Preferred	20
Unavailable	20

Chapter 3. IPv6 protocol	21
Extension headers	21
Fragmentation in an IPv6 network.	21
Fragmentation and UDP/RAW	22
Path MTU discovery	22
IPv6 routing	22
Router discovery	23
ICMPv6 redirects	23
Dynamic routing protocols	24
Considerations for route selection	26
Considerations for multipath routes	26
How does a VARY TCPIP,,OBEYFILE command affect routes?	26
ICMPv6	26
Multicasting	27
Multicast Listener Discovery (MLD)	27
Neighbor discovery (ND)	28
Router advertisements	28
Redirect processing	31
Duplicate Address Detection (DAD)	32
Address resolution	32
Neighbor unreachability detection	33
Assigning IP addresses to interfaces	33
Stateless address autoconfiguration	33
IP address takeover following an interface failure.	34
How to get addresses for VIPAs	35
Default address selection	36
Default destination address selection	36
Default source address selection	38
Migration and coexistence	39
Enabling IPv6 communication between IPv6 islands in an IPv4 environment	39
How to enable end-to-end communication between IPv4 and IPv6 applications	40
Considerations for configuring z/OS for IPv6	44
IPv6 stack support	44
INET considerations	45
IPv4-only stack	45
Dual-mode IPv4/IPv6 stack	46
Common INET considerations	46
Enabling AF_INET6 support in a Common INET environment	46
Disabling AF_INET6 support in a Common INET environment	46
Supporting a mixture of dual-mode stacks and IPv4-only stacks.	47
Configuring a common INET environment	48

Part 2. IPv6 enablement

49

 Chapter 4. Configuring support for z/OS	51
Ensure that important features are supported over IPv6	51
Assess automation and application impacts due to Netstat and message changes	51
Determine how remote sites connect to the local host	51
SNA access	52
Avoid using IP addresses for identifying remote hosts	52
Considerations when using the BIND parameter on the PORT statement	53
Security considerations	53
Application programming considerations	54
Enabling IPv6 support.	54
Configuring z/OS IPv6 support	55
Resolver processing.	56
Resolver configuration.	56
Resolver communications with the Domain Name System (DNS)	58
User exits	58
Which applications started with inetd are IPv6 enabled?	58

Modifying inetc.conf	58
How does IPv6 affect SMF records?	59
How does IPv6 affect the Policy Agent?	59
How does IPv6 affect SNMP?	60
Monitoring the TCP/IP network	60
How does IPv6 affect Netstat?	60
How does IPv6 affect Ping and Traceroute?	62
Diagnosing problems	62
How does IPv6 affect IPCS?	62
How does IPv6 affect packet and data tracing?	62
Chapter 5. Configuration guidelines	63
Connecting to an IPv6 Network	63
IPv6 address assignment guidelines	64
Updating DNS definitions	66
Including static VIPAs in DNS	66
Defining IPv4-only host names and IPv4/IPv6 host names	66
Using source VIPA	66
Using OMPROUTE or define static routes to improve network selection	67
Connecting to non-local IPv4 locations	68
IPv6-only application access to IPv4-only application	68
Part 3. Application enablement	69
Chapter 6. API support	71
UNIX socket APIs	71
z/OS UNIX Assembler Callable Services	71
z/OS C sockets	71
Native TCP/IP socket APIs	72
Sockets Extended macro API	72
Sockets Extended Call Instruction API	72
REXX sockets	72
CICS sockets	72
IMS sockets	72
Pascal API	72
TCP/IP C/C++ Sockets	73
Chapter 7. Basic socket API extensions for IPv6	75
Introduction	75
Design considerations	75
Protocol families	75
Address families	75
Special IP addresses	76
Name and address resolution functions	77
Protocol-independent nodename and service name translation	77
Socket address structure to host name and service name	82
Address conversion functions	83
Address testing macros	83
Interface identification	84
Socket options to support IPv6 (IPPROTO_IPV6 level)	84
Option to control sending of unicast packets	85
Options to control sending of multicast packets	85
Options to control receiving of multicast packets	86
Socket option to control IPv4 and IPv6 communications	87
Socket options for SOL_SOCKET, IPPROTO_TCP and IPPROTO_IP levels	87
Chapter 8. Enabling an application for IPv6	89
Changes to enable IPv6 support	89
Support for unmodified applications	89

Application awareness of whether system is IPv6 enabled	89
Socket address (sockaddr_in) structure changes	92
Address conversion functions	92
Resolver API processing	92
Special IPv6 addresses	93
Passing ownership of sockets across applications using givesocket and takesocket APIs	93
Using multicast and IPv6	94
IP addresses might not be permanent	94
Including IP addresses in the data stream	95
Example of an IPv4 TCP server program	95
Example of the simple TCP server program enabled for IPv6	96
Chapter 9. Advanced socket APIs	99
Controlling the content of the IPv6 packet header	99
Socket options and ancillary data to support IPv6 (IPPROTO_IPV6 level)	100
Socket option to support ICMPv6 (IPPROTO_ICMPV6 level)	112
Using ancillary data on sendmsg() and recvmsg()	113
Interactions between socket options and ancillary data	114
Understanding hop limit options	114
Understanding options for setting the source address	115
Understanding options for specifying the outgoing interface	115
Why use RAW sockets?	116
RAW protocol values	116
Application visibility of IP headers	116
ICMP considerations	117
Checksumming data	118
Part 4. Advanced topics	119
Chapter 10. Advanced concepts and topics	121
Tunneling	121
Configured tunnels	122
Automatic tunnels	123
6to4 addresses	123
6over4 tunnels	124
Application migration and coexistence overview	125
Application migration approaches	127
Translation mechanisms	127
Chapter 11. IPv6 support tables	131
Supported IPv6 standards	131
z/OS-specific features	131
Applications not enabled for IPv6	134
Part 5. Appendixes	137
Appendix A. Related protocol specifications (RFCs)	139
Internet drafts	154
Appendix B. Information APARs and technotes	155
Information APARs for IP documents	155
Information APARs for SNA documents	156
Other information APARs	156
Appendix C. Accessibility	159
Using assistive technologies	159
Keyboard navigation of the user interface	159
z/OS information	159

Notices	161
Trademarks	169
Bibliography	171
z/OS Communications Server information	171
z/OS Communications Server library	171
Index	177
Communicating Your Comments to IBM	181

Figures

1. IPv6 address space	4
2. Unicast address format	13
3. Global unicast address format	14
4. Link-local address format	14
5. Link-local scope zones	15
6. IPv4-mapped IPv6 address	16
7. OSA-Express QDIO interface ID format	16
8. Multicast address format	17
9. Flags in multicast address.	17
10. Communicating between IPv6 islands in an IPv4 world.	40
11. Communicating between IPv4 and IPv6 applications.	40
12. IPv6 application on dual-mode stack	42
13. IPv4-only application on a dual-mode stack.	43
14. Mixing dual-mode and IPv4-only stacks	48
15. z/OS socket APIs	71
16. Example of protocol-independent client application	91
17. IPv4 TCP server program	96
18. Simple TCP server program enabled for IPv6	97
19. Tunneling.	122
20. 6to4 address format	124
21. 6over4 address format	125
22. Dual-mode stack IP host	126
23. Application communication on a dual-mode host	127

Tables

1. IPv4/IPv6 comparison	6
2. Address types	9
3. Address type representation	11
4. Multicast scope field values	17
5. Source address selection	39
6. IPv6 support for different policy types	59
7. sockaddr format for AF_INET	76
8. sockaddr format for AF_INET6	76
9. Special IP addresses	76
10. Getaddrinfo application capabilities 1	79
11. Getaddrinfo application capabilities 2	80
12. Address conversion functions	83
13. Address testing macros	83
14. Function calls	84
15. Socket options for getsockopt() and setsockopt()	85
16. Using socket() to determine IPv6 enablement	90
17. sockaddr structure changes	92
18. Address conversion function changes	92
19. Resolver API changes	93
20. Special IPv6 address changes	93
21. givesocket() and takesocket() changes	93
22. Multicast options	94
23. Sockets options at the IPPROTO_IPV6 level	100
24. Ancillary data on sendmsg() (Level = IPPROTO_IPV6)	101
25. Ancillary data on recvmsg() (Level = IPPROTO_IPV6)	101
26. Sockets options at the IPPROTO_ICMPV6 level	112
27. Macros used to manipulate filter value	112
28. Supported IPv6 standards	131
29. Link-layer device support	132
30. Virtual IP Addressing support	132
31. Sysplex support	132
32. IP routing functions	133
33. Miscellaneous IP/IF-layer functions	133
34. Transport-layer functions	133
35. Network management and accounting functions	134
36. Security functions	134
37. Applications not enabled for IPv6	134
38. IP information APARs for z/OS Communications Server	155
39. SNA information APARs for z/OS Communications Server	156
40. Non-document information APARs	157

About this document

This document contains information relating to the IPv6 protocol and the implementation of the protocol on z/OS® Communications Server Version 1 Release 8.

This document supports z/OS.e.

Who should read this document

The reader of this document should be familiar with the IPv6 protocol.

Part 1, “IPv6 overview,” on page 1, Part 2, “IPv6 enablement,” on page 49, and Part 4, “Advanced topics,” on page 119 are intended for programmers and system administrators who are familiar with TCP/IP, MVS™, and z/OS UNIX®.

Part 3, “Application enablement,” on page 69 is intended for application programmers.

How this document is organized

This document contains the following parts and chapters:

- Part 1, “IPv6 overview,” on page 1, contains information about the IPv6 protocol. It contains the following chapters:
 - Chapter 1, “Introduction,” on page 3 provides an introduction to IPv6 for z/OS Communications Server Version 1 Release 8.
 - Chapter 2, “IPv6 addressing,” on page 9 contains a discussion of the IPv6 addressing model and the different IPv6 address types.
 - Chapter 3, “IPv6 protocol,” on page 21 provides a description of the z/OS Communications Server Version 1 Release 8 implementation of the IPv6 protocol.
- Part 2, “IPv6 enablement,” on page 49, contains information about functions specific to z/OS Communications Server Version 1 Release 8. It contains the following chapters:
 - Chapter 4, “Configuring support for z/OS,” on page 51 describes the IPv6 function provided in z/OS Communications Server Version 1 Release 8 and how to enable it.
 - Chapter 5, “Configuration guidelines,” on page 63 contains recommendations and guidance information for implementing the IPv6 functions provided in z/OS Communications Server Version 1 Release 8.
- Part 3, “Application enablement,” on page 69, contains information needed to IPv6-enable applications. It contains the following chapters:
 - Chapter 6, “API support,” on page 71 describes the various z/OS socket APIs and the level of IPv6 present for each API.
 - Chapter 7, “Basic socket API extensions for IPv6,” on page 75 describes basic socket API changes that most applications would use.
 - Chapter 8, “Enabling an application for IPv6,” on page 89 describes common issues and considerations involved in enabling existing IPv4 socket applications for IPv6 communications.
- Part 4, “Advanced topics,” on page 119 contains advanced topics and concepts.

- Part 5, “Appendices,” on page 137 “Appendices” provides information that you might find helpful. This document contains the following appendixes:
 - Appendix A, “Related protocol specifications (RFCs),” on page 139 lists the related protocol specifications for TCP/IP.
 - Appendix B, “Information APARs and technotes,” on page 155 lists information APARs for IP and SNA documents.
 - Appendix C, “Accessibility,” on page 159 describes accessibility features to help users with physical disabilities.
 - “Notices” on page 161 contains notices and trademarks used in this document.
 - “Bibliography” on page 171 contains descriptions of the documents in the z/OS Communications Server library.

How to use this document

To use this document, you should be familiar with z/OS TCP/IP Services and the TCP/IP suite of protocols.

Determining whether a publication is current

As needed, IBM® updates its publications with new and changed information. For a given publication, updates to the hardcopy and associated BookManager® softcopy are usually available at the same time. Sometimes, however, the updates to hardcopy and softcopy are available at different times. The following information describes how to determine if you are looking at the most current copy of a publication:

- At the end of a publication’s order number there is a dash followed by two digits, often referred to as the dash level. A publication with a higher dash level is more current than one with a lower dash level. For example, in the publication order number GC28-1747-07, the dash level 07 means that the publication is more current than previous levels, such as 05 or 04.
- If a hardcopy publication and a softcopy publication have the same dash level, it is possible that the softcopy publication is more current than the hardcopy publication. Check the dates shown in the Summary of Changes. The softcopy publication might have a more recently dated Summary of Changes than the hardcopy publication.
- To compare softcopy publications, you can check the last two characters of the publication’s file name (also called the book name). The higher the number, the more recent the publication. Also, next to the publication titles in the CD-ROM booklet and the readme files, there is an asterisk (*) that indicates whether a publication is new or changed.

How to contact IBM service

For immediate assistance, visit this Web site: <http://www.software.ibm.com/network/commserver/support/>

Most problems can be resolved at this Web site, where you can submit questions and problem reports electronically, as well as access a variety of diagnosis information.

For telephone assistance in problem diagnosis and resolution (in the United States or Puerto Rico), call the IBM Software Support Center anytime (1-800-IBM-SERV). You will receive a return call within 8 business hours (Monday – Friday, 8:00 a.m. – 5:00 p.m., local customer time).

Outside of the United States or Puerto Rico, contact your local IBM representative or your authorized IBM supplier.

If you would like to provide feedback on this publication, see “Communicating Your Comments to IBM” on page 181.

Conventions and terminology used in this document

Commands in this book that can be used in both TSO and z/OS UNIX environments use the following conventions:

- When describing how to use the command in a TSO environment, the command is presented in uppercase (for example, NETSTAT).
- When describing how to use the command in a z/OS UNIX environment, the command is presented in bold lowercase (for example, **netstat**).
- When referring to the command in a general way in text, the command is presented with an initial capital letter (for example, Netstat).

Samples used in this book might not be updated for each release. Evaluate a sample carefully before applying it to your system.

For definitions of the terms and abbreviations used in this document, you can view the latest IBM terminology at the IBM Terminology website.

Clarification of notes

Information traditionally qualified as **Notes** is further qualified as follows:

Note Supplemental detail

Tip Offers shortcuts or alternative ways of performing an action; a hint

Guideline

Customary way to perform a procedure

Rule Something you must do; limitations on your actions

Restriction

Indicates certain conditions are not supported; limitations on a product or facility

Requirement

Dependencies, prerequisites

Result Indicates the outcome

How to read a syntax diagram

This syntax information applies to all commands and statements included in this document that do not have their own syntax described elsewhere in this document.

The syntax diagram shows you how to specify a command so that the operating system can correctly interpret what you type. Read the syntax diagram from left to right and from top to bottom, following the horizontal line (the main path).

Symbols and punctuation

The following symbols are used in syntax diagrams:

Symbol	Description
►►	Marks the beginning of the command syntax.
►	Indicates that the command syntax is continued.
	Marks the beginning and end of a fragment or part of the command syntax.
◄◄	Marks the end of the command syntax.

You must include all punctuation such as colons, semicolons, commas, quotation marks, and minus signs that are shown in the syntax diagram.

Commands

Commands that can be used in both TSO and z/OS UNIX environments use the following conventions in syntax diagrams:

- When describing how to use the command in a TSO environment, the command is presented in uppercase (for example, NETSTAT).
- When describing how to use the command in a z/OS UNIX environment, the command is presented in bold lowercase (for example, **netstat**).

Parameters

The following types of parameters are used in syntax diagrams.

Required

Required parameters are displayed on the main path.

Optional

Optional parameters are displayed below the main path.

Default

Default parameters are displayed above the main path.

Parameters are classified as keywords or variables. For the TSO and MVS console commands, the keywords are not case sensitive. You can code them in uppercase or lowercase. If the keyword appears in the syntax diagram in both uppercase and lowercase, the uppercase portion is the abbreviation for the keyword (for example, OPERand).

For the z/OS UNIX commands, the keywords must be entered in the case indicated in the syntax diagram.

Variables are italicized, appear in lowercase letters, and represent names or values you supply. For example, a data set is a variable.


Syntax examples

In the following example, the USER command is a keyword. The required variable parameter is *user_id*, and the optional variable parameter is *password*. Replace the variable parameters with your own values.



Longer than one line

If a diagram is longer than one line, the first line ends with a single arrowhead and the second line begins with a single arrowhead.



Required operands

Required operands and values appear on the main path line.



You must code required operands and values.


Choose one required item from a stack

If there is more than one mutually exclusive required operand or value to choose from, they are stacked vertically.



Optional values

Optional operands and values appear below the main path line.



You can choose not to code optional operands and values.

Choose one optional operand from a stack

If there is more than one mutually exclusive optional operand or value to choose from, they are stacked vertically below the main path line.




Repeating an operand

An arrow returning to the left above an operand or value on the main path line means that the operand or value can be repeated. The comma means that each operand or value must be separated from the next by a comma. If no comma appears in the returning arrow, the operand or value must be separated from the next by a blank.



Selecting more than one operand

An arrow returning to the left above a group of operands or values means more than one can be selected, or a single one can be repeated.



Nonalphanumeric characters

If a diagram shows a character that is not alphanumeric (such as parentheses, periods, commas, and equal signs), you must code the character as part of the syntax. In this example, you must code OPERAND=(001,0.001).



Blank spaces in syntax diagrams

If a diagram shows a blank space, you must code the blank space as part of the syntax. In this example, you must code OPERAND=(001 FIXED).



Default operands

Default operands and values appear above the main path line. TCP/IP uses the default if you omit the operand entirely.



Variables

A word in all lowercase italics is a *variable*. Where you see a variable in the syntax, you must replace it with one of its allowable names or values, as defined in the text.



Syntax fragments

Some diagrams contain syntax fragments, which serve to break up diagrams that are too long, too complex, or too repetitious. Syntax fragment names are in mixed case and are shown in the diagram and in the heading of the fragment. The fragment is placed below the main diagram.



Syntax fragment:

—1ST_OPERAND,2ND_OPERAND,3RD_OPERAND—

Prerequisite and related information

z/OS Communications Server function is described in the z/OS Communications Server library. Descriptions of those documents are listed in “z/OS Communications Server information” on page 171, in the back of this document.

Required information

Before using this product, you should be familiar with TCP/IP, VTAM®, MVS, and UNIX System Services.

Related information

This section contains subsections on:

- “Softcopy information”
- “Other documents”
- “Redbooks” on page xx
- “Where to find related information on the Internet” on page xxi
- “Using LookAt to look up message explanations” on page xxii
- “Using IBM Health Checker for z/OS” on page xxiii

Softcopy information

Softcopy publications are available in the following collections:

Titles	Order Number	Description
<i>z/OS V1R7 Collection</i>	SK3T-4269	This is the CD collection shipped with the z/OS product. It includes the libraries for z/OS V1R7, in both BookManager and PDF formats.
<i>z/OS Software Products Collection</i>	SK3T-4270	This CD includes, in both BookManager and PDF formats, the libraries of z/OS software products that run on z/OS but are not elements and features, as well as the <i>Getting Started with Parallel Sysplex</i> ® bookshelf.
<i>z/OS V1R7 and Software Products DVD Collection</i>	SK3T-4271	This collection includes the libraries of z/OS (the element and feature libraries) and the libraries for z/OS software products in both BookManager and PDF format. This collection combines SK3T-4269 and SK3T-4270.
<i>z/OS Licensed Product Library</i>	SK3T-4307	This CD includes the licensed documents in both BookManager and PDF format.
<i>System Center Publication IBM S/390 Redbooks Collection</i>	SK2T-2177	This collection contains over 300 ITSO redbooks that apply to the S/390® platform and to host networking arranged into subject bookshelves.

Other documents

For information about z/OS products, refer to *z/OS Information Roadmap* (SA22-7500). The Roadmap describes what level of documents are supplied with each release of z/OS Communications Server, as well as describing each z/OS publication.

Relevant RFCs are listed in an appendix of the IP documents. Architectural specifications for the SNA protocol are listed in an appendix of the SNA documents.

The following table lists documents that might be helpful to readers.

Title	Number
<i>DNS and BIND</i> , Fourth Edition, O'Reilly and Associates, 2001	ISBN 0-596-00158-4
<i>Routing in the Internet</i> , Christian Huitema (Prentice Hall PTR, 1995)	ISBN 0-13-132192-7
<i>sendmail</i> , Bryan Costales and Eric Allman, O'Reilly and Associates, 2002	ISBN 1-56592-839-3
<i>SNA Formats</i>	GA27-3136
<i>TCP/IP Illustrated, Volume I: The Protocols</i> , W. Richard Stevens, Addison-Wesley Publishing, 1994	ISBN 0-201-63346-9
<i>TCP/IP Illustrated, Volume II: The Implementation</i> , Gary R. Wright and W. Richard Stevens, Addison-Wesley Publishing, 1995	ISBN 0-201-63354-X
<i>TCP/IP Illustrated, Volume III</i> , W. Richard Stevens, Addison-Wesley Publishing, 1995	ISBN 0-201-63495-3
<i>TCP/IP Tutorial and Technical Overview</i>	GG24-3376
<i>Understanding LDAP</i>	SG24-4986
<i>z/OS Cryptographic Service System Secure Sockets Layer Programming</i>	SC24-5901
<i>z/OS Integrated Security Services Firewall Technologies</i>	SC24-5922
<i>z/OS Integrated Security Services LDAP Client Programming</i>	SC24-5924
<i>z/OS Integrated Security Services LDAP Server Administration and Use</i>	SC24-5923
<i>z/OS JES2 Initialization and Tuning Guide</i>	SA22-7532
<i>z/OS Problem Management</i>	G325-2564
<i>z/OS MVS Diagnosis: Reference</i>	GA22-7588
<i>z/OS MVS Diagnosis: Tools and Service Aids</i>	GA22-7589
<i>z/OS MVS Using the Subsystem Interface</i>	SA22-7642
<i>z/OS Program Directory</i>	GI10-0670
<i>z/OS UNIX System Services Command Reference</i>	SA22-7802
<i>z/OS UNIX System Services Planning</i>	GA22-7800
<i>z/OS UNIX System Services Programming: Assembler Callable Services Reference</i>	SA22-7803
<i>z/OS UNIX System Services User's Guide</i>	SA22-7801
<i>z/OS XL C/C++ Run-Time Library Reference</i>	SA22-7821
<i>System z9 and zSeries OSA-Express Customer's Guide and Reference</i>	SA22-7935

Redbooks

The following Redbooks™ might help you as you implement z/OS Communications Server.

Title	Number
<i>Communications Server for z/OS V1R7 TCP/IP Implementation, Volume 1: Base Functions, Connectivity, and Routing</i>	SG24-7169
<i>Communications Server for z/OS V1R7 TCP/IP Implementation, Volume 2: Standard Applications</i>	SG24-7170
<i>Communications Server for z/OS V1R7 TCP/IP Implementation, Volume 3: High Availability, Scalability, and Performance</i>	SG24-7171

	Title	Number
	<i>Communications Server for z/OS V1R7 TCP/IP Implementation, Volume 4: Policy-Based Network Security</i>	SG24-7172
	<i>IBM Communication Controller Migration Guide</i>	SG24-6298
	<i>IP Network Design Guide</i>	SG24-2580
	<i>Managing OS/390® TCP/IP with SNMP</i>	SG24-5866
	<i>Migrating Subarea Networks to an IP Infrastructure</i>	SG24-5957
	<i>SecureWay Communications Server for OS/390 V2R8 TCP/IP: Guide to Enhancements</i>	SG24-5631
	<i>SNA and TCP/IP Integration</i>	SG24-5291
	<i>TCP/IP in a Sysplex</i>	SG24-5235
	<i>TCP/IP Tutorial and Technical Overview</i>	GG24-3376
	<i>Threadssafe Considerations for CICS</i>	SG24-6351

Where to find related information on the Internet

z/OS

This site provides information about z/OS Communications Server release availability, migration information, downloads, and links to information about z/OS technology

<http://www.ibm.com/servers/eserver/zseries/zos/>

z/OS Internet Library

Use this site to view and download z/OS Communications Server documentation

<http://www.ibm.com/servers/eserver/zseries/zos/bkserv/>

IBM Communications Server product

The primary home page for information about z/OS Communications Server

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

IBM Communications Server product support

Use this site to submit and track problems and search the z/OS Communications Server knowledge base for Technotes, FAQs, white papers, and other z/OS Communications Server information

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

IBM Systems Center publications

Use this site to view and order Redbooks, Redpapers, and Technotes

<http://www.redbooks.ibm.com/>

IBM Systems Center flashes

Search the Technical Sales Library for Techdocs (including Flashes, presentations, Technotes, FAQs, white papers, Customer Support Plans, and Skills Transfer information)

<http://www.ibm.com/support/techdocs/atstr.nsf>

RFCs

Search for and view Request for Comments documents in this section of the Internet Engineering Task Force Web site, with links to the RFC repository and the IETF Working Groups Web page

<http://www.ietf.org/rfc.html>

Internet drafts

View Internet-Drafts, which are working documents of the Internet Engineering Task Force (IETF) and other groups, in this section of the Internet Engineering Task Force Web site

<http://www.ietf.org/ID.html>

Information about Web addresses can also be found in information APAR II11334.

DNS Web sites: For more information about DNS, see the following USENET news groups and mailing addresses:

USENET news groups

comp.protocols.dns.bind

BIND mailing lists

<http://www.isc.org/ml-archives/>

BIND Users

- Subscribe by sending mail to bind-users-request@isc.org.
- Submit questions or answers to this forum by sending mail to bind-users@isc.org.

BIND 9 Users (This list might not be maintained indefinitely.)

- Subscribe by sending mail to bind9-users-request@isc.org.
- Submit questions or answers to this forum by sending mail to bind9-users@isc.org.

Note: Any pointers in this publication to Web sites are provided for convenience only and do not in any manner serve as an endorsement of these Web sites.

Using LookAt to look up message explanations

LookAt is an online facility that lets you look up explanations for most of the IBM messages you encounter, as well as for some system abends and codes. Using LookAt to find information is faster than a conventional search because in most cases LookAt goes directly to the message explanation.

You can use LookAt from these locations to find IBM message explanations for z/OS elements and features, z/VM[®], VSE/ESA[™], and Clusters for AIX[®] and Linux[™]:

- The Internet. You can access IBM message explanations directly from the LookAt Web site at <http://www.ibm.com/servers/eserver/zseries/zos/bkserv/lookat/>.
- Your z/OS TSO/E host system. You can install code on your z/OS or z/OS.e systems to access IBM message explanations using LookAt from a TSO/E command line (for example: TSO/E prompt, ISPF, or z/OS UNIX System Services).
- Your Microsoft[®] Windows[®] workstation. You can install LookAt directly from the z/OS Collection (SK3T-4269) or the z/OS and Software Products DVD Collection (SK3T-4271) and use it from the resulting Windows graphical user interface

(GUI). The command prompt (also known as the DOS > command line) version can still be used from the directory in which you install the Windows version of LookAt.

- Your wireless handheld device. You can use the LookAt Mobile Edition from <http://www.ibm.com/servers/eserver/zseries/zos/bkserv/lookat/lookatm.html> with a handheld device that has wireless access and an Internet browser (for example: Internet Explorer for Pocket PCs, Blazer or Eudora for Palm OS, or Opera for Linux handheld devices).

You can obtain code to install LookAt on your host system or Microsoft Windows workstation from:

- A CD-ROM in the *z/OS Collection* (SK3T-4269).
- The *z/OS and Software Products DVD Collection* (SK3T-4271).
- The LookAt Web site (click **Download** and then select the platform, release, collection, and location that suit your needs). More information is available in the LOOKAT.ME files available during the download process.

Using IBM Health Checker for z/OS

IBM Health Checker for z/OS is a z/OS component that installations can use to gather information about their system environment and system parameters to help identify potential configuration problems before they impact availability or cause outages. Individual products, z/OS components, or ISV software can provide checks that take advantage of the IBM Health Checker for z/OS framework. This book refers to checks or messages associated with this component.

For additional information about checks and about IBM Health Checker for z/OS, see *IBM Health Checker for z/OS: User's Guide*. z/OS V1R4, V1R5, and V1R6 users can obtain the IBM Health Checker for z/OS from the z/OS Downloads page at <http://www.ibm.com/servers/eserver/zseries/zos/downloads/>.

SDSF also provides functions to simplify the management of checks. See *z/OS SDSF Operation and Customization* for additional information.

How to send your comments

Your feedback is important in helping to provide the most accurate and high-quality information. If you have any comments about this document or any other z/OS Communications Server documentation:

- Go to the z/OS contact page at:
<http://www.ibm.com/servers/eserver/zseries/zos/webqs.html>
There you will find the feedback page where you can enter and submit your comments.
- Send your comments by e-mail to comsvrcf@us.ibm.com. Be sure to include the name of the document, the part number of the document, the version of z/OS Communications Server, and, if applicable, the specific location of the text you are commenting on (for example, a section number, a page number or a table number).

Summary of changes

Summary of changes for SC31-8885-04 z/OS Version 1 Release 8

This document contains information previously presented (SC31-8885-03), which supports z/OS Version 1 Release 7.

The information in this document includes descriptions of support for both IPv4 and IPv6 networking protocols. Unless explicitly noted, descriptions of IP protocol support concern IPv4. IPv6 support is qualified within the text.

New information

- IPv6 support for IPSec:
 - NAT traversal, see Table 36 on page 134
 - Release-specific support, see “Extension headers” on page 21
 - IP security for IPv6 configuration considerations, see “Multicast Listener Discovery (MLD)” on page 27 and “Neighbor discovery (ND)” on page 28
 - IPSec policies, “How does IPv6 affect the Policy Agent?” on page 59

Changed information

- Remove support for version 1 networking SLA MIB (pagtsnmp subagent):
 - Network SLAPM2 subagent comment, see Table 35 on page 134
- IPv6 support for IPSec:
 - Security function listings, see Table 36 on page 134
 - OSPF routing protocol, “Authentication with the IPv6 OSPF routing protocol” on page 25

Deleted information

- Remove support for version 1 networking SLA MIB (pagtsnmp subagent):
 - Removed pagtsnmp, see Table 37 on page 134
 - Removed RFC 2758, see Appendix A, “Related protocol specifications (RFCs),” on page 139
- AnyNet® function is removed from the z/OS V1R8 Communications Server product and therefore documentation describing AnyNet support has been deleted.
- IPv6 Support for IPSec:
 - TRMD, see Table 37 on page 134

This document contains terminology, maintenance, and editorial changes. Technical changes or additions to the text and illustrations are indicated by a vertical line to the left of the change.

You might notice changes in the style and structure of some content in this document—for example, headings that use uppercase for the first letter of initial words only, and procedures that have a different look and format. The changes are ongoing improvements to the consistency and retrievability of information in our documents.

**Summary of changes
for SC31-8885-03
z/OS Version 1 Release 7**

This document contains information previously presented in SC31-8885-02, which supports z/OS Version 1 Release 6.

This document refers to Communications Server data sets by their default SMP/E distribution library name. Your installation might, however, have different names for these data sets where allowed by SMP/E, your installation personnel, or administration staff. For instance, this document refers to samples in SEZAINST library as simply in SEZAINST. Your installation might choose a data set name of SYS1.SEZAINST, CS390.SEZAINST or other high level qualifiers for the data set name.

New information

- Enhancements to existing advanced socket APIs for z/OS UNIX Callable Services and Language Environment® C/C++ to support RFC3542, see Chapter 9, "Advanced socket APIs," on page 99.
- IPv6 support for HiperSockets™, see "Connecting to an IPv6 Network" on page 63.

Changed information

- SNMP IPv6 UDP MIB support, see "How does IPv6 affect SNMP?" on page 60.
- Server-specific WLM for sysplex distribution, see "Considerations when using the BIND parameter on the PORT statement" on page 53.
- Promotion of the use of IPv6 global unicast addresses

Site-local addresses were designed to use private address prefixes that could be used within a site without the need for a global prefix. Until recently, the full negative impacts of site-local address in the Internet were not fully understood. The Internet Engineering Task Force (IETF) has deprecated the special treatment given to this site-local prefix. Because of this, it is preferable to use global unicast addresses. This means addresses and prefixes that use the site-local prefix (fec0::/10) are being replaced with ones that use the global prefix for documentation (2001:0DB8::/32).

Deleted information

- All OROUTED information.
- The following tables from Chapter 4, "Configuring support for z/OS," on page 51:
 - IPv6 supported features table
 - IPv6 supported applications table.

This document contains terminology, maintenance, and editorial changes. Technical changes or additions to the text and illustrations are indicated by a vertical line to the left of the change.

You might notice changes in the style and structure of some content in this document—for example, headings that use uppercase for the first letter of initial words only, and procedures that have a different look and format. The changes are ongoing improvements to the consistency and retrievability of information in our documents.

**Summary of changes
for SC31-8885-02
z/OS Version 1 Release 6**

This document contains information previously presented in SC31-8885-01, which supports z/OS Version 1 Release 5.

New information

- IPv6 OSPF support for OMPROUTE , see “Dynamic routing protocols” on page 24
- Sysplex distributor policy performance monitoring, see “How does IPv6 affect the Policy Agent?” on page 59
- Sysplex enhancements, see Table 31 on page 132
- Select examples are enabled for z/OS library center advanced searches.

Changed information

- SNMP IPv6 MIBs, see “How does IPv6 affect SNMP?” on page 60
- OSPF support added to recommendations, see “Using OMPROUTE or define static routes to improve network selection” on page 67

This document contains terminology, maintenance, and editorial changes. Technical changes or additions to the text and illustrations are indicated by a vertical line to the left of the change.

Starting with z/OS V1R4, you may notice changes in the style and structure of some content in this document—for example, headings that use uppercase for the first letter of initial words only, and procedures that have a different look and format. The changes are ongoing improvements to the consistency and retrievability of information in our documents.

Part 1. IPv6 overview

This section contains the following chapters:

Chapter 1, “Introduction,” on page 3 provides an introduction to IPv6 for z/OS Communications Server.

Chapter 2, “IPv6 addressing,” on page 9 contains a discussion on the IPv6 addressing model and the different IPv6 address types.

Chapter 3, “IPv6 protocol,” on page 21 provides a description of the z/OS Communications Server implementation of the IPv6 protocol.

Chapter 1. Introduction

Internet Protocol Version 6 (IPv6) is the next generation of the Internet protocol designed to replace the current version, Internet Protocol Version 4 (IPv4). Most of today's internets use IPv4, which is approximately 20 years old and is approaching the end of its physical limits. The most significant issue surrounding IPv4 is the growing shortage of IPv4 addresses. In theory, 32 bits allow over 4 billion nodes, each with a globally unique address. In practice, the interaction between routing and addressing makes it impossible to exploit more than a small fraction of that number of nodes. Consequently, there is a growing concern that the continued growth of the Internet might lead to the exhaustion of IPv4 addresses early in the 21st century.

IPv6 fixes a number of problems in IPv4, such as the limited number of available IPv4 addresses. IPv6 uses 128-bit addresses, an address space large enough to last for the foreseeable future. It also adds many improvements to IPv4 in areas such as routing and network autoconfiguration. IPv6 is expected to gradually replace IPv4, with the two coexisting for a number of years during a transition period.

IPv6 is an evolutionary step from IPv4. Functions that work well in IPv4 have been kept in IPv6, and functions that did not work well in IPv4 have been removed.

z/OS Communications Server Version 1 Release 4 was the first release to incorporate IPv6 features. z/OS Communications Server enables you to do the following:

- Build an IPv6 network
- Start using IPv6-enabled applications
- Enable existing IPv4 applications to be IPv6 applications
- Access your SNA applications over an IPv6 network

Not all IPv6 features are supported. This document describes the support available and how to implement it. This chapter discusses some of the major differences between IPv4 and IPv6 and includes the following sections:

- “Expanded routing and addressing” on page 4
- “Hierarchical addressing and routing infrastructure” on page 4
- “Simplified IP header format” on page 4
- “Improved support for options” on page 4
- “Address autoconfiguration” on page 5
- “New protocol for neighbor node interaction” on page 5
- “Comparison of IPv6 and IPv4 characteristics” on page 6
- “Dual-mode stack support” on page 7

For more information about some of the features not yet supported, refer to Part 4, “Advanced topics,” on page 119.

Expanded routing and addressing

IPv6 uses a 128-bit address space, which has no practical limit on global addressability and provides 3.4×10^{50} unique addresses. This gives enough addresses so that every person could have a single IPv6 network with many nodes, and still the address space would be almost completely unused.

The greater availability of IPv6 addresses eliminates the need for private address spaces, which in turn eliminates one of the needs for network address translators (NATs) to be used between the private Intranet and the public Internet.

Hierarchical addressing and routing infrastructure

The use of hierarchical address formats is equally important as the expanded address space. The IPv4 addressing hierarchy includes network, subnet, and host components in an IPv4 address. With its 128-bit addresses, IPv6 provides globally unique and hierarchical addressing based on prefixes rather than address classes, which keeps routing tables small and backbone routing efficient.

The general format is as follows:

n bits	m bits	128-(n+m)bits
global routing prefix	subnet ID	interface ID

Figure 1. IPv6 address space

The global routing prefix is a value (typically hierarchically structured) assigned to a site; the subnet ID is an identifier of a link within the site; and the interface ID is a unique identifier for a network device on a given link (usually automatically assigned).

Simplified IP header format

The IPv6 header has a fixed size and its format is more simplified than the IPv4 header. Some fields in the IPv4 header were dropped in IPv6 or moved to optional IPv6 extension headers to reduce the common-case processing cost of packet handling, as well as keep the bandwidth cost of the IPv6 header as low as possible despite increasing the size of addresses. While the IPv6 address is four times the size of the IPv4 address, the total IPv6 header size is only twice as large as the IPv4 header size.

Improved support for options

Changes in the way IP header options are encoded allows for more efficient forwarding, less stringent limits on the length of options, and greater flexibility for introducing new options in the future. Optional IPv6 header information is conveyed in independent extension headers located after the IPv6 header and before the transport-layer header in each packet. In contrast to IPv4, most IPv6 extension headers are not examined or processed by intermediate nodes.

Address autoconfiguration

IPv6 provides for both stateless and stateful autoconfiguration. Stateless autoconfiguration allows a node to be configured in the absence of any configuration server. Stateless autoconfiguration also makes it possible for a node to configure its own globally routable addresses in cooperation with a local IPv6 router, by combining the 48- or 64-bit MAC address of the adapter with network prefixes that are learned from the neighboring router.

IPv6 allows the use of DHCPv6 for stateful autoconfiguration. DHCPv6 relies on a configuration server that maintains static tables to determine the addresses that are assigned to newly connected nodes. z/OS Communications Server does not support DHCPv6.

Tip: Manual configuration of addresses can be used in environments where complete local control is required (as with VIPA or additional LOOPBACK addresses).

New protocol for neighbor node interaction

Neighbor Discovery (ND) corresponds to a combination of the IPv4 protocols ARP, ICMP Router Discovery, and ICMP Redirect. Nodes (hosts and routers) use ND to determine the link-layer addresses for neighbors known to reside on attached links and to quickly purge cached values that become invalid. Hosts also use ND to find neighboring routers that are able to forward packets on their behalf. ND also defines a Neighbor Unreachability Detection algorithm. IPv4 does not contain a generally agreed upon protocol for performing Neighbor Unreachability Detection, although Dead Gateway Detection does address a subset of the problems that Neighbor Unreachability Detection solves.

Neighbor Discovery is used to do the following:

- Obtain configuration information including:

Router Discovery

Defines how hosts can automatically locate routers that reside on an attached link.

Prefix Discovery

Specifies how hosts discover the set of prefixes that are defined as being on-link (IPv6 address prefixes that reside on the shared link, such as an ethernet link), as well as those which are to be used when implementing Stateless Address Autoconfiguration.

Parameter Discovery

Allows a host to learn link parameters, such as the link MTU, and IP parameters, such as the hop limit to place in outgoing packets.

- Perform address resolution. Address resolution allows a node to determine the link-layer address of an on-link destination given the destination's IP address.
- Dynamically learn routes which can be used in next-hop determination. This specifies the algorithm for mapping the IP destination address into the IP address of the neighbor to which traffic should be sent. The next-hop can be either a router or the destination itself. Next-hop determination uses the on-link prefixes learned as part of Prefix Discovery to determine when the next hop is the destination itself.
- Determine when a neighbor is no longer reachable using Neighbor Unreachability Detection.

- Process Redirect messages. Routers use Redirect messages to notify a node that a better next-hop node should be used when forwarding packets to a particular destination. The new next-hop could be the actual destination, if the destination is on-link, or a different router, if the destination is off-link.

Comparison of IPv6 and IPv4 characteristics

There are major differences between IPv4 and IPv6. Table 1 lists these differences:

Table 1. IPv4/IPv6 comparison

IPv4	IPv6
Source and destination addresses are 32 bits (4 bytes) in length.	Source and destination addresses are 128 bits (16 bytes) in length. For more information, refer to Chapter 2, “IPv6 addressing,” on page 9.
Uses broadcast addresses to send traffic to all nodes on a subnet.	There are no IPv6 broadcast addresses. Instead, multicast scoped addresses are used. For more information refer to “Multicast scope” on page 17.
Fragmentation is supported at originating hosts and intermediate routers.	Fragmentation is not supported at routers. It is only supported at the originating host. For more information refer to “Fragmentation in an IPv6 network” on page 21.
IP header includes a checksum.	IP header does not include a checksum.
IP header includes options.	All optional data is moved to IPv6 extension headers. For more information refer to “Extension headers” on page 21.
IPSec support is optional.	IPSec support is required in a full IPv6 implementation.
No identification of payload for QoS handling by routers is present within the IPv4 header.	Payload identification for QoS handling by routers is included in the IPv6 header using the Flow Label field. For more information refer to “Option to provide QoS classification data” on page 111.
ICMP Router Discovery is used to determine the IPv4 address of the best default gateway and is optional.	Uses ICMPv6 Router Solicitation and Router Advertisement to determine the IPv6 address of the best default gateway and is a required function. For more information, refer to “Router advertisements” on page 28. z/OS sends Router Solicitations and processes Router Advertisements but does not send Router Advertisements.
Address Resolution Protocol (ARP) uses broadcast ARP Request frames to resolve an IPv4 address to a link layer address.	Uses multicast Neighbor Solicitation messages for address resolution. For more information refer to “Address resolution” on page 32.
Internet Group Management Protocol (IGMP) is used to manage local subnet group membership.	Uses Multicast Listener Discovery (MLD) messages to manage local subnet group membership. For more information refer to “Multicast Listener Discovery (MLD)” on page 27.

Table 1. IPv4/IPv6 comparison (continued)

IPv4	IPv6
Addresses must be configured either manually or through DHCP.	Addresses can be automatically assigned using stateless address autoconfiguration, assigned using DHCPv6, or manually configured. DHCPv6 is not supported in z/OS Communications Server V1R7.
Uses host address (A) resource records in the Domain Name System (DNS) to map host names to IPv4 addresses.	Uses host address (AAAA) resource records in the Domain Name System (DNS) to map host names to IPv6 addresses.
Uses pointer (PTR) resource records in the IN-ADDR.ARPA DNS domain to map IPv4 addresses to host names.	Uses pointer (PTR) resource records in the IP6.ARPA or IP6.INT DNS domain to map IPv6 addresses to host names.
For QoS, IPv4 supports both differentiated and integrated services.	Differentiated and integrated services are both supported. In addition, IPv6 provides flow label that can be used for more granular treatment of packets.

Dual-mode stack support

z/OS Communications Server can be an IPv4-only stack or a dual-mode stack. Dual-mode stack refers to a single TCP/IP stack supporting both IPv4 and IPv6 protocols at the same time.

Restriction: There is no support for an IPv6-only stack.

The following are several advantages of running in a dual-mode stack configuration:

- IPv4 and IPv6 applications can coexist on a single dual-mode stack.
- Unmodified applications can continue to send data over an IPv4 network.
- A single IPv6-enabled application can communicate using IPv4 and IPv6.
- IPv4 and IPv6 can coexist in the same devices and networks.

Chapter 2. IPv6 addressing

This chapter contains the following sections:

- “Textual representation of IPv6 addresses¹”
- “Textual representation of IPv6 prefixes¹” on page 10
- “Textual representation of IPv6 prefixes¹” on page 10
- “IPv6 address space” on page 11
- “Scope zones” on page 11
- “Categories of IPv6 addresses” on page 12
- “Typical IPv6 addresses assigned to a node” on page 19
- “IPv6 address states” on page 19

Textual representation of IPv6 addresses¹

IPv4 addresses are represented in dotted-decimal format. The 32-bit address is divided along 8-bit boundaries. Each set of 8 bits is converted to its decimal equivalent and separated by periods. In contrast, IPv6 addresses are 128 bits divided along 16-bit boundaries. Each 16-bit block is converted to a 4-digit hexadecimal number and separated by colons. The resulting representation is called colon-hexadecimal.

The following are the three conventional forms for representing IPv6 addresses as text strings:

- The preferred form is x:x:x:x:x:x:x:x, where the x's are the hexadecimal values of the eight 16-bit pieces of the address. For example:

FEDC:BA98:7654:3210:FEDC:BA98:7654:3210

1080:0:0:0:8:800:200C:417A

Guideline: It is not necessary to write the leading zeros in an individual field, but there must be at least one numeral in every field (except for the case described in the following bullet).

- As a result of some methods of allocating certain styles of IPv6 addresses, sometimes addresses contain long strings of zero bits. To make writing addresses containing zero bits easier, a special syntax is available to compress the zeros. A double colon (::) indicates multiple groups of 16 bits of zeros and can appear only once in an address. The :: can also be used to compress both leading and trailing zeros in an address.

For example the following addresses:

Table 2. Address types

Address type	Long form	Compressed form
Unicast	2001:DB8:0:0:8:800:200C:417A	2001:DB8::8:800:200C:417A
Multicast	FF01:0:0:0:0:0:101	FF01::101
Loopback	0:0:0:0:0:0:1	::1
Unspecified	0:0:0:0:0:0:0	::

- An alternative form that is sometimes more convenient when dealing with a mixed environment of IPv4 and IPv6 nodes is x:x:x:x:x:d.d.d.d, where the x's are the hexadecimal values of the 6 high-order 16-bit pieces of the address, and

the d's are the decimal values of the 4 low-order 8-bit pieces of the address (standard IPv4 representation). This is used for IPv4-compatible IPv6 addresses and IPv4-mapped IPv6 addresses. These types of addresses are used to hold embedded IPv4 addresses in order to carry IPv6 packets over IPv4 routing infrastructure. The address can be expressed in the following manner:

0:0:0:0:0:13.1.68.3
0:0:0:0:FFFF:129.144.52.38

or in compressed form:

::13.1.68.3
::FFFF:129.144.52.38

Textual representation of IPv6 prefixes¹

The text representation of IPv6 address prefixes is similar to the way IPv4 address prefixes are written in Classless Inter-Domain Routing (CIDR) notation. An IPv6 address prefix is represented by the notation:

ipv6-address/prefix-length

where

ipv6-address

An IPv6 address in any of the notations listed above.

prefix-length

A decimal value specifying how many of the leftmost contiguous bits of the address comprise the prefix.

For example, the following are legal representations of the 60-bit prefix 20010DB80000CD3 (hexadecimal):

2001:DB8:0000:CD30:0000:0000:0000:0000/60
2001:DB8::CD30:0:0:0:0/60
2001:DB8:0:CD30::/60

The following are not legal representations of the preceding prefix:

- 2001:DB8:0:CD3/60
Leading zeros might be dropped, but not trailing zeros, within any 16-bit chunk of the address.
- 2001:DB8::CD30/60
Address to the left of the forward slash (/) expands to 2001:DB8:0000:0000:0000:0000:CD30.
- 2001:DB8:0:CD3/60
Address to the left of the forward slash (/) expands to 2001:DB8:0000:0000:0000:0000:0CD3.

When writing both a node address and a prefix of that node address (for example, the node's subnet prefix), the two can be combined as follows:

- Node address - 2001:DB8:0:CD30:123:4567:89AB:CDEF

1. Copyright (C) The Internet Society (1998). All Rights Reserved. This document and translations of it can be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation can be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself cannot be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

- Subnet number - 2001:DB8:0:CD30::/60
- Combination of node address and subnet number - 2001:DB8:0:CD30:123:4567:89AB:CDEF/60

IPv6 address space

The type of a IPv6 address is identified by the high-order bits of the address as shown in Table 3.

Table 3. Address type representation

Address type	Binary prefix	IPv6 notation
Unspecified	00...0 (128 bits)	::/128
Loopback	00...1 (128 bits)	::1/128
Multicast	11111111	FF00::/8
Link-local unicast	1111111010	FE80::/10
Unassigned (formerly Site-local unicast)	1111111011	FEC0::/10
Global unicast aggregatable	(everything else)	

Anycast addresses are taken from the unicast address spaces (of any scope) and are not syntactically distinguishable from unicast addresses. Anycast is described as a cross between unicast and multicast. Like multicast, multiple nodes might be listening on an Anycast address. Like unicast, a packet sent to an Anycast address is delivered to one (and only one) of those nodes. The exact node to which it is delivered is based on the IP routing tables in the network.

For more information about different IPv6 addresses, refer to “Categories of IPv6 addresses” on page 12.

IPv6 addressing model

IPv6 unicast addresses of all types (excluding loopback and unspecified) can be assigned to a node’s interfaces.

All physical interfaces (excluding VIPA and loopback) are required to have at least one link-local unicast address. z/OS Communications Server only allows a single link-local address per interface. Other platforms might have more than one. A single interface can be assigned multiple unicast or anycast IPv6 addresses. Multiple IPv6 multicast groups of any scope can be joined on a single interface. A unicast address or a set of unicast addresses might be assigned to multiple physical interfaces if the implementation treats the multiple physical interfaces as one interface when presenting it to the Internet layer.

Currently, IPv6 continues the IPv4 model that a subnet prefix is associated with one link. Multiple subnet prefixes can be assigned to the same link.

Scope zones

Each IPv6 address has a specific scope in which it is defined. A scope is a topological area within which the IPv6 address can be used as a unique identifier for an interface or a set of interfaces. The scope for an IPv6 address is encoded as part of the address itself. A unicast address can have a link-local or global scope. A multicast address supports the following:

- Interface-local
- Link-local
- Subnet-local
- Admin-local
- Site-local (has been deprecated)
- Organization-local
- Global scopes

See “Unicast IPv6 addresses” on page 13 and “Multicast IPv6 Addresses” on page 17 for more discussions about unicast and multicast scopes.

A scope zone is an instance of a given scope. For instance, a link and all directly attached interfaces comprise a single link-local scope zone. A scope zone has the following properties:

- A scope zone is comprised of a contiguous set of interfaces and the links to which the interfaces are attached.
- An interface can belong to only one scope zone of each possible scope.
- A node can be connected to more than one scope zone of a given scope. For instance, a node can be connected to multiple link-local scope zones if it is attached to more than one LAN.
- The scope zone for an IPv6 address is not encoded within the address itself, but is instead determined by the interface over which the packet is sent or received.
- There is a single scope zone for IPv6 addresses of global scope which comprises all interfaces and links in the Internet.
- Packets that contain a source or destination address of a given scope can be routed only within the same scope zone, and cannot be routed between different scope zone instances.
- Addresses of a given scope can be reused in different scope zones.
- Scope zones associated with the inbound and intended outbound interfaces are compared to determine whether packets containing a limited scope address (for example, an address of scope other than global) can be successfully routed.
- Scope zone representations (zone indices) are valid only on the node where they are defined. The same zone can have separate representations in each node that belongs to that zone.

To identify a specific instance of a scope zone, a node assigns a unique scope zone index to each scope zone of the same scope to which it is attached.

Categories of IPv6 addresses

An IPv6 address is identified by the high-order bits of the address. The following categories of IP addresses are supported in IPv6:

Unicast

An identifier for a single interface. A packet sent to a unicast address is delivered to the interface identified by that address. It can be link-local scope, site-local scope, or global scope.

Guideline: Do not use site-local addresses.

Multicast

An identifier for a group of interfaces (typically belonging to different nodes). A packet sent to a multicast address is delivered to all interfaces identified by that address.

Anycast

An identifier for a group of interfaces (typically belonging to different nodes). A packet sent to an anycast address is delivered to the closest member of a group, according to the routing protocols' measure of distance.

Restriction: Although z/OS Communications Server can send or forward datagrams to an anycast address, z/OS Communications Server does not support functioning as an anycast endpoint.

There are no broadcast addresses in IPv6. Multicast addresses have superseded this function.

Unicast IPv6 addresses

IPv6 unicast addresses are aggregatable with prefixes of arbitrary bit-length similar to IPv4 addresses under Classless Interdomain Routing (CIDR).

There are several types of unicast addresses in IPv6:

- Global unicast
- Site-local unicast (has been deprecated)

Restriction: Although z/OS Communications Server can send or forward datagrams to an anycast address, z/OS Communications Server does not support functioning as an anycast endpoint.

- Link-local unicast

There are also some special-purpose subtypes of global unicast:

- IPv6 addresses with embedded IPv4 addresses

Additional address types or subtypes can be defined in the future.

A unicast address has the following format:

n bits	128-n bits
network prefix	interface ID

Figure 2. Unicast address format

Aggregatable global addresses

Aggregatable global unicast addresses are equivalent to public IPv4 addresses. They are globally routable and reachable on the IPv6 portion of the Internet.

A global unicast address has the following format:

Global routing prefix

Used to identify a specific customer site. The size of the field is 48 bits and allows an ISP to create multiple levels of addressing hierarchy within the network to both organize addressing and routing for downstream ISPs and identify sites.

Subnet ID

Used by an individual organization to identify subnets within its site. The organization can use these 16 bits to create 65 536 subnets or multiple levels of addressing hierarchy.

Interface ID

Indicates the interface on a specific subnet. The size of this field is 64 bits.

3 bits	45 bits	16 bits	64 bits
001	global routing prefix	subnet ID	interface ID

Figure 3. Global unicast address format

Local use addresses

There are two types of local-use unicast addresses defined:

- Link-local
- Site-local (has been deprecated)

The link-local address is for use on a single link.

Note: Site-local addresses were designed to use private address prefixes that could be used within a site without the need for a global prefix. The IETF has deprecated the special treatment given to the site-local prefix due to numerous problems in the actual use and deployment of site-local addresses. An IPv6 address constructed using a site-local prefix is now treated as global unicast address. The site-local prefix can be reassigned for other use by future IETF standards action.

Link-local addresses: Link-local addresses have the following format:

10 bits	54 bits	64 bits
1111111010	0	interface ID

Figure 4. Link-local address format

Requirement: A link-local address is required on each physical interface.

Link-local addresses are designed to be used for addressing on a single link for purposes such as automatic address configuration, neighbor discovery, or in the absence of routers. It also can be used to communicate with other nodes on the same link. A link-local address is automatically assigned.

Routers do not forward any packets with link-local source or destination addresses to other links.




Figure 5. Link-local scope zones

Figure 5 depicts two separate link-local scope zones. More than one interface can be connected to the same link for fault tolerance or extra bandwidth. Some nodes might allow the same link-local zone index to be assigned to each interface connected to the same physical link, while others might assign a unique link-local zone index to each interface even when more than one interface is connected to the same physical link. z/OS Communications Server takes the latter approach, assigning a unique link-local zone index to each physical interface.

Loopback address

The unicast address $0:0:0:0:0:0:1$ is called the loopback address. It cannot be assigned to any physical interface. It can be thought of as a link-local unicast address assigned to a virtual interface (typically called the loopback interface) that allows local applications to send messages to each other.

Restriction: The loopback address cannot be used as the source address in IPv6 packets that are sent outside of a node. An IPv6 packet with a destination address of loopback cannot be sent outside of a node and be forwarded by an IPv6 router. A packet received on an interface with destination address of loopback is dropped.

Unspecified address

The address $0:0:0:0:0:0:0$ is called the unspecified address. It is not assigned to any node. It indicates the absence of an address. One example of its use is in the Source Address field of any IPv6 packets sent by an initializing host before it has learned its own address.

Restriction: The unspecified address cannot be used as the destination address of IPv6 packets or in IPv6 routing headers. An IPv6 packet with a source address of unspecified cannot be forwarded by an IPv6 router.

IPv4-mapped IPv6 addresses

These addresses hold an embedded global IPv4 address. They are used to represent the addresses of IPv4 nodes as IPv6 addresses to applications that are enabled for IPv6 and are using AF_INET6 sockets. This allows IPv6-enabled applications to always deal with IP addresses in IPv6 format regardless of whether

the TCP/IP communications are occurring over IPv4 or IPv6 networks. The dual-mode TCP/IP stack performs the transformation of the IPv4-mapped addresses to and from native IPv4 format. IPv4-mapped addresses have the following format:

80 bits	16	32 bits
0000.....0000	FFFF	IPv4 address

Figure 6. IPv4-mapped IPv6 address

For example:

::FFFF:129.144.52.38

IPv6 interface identifiers

Interface identifiers in IPv6 unicast addresses are used to identify interfaces on a link. They are required to be unique on that link. In some cases, an interface's identifier is derived directly from that interface's link-layer address. z/OS Communications Server does not allow two links to have the same local address. Some implementations might allow the same interface identifier to be used on multiple interfaces on a single node, as long as they are attached to different links.

z/OS Communications Server allows the interface identifier to be generated (the default) or manually configured. When the interface ID is generated, then z/OS builds the interface ID when the interface becomes active based on the interface type as follows:

1. OSA-Express QDIO
2. OSA-Express returns the MAC address and a unique instance value during the start of an interface.
3. z/OS builds the interface identifier by inserting the unique instance value into the middle of the MAC address. This ensures that when multiple stacks share an OSA, each stack gets a unique interface ID. If a virtual MAC address is configured for this interface, then z/OS instead inserts the value 'FFFE'x into the middle of the MAC address.
4. HiperSockets

For HiperSockets interfaces, the interface ID generation works the same as for OSA-Express QDIO except that the HiperSockets device returns a 48-bit value that is unique for the HiperSockets CHPID rather than a MAC address. This ensures that when multiple stacks share a HiperSockets CHPID, each stack gets a unique interface ID.

5. MPCPTP6

For MPCPTP6 interfaces, z/OS randomly generates an interface ID.

24bits	16bits	24bits
MAC addr (bytes 1-3)	instance value	MAC addr (bytes 4-6)

Figure 7. OSA-Express QDIO interface ID format

A node can choose to use a different algorithm available for generation of interface identifiers for IPv6 addresses on a different platform.

Multicast IPv6 Addresses

An IPv6 multicast address is an identifier for a group of interfaces (typically on different nodes). It is identified with a prefix of 11111111 or FF in hexadecimal notation. It provides a way of sending packets to multiple destinations. An interface can belong to any number of multicast groups.

Multicast address format

Binary 11111111 at the start of the address identifies the address as being a multicast address. Multicast addresses have the following format:

8	4	4	112 bits
11111111	flgs	scope	group ID

Figure 8. Multicast address format

flgs is a set of 4 flags:

0	0	0	T
---	---	---	---

Figure 9. Flags in multicast address

- The 3 high-order flags are reserved, and must be initialized to 0.
- T = 0 indicates a permanently-assigned (well-known) multicast address, assigned by the Internet Assigned Number Authority (IANA).
- T = 1 indicates a non-permanently assigned (transient) multicast address.

Scope is a 4-bit multicast scope value used to limit the scope of the multicast group. Group ID identifies the multicast group, either permanent or transient, within the given scope.

Multicast scope

The scope field indicates the scope of the IPv6 internetwork for which the multicast traffic is intended. The size of this field is 4 bits. In addition to information provided by multicast routing protocols, routers use multicast scope to determine whether multicast traffic can be forwarded. For multicast addresses there are 14 possible scopes (some are still unassigned), ranging from interface-local to global (including both link-local and site-local).

Table 4 lists the defined values for the scope field:

Table 4. Multicast scope field values

Value	Scope
0	Reserved
1	Interface-local scope (same node)

Table 4. Multicast scope field values (continued)

Value	Scope
2	Link-local scope (same link)
3	Subnet-local scope
4	Admin-local scope
5	Site-local scope (same site)
8	Organization-local scope
E	Global scope
F	Reserved
Note: All other scope field values are currently undefined.	

For example, traffic with the multicast address of FF02::2 has a link-local scope. An IPv6 router never forwards this type of traffic beyond the local link.

Interface-local

The interface-local scope spans a single interface only. A multicast address of interface-local scope is useful only for loopback delivery of multicasts within a node, for example, as a form of interprocess communication within a computer. Unlike the unicast loopback address, interface-local multicast addresses can be joined on any interface.

Link-local

Link-local addresses are used by nodes when communicating with neighboring nodes on the same link. The scope of the link-local address is the local link.

Subnet-local

Subnet-local scope is given a different and larger value than link-local to enable possible support for subnets that span multiple links.

Admin-local

Admin-local scope is the smallest scope that must be administratively configured, that is, not automatically derived from physical connectivity or other, non-multicast-related configuration.

Site-local

The scope of a site-local address is the site or organization internetwork. Addresses must remain within their scope. A router must not forward packets outside of its scope.

Guideline: Site-local has been deprecated.

Organization-local

This scope is intended to span multiple sites belonging to a single organization.

Global

Global scope is used for uniquely identifying interfaces anywhere in the Internet.

Multicast groups

Group ID identifies the multicast group, either permanent or transient, within the given scope. The size of this field is 112 bits. Permanently assigned groups can use the group ID with any scope value and still refer to the same group. Transient assigned groups can use the group ID in different scopes to refer to different

groups. Multicast addresses from FF01:: through FF0F:: are reserved, well-known addresses. Use of these group IDs for any other scope values, with the T flag equal to 0, is not allowed.

All-nodes multicast groups: These groups identify all IPv6 nodes within a given scope. Defined groups include the following:

- Interface-local all-nodes group (FF01::1)
- Link-local all-nodes group (FF02::1)

All-routers multicast groups: These groups identify all IPv6 routers within a given scope. Defined groups include the following:

- Interface-local all-routers group (FF01::2)
- Link-local all-routers group (FF02::2)
- Site-local all-routers group (FF05::2)

Solicited-node multicast group: For each unicast address which is assigned to an interface, the associated solicited-node multicast group is joined on that interface. The solicited-node multicast address facilitates the efficient querying of network nodes during address resolution.

Anycast IPv6 Addresses

An IPv6 anycast address is an identifier for a set of interfaces (typically belonging to different nodes). A packet sent to an anycast address is delivered to one of the interfaces identified by that address (the nearest interface), according to the routing protocols' measure of distance. It uses the same formats as a unicast address, so one cannot differentiate between a unicast and an anycast address simply by examining the address. Instead, anycast addresses are defined administratively.

Typical IPv6 addresses assigned to a node

An IPv6 host is required to recognize the following addresses as identifying itself:

- Link-local address for each active IPv6 physical interface (cannot be manually defined)
- Assigned unicast addresses (autoconfigured or manually defined)
- IPv6 loopback address (::1)
- All-nodes multicast address (interface-local and link-local)
- Solicited node multicast addresses for each of its assigned unicast and anycast addresses
- Multicast addresses of all other groups to which the host belongs

IPv6 address states

An address state defines and controls how other algorithms work with a particular address.

Tentative

An address whose uniqueness on a link is being verified, prior to its assignment to an interface. A tentative address is not considered assigned to an interface in the usual sense. An interface discards received packets addressed to a tentative address, unless those packets are related to Duplicate Address Detection (DAD). For more information on DAD, refer to “Duplicate Address Detection (DAD)” on page 32.

Deprecated

An address assigned to an interface whose use is discouraged, but not forbidden. Packets sent from or to deprecated addresses are delivered as expected. A deprecated address continues to be used as a source address in existing communications where switching to a preferred address would be disruptive.

Preferred

An address assigned to an interface whose use is unrestricted. Preferred addresses can be used as the source or destination address of packets sent from or to the interface, respectively.

Unavailable

An unavailable address is one that is not yet assigned to the interface.

Chapter 3. IPv6 protocol

This chapter describes the IPv6 protocol implementation and contains the following sections:

- “Extension headers”
- “Fragmentation in an IPv6 network”
- “Path MTU discovery” on page 22
- “IPv6 routing” on page 22
- “ICMPv6” on page 26
- “Multicasting” on page 27
- “Neighbor discovery (ND)” on page 28
- “Assigning IP addresses to interfaces” on page 33
- “Default address selection” on page 36
- “Migration and coexistence” on page 39
- “Considerations for configuring z/OS for IPv6” on page 44
- “IPv6 stack support” on page 44
- “INET considerations” on page 45
- “Common INET considerations” on page 46

Guideline: You should be familiar with the IPv6 protocol in general.

Extension headers

In IPv6, IP-layer options within a packet are encapsulated in independent headers called extension headers. In contrast, IPv4 options are contained in the IP header itself.

| **Restriction:** Not all IPv6 extension headers are supported in z/OS Communications Server.

| The z/OS TCP/IP stack supports receipt of the following extension headers:

- Routing
- Fragmentation
- Hop-by-hop option
- Destination option
- Authentication (AH)
- Encapsulating Security Payload (ESP)

Fragmentation in an IPv6 network

Fragmentation is used by a source to send a packet larger than would fit in the path MTU to its destination. In order to send packets larger than the link minimum of 1280 bytes, a node must support determination of the minimum supported MTU along the path between the source and destination. This is accomplished by Path MTU discovery. For more information about path discovery, see “Path MTU discovery” on page 22.

The IPv6 IP header does not contain information about fragments. The fragmentation extension header carries this information. z/OS Communications Server allows for 2048 active IPv6 reassemblies in progress at any given time. z/OS Communications Server reassembly timeout for IPv6 reassemblies is 60 seconds. These two values are not configurable.

Fragmentation and UDP/RAW

Intermediate routers cannot fragment packets and UDP/RAW transports do not perform retransmission. To attempt to ensure that a UDP/RAW packet is not dropped due to fragmentation, one of the following conditions can occur:

- z/OS Communications Server always sends the packet using the minimum MTU (1280) unless the MTU for the destination is learned from an ICMPv6 Packet Too Big message.
- An application sends a packet using the IPV6_DONTFRAG socket option.

For example, a situation can occur where the MTU was learned by way of Path MTU discovery. In that case, the network topology changes, reducing the MTU to this particular destination. UDP/RAW sends with the original learned MTU and receives a Packet Too Big message. In this case, the packet is dropped, but subsequent sends learn the changed MTU and send with the appropriate size.

Path MTU discovery

When one IPv6 node has a large amount of data to send to another node, the data is transmitted in a series of IPv6 packets. It is preferable that these packets be of the largest size that can successfully traverse the path from the source node to the destination node. This packet size is referred to as the Path MTU (PMTU), and it is equal to the minimum link MTU of all the links in a path. IPv6 provides PMTU discovery as a standard mechanism for a node to discover the PMTU of an arbitrary path.

For IPv6, intermediate routers cannot fragment packets. An implementation must either support Path MTU discovery or send using IPv6 minimum link MTU. z/OS Communications Server supports path MTU discovery.

Path MTU discovery supports multicast as well as unicast destinations. When PMTU information is learned, it is cached for a period of time and then deleted in order to learn of increases in the MTU value.

IPv6 routing

Both replaceable and non-replaceable IPv6 static routes are supported by using BEGINROUTES profile statements.

Restriction: The GATEWAY statement in the TCP/IP profile does not support IPv6 static routes.

Dynamic routes for IPv6 are learned by:

- Router discovery
- ICMPv6 redirects
- Dynamic routing protocols

Replaceable static routes can be replaced by dynamic routes. If a replaceable static route is replaced by a dynamic route, and that dynamic route is later deleted, the replaceable static route is re-added.

Router discovery

Hosts can learn the network prefixes for all directly attached links from the router advertisements received from their routers. To determine if that host is on a directly attached link or on a remote link, check to see if another host's IPv6 address is constructed from a network prefix of one of the directly attached links. If it is on a directly attached link, data can be sent directly to that host without going through a router. Otherwise, it must be sent through some router using a default route that can also be learned from router advertisements.

Router advertisements are not a replacement for dynamic routing protocols such as IPv6 OSPF and IPv6 RIP. If a host is not using a dynamic routing protocol, some limitations apply.

If the host has multiple interfaces attached to more than one link, it must decide which interface to send the packet over. If there are multiple routers on the link attached to the interface, it must decide to which router it should send the packet. To make these decisions, it needs a route in its routing table. Without a dynamic routing protocol, the host uses the default route when selecting which router on which interface to send the packet. This behavior might not produce the desired results.

In the case where there are multiple default routers on the same physical link, the host might select a non-optimal router. This might not be a serious problem, because that router can send an ICMP Redirect, allowing the host to update its routing table and send subsequent packets to the correct router. If default routers are on multiple physical links, that is more serious. A router on one link is not able to redirect the host to use a different physical link. If the selected router cannot reach the destination, attempts to send data fails, even if the destination could be reached by a default router on another physical link. To resolve these limitations when not using a dynamic routing protocol, static routes might be needed to direct the traffic over the best interface using the appropriate router.

If a dynamic routing protocol is not used, routes to VIPAs cannot be advertised. For this reason, use a network prefix defined as being on-link for the interfaces which are associated with the VIPA. In this way, routers and hosts perceive that the VIPA is on a physical interface and sends Neighbor Discovery messages (the IPv6 equivalent of an ARP request) to get the MAC address of the interface. This is not the typical way to set up VIPAs and is not the typical way to set up VIPAs if a dynamic routing protocol is being used. Normally, they can be associated with interfaces on different LANs. But without a dynamic routing protocol, you must either take the suggested approach or define static routes at all routers on the same links as the z/OS system.

ICMPv6 redirects

ICMPv6 redirects replace static routes regardless of whether or not they are replaceable. Use the IGNOREREDIRECT keyword on the IPCONFIG6 statement in the TCP/IP profile to prevent the stack from adding routes learned by ICMPv6 redirects.

Rule: ICMPv6 redirects are always ignored when an IPv6 dynamic routing protocol is being used.

Dynamic routing protocols

The z/OS Communications Server OMPROUTE routing daemon supports the IPv6 OSPF and IPv6 RIP dynamic routing protocols. A host using one of these protocols can learn, from adjacent routers that are also using that protocol, the network prefixes and host addresses that can be reached.

IPv6 OSPF and IPv6 RIP can be used together with router discovery in the same network. IPv6 OSPF allows the host to learn the network prefixes and host addresses that can be reached indirectly by way of adjacent IPv6 OSPF routers (including default routes), as well as the network prefixes that can be reached directly on attached links in the IPv6 OSPF domain. IPv6 RIP allows the host to learn the network prefixes and host addresses that can be reached indirectly by way of adjacent IPv6 RIP routers (including default routes). Router discovery allows the host to learn default routes by way of adjacent routers participating in router discovery, as well as the network prefixes that can be reached directly on attached links.

In addition, the network prefixes that can be reached directly on attached links can be manually configured using the `Prefix` keyword on the `IPv6_Interface`, `IPv6 OSPF_Interface`, or `IPv6 RIP_Interface` statements in the OMPROUTE configuration file. When IPv6 OSPF or IPv6 RIP is used together with router discovery, certain routes can be learned from both methods. These routes consist of:

Default routes

Learned from both methods if adjacent routers are advertising themselves as default routers using both IPv6 OSPF or IPv6 RIP and router discovery. When this situation occurs, the default routes learned from IPv6 OSPF or IPv6 RIP takes precedence and generates the default routes in the TCPIP stack's IPv6 route table. Any default routes learned from router discovery are ignored as long as the default routes learned from IPv6 OSPF or IPv6 RIP exist.

Prefix routes

Learned from both router discovery and OMPROUTE under each of the following conditions:

- A router is advertising by way of router discovery that the prefix is on-link and the prefix is also manually configured to OMPROUTE using the `Prefix` keyword on an `IPv6_Interface`, `IPv6 OSPF_Interface`, or `IPv6 RIP_Interface` configuration statement.

Guideline: Use the `Prefix` keyword only when the prefix is not learned dynamically (using router discovery or a dynamic routing protocol).

For example, if there is a need to supplement the list of prefixes being advertised as on-link by the routers. If the same prefix is configured using the `Prefix` keyword and learned from router discovery, the route in the TCPIP stack's route table is the route added by OMPROUTE as a result of the `Prefix` keyword. Any route for the same prefix that is learned from router discovery is ignored as long as the OMPROUTE route exists.

Restriction: Prefixes learned from only OMPROUTE are not used for address autoconfiguration. If a prefix is learned from both OMPROUTE and router discovery, it can still be used for autoconfiguration even though the route learned from OMPROUTE is the one in the TCPIP stack route table.

- A router is advertising by way of router discovery that the prefix is on-link and a router is also advertising by way of IPv6 OSPF that the prefix is on-link.

In this case, the route in the TCPIP stack route table is the route added by OMPROUTE as a result of the information received by way of IPv6 OSPF. Any route for the same prefix that is learned from router discovery is ignored as long as the OMPROUTE route exists. As in the previous condition, the prefix learned from router discovery can still be used for address autoconfiguration.

- A router is advertising by way of router discovery that the prefix is on-link and it is also learned, by way of IPv6 OSPF or IPv6 RIP, that the prefix can be reached by way of an adjacent router.

In this case, the route in the TCPIP stack route table is the route added as the result of router discovery. This occurs because the router discovery information indicates that the prefix resides on a directly attached link, while the IPv6 OSPF or IPv6 RIP information indicates that the prefix can be reached indirectly, by way of the router from which the IPv6 OSPF or IPv6 RIP information was received. Any route for the prefix that is learned from IPv6 OSPF or IPv6 RIP is ignored as long as the router discovery route exists.

Tip for IPv6 OSPF routing protocol addressing conventions

IPv6 OSPF is based on IPv4 OSPF and has many similar concepts and controls. The primary difference between IPv6 OSPF and IPv4 OSPF is that for IPv6 OSPF, IP addresses are not used to communicate topology information. For example, in IPv4 OSPF, an interface is referred to by its IPv4 home address, but in IPv6 OSPF an interface is not referred to by any of its IPv6 home addresses. Instead, it is referred to by an integer interface ID. Similarly, IPv6 OSPF router IDs are not IPv6 home addresses; they are 32-bit integers written in IPv4-style dotted-decimal notation. Area IDs in IPv6 OSPF are also 32-bit integers written in IPv4-style dotted-decimal notation.

Guideline: Even though router IDs and area IDs in IPv6 OSPF are expressed similarly to the IPv4 equivalents, they are not the same constants. A router can have an IPv6 router ID which is different from its IPv4 router ID. If both IPv4 and IPv6 OSPF are running simultaneously, the area topology of each IP version can be completely different, with different area numbers and hierarchy.

Authentication with the IPv6 OSPF routing protocol

IPv4 OSPF includes authentication as part of the OSPF protocol. OMPROUTE supports both password authentication and MD5 cryptographic authentication for IPv4 OSPF. For IPv6 OSPF, authentication has been removed from OSPF itself. Instead, IPv6 OSPF relies on IPSec to ensure integrity and authentication of routing exchanges. As a result, OMPROUTE does not include any explicit authentication support, but instead relies on the underlying support provided by the z/OS TCP/IP stack.

To use IPSec to authenticate IPv6 OSPF routing exchanges on a link over which OMPROUTE establishes adjacencies, you must create a single manual security association (SA) for all traffic on that link, with corresponding filter definitions to permit the OSPF traffic. Use the interface SECCLASS to define different security associations for different links. This procedure is described in the IP security information in *z/OS Communications Server: IP Configuration Guide*.

Considerations for route selection

Route precedence is as follows:

- Host route to the destination.
- Route for a prefix of the destination. If there are routes to multiple prefixes of the destination, the route with the most specific prefix is chosen.
- Default route.

For IPv4, the concept exists of a special default multicast route with a destination of 224.0.0.0 and a netmask of 255.255.255.255. For IPv6, there is no special default multicast route. Because all IPv6 multicast addresses start with FF, the following prefix route serves the same function as the default multicast route:

```
destination = FF00::/8
```

Considerations for multipath routes

Multiple routes to the same destination are considered multipath routes. Multipath routes can be used for load balancing. Multipath route support for IPv6 is identical to multipath route support for IPv4. Define the MULTIPATH keyword on the IPCONFIG6 statement to control whether multiple routes are selected.

Guideline: If MULTIPATH is not enabled, the first active route added is selected.

When using a route that belongs to a multipath group, the MTU that is used is the minimum MTU of all routes in the multipath group.

How does a VARY TCPIP,,OBEYFILE command affect routes?

When a VARY TCPIP,,OBEYFILE command is issued and the profile contains a BEGINROUTES block, the following occurs:

- All static routes (both replaceable and non-replaceable) are deleted and replaced by any static routes defined in the BEGINROUTES block.
- All routes learned by way of ICMPv6 redirects are deleted.
- Routes learned by way of router advertisements or a dynamic routing daemon are not affected by the processing of the VARY TCPIP,,OBEYFILE command, with the following exception:
 - If the profile data set specified on the VARY TCPIP,,OBEYFILE command contains a non-replaceable static route to the same destination for which a route exists that was learned by way of router advertisements or a dynamic routing daemon, the existing route is deleted and replaced by the non-replaceable static route.

ICMPv6

The IP protocol moves data from one node to another. In order for IP to perform this task successfully, there are many other functions that need to be carried out as well, such as the following:

- Error reporting
- Route discovery
- Diagnostics
- Among others

In IPv6, all these tasks are carried out by the Internet Control Message Protocol (ICMPv6).

In addition, ICMPv6 provides a framework for Multicast Listener Discovery (MLD) and Neighbor Discovery (ND), which carry out the tasks of conveying multicast group membership information (the equivalent of the IGMP protocol in IPv4) and address resolution (performed by ARP in IPv4).

The following are types of ICMPv6 messages:

Error Report errors in the forwarding or delivery of IPv6 packets.

Informational

Provide diagnostic functions and additional host functionality such as MLD and ND.

The following ICMPv6 messages are supported:

- Destination unreachable
- Packet too big
- Time exceeded (hop limit exceeded)
- Echo request/reply
- Parameter problem
- Multicasting messages:
 - Group membership query
 - Report
 - Done
- Neighbor discovery:
 - Router solicitation and advertisement
 - Neighbor solicitation and advertisement
 - Redirect

Guideline: Not all ICMPv4 messages have equivalents in ICMPv6.

Multicasting

In early IP networks, a packet could be sent to either a single device (unicast) or to all devices (broadcast). A single transmission destined for a group of devices was not possible.

IPv6 uses multicast for those purposes for which IPv4 used broadcast; consequently, IPv6 does not support broadcast.

Applications can use multicast transmissions to enable efficient communication between groups of devices. Data is transmitted to a single multicast IP address and received by any device that needs to obtain the transmission.

Multicast Listener Discovery (MLD)

An IPv6 router uses MLD protocol to discover the following:

- The presence of multicast listeners (nodes wanting to receive multicast packets) on its directly attached links
- Which multicast addresses are of interest to those listeners

This information is provided to whichever multicast routing protocol is being used by the router. This ensures that multicast packets are delivered to all links where there are interested receivers. MLD is derived from IGMPv2.

Guideline: One important difference is that MLD uses ICMPv6 message types, rather than IGMP message types.

MLD has a router function and a listener function. The router function discovers the presence of multicast listeners and ensures delivery of multicast packets to listeners. The listener function informs routers when it starts and stops listening for a multicast address and responds to queries about multicast addresses. z/OS Communications Server V1R4 and above implement the listener function.

When a listener starts listening for a multicast address on an interface, it sends an MLD report message for that address on that interface.

When a listener stops listening for a multicast address on an interface, it sends a single MLD done message.

An MLD query message is sent by a router to query listeners about multicast addresses. A specific query is sent to listeners for a specific multicast address on a receiving interface. A general query is sent to listeners for all multicast addresses on a receiving interface. These query messages contain a maximum response delay (MRD) that causes listeners to delay report messages and not send them if another listener reports first. If no reports for the address are received from the link after the response delay of the last query has passed, the routers on the link assume that the address no longer has any listeners there; the address is therefore deleted from the list and its disappearance is made known to the multicast routing component.

If you configure IP security for IPv6, refer to special considerations in the *z/OS Communications Server: IP Configuration Guide* for information about filter rules for MLD packets.

Neighbor discovery (ND)

Neighbor discovery (ND) is an ICMPv6 function that enables a node to identify other hosts and routers on its links. It corresponds to a combination of IPv4 protocols:

- ARP
- ICMP Router Discovery
- ICMP Redirect

It maintains routes, MTU, retransmit times, reachability time, and prefix information based on information received from the routers. ND uses Duplicate Address Detection (DAD) to verify the host's home addresses are unique on the LAN.

ND uses Address Resolution to determine the link-layer addresses for neighbors on the LAN and Reachability Detection to determine neighbor reachability.

If you configure IP security for IPv6, refer to special considerations in the *z/OS Communications Server: IP Configuration Guide* for information about filter rules for neighbor discovery packets.

Router advertisements

Router advertisements are sent by routers to announce their availability. z/OS Communications Server receives router advertisements, but it does not originate them.

If the router advertisement indicates that the sending router should be used as a default router, a neighbor cache entry is created or updated for the sending router, and the following occurs:

- IPv6 dynamic default route is added (if not added by a previous advertisement)
- Next hop of default route is the advertisement's source address
- Interface of default route is the interface on which the advertisement was received
- Length of time that route remains valid is set or reset using the Lifetime value from the advertisement

A dynamic default route is not added due to the received router advertisement if the following exists:

- A non-replaceable static default route
- An IPv6 OSPF default route
- An IPv6 RIP default route

If a replaceable static default route exists, the dynamic default route is added due to the received router advertisement, replacing the replaceable route. The replaceable static default route is reinstated if the dynamic default route is later removed.

If the router advertisement indicates that the sending router should not be used as a default router, the following occurs:

- If an IPv6 dynamic default route exists with the advertisement's source as its next hop and the receiving interface as its interface and that route was added due to a received router advertisement (for example, not due to IPv6 OSPF or IPv6 RIP), it is deleted.
- Any IPv6 dynamic indirect routes with the advertisement's source as its next hop and the receiving interface as its interface are deleted. Exceptions to this are routes that were added due to a dynamic routing protocol such as IPv6 OSPF or IPv6 RIP.
- A neighbor cache entry is created or updated for the sending router. The neighbor cache entry contains data from the router advertisement such as the following:
 - Indication that neighbor is a router
 - Indication that neighbor is not a default router
 - Link-local and link-layer address of neighbor

A router advertisement can contain Prefix Information Options. These options inform nodes of additional specific routes that are available to them, and indicate prefixes for autoconfiguring addresses. A Prefix Information Option contains the on-link and autonomous flags.

The on-link flag, when set, indicates that on-link processing needs to be performed for the prefix on the shared link. When a prefix is on-link, the addresses in that prefix can be reached on that link without going through a router. The autonomous flag, when set, indicates that autoconfigure processing needs to be performed for the prefix on the shared link. A Prefix Information Option can have just the on-link flag set, just the autonomous flag set, or both flags set.

The sending router indicates that a prefix is on-link by setting the on-link flag and specifying a nonzero Valid Lifetime value for the prefix. If the Prefix Information Option indicates that the prefix is on-link, the following occurs:

- An IPv6 dynamic direct route is added (if not added by a previous advertisement)
- The destination of the route is the prefix being processed
- The interface of the route is the interface on which the advertisement was received
- The length of time that route remains valid is set or reset using the Valid Lifetime value from the Prefix Information Option

If a non-replaceable static direct route exists to this prefix or if a direct route to the prefix was added by OMPROUTE (due to the PREFIX parameter being specified on the IPV6_INTERFACE, IPV6 OSPF_INTERFACE, or IPV6 RIP_INTERFACE statement in the OMPROUTE configuration file or due to a router advertising by way of IPv6 OSPF that the prefix is on-link), then the dynamic direct route is not added. If a replaceable static direct route exists to this prefix, the dynamic direct route is added, replacing the replaceable route. The replaceable static direct route is reinstated if the dynamic direct route is later removed.

The sending router can indicate that a prefix is no longer on-link by setting the on-link flag and specifying a zero Valid Lifetime value for the prefix. In this case, if an IPv6 dynamic direct route exists with the prefix being processed as its destination and the receiving interface as its interface, and that route was added due to a received router advertisement (for example, not added by OMPROUTE), it is deleted.

The sending router can indicate that a prefix is to be used for address autoconfiguration by setting the autonomous flag and specifying a nonzero Valid Lifetime value for the prefix. If the Prefix Information Option indicates that the prefix should be used for address autoconfiguration, the following occurs:

- An IPv6 home address is added to the receiving interface for the autoconfigured address (if not added by a previous advertisement)
- An IPv6 implicit route is added for the receiving interface and the autoconfigured address (if not added by a previous advertisement)
- The length of time that home address and implicit route remain valid is set or reset using Valid Lifetime value from the Prefix Information Option
- The length of time that home address remains preferred (not deprecated) is set or reset using the Preferred Lifetime value from the Prefix Information Option

Restriction: Prefixes learned solely by using the Prefix parameter on the OMPROUTE IPV6_INTERFACE, IPV6 OSPF_INTERFACE, or IPV6 RIP_INTERFACE statement is never used for autoconfiguration.

If addresses are manually configured for an IPv6 interface by way of the INTERFACE statement, autoconfiguration of addresses for that interface is disabled. If a prefix is not 64 bits in length, it is not used for autoconfiguration of addresses. Unlike the prefix route and default route, the implicit route and home address cannot be deleted immediately. They must age out. If the Valid Lifetime value is set to infinity, the implicit route and home address do not time out. For more information about autoconfiguration, see “Stateless address autoconfiguration” on page 33.

Route timeouts

The valid lifetime for each type of route is updated (extending the life of the route) by the periodic receipt of router advertisements as long as the sending router is available and is not reconfigured relative to its defined prefixes or default router status.

When a Prefix Information Option contains a Valid Lifetime value of infinity, the associated implicit or prefix route is considered permanent and does not age unless a future Prefix Information Option for the prefix contains a non-infinity Valid Lifetime value.

Expiration of the valid lifetime for a default route is immediate if a future Router Advertisement indicates that the sending router is no longer a default router. Expiration of the valid lifetime for a prefix route is immediate if a future Prefix Information Option for the prefix contains a zero Valid Lifetime value. Expiration of the valid lifetime for an implicit route cannot be made immediate because the minimum lifetime allowed is two hours. It must age out naturally.

VARY TCPIP,,OBEYFILE command rules

Rules: Observe the following rules for the VARY TCPIP,,OBEYFILE command:

- If a non-replaceable static route in the profile data set specified on the VARY TCPIP,,OBEYFILE command has the same destination as an existing route that was added due to a received Router Advertisement, the existing route is replaced by the non-replaceable static route.
- If the profile data set specified on the VARY TCPIP,,OBEYFILE command specifies a manually configured home address for an interface that already has autoconfigured addresses, the autoconfigured addresses are deleted along with their associated implicit routes.

With the exception of the two preceding rules, all autoconfigured home addresses and routes added due to received Router Advertisements are maintained through VARY TCPIP,,OBEYFILE command processing.

Redirect processing

A node can receive a Redirect message from an on-link router if the router determines that the destination is on-link or if there is a better first-hop router for the given destination. z/OS Communications Server can be configured to ignore the IPv6 Redirects sent by routers by defining the IGNOREREDIRECT keyword on the IPCONFIG6 statement. In addition, IPv6 Redirects are ignored if the IPv6 OSPF or IPv6 RIP protocol of the OMPROUTE routing daemon is being used. If processing of Redirect messages is enabled, z/OS Communications Server begins using the new first-hop information which is identified in the Redirect message. A router must use its link-local address as the source address in Redirects that it originates. A received Redirect is only processed if the current route to the destination in the IPv6 route table has the source address of the Redirect as its next hop. Therefore, if Redirects are to be accepted, all static indirect routes must be configured using the next-hop router's link-local address. If the previous route to the destination was a host route, it is deleted from the route table to keep it from being used by Multipath processing.

If Redirect processing is disabled, z/OS Communications Server silently discards the Redirect message.

Duplicate Address Detection (DAD)

DAD is used to verify that an IPv6 home address is unique on the LAN before assigning the address to a physical interface (for example, QDIO). z/OS Communications Server responds to other nodes doing DAD for IP addresses assigned to the interface. DAD is not done for VIPAs or loopback addresses. DAD for local addresses is performed for physical interfaces when one of the following occurs:

- The interface is started (the autoconfigured link-local address and manually configured addresses/prefixes are checked).
- A VARY TCPIP,,OBEYFILE command is issued for a profile data set containing an INTERFACE ADDADDR for an already active interface.
- A Router Advertisement containing new prefix information and the autonomous bit set is received on an interface enabled for stateless autoconfiguration.

To disable DAD checking, specify DUPADDRDET 0 on the INTERFACE statement.

Duplicate Address Detection processing involves the following steps:

1. The host joins a link-local all-nodes multicast group at interface start processing.
2. The host joins a solicited-node group for the local address.
3. A neighbor solicitation is sent to the solicited-node multicast address with the tentative address for which DAD is being performed.
4. The host waits for a neighbor response (neighbor advertisement or neighbor solicitation) on the interface.
5. If no neighbor response is received within the specified retransmit time, the address is considered unique on the LAN.
6. If a neighbor response is received within the specified time, the address is not unique. The host leaves the solicited-node multicast group, issues a Duplicated Address Detected console message, and marks the address unavailable due to a duplicate address.

Unless DAD is disabled, the address is not considered assigned to an interface until DAD is successfully completed for the local address. Packets can be received for the all-nodes or solicited-node multicast groups, but there is no response because the address is not yet assigned to the interface. If the local address is a manually configured address, the addresses are displayed in a Netstat Home/-h report as Unavailable (if the interface has not been started or if DAD failed).

In situations where DAD is not done for the IPv6 home address (by specifying DUPADDRDET 0 on the INTERFACE statement or if it is a VIPA), the z/OS Communications Server host still responds if another node is doing DAD for an IPv6 address assigned to the interface or for IPv6 VIPAs when the interface is assigned to handle VIPAs. Note that responses are not sent for loopback addresses.

Address resolution

Address resolution in IPv6 is similar to ARP processing in IPv4, except ICMP neighbor solicitations, neighbor advertisements, router redirects, and router advertisements are used to obtain the link-layer (MAC) address. The host sends a neighbor solicitation to a solicited-node multicast address. It waits for a response for a period of time (retransmit time). If one is received, then the link-layer address contained in the neighbor advertisement is cached and any queued packets are sent to the address. If there is no response, the host repeats this process up to three times before it declares a neighbor unreachable.

A neighbor cache entry can also be built when a neighbor solicitation for a local address is received and the solicitation contains the sender's link-layer address (and the source address is not the unspecified address, that is, the sender is not performing DAD). The neighbor cache entry is built if it does not exist based on the assumption that a packet is soon sent to this neighbor. Building the cache entry reduces the overhead of having to perform the task of address resolution for the neighbor at a later time.

Issue the Netstat ND/-n command to display information for a specific neighbor or all neighbor cache entries. It displays the neighbor link-layer address, state, whether the neighbor is a router or host, and if a router is a default router. The following are possible neighbor states:

Incomplete

Address resolution is in progress.

Reachable

Positive confirmation of reachability was received.

Stale An unsolicited neighbor discovery message has updated the link-layer address. Reachability is verified the next time the entry is used.

Delay More than reachable time has elapsed since last positive confirmation of reachability. Default reachable time is 30 seconds. It can be overridden by data provided by neighbor advertisements. A small delay is experienced before starting a probe of neighbor (upper layers can provide confirmation).

Probe Neighbor solicitations are sent to verify neighbor reachability.

Neighbor unreachability detection

Neighbor unreachability detection verifies that two-way communication with a neighbor node exists. The host sends a neighbor solicitation to a node and waits for a solicited neighbor advertisement. If a solicited neighbor advertisement is received, the node is considered reachable. If there is no response, the host can repeat this process before it declares a neighbor unreachable. If a neighbor is found to be unreachable, the neighbor cache entry is deleted.

Assigning IP addresses to interfaces

Stateless address autoconfiguration is always used to generate and assign a link-local address to a physical IPv6 interface. If it cannot assign a link-local address, interface activation fails. No other addresses are assigned to the interface (whether they are assigned using stateless address autoconfiguration or manual configuration) until a link-local address has been successfully assigned. Link-local addresses are not aged out.

Stateless address autoconfiguration

The larger address field of IPv6 solves a number of problems inherent in IPv4, but the size of the address itself might be a potential problem for the TCP/IP administrator. As a result, IPv6 has the capability to automatically assign an address to an interface at initialization time. By doing this, a network can become operational with minimal action on the part of the TCP/IP administrator. Stateless autoconfiguration is supported for a physical interface (for example, QDIO) in z/OS Communications Server if no manually configured addresses are defined on

the interface. Manual configuration of the host's local addresses is not required except for VIPA interfaces. Stateless address autoconfiguration consists of the following steps:

1. During system startup, the host obtains an interface token from the interface hardware to create an interface ID. It generates its own addresses using a combination of router advertised prefixes and interface IDs.
2. Duplicate address detection is performed for the address. If a duplicate is not detected or DAD is disabled for the interface (DUPADDRDET 0 specified on the INTERFACE statement), the local address is added.
3. A stateless autoconfigured address is deleted when its valid lifetime expires or when a manually defined address is added to the interface.

An IPv6 address generated using stateless address autoconfiguration has two timers associated with it: a preferred lifetime and a valid lifetime. Router Advertisements contain the valid lifetime and preferred lifetime for a prefix. An IPv6 address goes through two phases to gracefully handle the address expiration:

Preferred

Use is unrestricted.

Deprecated

In anticipation of the expiration of the leased period, use of the address is discouraged.

When the preferred lifetime expires, the address created from the prefix is deprecated. When the valid lifetime expires, the address created from the prefix is deleted and an operator message is issued.

Autoconfiguration considerations

Consider the following during autoconfiguration:

- A manually configured address/prefix on an interface disables stateless autoconfiguration for the interface.
- INTERFACE *name* DELADDR *addr/prefix* and INTERFACE *name* DEPRADDR *addr/prefix* profile statements, activated by way of the VARY TCPIP,,OBEYFILE command, are not valid for autoconfigured addresses.
- A VARY TCPIP,,OBEYFILE command whose profile contains ADDADDR or DELADDR INTERFACE statements can affect stateless autoconfiguration:
 - An INTERFACE *name* ADDADDR *addr/prefix* profile statement, activated by way of the VARY TCPIP,,OBEYFILE, results in stateless autoconfigured addresses on the interface to be deleted. Stateless autoconfiguration capability is disabled.
 - If the DELADDR removes the last manually configured address/prefix, stateless autoconfiguration is enabled and subsequent router advertisements can generate autoconfigured addresses.
- Autoconfigured addresses are not automatically added to DNS. Consider using VIPA addresses in conjunction with autoconfigured addresses.

IP address takeover following an interface failure

The TCP/IP stack in z/OS Communications Server provides transparent fault-tolerance for failed (or stopped) IPv6 interfaces, when the stack is configured with redundant connectivity onto a LAN. This support is provided by the z/OS Communications Server interface-takeover function and applies to the IPv6 IPAQENET6 interface type.

At device or interface startup time, TCP/IP dynamically learns of redundant connectivity onto the LAN, and uses this information to select suitable backups in the case of a future failure of the device/interface. This support makes use of neighbor discovery flows for IPv6 interfaces, so upon failure (or stop) of an interface, TCP/IP immediately notifies stations on the LAN that the original IPv6 address is now reachable by way of the backup's link-layer (MAC) address. Users targeting the original IP address see no outage due to the failure, and they are unaware that any failure occurred.

Because this support is built upon neighbor discovery flows, no dynamic routing protocol in the IP layer is required to achieve this fault tolerance. To enable this support, you must configure redundancy onto the LAN by defining and activating multiple INTERFACEs onto the LAN. Note that an IPv4 device cannot back up an IPv6 interface, or vice versa.

The interface-layer fault-tolerance can be used in conjunction with VIPA addresses, where applications can target the VIPA address, and any failure of the real LAN hardware is handled by the interface-takeover function. This differs from traditional VIPA usage, where dynamic routing protocols are required to route around true hardware failures.

How to get addresses for VIPAs

Rule: All VIPAs must be manually configured.

VIPA interfaces are always active. IPv6 VIPAs can be site-local or global. Link-local VIPAs are not allowed because link-local addresses are for use only on the associated LAN and there is no VIPA LAN.

To globally enable SOURCEVIPA for IPv6, configure the SOURCEVIPA keyword on the IPCONFIG6 statement. Then, to enable SOURCEVIPA for particular interfaces, use the SOURCEVIPAINTERFACE parameter on the INTERFACE statement for those interfaces. The SOURCEVIPAINTERFACE parameter allows for the specification of the interface name of the VIRTUAL6 interface whose addresses should be used as SOURCEVIPA addresses.

Unlike IPv4, where the source VIPA selected is based upon the ordering of the HOME list, IPv6 SOURCEVIPA uses the addresses configured on the VIPA INTERFACE statement referenced by the SOURCEVIPAINTERFACE keyword on the INTERFACE statement for the outbound interface. When that VIPA interface has multiple addresses configured, the default source address selection algorithm selects among them. For detailed information about the algorithm, see "Default source address selection" on page 38.

Guidelines: Observe the following VIPA guidelines:

- Use different prefixes for IPv6 static VIPAs and for the IPv6 addresses assigned to real interfaces.
- Having static VIPAs configured with different prefixes than real addresses reduces the likelihood of address collisions between the manually configured VIPAs and the autoconfigured addresses of the real interfaces. This is also necessary as Duplicate Address Detection (DAD) is not performed for VIPA addresses.
- See 64 for information regarding static VIPAs.

Default address selection

IPv6 addressing architecture allows multiple unicast addresses to be assigned to interfaces. These addresses might have different reachability scopes (link-local, site-local, or global). These addresses can also be preferred or deprecated. Privacy considerations have introduced the concepts of public addresses and temporary addresses. The mobility architecture introduces home addresses and care-of addresses. In addition, multihoming situations result in more addresses per node. For example, a node can have multiple interfaces, some of them tunnels or virtual interfaces, or a site can have multiple ISP attachments with a global prefix per ISP.

The end result is that IPv6 implementations are often faced with multiple possible source and destination addresses when initiating communication. It is preferred to have default algorithms, common across all implementations, for selecting source and destination addresses so that developers and administrators can reason about and predict the behavior of their systems.

Furthermore, dual-mode stack implementations, which support both IPv6 and IPv4, very often need to choose between IPv6 and IPv4 when initiating communication. For example, DNS name resolution might yield both IPv6 and IPv4 addresses with the network protocol stack having both IPv6 and IPv4 source addresses available. In these cases, a policy that always prefers IPv6 or always prefers IPv4 might produce poor results. For example, if a DNS name resolves to a global IPv6 address and a global IPv4 address. If the node has assigned a global IPv6 address and a 169.254/16 autoconfigured IPv4 address, then IPv6 is the best choice for communication because the global address has a similar scope; therefore, a better chance of success. But if the node has assigned only a link-local IPv6 address and a global IPv4 address, then IPv4 is the best choice for communication because the scope more closely matches the scope of the destination to which you are communicating. The destination address selection algorithm solves this with a unified procedure for choosing among both IPv6 and IPv4 addresses.

Source address selection and destination address selection are discussed separately, but using a common framework enables the two algorithms together to yield useful results. The algorithms attempt to choose source and destination addresses of appropriate scope and configuration status (preferred or deprecated).

Default destination address selection

Resolver APIs have the capability to return multiple IP addresses as a result of a host name query. However, many applications use only the first address returned to attempt a connection or to send a UDP datagram. Therefore, sorting of these IP addresses is performed by the default destination address selection algorithm.

Establishing connectivity can depend on whether an IPv6 address or an IPv4 address is selected, which makes this sorting function even more important.

Default destination address selection occurs only when the system is enabled for IPv6 and the application is using the `getaddrinfo()` API to retrieve IPv6 and/or IPv4 addresses.

The default destination address selection algorithm sorts a list of destination addresses and generates a new list. The algorithm sorts together both IPv6 and IPv4 addresses by a set of rules. Rules are applied, in order, to the first and second address, choosing a best address. Rules are then applied to this best address and the third address. This continues until rules have been applied to all addresses and

the entire list of addresses has been sorted. If one of the rules is able to select the best address between two addresses, remaining rules are bypassed for those two addresses. Subsequent rules act as tie-breakers for earlier rules. The destination address selection algorithm attempts to predict what source address is selected by TCP/IP when the application initiates an outbound connection or sends a datagram using the destination address. This source address is used for some of the destination address selection criteria rules. Source address prediction processing assumes that the application itself does not explicitly specify a source IP address (using bind or ipv6_pktnfo) when initiating a connection or sending a datagram. If the application does explicitly specify a source address, then the destination address selected by this algorithm might not be optimal. The decision the application makes might assume that a different source address is used.

Rules: Observe the following rules:

Rule 1: Avoid unusable destinations.

If one address is reachable (the stack has a route to the particular address) and the other is unreachable, then place the reachable destination address prior to the unreachable address.

Rule 2: Prefer matching scope.

If the scope of one address matches the scope of its source address and the other address does not meet this criteria, then the address with the matching scope is placed before the other destination address.

The scopes of the destination addresses and their associated source addresses are determined by interrogating the high order bits of the address. The destination address can be a multicast or unicast address. Unicast Link-Local is mapped to multicast Link-Local, unicast Site-Local to multicast Site-Local, and unicast Global scope to multicast Global scope.

Rule 3: Avoid deprecated addresses.

If one address is deprecated and the other is non-deprecated, then the non-deprecated address is placed prior to the other address.

Rule 4: Prefer matching address formats.

If one address format matches its associated source address format and the other destination does not meet this criteria, then place the destination with the matching format prior to the other address.

Rule 5: Prefer higher precedence.

If the precedence of one address is higher than the precedence of the other address, then the address with the higher precedence is placed before the other destination address.

Rule 6: Use longest matching prefix.

If one destination address has a longer CommonPrefixLength with its associated source address than the other destination address has with its source address, then the address with the longer CommonPrefixLength is placed before the other address.

Rule 7: Leave the order unchanged.

No rule selected a better address of these two; they are equally good. Choose the first address as the better address of these two and the order is not changed.

Default source address selection

When the application or upper-layer protocol has not selected a source address for an outbound IPv6 packet (using bind or ipv6_pktinfo), the default source address selection algorithm selects one.

The goal of default source address selection is to select the address that is most likely to allow the packet to reach its destination and to support site renumbering. The group of candidate addresses consists of the addresses assigned to the outbound interface (both configured, dynamically generated, or both) or the addresses configured for the outbound interface's SOURCEVIPA interface. Any address which is preferred or deprecated is included in the candidate list. The algorithm is applied to the candidate address list to select the best source address for the packet. If there is only one address in the list of candidate source addresses, then that address is used. If there is more than one address in the candidate list, one is selected by applying the algorithm's rules to the addresses. Rules are applied, in order, to the first and second address, choosing a best address. Rules are then applied to this best address and the third address. This continues until rules have been applied to all addresses. If one of the rules is able to select the best address between two addresses, remaining rules are bypassed for those two addresses. Subsequent rules act as tie-breakers for earlier rules.

Rules: Observe the following rules:

Rule 1: Prefer same address.

If either address is the destination address, choose that address as the source address and terminate the entire algorithm.

Rule 2: Prefer appropriate scope.

If the scope of one address is preferable to the scope of the other address, then the address with better scope is the better address of these two.

As an example, how is the scope of one source address (SA) preferable to the scope of another source address (SB) for the given destination address (D)?

- If scope of SA < scope of SB: If scope of SA < scope of D then SB is the best address of SA and SB; otherwise SA is the best address.
- If scope of SB < scope of SA: If scope of SB < scope of D then SA is the best address of SA and SB; otherwise SB is the best address.

Rule 3: Avoid deprecated addresses.

If one address is deprecated and the other is preferred, then the preferred address is the better address of these two.

Rule 4: Use longest matching prefix.

If one address has a longer CommonPrefixLength with the destination than the other address, then the address with the longer CommonPrefixLength is the better address of these two.

Rule 5: Leave the order unchanged.

No rule selected a better address of these two; they are equally good. Choose the first address as the better address of these two.

VIPA considerations with source address selection

If SOURCEVIPA is configured for the outbound interface and the application has not requested that SOURCEVIPA be ignored (by way of Ignore Source VIPA socket option), the source address is selected from the SOURCEVIPA interface's addresses. Otherwise, source address is selected from the outbound interface's addresses.

Note that selection of a Source VIPA address for IPv6 is done differently from IPv4.

It is determined by the SOURCEVIPAINTERFACE parameter configured on the outbound interface, rather than the order of the HOME list.

When a socket is used to establish a TCP connection to an IPv6 destination or to send a UDP or RAW IP datagram to an IPv6 destination, the local address of the socket is determined based on the set of rules listed in Table 5:

Table 5. Source address selection

Source address selection for communication to IPv6 destinations		TCP, UDP, and RAW
IPCONFIG6 NOSOURCEVIPA	1. Is the socket already bound to a local IPv6 address?	Do not change the local address, use it as it is.
	2. Is the socket unbound (bound to the unspecified IP address)?	Use the IPv6 default source address selection algorithm (selecting an IPv6 address on the physical interface over which the IP packet is about to be sent).
IPCONFIG6 SOURCEVIPA	1. Is the socket already bound to a local IPv6 address?	Do not change the local address, use it as it is.
	2. Has setsockopt() with the NOSOURCEVIPA option been issued for the socket?	Use the IPv6 default source address selection algorithm (selecting an IPv6 address on the physical interface over which the IP packet is about to be sent).
	3. Is there a SOURCEVIPAINTERFACE option on the IPv6 INTERFACE definition over which the IP packet is about to be sent?	Use the IPv6 source address selection algorithm to select an IPv6 VIPA address from the IPv6 virtual interface pointed to by the SOURCEVIPAINTERFACE option.
	4. Is there no SOURCEVIPAINTERFACE option on the IPv6 INTERFACE definition over which the IP packet is about to be sent?	Use the IPv6 default source address selection algorithm (selecting an IPv6 address on the physical interface over which the IP packet is about to be sent).

Migration and coexistence

Enabling IPv6 communication between IPv6 islands in an IPv4 environment

Figure 10 on page 40 illustrates communication between IPv6 islands in an IPv4 environment:




Figure 10. Communicating between IPv6 islands in an IPv4 world

Tunneling

Tunneling provides a way to utilize an existing IPv4 routing infrastructure to carry IPv6 traffic. IPv6 nodes (or networks) that are separated by IPv4 infrastructure can build a virtual link by configuring a tunnel. IPv6-over-IPv4 tunnels are modeled as single-hop. In other words, the IPv6 hop limit is decremented by 1 when an IPv6 packet traverses the tunnel. The single-hop model serves to hide the existence of a tunnel. The tunnel is opaque to network and is not detectable by network diagnostic tools such as traceroute.

z/OS Communications Server does not support being a tunnel endpoint. This means that the z/OS Communications Server stack must have an IPv6 interface connected to an IPv6 capable router. The router is relied upon to handle all tunneling issues.

For more information, see “Tunneling” on page 121.

How to enable end-to-end communication between IPv4 and IPv6 applications

Figure 11 illustrates communication between IPv4 and IPv6 applications:




Figure 11. Communicating between IPv4 and IPv6 applications

Dual-mode stack

z/OS Communications Server can be an IPv4-only or dual-mode stack.

Restriction: There is no support for an IPv6-only stack.

By default, IPv6-enabled applications can communicate with both IPv4 and IPv6 peers. A socket option makes an IPv6-enabled application require all peers to be IPv6. See “Socket option to control IPv4 and IPv6 communications” on page 87 for detailed information about the IPV6_V6ONLY socket option.

IPv6 application on a dual-mode stack: An IPv6 application on a dual-mode stack can communicate with IPv4 and IPv6 partners as long as it does not bind to a native IPv6 address. If it binds to a native IPv6 address, it cannot communicate with an IPv4 partner because the native IPv6 address cannot be converted to an IPv4 address.

If a partner is IPv6, all communication uses IPv6 packets.

If a partner is IPv4, the following occurs:

- Both source and destination are IPv4-mapped IPv6 addresses.
- On inbound, the transport protocol layer maps the IPv4 address to its corresponding IPv4-mapped IPv6 address before returning to the application with AF_INET6 addresses.
- On outbound the transport protocol layer converts the IPv4-mapped addresses to native IPv4 addresses and send IPv4 packets.




Figure 12. IPv6 application on dual-mode stack

IPv4 application on a dual-mode stack: An IPv4 application running on a dual-mode stack can communicate with an IPv4 partner. The source and destination addresses are native IPv4 addresses and the packet is an IPv4 packet.

If a partner is IPv6 enabled and running on an IPv6-only stack, then communication fails. The partner only has a native IPv6 address (not an IPv4-mapped IPv6 address). The native IPv6 address for the partner cannot be converted into a form that the AF_INET application understands.




Figure 13. IPv4-only application on a dual-mode stack

Application Layer Gateways (ALG) and protocol translation

When IPv6-only nodes begin to appear in the network, AF_INET6 applications on these nodes might need to communicate with AF_INET applications. For a multihomed dual-mode IP host, it is likely that the host has both IPv4 and IPv6 interfaces over which requests for host-resident applications are received or sent. IPv4-only (AF_INET sockets) applications are not generally able to communicate with IPv6 partners, which means that only the IPv4 partners in the IPv4 network can communicate with those applications; an IPv6 partner cannot.

As soon as IPv6-only hosts are being deployed in a network, applications on those IPv6-only nodes cannot communicate with the IPv4-only applications on the dual mode hosts, unless one of multiple migration technologies are implemented either on intermediate nodes in the network or directly on the dual mode hosts.

Numerous RFCs describe solutions in this area. One solution is a SOCKS64 implementation that works as a SOCKS server that relays communication between IPv4 and IPv6 flows. SOCKS is a well-known technology, and the issues around it

are familiar. Servers do not require any changes, but client applications (or the stack on which the client applications reside) need to be socksified to be able to reach out through a SOCKS64 server to an IPv4-only partner.

Other solutions are based on a combination of network address translation, IP-level protocol translation, and DNS-flow catcher/interpreter. These solutions all have problems with application-level IP address awareness and end-to-end security.

Network address translation: IPv4 NAT translates one IPv4 (private) address into another IPv4 (external) address. IPv6 NAT-PT translates an IPv4 address into an IPv6 address.

Rules: There are several limitations with NAT-PT:

- All requests and responses pertaining to a session must be routed through the same NAT-PT translator.
- There is a protocol translation limitation because a number of IPv4 fields have changed meaning in IPv6. Details of IPv4 to IPv6 protocol translation can be found in the Stateless IP/ICMP Translation Algorithm (SIIT) RFC.
- If an application carries the IP address in the payload, ALGs must be incorporated.
- Lack of end-to-end security. The two end nodes that seek IPSec network level security must both use IPv4 or IPv6.
- DNS messages and DNSSEC translation. An IPv4 end-node that demands DNS replies be signed rejects replies that have been tampered with by NAT-PT.

Restriction: z/OS Communications Server TCP/IP does not provide a SOCKS64 server and does not contain NAT-PT functionality. If an IPv6-only client requires access to an IPv4-only server running on z/OS, an external SOCKS64 or NAT-PT node is required to translate the IPv6 packet to a corresponding IPv4 packet and vice versa.

Considerations for configuring z/OS for IPv6

This section describes some general considerations for configuring IPv6 on z/OS, including cases where multiple types of TCP/IP stacks are present.

Guideline: In this section, stack or TCP/IP stack is used as a generic term to describe a protocol stack that can be defined as a UNIX System Services AF_INET Physical File System (PFS) in the BPXPRMxx parmlib member (for example, z/OS CS TCP/IP).

IPv6 stack support

IPv4-only stack

Some TCP/IP stacks only support IPv4 interfaces and are only capable of sending or receiving IPv4 packets. These TCP/IP stacks are generally referred to as IPv4-only stacks, as they support IPv4 but do not support communication over IPv6 networks.

An IPv4-only stack supports AF_INET socket applications, but does not support AF_INET6 socket applications.

Restriction: z/OS Communications Server TCP/IP can be started as IPv4-only stack.

IPv6-only stack

An IPv6-only stack supports IPv6 interfaces, but it does not support IPv4 interfaces. These TCP/IP stacks support AF_INET6 sockets and applications that use them, as long as the IP addresses that are used are not IPv4-mapped IPv6 addresses. They do not support AF_INET sockets. Applications can send and receive IPv6 packets by way of an IPv6-only stack, but they cannot send and receive IPv4 packets.

Restriction: z/OS Communications Server TCP/IP cannot be started as an IPv6-only stack.

Dual-mode stack

Many IPv6 TCP/IP stacks support both IPv4 and IPv6 interfaces and are capable of receiving and sending IPv4 and IPv6 packets over the corresponding interfaces. These TCP/IP stacks are generally referred to as a dual-mode stack IP stacks. This does not indicate that there are two separate TCP/IP stacks running on such a node, but it does indicate that the TCP/IP stack has built-in support for both IPv4 and IPv6.

A dual-mode stack supports AF_INET and AF_INET6 socket applications. AF_INET applications can communicate using IPv4 addresses. IPv6-enabled applications that use AF_INET6 sockets can communicate using both IPv6 addresses and IPv4 addresses (using the IPv4-mapped IPv6 address format).

Guideline: z/OS Communications Server TCP/IP can be started as a dual-mode stack.

INET considerations

This section describes the INET considerations for IPv4-only and dual-mode IPv4/IPv6 stacks.

IPv4-only stack

An IPv4-only stack supports AF_INET applications, but it does not support AF_INET6 applications. Start an IPv4-only stack in an integrated sockets environment in one of the following ways:

- Do not code an AF_INET6 statement in BPXPRMxx. This method is the easier of the two. When AF_INET6 is not enabled, the underlying TCP/IP stack is started as an IPv4-only stack, even if it is capable of supporting IPv6.
Restriction: This is the only way to start z/OS Communications Server TCP/IP as an IPv4-only stack in an integrated sockets environment.
- Run a TCP/IP stack that is not capable of supporting IPv6. When starting a TCP/IP stack that does not support IPv6, the stack ignores any AF_INET6 definitions that might appear in BPXPRMxx. As a result, the stack is started as an IPv4-only stack, even when AF_INET6 is coded in BPXPRMxx.

When a TCP/IP stack is started as an IPv4-only stack in an Integrated Sockets environment, applications can open AF_INET sockets and can only send and receive IPv4 packets over IPv4 interfaces. However, applications are unable to open AF_INET6 sockets.

Dual-mode IPv4/IPv6 stack

When both AF_INET and AF_INET6 are coded in BPXPRMxx and a dual-mode capable stack is started, both AF_INET and AF_INET6 sockets are supported by the stack, and applications can send and receive IPv4 and IPv6 packets.

Requirements: To enable AF_INET6 support in an integrated sockets environment, the following two conditions must exist:

- AF_INET6 must be configured in BPXPRMxx. Note that AF_INET6 support can be dynamically enabled by configuring AF_INET6 in BPXPRMxx and then issuing the SETOMVS RESET= command to activate the new configuration.
- A dual-mode capable stack must be started after AF_INET6 is configured in BPXPRMxx. Note that if a dual-mode capable TCP/IP stack is started before configuring BPXPRMxx, it remains an IPv4-only stack as long as it remains active. However, if it is stopped and then restarted, it restarts as a dual-mode TCP/IP stack if AF_INET6 is configured in BPXPRMxx at the time it is restarted.

Requirement: To enable AF_INET6 support for z/OS Communications Server TCP/IP, z/OS Communications Server TCP/IP must be started as a dual-mode stack. z/OS Communications Server TCP/IP does not support being started as an IPv6-only stack. In other words, if AF_INET6 is coded in BPXPRMxx, AF_INET must also be coded. If it is not, then the z/OS Communications Server TCP/IP stack fails to initialize.

Common INET considerations

This section describes additional INET considerations.

Enabling AF_INET6 support in a Common INET environment

Requirements: To enable AF_INET6 support in a Common INET environment, the following conditions must exist:

- AF_INET6 must be configured in BPXPRMxx. Note that AF_INET6 support can be dynamically enabled by configuring AF_INET6 in BPXPRMxx and then issuing the SETOMVS RESET= command to activate the new configuration.
- At least one dual-mode capable stack must be started after AF_INET6 is configured in BPXPRMxx. Note that any dual-mode capable TCP/IP stack started before configuring BPXPRMxx remains an IPv4-only stack as long as it remains active. However, if it is stopped and then restarted, it restarts as a dual-mode TCP/IP stack if AF_INET6 is configured in BPXPRMxx at the time it is restarted.

Guideline: Do not start some z/OS CS TCP/IP stacks with AF_INET6 support and some without AF_INET6 support. If AF_INET6 support is dynamically enabled, you should stop and restart all TCP/IP stacks which were active when AF_INET6 support was enabled. This allows these TCP/IP stacks to become dual-mode stacks. After this occurs, all applications which are capable of opening AF_INET6 sockets should be stopped and restarted. This allows the restarted applications to communicate over IPv4 and IPv6 networks.

Disabling AF_INET6 support in a Common INET environment

Disable AF_INET6 support in a Common INET environment in one of the following ways:

- Stop all active dual-mode TCP/IP stacks while IPv4-only stacks remain active. Applications are no longer be able to open AF_INET6 sockets, although they can

continue to use any AF_INET6 sockets that are already open and not bound to one of the stopped dual-mode TCP/IP stacks. However, applications are able to open AF_INET sockets.

- Dynamically disable AF_INET6 in BPXPRMxx and stop all active dual-mode TCP/IP stacks. When restarted, the dual-mode capable TCP/IP stacks start as IPv4-only stacks. In effect, this is a subset of the previous case. To disable AF_INET6 support, issue the SETOMVS RESET= command to set the AF_INET6 MAXSOCKETS value to 0.

Supporting a mixture of dual-mode stacks and IPv4-only stacks

When AF_INET6 sockets are supported, an IPv6-enabled application can use an AF_INET6 socket to send and receive data with both IPv4 and IPv6 partners.

When communicating with an IPv6 partner, a native IPv6 address is used. When communicating with an IPv4 partner, the IPv4 address is encoded as an IPv4-mapped IPv6 address. When an IPv4-mapped IPv6 address is used on an AF_INET6 socket, a dual-mode TCP/IP stack realizes the partner is attached to the IPv4 network and routes packets over IPv4 interfaces.

As long as all TCP/IP stacks started in a Common INET environment provide native support AF_INET6 sockets, socket calls can be passed directly to the underlying TCP/IP stack. However, when both dual-mode stacks and IPv4-only stacks are started in a Common INET environment, the IPv4-only stacks are not able to process the native AF_INET6 socket calls. As a result, an application which uses IPv4-mapped IPv6 addresses on an AF_INET6 socket needs transformations done by Common INET to communicate with partners over any active IPv4-only stack.

Common INET provides AF_INET6 transformations that allow AF_INET6 applications to communicate with an IPv4 peer over IPv4-only stack. The AF_INET6 transformations convert AF_INET6 socket calls to the corresponding AF_INET socket calls before sending them to an IPv4-only stack and converts AF_INET responses received from the IPv4-only stack to the corresponding AF_INET6 responses before making them available to the AF_INET6 application. Note that even with this transformation, AF_INET6 applications must use IPv4-mapped IPv6 addresses to communicate with IPv4 applications.

Figure 14 on page 48 shows a mixture of dual-mode stacks and IPv4-only stacks:




Figure 14. Mixing dual-mode and IPv4-only stacks

Configuring a common INET environment

If a mixture of dual-mode capable stacks and IPv4-only stacks are started in a Common INET environment, the default stack should be one of the dual-mode capable stacks. Common INET routes certain requests to the default stack, and this enables the stack with more functional capability to process these requests.

If AF_INET6 support is dynamically configured in BPXPRMxx, stop and restart all dual-mode-capable TCP/IP stacks. After the TCP/IP stacks have been stopped and restarted, stopped and restarted all IPv6-enabled applications.

Part 2. IPv6 enablement

This section contains the following chapters:

Chapter 4, “Configuring support for z/OS,” on page 51 describes the IPv6 function provided in z/OS Communications Server and how to enable it.

Chapter 5, “Configuration guidelines,” on page 63 contains recommendations and guidance information for implementing the IPv6 functions provided in z/OS Communications Server.

Chapter 4. Configuring support for z/OS

This chapter describes the configuration support needed for z/OS and contains the following sections:

- “Ensure that important features are supported over IPv6”
- “Assess automation and application impacts due to Netstat and message changes”
- “Determine how remote sites connect to the local host”
- “Determine how remote sites connect to the local host”
- “SNA access” on page 52
- “Avoid using IP addresses for identifying remote hosts” on page 52
- “Considerations when using the BIND parameter on the PORT statement” on page 53
- “Security considerations” on page 53
- “Application programming considerations” on page 54
- “Enabling IPv6 support” on page 54
- “Resolver processing” on page 56
- “User exits” on page 58
- “Which applications started with inetd are IPv6 enabled?” on page 58
- “How does IPv6 affect SMF records?” on page 59
- “How does IPv6 affect the Policy Agent?” on page 59
- “How does IPv6 affect SNMP?” on page 60
- “Monitoring the TCP/IP network” on page 60
- “Diagnosing problems” on page 62

Ensure that important features are supported over IPv6

See Chapter 11, “IPv6 support tables,” on page 131 to ensure all needed features are supported over IPv6.

Assess automation and application impacts due to Netstat and message changes

Netstat output for stacks that are IPv6-enabled has a different format in order to accommodate the longer IPv6 address. This becomes an issue when applications that parse Netstat output are used. The same considerations also apply to applications which use IP addresses in their automation because IP addresses now have a longer format.

Determine how remote sites connect to the local host

It is likely that clients that are not connected to a link that is directly attached to a z/OS image require access to servers that run on that z/OS image. Because z/OS provides a dual-stack implementation, z/OS can send IPv4 packets to partner nodes that are connected to the IPv4 network and IPv6 packets to partner nodes that are connected to the IPv6 network. If the client node is connected to the same routing infrastructure as the z/OS node, traffic is routed between z/OS and the client node by way of the native network transport.

In some cases, the two nodes might not be connected to the same routing infrastructure. For instance, each node might be attached to distinct IPv6 networks that are separated by an intermediate IPv4 network. When this occurs, tunneling might be used to transmit the native IPv6 packets across the IPv4 network. This allows nodes in the disjoint IPv6 networks to send packets to one another.

z/OS does not support functioning as an endpoint for this type of tunnel. However, z/OS might route traffic over a tunnel in the intermediate network. In this case, the tunnel endpoint used by z/OS would be an IPv6/IPv4 router in the network that supports one of several tunneling protocols. The tunnel endpoint used by z/OS might be attached to the same LAN to which z/OS attaches or might be attached to a remote network link. In either case, the presence of the tunnel endpoint is transparent to z/OS; from the z/OS perspective, traffic is routed over the native IPv6 network.

SNA access

Both Enterprise Extender and TN3270 allow access to SNA applications over an IPv6 network as well as an IPv4 network. For both protocols, it is possible to simultaneously support connectivity over IPv4 and IPv6 networks. Enterprise Extender uses separate path statements and connection networks for each protocol. By assigning different weights to Transmission Groups that use different network protocols, it is possible to have SNA traffic prefer being routed over the IPv6 network or the IPv4 network. For TN3270, the network protocol used is determined by the remote TN3270 client.

Guideline: For Enterprise Extender and TN3270, use global unicast addresses. While link-local addresses might work in certain configurations, they are not suitable for use when connecting between partner companies. There are few, if any, IPv6 NAT devices which can perform the necessary mappings between limited scope addresses and globally routable addresses and, given the vast number of globally unique IPv6 addresses available, are not necessary.

Avoid using IP addresses for identifying remote hosts

In IPv4 networks, some sites and applications attempt to use the remote IP address to identify the client node which is connecting. In general, do not do this for IPv4, because the client address can often be unpredictable, either due to the client using DHCP to obtain its address or due to the client accessing the server from behind a NAT (Network Address Translator) device.

In IPv6, the client address is likely to become even more volatile than it is in IPv4 networks. Using Stateless Address Autoconfiguration, a client's address is dynamically derived from the MAC address of the network adapter used for connectivity. IPv6 also allows clients to pseudo-randomly generate IP addresses, referred to as temporary addresses, which can be used for one or more connections. These temporary addresses can be generated as frequently as the client desires- once a day, once an hour, or even more frequently. In general, the temporary addresses are not placed in the DNS, making it impossible to use DNS to map the IP address to a host name.

Result: The client IP addresses are unpredictable and subject to frequent change. In addition, it is possible, and even likely, that a server is unable to map the client address to a host name. If a mechanism to identify the remote host is required, then a different mechanism (client certificate, password, and so on) should be used to identify the remote host. For example, this approach is used by Enterprise

Extender. For IPv6, Enterprise Extender does not support configuring or passing IPv6 addresses. Instead, it uses hostnames to identify Enterprise Extender nodes.

Considerations when using the BIND parameter on the PORT statement

The PORT statement reserves a port for the use of a particular server. It normally does not distinguish between IPv4 and IPv6; the port is reserved regardless of which type of address the application uses. The BIND keyword on the PORT statement allows you to force an INADDR_ANY listener to listen on a particular IP address. You can now specify an IPv6 address on this keyword. INADDR_ANY listeners are converted to an IPv4 address, but ignores an IPv6 address on the BIND keyword. IN6ADDR_ANY listeners are converted to either an IPv4 address (the IPv4-mapped form of that address) or an IPv6 address, depending on what is specified with the BIND keyword.

If you use the BIND option, your server can listen only for IPv4 connections or IPv6 connections, but not both. To have the same service serve both IPv4 and IPv6 clients, you might need to start two instances of it, one bound to an IPv4 address and one to an IPv6 address.

With SHAREPORT or SHAREPORTWLM keyword, you can start multiple instances of the server and have connections automatically load balanced between them. This function is supported for TCP listeners only. All IPv4 connection requests are load balanced between the set of IPv4 listeners (including AF_INET6 IN6ADDR_ANY listeners), while all IPv6 connection requests are load balanced between the set of IPv6 listeners. See the *z/OS Communications Server: IP Configuration Reference* for information about the load balancing algorithms used by each of these parameters.

Security considerations

On z/OS Communications Server, not all security features that are supported over an IPv4 transport are enabled when communicating by way of an IPv6 transport. For instance, IPSec, Network Access Control, Stack and Port Access Control, TLS, SSL, and Kerberos (Kerberos Version 5 and GSSAPIs) are enabled for both IPv4 and IPv6, whereas Intrusion Detection is enabled for IPv4 but not for IPv6. Refer to Table 36 on page 134 for a list of features supported for IPv4 or IPv6.

When a security function is supported over IPv4 but not over IPv6, the security feature is exercised when data is transmitted over the IPv4 transport. This is true whether the application uses AF_INET or AF_INET6 sockets. However, when an AF_INET6 socket application communicates over the IPv6 transport, security features that are supported over IPv4 only are not exercised.

Result: For the same local application, some security features can be exercised when communicating by way of IPv4, but not when communicating by way of IPv6.

To avoid creating a potential security exposure, it is important to determine if any important security features are supported over IPv4 but not over IPv6 prior to enabling AF_INET6 on a given LPAR. If only a subset of applications utilize such a security feature, then it is sufficient to ensure that those applications communicate only over the IPv4 transport.

To ensure that the IPv4 transport is used, the following methods are available:

- Verify that the application uses AF_INET sockets. Applications that use AF_INET sockets are able to communicate only by way of the IPv4 transport.
- Configure the application to bind to an IPv4 address. Applications that bind to an IPv4 address are able to communicate using the IPv4 transport only.
- Use the BIND parameter on the PORT statement to cause the application to bind to an IPv4 address.

Application programming considerations

Refer to Part 3, "Application enablement," on page 69 for information about application programming considerations.

Enabling IPv6 support

z/OS Communications Server can be run as an IPv4-only stack or as a dual-mode stack (IPv4 and IPv6). The BPXPRMxx parmlib member determines which mode is used. The following configurations are possible:

- INET IPv4 only
- INET IPv4/IPv6 dual-mode stack
- CINET IPv4 only
- CINET IPv4/IPv6 dual-mode stack

Restriction: After a stack has been started, you cannot change its mode without stopping and restarting the stack.

You can configure either a single AF_INET or both AF_INET and AF_INET6. Although coding AF_INET6 alone is not prohibited, TCP/IP does not start because the master socket is AF_INET and the call to open it fails.

IPv4-only BPXPRMxx sample definition:

```
FILESYSTYPE Type(INET) Entrypoint(EZBPFINI)
NETWORK DOMAINNAME(AF_INET)
    DOMAINNUMBER(2)
    MAXSOCKETS(2000)
    TYPE(INET)
```

INET IPv4/IPv6 dual-mode stack BPXPRMxx sample definition:

Dual-mode stack support is defined by using two NETWORK statements (one for AF_INET and one for AF_INET6) in the BPXPRMxx parmlib member. For example:

```
FILESYSTYPE Type(INET) Entrypoint(EZBPFINI)
NETWORK DOMAINNAME(AF_INET)
    DOMAINNUMBER(2)
    MAXSOCKETS(2000)
    TYPE(INET)
NETWORK DOMAINNAME(AF_INET6)
    DOMAINNUMBER(19)
    MAXSOCKETS(3000)
    TYPE(INET)
```

Separate MAXSOCKETS values are supported. The IPv6 default is the IPv4 specified value.

CINET IPv4-only BPXPRMxx sample definition:

Multiple TCP/IP stacks in one MVS image or LPAR are only supported by using Common INET (CINET). Each TCP/IP stack is defined in the BPXPRMxx parmlib member using a SUBFILESYSTYPE statement. These definitions are identical to what was used prior to IPv6 support. The following example shows the definitions for three IPv4-only stacks:

```
FILESYSTYPE TYPE(CINET) ENTRYPOINT(BPXTCINT)
NETWORK DOMAINNAME(AF_INET)
  DOMAINNUMBER(2)
  MAXSOCKETS(2000)
  TYPE(CINET)
  INADDRANYPORT(20000)
  INADDRANYCOUNT(100)
SUBFILESYSTYPE NAME(TCPCS) TYPE(CINET) ENTRYPOINT(EZBPFINI)
SUBFILESYSTYPE NAME(TCPCS2) TYPE(CINET) ENTRYPOINT(EZBPFINI)
SUBFILESYSTYPE NAME(TCPCS3) TYPE(CINET) ENTRYPOINT(EZBPFINI)
```

CINET IPv4/IPv6 dual-mode stack BPXPRMxx sample definition:

Dual-mode stack (IPv4/IPv6) support is defined by using two NETWORK statements in the BPXPRMxx member. Each TCP/IP stack is defined in the BPXPRMxx parmlib member with SUBFILESYSTYPE. All z/OS Communications Server stacks defined under the two NETWORK statements are IPv4 or IPv6 stacks. The following example shows the definitions for three dual (IPv4/IPv6) stacks:

```
FILESYSTYPE TYPE(CINET) ENTRYPOINT(BPXTCINT)
NETWORK DOMAINNAME(AF_INET)
  DOMAINNUMBER(2)
  MAXSOCKETS(2000)
  TYPE(CINET)
  INADDRANYPORT(20000)
  INADDRANYCOUNT(100)
NETWORK DOMAINNAME(AF_INET6)
  DOMAINNUMBER(19)
  MAXSOCKETS(3000)
  TYPE(CINET)
SUBFILESYSTYPE NAME(TCPCS) TYPE(CINET) ENTRYPOINT(EZBPFINI)
SUBFILESYSTYPE NAME(TCPCS2) TYPE(CINET) ENTRYPOINT(EZBPFINI)
SUBFILESYSTYPE NAME(TCPCS3) TYPE(CINET) ENTRYPOINT(EZBPFINI)
```

Configuring z/OS IPv6 support

The following configuration statements enable IPv6 addresses to be configured. Refer to the *z/OS Communications Server: IP Configuration Reference* for detailed information on each of these statements.

BEGINROUTES

Code this statement to add static IPv6 routes to the IP routing table. BEGINROUTES with IPv6 addresses coded is rejected if the stack is not enabled for IPv6. The GATEWAY statement does not support IPv6 routes.

DELETE PORT (BIND IP address)

IPv6 must be enabled for IPv6 addresses to be coded on these configuration statements.

INTERFACE

An IPv6-enabled stack still uses DEVICE and LINK to define IPv4 interfaces. However, you cannot use DEVICE and LINK to define IPv6 interfaces. You must use the INTERFACE statement to define IPv6 interfaces. The stack must be enabled for IPv6 to use this statement.

IPCONFIG

A FORMAT keyword has been added to control the format of the command output if the stack is not enabled for IPv6.

IPCONFIG6

This statement is rejected if the stack is not enabled for IPv6. However, the SOURCEVIPA option has a dependency on the INTERFACE statement. You must specify the SOURCEVIPAINTERFACE keyword on the INTERFACE statement for each interface on which you desire that SOURCEVIPA take effect.

PKTTRACE

IPv6 must be enabled for IPv6 addresses to be coded on these configuration statements.

PORt (BIND IP address)

IPv6 must be enabled for IPv6 addresses to be coded on these configuration statements.

Resolver processing

IPv6 support introduces several changes to how host name and IP address resolution is performed. These changes affect several areas of resolver processing, including:

- New resolver APIs are introduced for IPv6 enabled applications. See “Name and address resolution functions” on page 77 for more details.
- New DNS resource records are defined to represent hosts with IPv6 addresses; therefore new, network flows between resolvers and name servers (in place of DNS IPv4 A records).
- A new algorithm is defined to describe how a resolver needs to sort a list of IP addresses returned for a multihomed host. See “Default destination address selection” on page 36 for more information.
- New statements in the resolver configuration files are defined, and new search orders are implemented for local host tables processing.

Resolver configuration

In order to avoid impacting existing IPv4 queries, the use of /etc/hosts, HOSTS.LOCAL, HOSTS.SITEINFO, and HOSTS.ADDINFO files continue to be supported for IPv4 addresses only. The HOSTS.SITEINFO and HOSTS.ADDRINFO files continue to be generated from HOSTS.LOCAL file by way of the MAKESITE utility.

ETC.IPNODES is a new local host file (in the style of /etc/hosts) that might contain both IPv4 and IPv6 addresses. IPv6 addresses can be defined in ETC.IPNODES only. The introduction of this file allows the administration of local host files to more closely resemble that of other TCP/IP platforms and eliminates the requirement of post-processing the files (specifically, MAKESITE).

The following new search order is used for selecting new ETC.IPNODES local host files for IPv6 searches in MVS and UNIX environments:

1. GLOBALIPNODES
2. RESOLVER_IPNODES environment variable (UNIX only)
3. userid/jobname.ETC.IPNODES
4. *hlq.ETC.IPNODES*
5. DEFAULTIPNODES

6. /etc/iphnodes

IPv6 search order is simplified, but to minimize migration concerns, the IPv4 search order continues to be supported as in previous releases. The side effect of this is that by default, you would be required to maintain two different local host files (for example, IPv4 addresses in HOSTS.LOCAL, IPv6 and IPv4 addresses in ETC.IPNODES) for your system.

An easier approach is to use the new COMMONSEARCH statement in the resolver setup file. By specifying COMMONSEARCH, you indicate that only the new IPv6 search order should be used, regardless of whether the search is for IPv6 or IPv4 resources. This means that only one file (ETC.IPNODES) has to be managed for the system, and that all the APIs utilize the same single file. The use of COMMONSEARCH reduces IPv6 and IPv4 searching to a single search order, and also reduces the z/OS UNIX and native MVS environments to a single search order.

For detailed information about search orders, refer to *z/OS Communications Server: IP Configuration Guide*.

IPv4-only configuration statements

Only IPv4 addresses can be specified on the NAMESERVER and NSINTERADDR TCPIP.DATA statements. This implies that all resolver communications with a name server occurs using AF_INET sockets, even when resource records related to IPv6 addresses are being queried.

The other statement in the TCPIP.DATA data set that currently supports IP address specification is the SORTLIST directive. SORTLIST is used for sorting IPv4 addresses only; the default destination address selection algorithm is used to sort IPv6 addresses.

IPv6/IPv4 configuration statements

Use the following statements for IPv6/IPv4 configuration:

COMMONSEARCH/NOCOMMONSEARCH resolver setup statement

Use these statements when a common local host file search order is to be used or not used. The COMMONSEARCH statement allows the same search order of local host files be used for an IPv4 or a IPv6 query. It also allows the same search order to be used in both the native MVS and z/OS UNIX environments.

GLOBALIPNODES resolver setup statement

Use this statement to specify the global local host file.

DEFAULTIPNODES resolver setup statement

Use this statement to specify the default local host file.

Steps for implementing the resolver functions

Perform the following steps to implement the resolver functions

1. Add new resolver setup statements.

2. Create the IPNODES local host files.

3. Add IPv6 resource records to DNS.

For detailed information, refer to Understanding resolvers in the *z/OS Communications Server: IP Configuration Guide*.

Resolver communications with the Domain Name System (DNS)

To retrieve IPv6 data from the proper name server, you must ensure that the resolver configuration data set points to name servers that can resolve the IPv6 queries. A resolver does not have to communicate with a name server over an IPv6 network in order to retrieve IPv6 data. The z/OS resolver can use only IPv4 to communicate with a name server.

User exits

Several TCP/IP applications provide exit facilities that can be used for a variety of purposes. Several of these exits include IP addresses or SOCKADDR structures as part of the parameters passed to the exits.

The following exits are available to support IPv6 addresses or SOCKADDR structures:

- FTP - All FTP exits have been enhanced to support IPv6 addresses except for FTPSMFEX. Samples for these exits are provided in SEZAINST. Refer to *z/OS Communications Server: New Function Summary* for more information on changes to these exits:
 - FTCHKCMD
 - FTCHKCM1
 - FTCHKCM2
 - FTCHKJES
 - FTCHKPWD
 - FTPOSTPA
 - FTPOSTPR
- The TSO remote execution server user exit - RXEXIT.

Which applications started with inetc are IPv6 enabled?

The following z/OS UNIX applications support IPv6 addresses:

- Internet daemon (inetd) server
- Remote execution (orexecd) server
- Remote shell (orshd) server
- Telnet server (otelnetd)

Modifying inetc.conf

The inetc.conf file must be modified to support the IPv6-enabled applications. In order for the z/OS UNIX servers to support IPv6 connections, tcp6 must be specified for the protocol of the service name in the inetc.conf file. When tcp6 is defined, IPv4 clients are also supported.

The z/OS UNIX rsh server and Telnet server support Kerberos for IPv4 connections, but not for IPv6 connections.

How does IPv6 affect SMF records?

Most of the TCP/IP SMF records currently contain IP addresses as part of their content. The data in these records is typically processed by programs, some of which are real-time SMF exits and others that post-process the SMF records after the records are created. In z/OS V1R2, a new type of TCP/IP SMF record, type 119, was introduced. The type 119 SMF records were created to provide a standardized structure for all SMF records provided by TCP/IP. This included a standard representation of IP addresses appearing across all type 119 records in which IPv4 addresses appear in IPv4-mapped form and IPv6 addresses appear as is.

Guideline: The type 119 records constitute a superset of the older type 118 records in terms of data that is available. Users exploiting IPv6 should migrate to the SMF 119 record.

Type 118 FTP client and server transfer completion records are generated for IPv6 connections. In this case, the FTP records use IP addresses of 255.255.255.255 to indicate that the address cannot be included. All other type 118 SMF records are not generated for IPv6 connections.

For more information about SMF records, see the *z/OS Communications Server: IP Configuration Guide*.

How does IPv6 affect the Policy Agent?

The Policy Agent supports IPv6 in the following ways:

- Table 6 lists the policy types that support IPv6.
- IPv6 XCF addresses can be specified in a sysplex distributor environment.

Table 6. IPv6 support for different policy types

Policy type	IPv6 supported?
IDS	No
IPSec	Yes
QoS	Yes
AT-TLS	Yes

When IPv6 addresses are used in policies for a given stack, as configured to Policy Agent using the `TcpImage` configuration statement, the stack must be IPv6 enabled. IPv6 policy is installed but is not enforceable in a stack that is not IPv6 enabled. If the corresponding stack is recycled later with IPv6 enabled, all policies are read and parsed again. At this point, any policies with IPv6 addresses are enforced.

The use of IPv6 interfaces in QoS policies is problematic, because such interfaces can be assigned multiple IP addresses. As a result, the only way to specify IPv6 addresses in policies is by interface name. The interface name can also be used for IPv4 interfaces, as well as the IPv4 address. The name specified in the policies for IPv4 interfaces is the name specified on the `LINK` statement in the TCP/IP profile. For IPv6 interfaces, it is the name specified on the `INTERFACE` statement. IPv6 interfaces can be specified for QoS policies and also for the `SetSubnetPrioTosMask` statement or `LDAP` object.

To support sysplex distributor policy performance monitoring, as specified using the PolicyPerfMonitorForSDR configuration statement, the Policy Agent needs to establish TCP connections between the qosCollector threads that run on the distributing stacks and the qosListener threads that run on the target stacks. Depending on the sysplex configuration, either one or two connections between these threads are established. One connection is established for all target stacks that are configured using IPv4, and one connection is established for all target stacks configured using IPv6. Because a given target can be configured using both IPv4 and IPv6, it is possible that two connections are established between a given qosCollector and qosListener thread. When this occurs, only information related to distributed IPv4 DVIPAs flows over the IPv4 connection and likewise for the IPv6 connection.

How does IPv6 affect SNMP?

The following SNMP components operate over IPv6 networks and handle IPv6-related management data.

Requirement: The TCP/IP stack on your system must support IPv6 networking to take advantage of the IPv6 support offered by these components. If not, these applications operate in IPv4 mode.

- SNMP agent
- z/OS UNIX `snmp/osnmp` command
- Trap Forwarder daemon
- Distributed Protocol Interface (DPI[®])
- TN3270 Telnet subagent

The TCP/IP subagent supports IPv6 management data in the following MIB modules:

- IF-MIB from RFC 2233 - Interface data
- IP-MIB from draft-ietf-ipv6-rfc2011-update-04.txt - IP and ICMP data
- IP-FORWARD-MIB from draft-ietf-ipv6-rfc2096-update-05.txt - Route data
- TCP-MIB from draft-ietf-ipv6-rfc2012-update-04.txt - TCP connection data
- UDP-MIB from draft-ietf-ipv6-rfc2013-update-03.txt - UDP endpoint data
- TCP/IP Enterprise-specific MIB (IBMTCPIMVS-MIB)

Refer to Managing TCP/IP Network Resources with SNMP in *z/OS Communications Server: IP System Administrator's Commands* for more details regarding the TCP/IP Subagent support.

Monitoring the TCP/IP network

This section describes how IPv6 affects reports.

How does IPv6 affect Netstat?

- In order to accommodate full IPv6 address information, Netstat reports have been redesigned. If the TCP/IP stack is IPv6 enabled, reports are displayed in a different format than with IPv4. This might impact applications that are used to parse Netstat output. The same considerations apply to applications which use IP addresses in their automation since IP addresses now have a longer size. If

the TCP/IP stack is not IPv6 enabled, the report format is unchanged unless the FORMAT LONG parameter is specified on the Netstat command or on the IPCONFIG PROFILE statement.

- IPv6 statistic information is added to the Netstat STATS/-S report.
- Information regarding whether the stack is IPv6 enabled or not is added to the Netstat UP/-u report.
- For a server that opens an AF_INET6 socket, binds to IN6ADDR_ANY, and does a socketopt with IPv6_V6ONLY against the socket, the local address information in the connection related reports are contained the text (IPV6_ONLY).

```
Netstat ALLCONN/-a example on an IPv6 enabled stack:  
MVS TCP/IP NETSTAT CS V1R6          TCPIP NAME: TCPCS      17:40:36  
User Id Conn State  
-----  ----  -----  
FTPABC1 00000021 Listen  
  Local Socket: 0.0.0.0.21  
  Foreign Socket: 0.0.0.0.0  
FTPDV6 00000086 Listen  
  Local Socket: ::.21 (IPV6_ONLY)  
  Foreign Socket: ::.0
```

Control of output format

When the stack is IPv6-enabled, the report output is displayed in the new format, which is referred to as long format.

In order to allow the stack to be configured for IPv4-only operation (not IPv6 enabled and short format displays), but still allow a developer who needs to modify programs that rely on Netstat output to update and test new versions of these programs with long format output from Netstat, the following output format control options are available:

FORMAT SHORT

The output is displayed in the existing IPv4 format.

FORMAT LONG

The output is displayed in the format which supports IPv6 addresses.

A stack-wide output format parameter (FORMAT SHORT/LONG) can be specified on the IPCONFIG profile statement. It instructs Netstat to produce output in one of the above formats. FORMAT SHORT is only applicable when the stack is not IPv6 enabled.

In addition to the stack-wide FORMAT parameter, a Netstat command line option FORMAT/-M with keyword SHORT/LONG is supported to override the stack-wide parameter. When a user specifies the Netstat command line format option, it overrides the stack-wide format parameter on an IPv4-only stack.

What has changed?

All Netstat reports have been modified to support IPv6.

The following Netstat report is added to display Neighbor Discovery cache information:

- Netstat ND/-n

Guideline: The Netstat GATE/-g is not enhanced to support IPv6 routes. Netstat ROUTE/-r is the suggested alternative.

For more detailed information, refer to Netstat in *z/OS Communications Server: IP System Administrator's Commands*.

How does IPv6 affect Ping and Traceroute?

Ping and Traceroute provide the following support for IPv6:

- IPv6 IP addresses or host names that resolve to IPv6 IP addresses, can be used for destinations.
- IPv6 IP addresses can be used as the source IP address for the command's outbound packets.
- IPv6 IP addresses or interface names can be used as the outbound interface.
- A new ADDRTYPE/-A command option can be specified to indicate whether an IPv4 or IPv6 IP address should be returned from host name resolution.
- IPv4-mapped IPv6 IP addresses are not supported for any option value.

Diagnosing problems

This section describes IPv6 problem diagnosis considerations.

How does IPv6 affect IPCS?

IPCS formatting has been enhanced for IPv6 for TCPIPCS dump analysis and CTRACE components SYSTCPIP and SYSTCPDA. For detailed information about IPCS, refer to TCP/IP services traces and IPCS support in the *z/OS Communications Server: IP Diagnosis Guide*.

How does IPv6 affect packet and data tracing?

Packet and data trace functions have been enhanced for IPv6 to allow tracing of IPv6 addresses. For detailed information about trace functions, refer to TCP/IP services traces and IPCS support in *z/OS Communications Server: IP Diagnosis Guide*.

Chapter 5. Configuration guidelines

This chapter describes IPv6 configuration guidelines and contains the following sections:

- “Connecting to an IPv6 Network”
- “IPv6 address assignment guidelines” on page 64
- “Updating DNS definitions” on page 66
- “Using source VIPA” on page 66
- “Using OMPROUTE or define static routes to improve network selection” on page 67
- “Connecting to non-local IPv4 locations” on page 68
- “IPv6-only application access to IPv4-only application” on page 68

Connecting to an IPv6 Network

z/OS Communications Server TCP/IP supports direct attachment to IPv6 networks in the following ways:

IPAQENET6 interface type

TCP/IP attaches to an IPv6 LAN by way of OSA-Express in QDIO mode, using either Fast Ethernet or Gigabit Ethernet. A single physical LAN can carry both IPv4 and IPv6 packets over the same media. While the physical network is shared, from a logical view there are two separate LANs, one carrying IPv4 traffic and one carrying IPv6 traffic. A single OSA-Express port can be used to carry both IPv4 and IPv6 traffic simultaneously.

MPCPTP6 interface type

TCP/IP can directly communicate with other IPv6 z/OS Communications Server TCP/IP V1R5 (or later) images, using ESCON® Channel-to-Channel Adapters, XCF connectivity (if the stacks are in the same sysplex), or the IUTSAMEH facility (if the stacks are on the same LPAR).

IPAQIDIO6 interface type

TCP/IP can directly communicate with other IPv6 z/OS Communications Server TCP/IP V1R7 (or later) images and z/Linux images using HiperSockets connectivity. This applies only to stacks running on the same central processor complex and running on a zSeries® server that supports IPv6 HiperSockets .

Guideline: All three of these interface types can be used for LPAR-to-LPAR IPv6 communication, best performance is achieved by using the IPAQIDIO6 interface type (if both stacks meet the criteria previously listed). The performance of the other interface types varies with the speed of the underlying media.

For stack-to-stack communications within a single LPAR, the MPCPTP6 interface type (using IUTSAMEH) provides the best performance.

To transport IPv6 traffic to another host, z/OS TCP/IP must send traffic using native IPv6 packets. Note that when communicating with another IPv6 host, a router within the network might tunnel the IPv6 packet across an IPv4 network to a remote IPv6 LAN or host. However, z/OS Communications Server TCP/IP

cannot be the tunnel endpoint, and the tunneling by an intermediate router is transparent to z/OS Communications Server TCP/IP.

IPv6 address assignment guidelines

This section provides IPv6 address assignment guidelines.

Avoid using site-local addresses

Site-local addresses were designed to use private address prefixes that could be used within a site without the need for a global prefix. Until recently, the full negative impacts of site-local addresses in the Internet were not fully understood. Due to problems in the use and deployment of addresses constructed using a site-local prefix, the IETF has deprecated the special treatment given to the site-local prefix. An IPv6 address constructed using a site-local prefix is now treated as a global unicast address. The site-local prefix can be reassigned for other use by future IETF standards action.

Guideline: Because of this, site-local unicast addresses should not be used. Use global unicast addresses instead of site-local addresses.

Defining the interface ID for physical interfaces

If you do not manually configure the interface ID, the system selects an interface ID for you, using a random value (on an MPCPTP6 interface), a value derived from the MAC address (on an IPAQENET6 interface), or a value derived from the IQD CHPID (on an IPAQIDIO6 interface). To simplify the configuration effort, let the system select the interface ID. In some cases, though, it is necessary or desirable to control all IPv6 addresses which are assigned to a physical adapter. This might be useful if other IPv6 nodes need to define static routes to this host, or if you use IPv6 addresses in Multi-Level Security policies.

Use stateless address autoconfiguration for physical interfaces

IPv6 addresses for physical interfaces can be manually defined or can be automatically assigned by stateless address autoconfiguration. Use the stateless address autoconfiguration for this assignment. Using stateless address autoconfiguration reduces the amount of definition required to enable IPv6 support, while making future site renumbering easier.

Use VIPAs

Using static VIPAs removes hardware as a single point of failure for connections being routed over the failed hardware. If you are not using dynamic routing, configure at least one static VIPA for each LAN to which z/OS Communications Server TCP/IP is connected. Each VIPA configured this way should be associated with all physical adapters connected to that same LAN.

Requirement: Static VIPAs must be manually configured; z/OS Communications Server TCP/IP does not support stateless address autoconfiguration for VIPAs.

Dynamic VIPAs (DVIPAs) can also be used in an IPv6 network. The decision to use DVIPAs in an IPv6 network is similar to the decision to use DVIPAs in an IPv4 network. For detailed information, refer to *Using Dynamic VIPAs (DVIPAs) in the z/OS Communications Server: IP Configuration Guide*.

Selecting the network prefix

z/OS Communications Server TCP/IP does not perform duplicate address detection for VIPAs, because they are not assigned to a physical interface attached to the LAN.

Guideline: To avoid possible address collisions, the network prefix used for static VIPAs should be different from the network prefix used for physical interfaces (either manually configured or autoconfigured using stateless address autoconfiguration).

If either the IPv6 OSPF or IPv6 RIP dynamic routing protocol of OMPROUTE is being used, the network prefix for a static VIPA should not be the same as any prefix defined as on-link on a physical link. The VIPA can then be associated with interfaces attached to any physical link, thus enabling maximum redundancy. This association between VIPAs and interfaces attached to physical links is accomplished using the SOURCEVIPAINTERFACE parameter of the INTERFACE statement for the interface attached to the physical link.

If IPv6 OSPF or IPv6 RIP dynamic routing protocol of OMPROUTE is not being used, the network prefix for a static VIPA should be selected from the set of prefixes which are advertised by way of router discovery by one or more routers attached to the LAN. The prefix should be advertised as on-link and not to be used for address autoconfiguration. By using an on-link prefix, hosts and routers attached to the LAN use neighbor discovery address resolution to obtain a link-layer address for the VIPA. z/OS Communications Server TCP/IP selects a link-layer address of an attached physical interface when responding to the query, and the attached host or router forwards the packet to z/OS Communications Server TCP/IP. This eliminates the need to define static routes for VIPAs at hosts and routers attached to the same LAN as z/OS Communications Server TCP/IP. By using a prefix that is not being used for address autoconfiguration, the network prefix is not used by hosts for autoconfiguring addresses for physical interfaces.

Selecting the interface identifier

The VIPA interface identifier must be unique among all IP addresses that are created using the combination of network prefix and interface identifier. Any scheme can be used in generating the interface identifiers, as long as they are unique. By using a network prefix that is not used by stateless address autoconfiguration, it is only necessary to ensure the interface identifier is unique among all VIPAs that are sharing the same network prefix.

Effects of site renumbering on static VIPAs

When renumbering a site, new network prefixes are assigned to subnetworks. The existing network prefixes are marked as deprecated, during which time either the new prefixes or the old, deprecated prefixes can be used. After some time period, the deprecated network prefixes are deleted, along with all IPv6 addresses which use the network prefix.

For autoconfigured addresses, this process is automatically managed by stateless address autoconfiguration algorithms. For manually defined addresses, including all VIPAs, the process must be managed manually. When a prefix is to be deprecated, addresses that use the prefix should be deprecated using the INTERFACE DEPRADDR statement. After the prefix has expired, addresses that use the prefix should be deleted using the INTERFACE DELADDR statement.

Updating DNS definitions

This section describes considerations for updating DNS definitions.

Including static VIPAs in DNS

Include static VIPAs in DNS, in both the forward and reverse zones. If VIPAs are used, it is unnecessary to include IPv6 addresses assigned to interfaces.

Requirement: IPv6 Enterprise Extender requires that hostname resolution be used for the static VIPA. This hostname resolution can be from a DNS or a local hosts file (/etc/hosts).

Defining IPv4-only host names and IPv4/IPv6 host names

In general, IPv6 connectivity between two hosts is preferred over IPv4 connectivity. In many cases, IPv4 is used only if one of the nodes does not support IPv6. This can lead to undesirable paths in the network being used for communication between two hosts. For instance, when a native IPv6 path does not exist, data can be tunneled over the IPv4 network, even when a native IPv4 path exists.

This can lead to longer connection establishment to an AF_INET application which resides on a dual-stack host. The client first attempts to connect using each IPv6 address defined for the dual-stack host before attempting to connect with IPv4. A well-behaved client cycles through all the addresses returned and ultimately, connects using IPv4. However, this takes both time and network resources to accomplish, and not all clients are well-behaved or bug-free.

To avoid undesirable tunneling, as well as other potential problems, configure two host names in DNS. The existing host name should continue to be used for IPv4 connectivity, so as to minimize disruption when connecting to unmodified AF_INET server applications. A new host name should also be defined, for which both IPv4 and IPv6 should be configured. When connecting using the old host name, AF_INET6 clients connect using IPv4. When connecting using the new host name, AF_INET6 clients attempt to connect using IPv6 and, if that fails, falls back and connects using IPv4.

Using two host names allows the client to choose the network path that is taken. The client can route over IPv6 when the destination application is IPv6 enabled and a native IPv6 path exists, or take an IPv4 path.

The use of distinct host names for IPv4 and IPv4/IPv6 addresses is not strictly required. A single host name can be used to resolve to both IPv4 and IPv6 addresses. In addition, the use of distinct host names is only necessary during the initial transition phase when native IPv6 connectivity does not exist and applications have not yet been enabled for IPv6. After both of these occur, a single host name can be used.

Using source VIPA

Use a VIPA, either static or dynamic, be used as the source IP address on IPv6 hosts. Using a VIPA allows an IPv6 address to be resolved to a host name, assuming the guidelines in “Updating DNS definitions” are implemented. Define the VIPA using any of the following available configuration statements:

- SOURCEVIPAINT parameter on the INTERFACE statement
- TCPSTACKSOURCEVIPA parameter on the IPCONFIG6 statement

- SRCIP statement

Refer to Virtual IP Addressing in the *z/OS Communications Server: IP Configuration Guide* for additional information on choosing an appropriate method for specifying a source VIPA.

Using OMPROUTE or define static routes to improve network selection

The IPv6 OSPF or IPv6 RIP dynamic routing protocol provided by the OMPROUTE routing daemon should be used to provide information about the IPv6 prefixes and hosts that can be accessed indirectly by way of adjacent routers. IPv6 OSPF or IPv6 RIP can be used, either alone or together with IPv6 router discovery, to provide complete routing information.

If the IPv6 OSPF or IPv6 RIP dynamic routing protocol of OMPROUTE is not being used, the only routes that are learned (by way of router discovery) that can be used to access hosts that are not on directly attached links are default routes. Hosts can then use the default routes when sending packets to remote hosts. If a host selects a non-optimal router when sending data, the router can redirect the host to use a more optimal router when sending data to the remote host, as long as the optimal router is on the same LAN as the original router.

When a host is connected to multiple LANs, this processing might result in the following situations:

- A non-optimal router is used
- A router is used that cannot reach the final destination

For instance, if a host selects a router on one LAN, but the optimal router is on another LAN, the router on the first LAN cannot redirect the host to the second LAN. In this case, configure a static route to allow the host to initially select the optimal network path.

Guidelines: When defining static routes, use the following guidelines:

Use subnet routes instead of host routes

Remote IP addresses are difficult to predict. When using extensions to stateless address autoconfiguration, some clients can change their IP addresses on a routine basis, such as once an hour or once a day. In addition, these addresses can be created using cryptographic algorithms, making it difficult to impossible to predict which IP address a client might use. Defining static host routes to be used when communicating with such a client is equally as difficult or impossible.

Instead of defining a host route, define subnet routes. The network prefixes used in generating IPv6 addresses are much more stable than the interface identifiers used by hosts, typically changing only when a site is renumbered.

Use the link-local address of gateway router

When defining the gateway router for a static route, use the link-local address for the router. Link-local addresses do not change as the result of site renumbering, minimizing potential updates to the static routes. This is required in order to honor and process an ICMPv6 redirect message.

Effects of site renumbering on static routes

When a remote site is renumbered, new network prefixes are defined for the remote site and the old network prefixes are deprecated. After a time period, the old network prefixes are deleted.

A static route to a remote subnet should be created when a prefix is defined and should remain as long as the prefix is either preferred or deprecated. Only when the remote prefix is deleted should the static route be deleted.

Connecting to non-local IPv4 locations

If native IPv6 connectivity does not exist between two IPv6 sites, IPv6 over IPv4 tunneling can be used to provide IPv6 connectivity to the two sites. z/OS Communications Server TCP/IP can make use of an IPv6 over IPv4 tunnel to send packets to a remote site, but cannot be used as a tunnel endpoint itself. Instead, an intermediate router which supports IPv6 over IPv4 tunneling must act as the tunnel endpoint.

See “Enabling IPv6 communication between IPv6 islands in an IPv4 environment” on page 39 for more information on IPv6 over IPv4 tunnels.

IPv6-only application access to IPv4-only application

When an IPv6-only application needs to communicate with an IPv4-only host or application, some form of IPv6-to-IPv4 translation or application-layer gateway must occur. If needed, an outboard protocol translator or application-layer gateway component must be used, as z/OS Communications Server TCP/IP does not include such support. There are various technologies which can be used, such as NAT-PT or SOCKS64. See “Application Layer Gateways (ALG) and protocol translation” on page 43 for more information.

Part 3. Application enablement

Before reading this part, you should have a good understanding of the information presented in Part 1, “IPv6 overview,” on page 1.

This section contains the following chapters:

Chapter 6, “API support,” on page 71 describes the various z/OS socket APIs and the level of IPv6 present for each API.

Chapter 7, “Basic socket API extensions for IPv6,” on page 75 describes basic socket API changes that most applications would use.

Chapter 8, “Enabling an application for IPv6,” on page 89 describes common issues and considerations involved in enabling existing IPv4 socket applications for IPv6 communications.

Chapter 9, “Advanced socket APIs,” on page 99 discusses advanced IPv6 API functions that can be used by specialized IP applications.

For detailed information on specific APIs, refer to the following documentation:

- TCP/IP socket APIs are defined in the *z/OS Communications Server: IP Sockets Application Programming Interface Guide and Reference*.
- UNIX Language Environment C/C++ socket APIs are defined in the *z/OS XL C/C++ Run-Time Library Reference*.
- UNIX System Services Callable APIs are defined in the *z/OS UNIX System Services Programming: Assembler Callable Services Reference*.

Chapter 6. API support

This chapter describes API support and contains the following sections:

- “Native TCP/IP socket APIs” on page 72
- “Native TCP/IP socket APIs” on page 72

z/OS provides a versatile and diverse set of socket API libraries to support the various z/OS application environments. Figure 15 illustrates the relationship of the various z/OS socket APIs and the level of IPv6 present for each API.




Figure 15. z/OS socket APIs

The following are the two main socket API execution environments in z/OS:

- UNIX [implemented by UNIX System Services (Language Environment)]
- Native TCP/IP (implemented by TCP/IP in z/OS CS)

UNIX socket APIs

This section contains information about UNIX socket APIs.

z/OS UNIX Assembler Callable Services

z/OS UNIX Assembler Callable Services is a generalized call-based interface to z/OS UNIX IP sockets programming. This API supports both IPv4 and IPv6 communications. It includes support for the basic IPv6 API features and for a subset of the advanced IPv6 API features. For more information, refer to the *z/OS UNIX System Services Programming: Assembler Callable Services Reference*.

z/OS C sockets

z/OS UNIX C sockets is used in the z/OS UNIX environment. Programmers use this API to create applications that conform to the POSIX or XPG4 standard (a

UNIX specification). This API supports both IPv4 and IPv6 communications. It includes support for the basic IPv6 API features and for a subset of the advanced IPv6 API features. For more information on this API, refer to the *z/OS XL C/C++ Run-Time Library Reference*.

Native TCP/IP socket APIs

The following TCP/IP Services APIs are included in this library. For more information on these APIs (excluding CICS[®]), refer to *z/OS Communications Server: IP Sockets Application Programming Interface Guide and Reference*.

Sockets Extended macro API

The Sockets Extended macro API is a generalized assembler macro-based interface to IP socket programming. It includes support for IPv4 and for the basic IPv6 socket API functions.

Sockets Extended Call Instruction API

The Sockets Extended Call Instruction API is a generalized call-based interface to IP sockets programming. It includes support for IPv4 and for the basic IPv6 socket API functions.

REXX sockets

The REXX sockets programming interface implements facilities for IP socket communication directly from REXX programs by way of an address rxsocket function. It includes support for IPv4 and for the basic IPv6 socket API functions.

CICS sockets

The CICS socket interface enables you to write CICS applications that act as clients or servers in a TCP/IP-based network. Applications can be written in C language, using the C sockets programming interface, or they can be written in COBOL, PL/I, or assembler, using the Extended Sockets programming interface. This API supports TCP/IP communications over IPv4 and basic IPv6 socket API functions. For more information, refer to the *z/OS Communications Server: IP CICS Sockets Guide*.

IMS sockets

The Information Management System (IMSTM) socket interface supports development of client/server applications in which one part of the application executes on a TCP/IP-connected host and the other part executes as an IMS application program. The programming interface used by both application parts is the socket programming interface. This API currently supports TCP/IP communications over IPv4 only, but will probably support IPv6 communications in a future release. For more information, refer to *z/OS Communications Server: IP IMS Sockets Guide*.

Pascal API

The Pascal socket application programming interface enables you to develop TCP/IP applications in the Pascal language. It only supports TCP/IP communications over IPv4. It is unlikely that this API will be enhanced to support IPv6 in the future. Applications using this API are encouraged to migrate their application to one of the other socket APIs that are IPv6 enabled.

TCP/IP C/C++ Sockets

The C/C++ Socket interface supports IPv4 socket function calls that can be invoked from C/C++ programs. This API is very similar to the UNIX C socket API that is the recommended socket API for C/C++ application development on z/OS. The TCP/IP C/C++ sockets API will not be enhanced for IPv6 support. Existing applications that will be enabled for IPv6 should consider migrating to the UNIX C socket API.

There are several higher level C/C++ APIs that rely on the TCP/IP sockets for communications over an IP network, including the following:

- Resource Reservation Setup Protocol API (RAPI)
- Sun and NCS Remote Procedure Call (RPC)
- X Window System and Motif
- X/Open Transport Interface (XTI)

These APIs do not support IPv6 communications.

Guideline: CICS programs written to use the IP CICS C Sockets API must use the TCP/IP C headers. Include the following definition to expose the required IPv6 structures, macros, and definitions in the header files:

```
#define __CICS_IPV6
```

Refer to C Language application programming in the *z/OS Communications Server: IP CICS Sockets Guide* for guidance on using the IP CICS C Sockets API.

Chapter 7. Basic socket API extensions for IPv6

This chapter describes the basic extensions to the socket interface and new features of IPv6 as described in the Internet Engineering Task Force (IETF) RFC 3493, *Basic Socket Interface Extensions for IPv6*, and contains the following sections:

- “Introduction”
- “Design considerations”
- “Name and address resolution functions” on page 77
- “Interface identification” on page 84
- “Socket options to support IPv6 (IPPROTO_IPV6 level)” on page 84

Note: All examples in this chapter are shown using UNIX Language Environment C; see *z/OS XL C/C++ Run-Time Library Reference* for details.

Introduction

IPv4 addresses are 32 bits long, but IPv6 interfaces are identified by 128-bit addresses. The socket interface makes the size of an IP address visible to an application; virtually all TCP/IP applications using sockets have knowledge of the size of an IP address. Those parts of the API that expose the addresses must be changed to accommodate the larger IPv6 address size. IPv6 also introduces new features, some of which must be made visible to applications by way of the API.

Design considerations

The two main programming tasks associated with IPv6 exploitation involve migrating existing application programs to support IPv6 and designing new programs for IPv6. In both cases, the changed or new code should be designed so that it is capable of using IPv4 or IPv6 addresses. Servers should be designed so that they can communicate with both IPv4 and IPv6 clients. Existing IPv4 client and server programs should continue to operate properly as long as only IPv4 connectivity is required between clients and servers.

The following sections describe key differences between IPv4 and IPv6.

Requirement: It is assumed that you have a basic knowledge of IPv4 socket programming for clients and servers.

Protocol families

IPv4 socket applications use a AF_INET (equivalent to PF_INET) protocol family. For IPv6, a new protocol family of AF_INET6 (equivalent to PF_INET6) has been defined. The protocol family is the first parameter to the socket() function that is used to obtain a socket descriptor. For most applications, an AF_INET6 socket can be used to communicate with IPv4 and IPv6 clients.

Address families

Most socket functions require a socket descriptor and a generic socket address structure called a sockaddr. The exact format of the sockaddr structure depends on the address family. For IPv4 sockets, the sockaddr structure is sockaddr_in. For IPv6, the sockaddr structure sockaddr_in6 is used.

The following socket functions have a sockaddr as one of their parameters:

```
bind()  
connect()  
sendmsg()  
sendto()  
accept()  
recvfrom()  
recvmsg()  
getpeername()  
getsockname()
```

The sockaddr structure that is used in these functions must be the proper structure for the socket family.

For IPv4 (AF_INET), the sockaddr (sockaddr_in) contains the information shown in Table 7.

Table 7. sockaddr format for AF_INET

sockaddr length	1 byte	Not used, should be set to 0
family	1 byte	AF_INET
port	2 bytes	TCP or UDP port number
IP address	4 bytes	IPv4 IP address
reserved	8 bytes	Not used

For IPv6 (AF_INET6), the sockaddr (sockaddr_in6) contains additional information. Also, note that the IP address for IPv6 is 16 bytes long instead of 4 bytes long as in IPv4.

Table 8. sockaddr format for AF_INET6

sockaddr length	1 byte	Not used, should be set to 0
family	1 byte	AF_INET6
port	2 bytes	TCP or UDP port number (same as v4)
flowinfo	4 bytes	Flow information
IP address	16 bytes	IPv6 IPaddress
scope ID	4 bytes	Used to determine IP address scope

Special IP addresses

Like IPv4, IPv6 also defines loopback and wildcard (INADDR_ANY) addresses. The differences are shown in Table 9.

Table 9. Special IP addresses

	IPv4	IPv6
Loopback address	127.0.0.1	::1 (15 bytes of zeros, 1 byte of 1)
Wildcard address	0.0.0.0	:: (16 bytes of zeros)
Multicast address	224.0.0.1 - 239.255.255.255	Refer to "Multicast IPv6 Addresses" on page 17

Name and address resolution functions

IPv6 introduces new APIs for the Resolver function. These APIs allow applications to resolve host names to IP addresses and vice versa. The primary new APIs are `getaddrinfo`, `getnameinfo`, and `freeaddrinfo`. The APIs are designed to work with both IPv4 and IPv6 addressing. The use of these new APIs should be considered if an application is being designed for eventual use in an IPv6 environment.

The way in which hostname (`getaddrinfo`) or IP address (`getnameinfo`) resolution is performed depends on the Resolver specifications contained in the Resolver setup files and TCPIP.DATA configuration files. These specifications determine whether the APIs query a name server first, then search the local host tables, or whether the order is reversed, or even if one of the steps is eliminated completely. The specifications also control, if local host tables have to be searched, which tables that are accessed. For detailed information about Resolver setup, see “Resolver configuration” on page 56.

Protocol-independent nodename and service name translation

`Getaddrinfo` is conceptually a replacement for the existing `gethostbyname` and `getservbyname` APIs. `Getaddrinfo` takes an input hostname, or an input servicename, or both, and returns (when resolution is successful) one or more `addrinfo` structures. `Getaddrinfo` can also accept as input, a hostname or a servicename in numeric form, and returns the same value in presentation form using the `addrinfo` structure. An `addrinfo` structure contains the following output information:

- Pointer to `sockaddr_in` or `sockaddr_in6` structure containing an IP address and service port
- Length of `sockaddr` structure and family type (`AF_INET`, `AF_INET6`) of the `sockaddr` structure
- Socktype and protocol values usable with this `sockaddr` structure
- Pointer to canonical name associated with the input hostname (applicable only in the first `addrinfo` structure)
- Pointer to next `addrinfo` structure (set to 0 in the last element of the chain)

The storage for the `addrinfo` structures is allocated by the Resolver from the application’s address space, and the application should use the `freeaddrinfo` API to release the `addrinfo` structures when the information is no longer required. The application should not manipulate the chain of `addrinfo` structures returned by way of `getaddrinfo`, but rather that the application simply return the entire chain, as received, back to the Resolver by way of `freeaddrinfo`.

In addition to hostname or servicename, one of which must be present on a valid `getaddrinfo` invocation, the application can specify additional input to the Resolver on the `getaddrinfo` invocation. This input is optional, and if specified, is passed by way of an input `addrinfo` structure. The input settings include the following possibilities:

- Family type of `sockaddr` structure required on output.
- Socktype and protocol values for which the returned IP address and port number must work. This would be used primarily for cases where a *servicename* was being resolved, as might typically have been done previously by way of `getservbyname`.
- Various input flag settings include the following:
 - `AI_ADDRCONFIG`

- AI_ALL
- AI_CANONNAME
- AI_NUMERICHOST
- AI_NUMERICSERV
- AI_PASSIVE
- AI_V4MAPPED

In the absence of any specific input from the application, the Resolver assumes that any sockaddr type is acceptable (that is, both IPv4 and IPv6 addresses) as output. Thus, by default, the Resolver searches for both IPv6 and IPv4 address by way of DNS or by way of local host files (such as /etc/hosts). Obviously, this might not always be the best choice for the application issuing `getaddrinfo`. By using the above input fields, an application issuing `getaddrinfo()` can influence the processing performed by the Resolver function for that given request in the following ways:

- The application can specify that the sockaddr returned by `getaddrinfo` should be of family type AF_INET, AF_INET6 or AF_UNSPEC (meaning either family type would be acceptable). For example, if AF_INET is specified, the Resolver does not perform any searches for IPv6 addresses for *hostname*, because the output requested must be an IPv4 address.
- The application can specify the following:
 - Both IPv6 and IPv4 addresses should be returned
 - IPv4 should be returned only if there are no IPv6 addresses resolved
 - Only IPv6 addresses should be returned
 - Only IPv4 addresses should be returned.

This information, indicated by the input combination of family type and the AI_ALL and AI_V4MAPPED flags, to a large extent controls the types of searches performed by the Resolver during the course of the processing.

- The application can specify that IPv6 addresses should be returned only when the system has IPv6 interfaces defined and can specify that IPv4 addresses should be returned only when IPv4 interfaces are defined. This preference, indicated by way of the AI_ADDRCONFIG flag, allows the application to eliminate resolution searches looking for addresses that cannot be used by the application.
- The application can specify whether the sockaddr returned should contain an address for passive (that is, the INADDR_ANY address) or active (that is, the loopback address) socket activation. This choice is indicated by way of the AI_PASSIVE flag, and is applicable only in the absence of an input *hostname*.
- The application can specify that only translation from presentation to numeric format should be performed for *hostname*, or service name, or both. This option is indicated by setting the AI_NUMERICHOST (for *hostname*) or AI_NUMERICSERV (for *servicename*) flags, which indicate that the associated input value must be in numeric format or the `Getaddrinfo` request should be failed.
- The application can specify that only a given socktype or protocol value should be used for looking up the port number associated with the input *servicename*, or can request that all valid socktypes and protocols (TCP and UDP) be used for the `getservbyname` processing. This preference is indicated by way of the socktype and protocol settings.

With such a flexible interface, the application programmer must decide what inputs are reasonable for the capabilities of the application being created or modified. The most likely application uses are the following.

Table 10 shows the two most likely application usages and the suggested getaddrinfo input settings that coincide with that functionality:

- IPv6-capable when the underlying system is IPv6 capable
- IPv4-capable only

Table 10. Getaddrinfo application capabilities 1

Application capabilities	Sockaddr family to request	Additional flags to set	Expected outputs
(IPv4 only) Application is pure IPv4 and cannot handle any IPv6 addresses.	AF_INET	AI_ADDRCONFIG	Getaddrinfo returns one or more addrinfo structures, each pointing to an IPv4 address saved in an AF_INET sockaddr. No addrinfos are returned if there is no IPv4 interfaces defined on the system. No searches of any kind are performed for IPv6 addresses as part of this request.
(IPv6 capable) Application wants all known addresses for hostname, in IPv6 format when the system supports IPv6, or in IPv4 format otherwise.	AF_UNSPEC	AI_ADDRCONFIG, AI_ALL -or -AI_ADDRCONFIG, AI_V4MAPPED, AI_ALL	Getaddrinfo returns one or more addrinfo structures, each pointing to a sockaddr structure. The sockaddrs consists of one of the following sets: <ul style="list-style-type: none"> • All AF_INET6 sockaddrs, containing IPv6 or mapped IPv4 addresses, if the system supports IPv6 processing (only when AI_V4MAPPED coded). • AF_INET6 sockaddrs, containing IPv6 addresses, and AF_INET sockaddrs, containing IPv4 addresses, if the system supports IPv6 processing (only when AI_V4MAPPED is NOT coded). • All AF_INET sockaddrs, containing IPv4 addresses, if the system does not support IPv6 processing. In all cases, the IPv6 addresses are returned only if there is an IPv6 interface defined on the system, and the IPv4 addresses are returned only if there is an IPv4 interface defined.

An application with no interest in utilizing IPv6 wants to utilize the first entry in Table 10. Otherwise, if there is some interest in utilizing IPv6 functionality, an application would achieve the greatest flexibility by using the second table entry. Using the IPv6 entry approach, the application places the burden of supplying a workable sockaddr structure on the Resolver logic. If IPv6 is supported on the system, the Resolver endeavors to return AF_INET6 sockaddrs to the application; otherwise, the Resolver returns AF_INET sockaddrs to the application. The choice of coding or not coding AI_V4MAPPED in this situation depends on the

application's preference regarding receiving AF_INET6 sockaddrs: the more the application wants to deal exclusively with AF_INET6 sockaddrs, the more reason to code AI_V4MAPPED.

Table 10 on page 79 should be sufficient for most application usages. However, there are other likely application capability models possible, and Table 11 provides some guidance on how to code the Getaddrinfo invocations for those applications.

Table 11. Getaddrinfo application capabilities 2

Application capabilities	Sockaddr family to request	Additional flags to set	Expected outputs
Application is pure IPv6 and cannot handle any mapped IPv4 addresses.	AF_INET6	AI_ADDRCONFIG	Getaddrinfo returns one or more addrinfo structures, each pointing to an IPv6 address saved in an AF_INET6 sockaddr. No addrinfos is returned if there is no IPv6 interfaces defined on the system. No searches of any kind are performed for IPv4 addresses as part of this request.
Application prefers IPv6 addresses, requires IPv6 address format, but can handle mapped IPv4 addresses if necessary.	AF_INET6	AI_ADDRCONFIG, AI_V4MAPPED	Getaddrinfo returns one or more addrinfo structures, each pointing to an AF_INET6 sockaddr. The addresses within the sockaddrs consists of one of the following sets: <ul style="list-style-type: none"> • All IPv6 addresses, if there is an IPv6 interface defined on the system and IPv6 addresses exist for hostname • All mapped IPv4 addresses, if there were no IPv6 addresses to be returned for hostname and there was an IPv4 interface defined for the system
Application prefers IPv6 addresses, but can handle native IPv4 addresses if necessary.	AF_UNSPEC	AI_ADDRCONFIG	Getaddrinfo returns one or more addrinfo structures, each pointing to a sockaddr structure. The sockaddrs consists of one of the following sets: <ul style="list-style-type: none"> • All AF_INET6 sockaddrs, containing IPv6 addresses, if there is an IPv6 interface defined on the system and IPv6 addresses exist for hostname • All AF_INET sockaddrs containing IPv4 addresses, if there were no IPv6 addresses to be returned for hostname and there was an IPv4 interface defined for the system

Table 11. *Getaddrinfo* application capabilities 2 (continued)

Application capabilities	Sockaddr family to request	Additional flags to set	Expected outputs
Application wants all known addresses for hostname, in IPv6 format.	AF_INET6	AI_ADDRCONFIG, AI_V4MAPPED, AI_ALL	Getaddrinfo returns one or more addrinfo structures, each pointing to an AF_INET6 sockaddr. The addresses within the sockaddrs consists of all IPv6 addresses, if there is an IPv6 interface defined on the system and mapped IPv4 addresses, if there is an IPv4 interface defined for the system, associated with hostname.
Application wants all known addresses for hostname, in native (IPv6 or IPv4) format.	AF_UNSPEC	AI_ADDRCONFIG, AI_ALL	Getaddrinfo returns one or more addrinfo structures, each pointing to a sockaddr structure. The sockaddr structures are a mixture of AF_INET6 sockaddrs (each containing an IPv6 address) and AF_INET sockaddrs (each containing an IPv4 address). The IPv6 addresses are returned only if there is an IPv6 interface defined on the system, and the IPv4 addresses are returned only if there was an IPv4 interface defined for the system.
Application wants all known addresses for hostname, regardless of system connectivity, in native format.	AF_UNSPEC	AI_ALL	Getaddrinfo returns one or more addrinfo structures, each pointing to a sockaddr structure. The sockaddr structures can be a mixture of AF_INET6 sockaddrs (each containing an IPv6 address) or AF_INET sockaddrs (each containing an IPv4 address), depending on the address resolution.
Default settings when IPv6 is enabled on the system.	AF_UNSPEC	NONE	<p>Getaddrinfo returns one or more addrinfo structures, each pointing to a sockaddr structure. The sockaddrs consists of one of the following sets:</p> <ul style="list-style-type: none"> • All AF_INET6 sockaddrs, containing IPv6 addresses, if there is an IPv6 address defined for hostname in any queried domain name server or defined in a local hosts table. No searches for IPv4 addresses are performed for hostname. • All AF_INET sockaddrs, containing IPv4 addresses, if there are no IPv6 addresses found for hostname. <p>In either case, the actual availability of IPv6 or IPv4 interfaces on the system is not taken into consideration.</p>

Table 11. Getaddrinfo application capabilities 2 (continued)

Application capabilities	Sockaddr family to request	Additional flags to set	Expected outputs
Default settings when IPv6 is not enabled on the system.	AF_UNSPEC	NONE	Getaddrinfo returns one or more addrinfo structures, each pointing to a sockaddr structure. The sockaddr structures can be a mixture of AF_INET6 sockaddrs (each containing an IPv6 address) or AF_INET sockaddrs (each containing an IPv4 address), depending on the address resolution performed. The actual availability of IPv6 or IPv4 interfaces on the system is not taken into consideration.

Regardless of the application model in use, and because output from getaddrinfo can be a chain of addrinfo structures, the application should attempt to use each address, in the order received, to open a socket and connect or send a datagram to the target host name until it is successful, versus simply using the first address and stopping if a failure is encountered.

The application is now responsible for freeing the storage (addrinfo and sockaddr structures, and so on) associated with the new resolver APIs. The new freeaddrinfo API should be used to free this storage. If the application neglects to perform this step, the resolver cleans up the storage when the process terminates, but storage constraints might occur before termination if a large number of getaddrinfo APIs are performed.

Socket address structure to host name and service name

Conceptually, Getnameinfo is a replacement for the existing gethostbyaddr and getservbyport APIs. Getnameinfo takes an input *IP address* or an input *port number*, or both, and returns (when resolution is successful) the host name or the service location. These parameters are passed in a sockaddr structure that also contains the address family.

In addition to *IP address* or *port number*, one of which must be present on a valid getnameinfo invocation, the application can specify additional input to the Resolver on the getnameinfo invocation. This input is optional. The input settings include the following (various input flag settings can be specified):

NI_NOFQDN

Specifies that only the host name portion of the fully qualified domain name (FQDN) is returned for local hosts.

NI_NUMERICHOST

Specifies that the numeric form of the host name, its IP address, is returned instead of its name. No resolution takes place for the specified input if the NI_NUMERICHOST flag is on.

NI_NUMERICSERV

Specifies that the numeric form of the service name, the port number, is returned instead of the service name. No resolution takes place for the specified input if the NI_NUMERICSERV flag is on.

NI_NAMEREQD

Specifies that an error is returned if the host name cannot be located. (If NI_NAMEREQD is not specified, the numeric form of the host name, the IP address, is returned).

NI_DGRAM

Specifies that the service is a datagram service (SOCK_DGRAM). The default behavior assumes that the service is a stream service.

Address conversion functions

IP addresses often need to be given to a socket application in character (string) format. It is also common for socket applications to need to display IP addresses in string format. The following functions work for IPv4 and IPv6 addresses:

inet_ntop

Convert a binary IP address (either v4 or v6) into string format.

inet_pton

Convert an IP address in string format to binary format.

The functions `inet_ntoa` and `inet_addr` are still available, but they cannot be used for IPv6 addresses.

Table 12. Address conversion functions

Function	z/OS UNIX Assembler Callable services	C/C++ using Language Environment	IP CICS C sockets	REXX	Socket Extended macro/call (includes CICS EZASOKET)
inet_pton	No	Yes	Yes	No	No
inet_ntop	No	Yes	Yes	No	No
PTON	No	No	No	No	Yes
NTOP	No	No	No	No	Yes

Address testing macros

The macros listed in Table 13 can be used to test for special IPv6 addresses.

Table 13. Address testing macros

Macros	Assembler Callable services	C/C++ using Language Environment	IP CICS C sockets	REXX	Socket Extended macro/call (includes CICS EZASOKET)
IN6_IS_ADDR_UNSPECIFIED	No	Yes	Yes	No	No
IN6_IS_ADDR_LOOPBACK	No	Yes	Yes	No	No
IN6_IS_ADDR_MULTICAST	No	Yes	Yes	No	No
IN6_IS_ADDR_LINKLOCAL	No	Yes	Yes	No	No
IN6_IS_ADDR_SITELOCAL	No	Yes	Yes	No	No
IN6_IS_ADDR_V4MAPPED	No	Yes	Yes	No	No
IN6_IS_ADDR_V4COMPAT	No	Yes	Yes	No	No
IN6_IS_ADDR_MC_NODELOCAL	No	Yes	Yes	No	No
IN6_IS_ADDR_MC_LINKLOCAL	No	Yes	Yes	No	No
IN6_IS_ADDR_MC_SITELOCAL	No	Yes	Yes	No	No
IN6_IS_ADDR_MC_ORGLOCAL	No	Yes	Yes	No	No
IN6_IS_ADDR_MC_GLOBAL	No	Yes	Yes	No	No

The macros function in the following ways:

- The first seven macros return true if the address is of the specified type, or false otherwise.
- The last five macros test the scope of a multicast address and return true if the address is a multicast address of the specified scope, or false if the address is either not a multicast address or not of the specified scope.
- IN6_IS_ADDR_LINKLOCAL and IN6_IS_ADDR_SITELOCAL return true only for the two types of local-use IPv6 unicast addresses (link-local and site-local), and that by this definition, the IN6_IS_ADDR_LINKLOCAL macro returns false for the IPv6 loopback address (::1). These two macros do not return true for IPv6 multicast addresses of either link-local scope or site-local scope.

Interface identification

IPv6 interfaces can have many different IP addresses. IPv6 allows a socket application to specify an interface to use for sending data by specifying an interface index. Certain socket options allow specification an interface index. Also, socket options for IPv6 multicast join group and IPv6 multicast leave group allow optional specification of an interface index.

The function, `if_nameindex()`, allows socket applications to obtain a list of interface names and their corresponding index. Also, two functions, `if_nametoindex()` and `if_indextoname()` allow translation of an interface name to its index and translation of an interface index to an interface name. The function, `if_freenameindex()`, is used to free dynamic storage allocated by the `if_nameindex()` function.

For non-C/C++ (Language Environment applications), a new ioctl function code (SIOCGIFNAMEINDEX) is provided. Use Table 14 to determine which APIs support this new ioctl.

Table 14. Function calls

Function/IOCTL	z/OS UNIX Assembler Callable services	C/C++ using Language Environment	IP CICS C sockets	REXX	Socket Extended macro/call (includes CICS EZASOKET)
<code>if_nametoindex</code>	No	Yes	Yes	No	No
<code>if_indextoname</code>	No	Yes	Yes	No	No
<code>if_nameindex</code>	No	Yes	Yes	No	No
<code>SIOCGIFNAMEINDEX</code>	Yes	No	No	Yes	Yes
<code>if_freenameindex</code>	No	Yes	Yes	No	No

Socket options to support IPv6 (IPPROTO_IPV6 level)

A group of socket options is defined to support IPv6. They are defined with a level of IPPROTO_IPV6. The individual options begin with IPV6_.

Restriction: These options are only allowed on AF_INET6 sockets.

In most cases, an IPV6_xxx option can be set on an AF_INET6 socket that is using IPv4-mapped IPv6 addresses but have no effect. For example, the

IPV6_UNICAST_HOPS socket option is used to set a hop limit value in the IPv6 header. Because IPv4 packets are used with IPv4-mapped IPv6 addresses, the hop limit value is not used.

Guideline: The Sockets Extended macro/call APIs do not use level as an input to `getsockopt()` and `setsockopt()`. However, other IPv6-enabled APIs do use level as input. For detailed information about `setsockopt()` and `getsockopt()` input and output, refer to the API-specific documentation.

Table 15. Socket options for `getsockopt()` and `setsockopt()`

Socket options <code>getsockopt()</code> <code>setsockopt()</code>	z/OS UNIX Assembler Callable services	C/C++ using Language Environment	IP CICS C sockets	REXX	Sockets Extended macro/call (includes CICS EZASOKET)
IPV6_UNICAST_HOPS	Yes	Yes	Yes	Yes	Yes
IPV6_MULTICAST_IF	Yes	Yes	Yes	Yes	Yes
IPV6_MULTICAST_LOOP	Yes	Yes	Yes	Yes	Yes
IPV6_MULTICAST_HOPS	Yes	Yes	Yes	Yes	Yes
IPV6_JOIN_GROUP	Yes	Yes	Yes	Yes	Yes
IPV6_LEAVE_GROUP	Yes	Yes	Yes	Yes	Yes
IPV6_V6ONLY	Yes	Yes	Yes	Yes	Yes

Option to control sending of unicast packets

Use the following option to control sending of unicast packets:

IPV6_UNICAST_HOPS

The IPv6 header contains a hop limit field that controls the number of hops over which a datagram can be sent before being discarded. This is similar to the TTL field in the IPv4 header. The IPV6_UNICAST_HOPS socket option can be used to set the default hop limit value for an outgoing unicast packet. The socket option value should be between 0 and 255 inclusive. A socket option value of -1 is used to clear the socket option. This causes the stack default to be used.

A `getsockopt()` with this option returns the value set by a `setsockopt()`. If a `setsockopt()` has not been performed, the stack's default value is returned.

The HOPLIMIT parameter on the IPCONFIG6 statement influences the default hop limit when this socket option is not set. An application must be APF-authorized or have superuser authority to set this option to a value greater than the value of HOPLIMIT on the IPCONFIG6 statement. Refer to the *z/OS Communications Server: IP Configuration Reference* for more information about the IPCONFIG6 statement.

Tip: This function is similar to the IPv4 socket option IP_TTL.

Options to control sending of multicast packets

The following three options allow an application to control certain features in the transmission of IPv6 multicast packets. These socket options do not have to be set to send multicast packets. Supplying a multicast address as the destination address is the only thing required to send an IPv6 multicast packet.

IPV6_MULTICAST_IF

This socket option allows an application to control the outgoing interface used for a multicast packet. The socket option value is the interface index of the interface to be used.

A getsockopt() with this option returns the value set by setsockopt(). If a setsockopt() has not been done, a value of 0 is returned.

Tip: This function is similar to the IPv4 socket option IP_MULTICAST_IF.

IPV6_MULTICAST_HOPS

The IPv6 header contains a hop limit field that controls the number of hops over which a datagram can be sent before being discarded. This is similar to the TTL field in the IPv4 header. The IPV6_MULTICAST_HOPS socket option can be used to set the default hop limit value for an outgoing multicast packet. The socket option value should be between 0 and 255 inclusive. A socket option value of -1 is used to clear the socket option. This causes the default value of 1 to be used.

A getsockopt() with this option returns the value set by a setsockopt(). If a setsockopt() has not been done, the default value of 1 is returned.

The default value is 1. An application must be APF-authorized or have superuser authority to set this option to a value greater than the value of HOPLIMIT on the IPCONFIG6 statement. Refer to the *z/OS Communications Server: IP Configuration Reference* for more information on the IPCONFIG6 statement.

Tip: This function is similar to the IPv4 socket option IP_MULTICAST_TTL.

IPV6_MULTICAST_LOOP

When a multicast packet is sent, if the sender belongs to the multicast group to which the packet was sent, then this option controls whether the sender receives a copy of the packet or not. If this option is enabled, then the sender receives a copy of the packet. The socket option value should be 1 to enable the option, or 0 to disable the option.

A getsockopt() with this option returns the value set by a setsockopt(). If a setsockopt() has not been done, the default value of 1 (enabled) is returned.

Tip: This function is similar to the IPv4 socket option IP_MULTICAST_LOOP.

Options to control receiving of multicast packets

Use the following option to control receiving of multicast packets:

IPV6_JOIN_GROUP

This socket option allows an application to join a multicast group on a specific local interface. The socket option data specifies an IPv6 multicast address and an IPv6 interface index. IPv4-mapped IPv6 multicast addresses are not supported. If an interface index of 0 is specified, the stack selects a local interface. An application that wants to receive multicast packets destined for a multicast group needs to join that group. It is not necessary to join a multicast group to send multicast packets.

Restriction: Getsockopt() does not support this option.

Tip: This function is similar to the IPv4 socket option IP_ADD_MEMBERSHIP.

IPV6_LEAVE_GROUP

This socket option is used by an application to leave a multicast group it previously joined. The socket option data specifies an IPv6 multicast address and an IPv6 interface index. If an interface index of 0 is used to join a multicast group, an interface index of 0 must be used to leave the group.

Restriction: Getsockopt() does not support this option.

Tip: This function is similar to the IPv4 socket option IP_DROP_MEMBERSHIP.

Socket option to control IPv4 and IPv6 communications

Use the following option to control IPv4 and IPv6 communications:

IPV6_V6ONLY

An AF_INET6 socket can be used for IPv6 communications, IPv4 communications, or a mix of IPv6 and IPv4 communications. The IPV6_V6ONLY socket option allows an application to limit an AF_INET6 socket to IPv6 communications only. A nonzero socket option value enables the option; a value of 0 disables the option.

A getsockopt() with this option returns the value set by a setsockopt(). If a setsockopt() has not been done, the default value of 0 (disabled) is returned.

If an application wants to enable this option, the setsockopt() must be set prior to binding the socket, connecting the socket, or sending data over the socket. This option cannot be changed (either enabled or disabled) after the socket has been bound. (An implicit bind is done for datagram sockets on connect or send operations if the socket is not already bound.)

Socket options for SOL_SOCKET, IPPROTO_TCP and IPPROTO_IP levels

Socket options at the SOL_SOCKET and IPPROTO_TCP levels are not dependent on the IP layer being used. They are supported for both AF_INET and AF_INET6 sockets.

Socket options at the IPPROTO_IP level support IPv4. They are not supported on AF_INET6 sockets.

Not all socket options at these levels are supported by all APIs. Check the API specific documentation for information on a specific socket option.

Chapter 8. Enabling an application for IPv6

This chapter describes how to enable an application for IPv6 and contains the following sections:

- “Changes to enable IPv6 support”
- “Support for unmodified applications”

Changes to enable IPv6 support

Several coding changes are needed to enable an application for IPv6 communications. Chapter 7, “Basic socket API extensions for IPv6,” on page 75 describes the changes to the basic Socket APIs that most applications use. Chapter 9, “Advanced socket APIs,” on page 99 describes the changes to advanced functions (which are typically used by a small number of TCP/IP applications) of the socket APIs that facilitate IPv6 communications. The sections in this chapter describe some of the general considerations involved in enabling an application for IPv6. Note that while many of the examples and references in this chapter assume the use of C/C++ sockets supported by the Language Environment (LE), most of the concepts (unless explicitly noted) apply to the other Socket API libraries that support IPv6. For a more detailed description of the actual APIs, see Chapter 7, “Basic socket API extensions for IPv6,” on page 75 and Chapter 9, “Advanced socket APIs,” on page 99 and the documentation for the specific API you are using.

Guideline: You should be familiar with IPv6 in general and IPv6 support on z/OS Communications Server.

Support for unmodified applications

During the transition period where networks, routers, and hosts are upgraded to support IPv6, it is expected that most IPv6-enabled hosts also continue to have IPv4 connectivity. This is accomplished with dual-mode stack support that allows a single TCP/IP protocol stack to support both IPv4 and IPv6 communications. TCP/IP on z/OS supports dual-mode stack operation. As a result, applications that are not IPv6-enabled continue to function over an IPv4 network, without any changes. However, at some point during the IPv6 deployment process, some IP hosts might only have connectivity to IPv6 networks or have a TCP/IP protocol stack that is capable of IPv6 communications only. Various migration and coexistence techniques can be employed to allow IPv6-only hosts to communicate with IPv4-only applications as described in “Migration and coexistence” on page 39. However, in the absence of these mechanisms, an application needs to be enabled for IPv6 in order to allow for communications with IPv6-only hosts or applications.

Application awareness of whether system is IPv6 enabled

A z/OS system might or might not be enabled for IPv6 communications. Enabling a z/OS system for IPv6 support requires explicit configuration by the system administrator to allow AF_INET6 sockets to be created. As a result, an application cannot typically assume that IPv6 is enabled on the systems where the application is running. Some exceptions do exist. For example, applications can run on a limited number of systems that are known to be IPv6 enabled. However, in general, most applications that are being enhanced to support IPv6 must first perform a run-time test to determine whether IPv6 is enabled on the system where

they are executing. If the system is not enabled for IPv6, the application should proceed with its existing IPv4 logic. If the system is enabled for IPv6, the application can now use AF_INET6 sockets and features to communicate with both IPv4 and IPv6 applications.

Determine if a system is enabled for IPv6 by attempting to create an AF_INET6 socket. If this operation is successful, the application can assume that IPv6 is enabled. If the operation fails (with return code EAFNOSUPPORT) the application should revert to its IPv4 logic and create an AF_INET socket.

Table 16. Using socket() to determine IPv6 enablement

Affected socket API call	Changes required
socket()	Specify AF_INET6 as the Address Family (or domain) parameter. This API call fails if the system is not enabled for IPv6.

The getaddrinfo() API is an alternative mechanism that can be used by TCP/IP client applications to determine whether IPv6 is enabled. This API is a replacement for the gethostbyname() API and is typically used by TCP/IP client programs to resolve a host name to an IP address. For example, a client application that receives the server application's host name or IP address (such as FTP) as input can invoke the getaddrinfo() function prior to opening up a socket with a selected set of options. This allows the application to receive a list of addrinfo structures (one for each IP address of the destination host) that contain the following information:

- The address family of the IP address (AF_INET or AF_INET6)
- A pointer to a socket address structure of the appropriate type (sockaddr_in or sockaddr_in6) that is fully initialized (including the IP address and Port fields)
- The length of the socket address structure

A client application can be coded with this information in a manner that allows it to be protocol-independent without having to perform specific run-time checks to determine whether IPv6 is enabled or not and without having to have dual-path logic (IPv4 versus IPv6). The following is an example of this approach:

```

int
myconnect(char *hostname)
{
    struct addrinfo *res, *aip;
    struct addrinfo hints;
    char buf[INET6_ADDRSTRLEN];
    static char *servicename = "21";
    int sock = -1;
    int error;

    /* Initialize the hints structure for getaddrinfo() call.
       This application can deal with either IPv4 or IPv6 addresses.
       It relies on getaddrinfo to return the most appropriate IP address
       and socket address structure based on the current configuration */

    bzero(&hints, sizeof (hints));
    hints.ai_socktype = SOCK_STREAM; /* Interested in streams sockets
                                      only */
    /* Note that we are asking for all IP addresses to be returned (IPv4
       or IPv6) based on the system connectivity. Also, note that we
       would prefer all addresses to be returned in sockaddr_in6 format
       if the system is enabled for IPv6. In addition, we also specify
       a numeric port using AI_NUMERICSERV so that the returned socket
       address structures are primed with our port number. */

    hints.ai_flags = AI_ALL | AI_V4MAPPED | AI_ADDRCONFIG |
                     AI_NUMERICSERV;
    hints.ai_family = AF_UNSPEC;
    error = getaddrinfo(hostname, servicename, &hints, &res);
    if (error != 0) {
        (void) fprintf(stderr,
                      "getaddrinfo: %s for host %s service %s\n",
                      gai_strerror(error), hostname, servicename);
        return (-1);
    }
    for (aip = res; aip != NULL; aip = aip->ai_next) {
    /*
     * Loop through list of addresses returned, opening sockets
     * and attempting to connect() until successful. The
     * The address type depends on what getaddrinfo()
     * gave us.
     */
        sock = socket(aip->ai_family, aip->ai_socktype,
                      aip->ai_protocol);
        if (sock == -1) {
            printf("Socket failed: %d\n", sock);
            freeaddrinfo(res);
            return (-1);
        }
        /* Connect to the host. */
        if (connect(sock, aip->ai_addr, aip->ai_addrlen) == -1) {
            printf("Connect failed, errno=%d, errno2=%08x\n",
                   errno, __errno2());
            (void) close(sock);
            sock = -1;
            continue;
        }
        break;
    }
    freeaddrinfo(res);
    return (sock);
}

```

Figure 16. Example of protocol-independent client application

When this example executes on a system where IPv6 is not enabled, only IPv4 addresses are returned in AF_INET format (in sockaddr_in structures). When this

identical example executes on a IPv6-enabled system, both IPv4 and IPv6 addresses are returned, and the IPv4 addresses are returned in IPv4-mapped IPv6 address format (in sockaddr_in6 structures). Note that an AF_INET6 socket can be used for the connection even when the address returned by getaddrinfo() is an IPv4-mapped IPv6 address.

Socket address (sockaddr_in) structure changes

As mentioned in Chapter 7, “Basic socket API extensions for IPv6,” on page 75, the socket address structure (sockaddr) is larger for IPv6 and has a slightly different format. This structure is passed as input or output on several socket API calls. The type of structure passed must match the address family of the socket being used on the socket API call. As a result, application changes are necessary. Table 17 describes the necessary changes:

Table 17. sockaddr structure changes

Affected Socket API calls	Changes required
Bind(), connect(), sendmsg(), sendto()	The length and type of sockaddr structure passed must match the address family of the socket being used (structure sockaddr_in or sockaddr_in6).
accept(), recvmsg(), recvfrom(), getpeername(), getsockname()	The sockaddr structure passed needs to be sufficiently large for the address family of the socket being used on these APIs. Note that the larger sockaddr_in6 structure can be passed even for AF_INET sockets. However, the application needs to be aware that the format of the sockaddr structure returned depends on the address family of the input socket.
UNIX System Services BPX1SRX (Send/Recv CSM buffers using sockets)	The length and type of sockaddr structure passed must match the address family of the socket being used (structure sockaddr_in or sockaddr_in6).

Address conversion functions

Because IPv6 and IPv4 addresses have a different format and size, changes are required when formatting these addresses for presentation purposes. Two utility functions have been introduced for a selected set of socket APIs to help applications perform this processing. A formatted IPv6 address uses significantly more space than a formatted IPv4 address (46 bytes versus 16 bytes) and this might affect the layout of any messages and displays that include an IP address.

Table 18. Address conversion function changes

Affected API call	Changes required
Translating an IP address from numeric form to presentation form using inet_ntoa()	Convert to use inet_ntop() function. This function can be used for both IPv4 and IPv6 addresses.
Translating a presentation form IP address to numeric form using inet_addr()	Convert to use inet_ntop() function. This function can be used for both IPv4 and IPv6 addresses.

Resolver API processing

TCP/IP applications typically need to resolve a host name to an IP address and sometimes need to resolve an IP address to a host name. Applications perform this processing by invoking resolver APIs, such as gethostbyname() and gethostbyaddr(). A new set of resolver APIs was introduced to support IPv6. Applications that currently use resolver APIs need to be modified to use the new

APIs in order to be enabled for IPv6. The older resolver APIs continue to be supported for IPv4 communications. For more information about resolver APIs, refer to “Name and address resolution functions” on page 77.

Table 19. Resolver API changes

Affected API call	Changes required
gethostbyname()	Use new getaddrinfo() API. These APIs can be used even if the system is not IPv6 enabled. Note that the freeaddrinfo() API needs to be issued to free up storage areas returned by the getaddrinfo() API.
gethostbyaddr()	Use the new getnameinfo() API. This API can also be used on a system that is not IPv6 enabled.

Special IPv6 addresses

IPv4 provides two IP addresses that have the following special meaning in the context of socket programs:

- The Loopback Address, typically 127.0.0.1, allows applications to connect() to or send datagrams to other applications on the same host.
- The INADDR_ANY address (0.0.0.0) allows TCP/IP server applications that specify it on a bind() call to accept incoming connections or datagrams across any network interface configured on the local host.

The concept of these special IPv4 addresses is also available in IPv6. The changes are described in Table 20.

Table 20. Special IPv6 address changes

Socket API calls	Changes required
Binding a socket to the IPv4 wildcard address (INADDR_ANY - 0.0.0.0)	Specify IPv6 IN6ADDR_ANY (::) in the sockaddr_in6 structure.
Using LOOPBACK (127.0.0.1) on bind(), connect(), sendto(), sendmsg()	Specify IPv6 Loopback address (::1) in the sockaddr_in6 structure.

See Chapter 7, “Basic socket API extensions for IPv6,” on page 75 for details about any constant definitions available for these special IPv6 addresses and the socket API that you are using.

Passing ownership of sockets across applications using givesocket and takesocket APIs

If your application is using the givesocket() and takesocket() APIs to pass ownership of a socket from one program to another, some changes are necessary for IPv6 enablement. The givesocket() and takesocket() APIs now support an address family of AF_INET6 for the socket being given or taken. The address family specified by the program performing the takesocket() must match the address family specified by the program that performed the givesocket(). As a result, care should be taken in coordinating the updates for IPv6 support across the partner applications performing givesocket and takesocket processing.

Table 21. givesocket() and takesocket() changes

Affected API call	Changes required
givesocket()	Specify AF_INET6 (Decimal 19) as the domain when giving an AF_INET6 socket.

Table 21. *givesocket() and takesocket() changes (continued)*

Affected API call	Changes required
getclientid()	Specify AF_INET6 as the domain when dealing with an AF_INET6 socket.
takesocket()	Specify AF_INET6 as the domain when taking an AF_INET6 socket.

Using multicast and IPv6

IPv6 provides enhanced support for multicast applications, including more granularity in the scope of multicast addressing and new socket options to allow an application to exploit this support. Table 22 lists IPv4 multicast setsockopt() and getsockopt() options and the equivalent IPv6 multicast options.

Table 22. *Multicast options*

Multicast function	IPv4	IPv6
Level specified on setsockopt()/getsockopt()	IP_PROTO	IPPROTO_IPV6
Joining a multicast group	IP_ADD_MEMBERSHIP	IPV6_JOIN_GROUP
Leaving a multicast group	IP_DROP_MEMBERSHIP	IPV6_LEAVE_GROUP
Select outbound interface for sending multicast datagrams	IP_MULTICAST_IF	IPV6_MULTICAST_IF
Set maximum hop count	IP_MULTICAST_TTL	IPV6_MULTICAST_HOPS
Enabling multicast loopback	IP_MULTICAST_LOOP	IPV6_MULTICAST_LOOP

In addition to the changes in the setsockopt() and getsockopt() options, the input and output parameters specified for these options are also changed when compared to IPv4. For example, selecting an outgoing interface for sending multicast IPv6 datagram involves passing an interface index that identifies the interface versus passing the IP address of the interface. For a detailed description of the IPv6 multicast options see “Options to control sending of multicast packets” on page 85.

An important consideration in updating your multicast application for IPv6 is how these changes are provided to the other partner applications participating in these multicast operations. For example, if a partner application in the network that is receiving these multicast packets is not updated, then the application sending the multicast datagrams might need to send them twice, once to an IPv4 multicast address and once to an IPv6 multicast address. Also, in order to perform this type of processing the application needs to create two separate sockets, an AF_INET socket and a AF_INET6 socket. There is no support equivalent to IPv4-mapped IPv6 addresses that would allow an AF_INET6 socket to be used in sending IPv4 multicast packets. As an alternative solution, first enable all the receiver applications for IPv6 and then enable the sender applications.

IP addresses might not be permanent

Long-term use of an address is discouraged as IPv6 allows for IP addresses to be dynamically renumbered. Applications should rely on DNS resolvers to cache the appropriate IP addresses and should avoid having IP addresses in configuration files.

Including IP addresses in the data stream

Applications that include IP addresses in the data they transmit over TCP/IP require changes when enabling for IPv6, as the IPv6 addresses have a different format from IPv4 addresses. The following options can be considered in dealing with these changes:

- Determine whether IP addresses are really needed in the data exchanged by the applications.
- Change the partner applications processing to always send IP addresses encoded using IPv6 format. In the case where IPv4 addresses are being used, they can be represented as IPv4-mapped IPv6 addresses.
- Include a version identifier that describes the format of the IP address being sent (IPv4 or IPv6).
- Modify applications to use host names instead of IP addresses in the data stream. This approach requires that the partner receiving the host name is able to resolve it to an IP address. Also note that a single IP host can have multiple IP addresses.
- In many cases, you might not be able to change all partner applications in your network at the same time. As a result, determining the type of IP address to send is a key consideration. Consider the following options when making this decision:
 - Determine the level of support when the connection is established by exchanging version or supported functions.
 - Encode the IPv6 addresses using new options. If the option is rejected by the peer, then it does not support IPv6.
 - Base the decision on the partner application's IP address. If the partner's source IP address is an IPv4 address then only use IPv4 addresses; otherwise, use an IPv6 address. This option can cause an IPv6-enabled partner application to be treated as an IPv4 partner if that application uses an IPv4-mapped IPv6 address to connect.

Example of an IPv4 TCP server program

The following example shows a simple IPv4 TCP server program written in C. The program opens a TCP socket, binds it to port 5000, and then performs a `listen()` followed by an `accept()` call. When a connection is accepted the server sends a `Hello` text string back to the client and closes the socket. This sample program is later shown with the changes required to make it IPv6 enabled.

```

/* simpleserver.c
   A very simple TCP socket server
 */
#include <sys/socket.h>
#include <netinet/in.h>
#include <arpa/inet.h>
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
int main(int argc,const char **argv)
{
    int serverPort = 5000;
    int rc;
    struct sockaddr_in serverSa;
    struct sockaddr_in clientSa;
    int clientSaSize;
    int on = 1;
    int c;
    int s = socket(PF_INET,SOCK_STREAM,0);
    rc = setsockopt(s,SOL_SOCKET,SO_REUSEADDR,&on,sizeof on);
    /* initialize the server's sockaddr */
    memset(&serverSa,0,sizeof(serverSa));
    serverSa.sin_family = AF_INET;
    serverSa.sin_addr.s_addr = htonl(INADDR_ANY);
    serverSa.sin_port = htons(serverPort);
    rc = bind(s,(struct sockaddr *)&serverSa,sizeof(serverSa));
    if (rc < 0)
    {
        perror("bind failed");
        exit(1);
    }
    rc = listen(s,10);
    if (rc < 0)
    {
        perror("listen failed");
        exit(1);
    }
    rc = accept(s,(struct sockaddr *)&clientSa,&clientSaSize);
    if (rc < 0)
    {
        perror("accept failed");
        exit(1);
    }
    printf("Client address is: %s\n",inet_ntoa(clientSa.sin_addr));
    c = rc;
    rc = write(c,"hello\n",6);
    close (s);
    close (c);
    return 0;
}

```

Figure 17. IPv4 TCP server program

Example of the simple TCP server program enabled for IPv6

The simple TCP server program is now shown with the changes (in bold) that are required to allow it to accept connections from IPv6 clients.

```

/*
 * A very simple TCP socket server for v4 or v6
 */
#include <sys/socket.h>
#include <netinet/in.h>
#include <arpa/inet.h>
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <string.h>
int main(int argc,const char **argv)
{
    int serverPort = 5000;
    int rc;
    union {
        struct sockaddr_in sin;
        struct sockaddr_in6 sin6;
    } serverSa;
    union {
        struct sockaddr_in sin;
        struct sockaddr_in6 sin6;
    } clientSa;

    int clientSaSize = sizeof(clientSa);
    int on = 1;
    int family;
    socklen_t serverSaSize;
    int c;
    char buf[INET6_ADDRSTRLEN];

    int s = socket(PF_INET6,SOCK_STREAM,0);
    if (s < 0)
    {
        fprintf(stderr, "IPv6 not active, falling back to IPv4...\n");
        s = socket(PF_INET,SOCK_STREAM,0);
        if (s < 0)
        {
            perror("socket failed");
            exit (1);
        }
        family = AF_INET;
        serverSaSize = sizeof(struct sockaddr_in);
    }
    else /* got a v6 socket */
    {
        family = AF_INET6;
        serverSaSize = sizeof(struct sockaddr_in6);
    }
    printf("socket descriptor is %d, family is %d\n",s,family);

```

Figure 18. Simple TCP server program enabled for IPv6 (Part 1 of 2)

```

rc = setsockopt(s,SOL_SOCKET,SO_REUSEADDR,&on,sizeof on);

/* initialize the server's sockaddr */
memset(&serverSa,0,sizeof(serverSa));
switch(family)
{
    case AF_INET:
        serverSa.sin.sin_family = AF_INET;
        serverSa.sin.sin_addr.s_addr = htonl(INADDR_ANY);
        serverSa.sin.sin_port = htons(serverPort);
        break;
    case AF_INET6:
        serverSa.sin6.sin6_family = AF_INET6;
        serverSa.sin6.sin6_addr = in6addr_any;
        serverSa.sin6.sin6_port = htons(serverPort);
}
rc = bind(s,(struct sockaddr *)&serverSa,serverSaSize);
if (rc < 0)
{
    perror("bind failed");
    exit(1);
}
rc = listen(s,10);
if (rc < 0)
{
    perror("listen failed");
    exit(1);
}
rc = accept(s,(struct sockaddr *)&clientSa,&clientSaSize);
if (rc < 0)
{
    perror("accept failed");
    exit(1);
}
c = rc;
printf("Client address is: %s\n",
       inet_ntop(clientSa.sin.sin_family,
                 clientSa.sin.sin_family == AF_INET
                 ? &clientSa.sin.sin_addr
                 : &clientSa.sin6.sin6_addr,
                 buf, sizeof(buf)));

if(clientSa.sin.sin_family == AF_INET6
    && ! IN6_IS_ADDR_V4MAPPED(&clientSa.sin6.sin6_addr))
    printf("Client is v6\n");
else
    printf("Client is v4\n");

rc = write(c,"hello\n",6);
close (s);
close (c);
return 0;
}

```

Figure 18. Simple TCP server program enabled for IPv6 (Part 2 of 2)

Chapter 9. Advanced socket APIs

This chapter describes the advanced socket APIs and includes the following sections:

- “Controlling the content of the IPv6 packet header”
- “Using ancillary data on sendmsg() and recvmsg()” on page 113
- “Interactions between socket options and ancillary data” on page 114
- “Why use RAW sockets?” on page 116

Before using advanced socket APIs in a multilevel security environment, refer to *Preparing for TCP/IP networking in a multilevel secure environment in z/OS Communications Server: IP Configuration Guide*. The advanced socket API for IPv6 support includes the following:

- IPv6 RAW socket support
- New socket options
- New ancillary data objects on sendmsg/recvmsg
- The ability to receive inbound packet information, including the following:
 - Arriving interface index
 - Destination IP address
 - Hop limit
 - Routing headers
 - Hop-by-hop option
 - Destination options
 - Traffic class by way of ancillary data
- The ability to set outgoing packet information, including the following:
 - Interface to use
 - Source IP address
 - Hop limit
 - Next hop address
 - Routing headers
 - Hop-by-hop options
 - Destination option
 - Traffic class (This can be set by socket options or ancillary data with some restrictions.)

z/OS UNIX C/C++ and z/OS UNIX Assembler Callable APIs support the advanced socket API for IPv6. The advanced socket API for IPv6 is not implemented in native TCP/IP socket APIs.

Controlling the content of the IPv6 packet header

This section contains information about socket options and how to control the content of the IPv6 packet header.

Socket options and ancillary data to support IPv6 (IPPROTO_IPV6 level)

An application can use socket options to enable or disable a function for a socket. An application can also provide a value to be used for a function with a socket option. After an option is enabled, it remains in effect for the socket until it is disabled.

An application can also use ancillary data on the sendmsg() API to enable a function or provide a value for the packet being sent by way of sendmsg(). The value of the ancillary data is in effect for that packet only. Note that the value of the ancillary data can override a socket option value. For a detailed explanation of ancillary data, see "Using ancillary data on sendmsg() and recvmsg()" on page 113.

An application can also receive ancillary data on the recvmsg() API. The returned Ancillary data is enabled for any socket options that return data on recvmsg.

A group of advanced socket options and ancillary data is defined to support IPv6. They are defined with a level of IPPROTO_IPV6 or IPPROTO_ICMPV6. The individual options begin with IPV6_ and ICMP6_ respectively. These options are only allowed on AF_INET6 sockets. In most cases, these options can be set on an AF_INET6 socket that is using IPv4-mapped IPv6 addresses, but have no effect. For example, the IPV6_HOPLIMIT ancillary data option is used to set a hop limit value in the IPv6 header. Because IPv4 packets are used with IPv4-mapped IPv6 addresses, the hop limit value is not used. The following are the only advanced socket options that have an effect on an AF_INET6 socket that is using IPv4-mapped IPv6 addresses:

- IPV6_PKTINFO
- IPV6_RECVPKTINFO
- IPV6_TCLASS
- IPV6_RECVTCLASS

Table 23. Sockets options at the IPPROTO_IPV6 level

Socket options getsockopt() setsockopt()	Assembler Callable Services	C/C++ using Language Environment	REXX	Sockets Extended macro/call
IPV6_CHECKSUM	Y	Y	N	N
IPV6_DONTFRAG	Y	Y	N	N
IPV6_DSTOPTS	Y	Y	N	N
IPV6_HOPOPTS	Y	Y	N	N
IPV6_NEXTHOP	Y	Y	N	N
IPV6_PATHMTU[valid only on getsockopt()]	Y	Y	N	N
IPV6_PKTINFO	Y	Y	N	N
IPV6_RECVDSTOPTS	Y	Y	N	N
IPV6_RECVHOPLIMIT	Y	Y	N	N
IPV6_RECVHOPOPTS	Y	Y	N	N
IPV6_RECVPATHMTU	Y	Y	N	N
IPV6_RECVPKTINFO	Y	Y	N	N
IPV6_RECVRTHDR	Y	Y	N	N

Table 23. Sockets options at the IPPROTO_IPV6 level (continued)

Socket options getsockopt() setsockopt()	Assembler Callable Services	C/C++ using Language Environment	REXX	Sockets Extended macro/call
IPV6_RECVTCLASS	Y	Y	N	N
IPV6_RTHDR	Y	Y	N	N
IPV6_RTHDRDSTOPTS	Y	Y	N	N
IPV6_TCLASS	Y	Y	N	N
IPV6_USE_MIN_MTU	Y using BPX1	Y	N	N

Table 24. Ancillary data on sendmsg() (Level = IPPROTO_IPV6)

Ancillary data on sendmsg()	Assembler Callable Services	C/C++ using Language Environment	REXX	Sockets Extended macro/call
IP_QOS_CLASSIFICATION ²	Y	Y	N	N
IPV6_DONTFRAG	Y	Y	N	N
IPV6_DSTOPTS	Y	Y	N	N
IPV6_HOPLIMIT ²	Y	Y	N	N
IPV6_HOPOPTS	Y	Y	N	N
IPV6_NEXTHOP	Y	Y	N	N
IPV6_PKTINFO ²	Y	Y	N	N
IPV6_RTHDR	Y	Y	N	N
IPV6_RTHDRDSTOPTS	Y	Y	N	N
IPV6_TCLASS	Y	Y	N	N
IPV6_USE_MIN_MTU	Y	Y	N	N

Table 25. Ancillary data on recvmsg() (Level = IPPROTO_IPV6)

Ancillary data on recvmsg()	Assembler Callable Services	C/C++ using Language Environment	REXX	Sockets Extended macro/call
IPV6_DSTOPTS	Y	Y	N	N
IPV6_HOPLIMIT	Y	Y	N	N
IPV6_HOPOPTS	Y	Y	N	N
IPV6_PATHMTU	Y	Y	N	N
IPV6_PKTINFO	Y	Y	N	N
IPV6_RTHDR	Y	Y	N	N
IPV6_TCLASS	Y	Y	N	N

Options for path MTU discovery

Use the following options for MTU discovery:

2. This option is supported as ancillary data for UDP and RAW protocols. It is not possible to use ancillary data to transmit options for TCP because there is not a one-to-one mapping between send operations and the TCP segments being transmitted.

IPV6_USE_MIN_MTU (used with TCP, UDP and RAW applications)

For IPv6, only the endpoint nodes can fragment a packet. Path MTU discovery determines the largest packet that can be sent to a destination without requiring fragmentation by an intermediate node (because that is not supported). In some cases, an application might not want to have the overhead of path MTU discovery. All nodes in an IPv6 network are required to support a minimum MTU of 1280 bytes. When an application enables this option, path MTU discovery is bypassed. If a direct route to the destination is not available, the minimum MTU size (1280 bytes) is used to send packets that otherwise might require fragmentation. If a direct route is available, the link's MTU size is used, because path MTU discovery is not needed when there are no intermediate nodes in the path.

For unicast destinations, this option disabled (this is the default). This avoids sending packets with the minimum MTU size. Instead, path MTU discovery processing information is used.

For multicast destinations, this option enabled (this is the default). This prevents path MTU discovery information from being used. If a direct route is not available, packets are sent with the minimum MTU size. If a direct route is available, packets are sent using the link's MTU, because no intermediate nodes are in the path.

This option can be enabled or disabled for the following:

- A socket with a setsockopt()
- A single send operation with ancillary data on the sendmsg()

A value of -1 passed on the set socket option causes the default values for unicast and multicast destinations to be used.

A value of 0 disables this option for both unicast and multicast destinations. Path MTU discovery information is used to send packets greater than the minimum MTU size.

A value of 1 enables this option for unicast and multicast destinations. All packets are sent without using path MTU discovery information, using the minimum MTU size, unless a direct route is available to the destination.

A getsockopt() with this option returns the value set by a setsockopt(). If a setsockopt() has not been done, the default value of -1 (disabled for unicast, enabled for multicast) is returned.

IPV6_DONTFRAG (used with UDP and RAW applications)

The IPV6_DONTFRAG option enables the application to indicate that the packet should not be fragmented by the local z/OS host.

This option is useful for applications that want to discover the actual path MTU.

Guideline: When using the IPV6_DONTFRAG socket option, use the IPV6_RECVPATHMTU socket option also. Otherwise, packets are silently discarded without any notification to the application.

This option can be enabled or disabled for the following:

- A socket with a setsockopt()
- A single send operation with ancillary data on the sendmsg()

A value of 1 enables this option for unicast or multicast destinations.

A `getsockopt()` with this option returns the value set by a `setsockopt()`. If a `setsockopt()` has not been performed, then `getsockopt()` returns a value of 0.

If `IPV6_DONTFRAG` is specified along with `IPV6_USE_MIN_MTU`, the `IPV6_DONTFRAG` setting is ignored, resulting in selection of the minimum architected IPv6 MTU size (1280 bytes).

IPV6_RECVPATHMTU (used with UDP and RAW applications)

The `IPV6_RECVPATHMTU` option enables the application to receive notifications about changes to the path MTU. This option notifies the application about all path MTU changes for all destinations, not only the ones initiated by this socket.

When the `IPV6_RECVPATHMTU` socket option is enabled, the path MTU is returned as ancillary data on the `recvmsg()` API (for an empty message) whenever the path MTU changes. The path MTU can change if the application sends a packet with the `IPV6_DONTFRAG` option and the packet is larger than the current path MTU. The path MTU can also change if the stack receives a corresponding ICMPv6 packet too big error. The ancillary data level is `IPPROTO_IPV6`. The option name is `IPV6_PATHMTU`. For a detailed explanation of ancillary data, see “Using ancillary data on `sendmsg()` and `recvmsg()`” on page 113.

This option can be enabled or disabled for a socket with a `setsockopt()`.

A value of 1 enables this option.

A `getsockopt()` with this option returns the value set by a `setsockopt()`. If a `setsockopt()` has not been performed, then `getsockopt()` returns a value of 0.

IPV6_PATHMTU (used with UDP and RAW applications)

The `IPV6_PATHMTU` option enables the application to retrieve the current path MTU to a given destination for which it has done a `connect()`.

This option is useful for applications also using `IPV6_RECVPATHMTU` that want to pick a good starting value.

This option is valid only on a `getsockopt()`. It returns the MTU that the stack uses on this connected socket.

Options to control the sending of packets

Some of these options add extension headers to outbound packets. z/OS TCP/IP allows the application to specify a maximum of 512 bytes of extension headers for an outbound packet. Additionally, for `IPV6_RTHDR`, z/OS TCP/IP allows the application to specify a maximum of 8 intermediate addresses in the routing header.

Use the following options to control the sending of packets:

IPV6_PKTINFO (used with UDP and RAW applications)

The `IPV6_PKTINFO` option enables the application to provide the following pieces of information:

- The source IP address for an outgoing packet
- The outgoing interface for a packet

The option value contains a 16-byte IPv6 address and a 4-byte interface index. An application can provide a nonzero value for one or both pieces of information.

To perform this operation, an application must meet one of the following criteria:

- Be APF authorized
- Have superuser authority
- The SERVAUTH resource
EZB.SOCKOPT.*sysname.tcpname.IPV6_PKTINFO* must be defined and the application must at least have READ access to it.

This option can be enabled or disabled for the following:

- A socket with a setsockopt()
- A single send operation with ancillary data on the sendmsg()

To disable the option, specify both the IPv6 address and the interface index as 0 in the option value.

A getsockopt() with this option returns the value set by setsockopt(). If a setsockopt() has not been done, a value of 0 is returned.

See “Understanding options for setting the source address” on page 115 for a discussion of the interaction of socket options and ancillary data for the setting of the source address. See “Understanding options for specifying the outgoing interface” on page 115 for a discussion of the interaction of socket options and ancillary data for determining the outgoing interface.

IPV6_HOPLIMIT (used with UDP and RAW applications)

The IPv6 header contains a hop limit field that controls the number of hops over which a datagram can be sent before being discarded. This is similar to the TTL field in the IPv4 header. The IPV6_HOPLIMIT option can be used to set the hop limit value for an outgoing packet. The option value should be between 0 and 255 inclusive. A value of -1 causes the TCP/IP protocol stack default to be used.

To perform this operation, an application must meet one of the following criteria:

- Be APF authorized
- Have superuser authority
- The SERVAUTH resource
EZB.SOCKOPT.*sysname.tcpname.IPV6_HOPLIMIT* must be defined and the application must at least have READ access to it

Note that the IPV6_UNICAST_HOPS socket option and the IPV6_MULTICAST_HOPS socket option are available to set a hop limit value also. See “Understanding hop limit options” on page 114 for a discussion of the interaction of IPV6_UNICAST_HOPS, IPV6_MULTICAST_HOPS and IPV6_HOPLIMIT.

IPV6_NEXTHOP (used with UDP and RAW applications)

The IPV6_NEXTHOP enables the application to specify the next hop address for an outgoing packet. The option value contains a sockaddr_in6 socket address structure and must contain an IPv6 address.

Restriction: This option does not support IPv4 mapped addresses.

To perform this operation, an application must meet one of the following criteria:

- Be APF authorized

- Have superuser authority
- The SERVAUTH resource EZB.SOCKOPT.sysname.tcpname.IPV6_NEXTHOP must be defined and the application must at least have READ access to it

This option can be enabled or disabled for the following

- A socket with a setsockopt()
- A single send operation with ancillary data on the sendmsg()

Restriction: IPV6_NEXTHOP is valid only for unicast destinations.

An option value with an *optlen* value of 0 disables IPV6_NEXTHOP. This option does not have any meaning for multicast destinations and is ignored for multicast.

A getsockopt() with this option returns the value set by a setsockopt(). If a setsockopt() has not been performed, then getsockopt() returns a value of 0 in *optlen*.

See “Understanding options for specifying the outgoing interface” on page 115 for a discussion of the interaction of socket options and ancillary data for determining the outgoing interface.

Tips:

- If you use this socket option in a Common INET environment, establish affinity to the desired stack to ensure predictable results (as not all stacks might have a route to the specified next hop address).
- If you specify a link-local address as the next hop address, specify the outgoing interface either on IPV6_PKTINFO or by using the scope portion of the socket address structure.

Rule: The next hop address cannot be a multicast address and must be a neighbor (for example, the stack must have a direct route to the next hop address).

IPV6_RTHDR (used with UDP and RAW applications)

The IPV6_RTHDR option enables the application to specify an IPv6 routing header (as an extension header) for an outgoing packet. The option value contains a type 0 routing header. An application can specify at most one routing header. z/OS TCP/IP allows the application to specify a maximum of eight IPv6 addresses in the routing header.

To perform this operation, an application must meet one of the following criteria:

- Be APF authorized
- Have superuser authority
- The SERVAUTH resource EZB.SOCKOPT.sysname.tcpname.IPV6_RTHDR must be defined and the application must at least have READ access to it

This option can be enabled or disabled for the following:

- A socket with a setsockopt()
- A single send operation with ancillary data on the sendmsg()

A `getsockopt()` with this option returns the value set by a `setsockopt()`. If a `setsockopt()` has not been performed, then `getsockopt()` returns a value of 0 in `optlen`.

Tip: If you use this socket option in a Common INET environment, establish affinity to the desired stack to ensure predictable results (as not all stacks might have a path to the destination starting with the first entry in the specified routing header).

A z/OS UNIX C/C++ application can use the following utilities to build routing headers:

- `inet6_rth_space()` - return number of bytes required for routing header
- `inet6_rth_init()` - initialize buffer data for routing header
- `inet6_rth_add()` - add one IPv6 address to the routing header

Refer to *z/OS XL C/C++ Run-Time Library Reference* for a description of these utilities.

A z/OS UNIX Assembler Callable Services application needs to build the routing headers explicitly. Refer to *z/OS UNIX System Services Programming: Assembler Callable Services Reference* for information about z/OS UNIX Assembler Callable Services and the data structures defined in the `BPXYSOCK` macro.

IPV6_DSTOPTS (used with UDP and RAW applications)

The `IPV6_DSTOPTS` option enables the application to specify destination options that get examined by the host at the final destination.

The `IPV6_DSTOPTS` option can be used to set a destination options header (as an extension header) for an outgoing packet. The option value contains a destination options header.

To perform this operation, an application must meet one of the following criteria:

- Be APF authorized
- Have superuser authority
- The SERVAUTH resource

`EZB.SOCKOPT.sysname.tcpname.IPV6_DSTOPTS` must be defined and the application must at least have READ access to it

This option can be enabled or disabled for the following:

- A socket with a `setsockopt()`
- A single send operation with ancillary data on the `sendmsg()`

A `getsockopt()` with this option returns the value set by a `setsockopt()`. If a `setsockopt()` has not been performed, then `getsockopt()` returns a value of 0 in `optlen`.

A z/OS UNIX C/C++ application can use the following utilities to build the following destination options headers:

- `inet6_opt_init()` - initialize buffer data for options header
- `inet6_opt_append()` - add one TLV option to the options header
- `inet6_opt_finish()` - finish adding TLV options to the option header

- `inet6_opt_set_val()` - add one component of the option content to the option

Refer to *z/OS XL C/C++ Run-Time Library Reference* for a description of these utilities.

A z/OS UNIX Assembler Callable Services application needs to build the options headers explicitly. Refer to *z/OS UNIX System Services Programming: Assembler Callable Services Reference* for information about z/OS UNIX Assembler Callable Services and the data structures defined in the BPXYSOCK macro.

IPV6_RTHDRDSTOPTS (used with UDP and RAW applications)

The IPV6_RTHDRDSTOPTS option enables the application to specify destination options that get examined by every IP host that appears in the routing header.

The IPV6_RTHDRDSTOPTS option can be used to set a destination options header (as an extension header) for an outgoing packet. The option value contains a destination options header. This option is ignored if the application does not also use the IPV6_RTHDR option to specify a routing header.

To perform this operation, an application must meet one of the following criteria:

- Be APF authorized
- Have superuser authority
- The SERVAUTH resource

`EZB.SOCKOPT.sysname.tcpname.IPV6_RTHDRDSTOPTS` must be defined and the application must at least have READ access to it

This option can be enabled or disabled for the following:

- A socket with a `setsockopt()`
- A single send operation with ancillary data on the `sendmsg()`

A `getsockopt()` with this option returns the value set by a `setsockopt()`. If a `setsockopt()` has not been performed, then `getsockopt()` returns a value of 0 in optlen.

A z/OS UNIX C/C++ application can use the following utilities to build Destination options headers:

- `inet6_opt_init()` - initialize buffer data for options header
- `inet6_opt_append()` - add one TLV option to the options header
- `inet6_opt_finish()` - finish adding TLV options to the option header
- `inet6_opt_set_val()` - add one component of the option content to the option

Refer to *z/OS XL C/C++ Run-Time Library Reference* for a description of these utilities.

A z/OS UNIX Assembler Callable Services application needs to build the options headers explicitly. Refer to *z/OS UNIX System Services Programming: Assembler Callable Services Reference* for information about z/OS UNIX Assembler Callable Services and the data structures defined in the BPXYSOCK macro.

IPV6_TCLASS (used with TCP, UDP and RAW applications)

The IPv6 header contains a traffic class field that can be used to identify and distinguish between different classes or priorities of IPv6 packets. This is similar to the type of service (ToS) field in the IPv4 header. The IPV6_TCLASS option can be used to set the traffic class value for an outgoing packet. However, if a QoS policy that specifies a traffic class for the packet is also in effect, then the stack ignores the value specified with the IPV6_TCLASS option and uses the value specified by the QoS policy.

To perform this operation, an application must meet one of the following criteria:

- Be APF authorized
- Have superuser authority
- The SERVAUTH resource `EZB.SOCKOPT.sysname.tcpname.IPV6_TCLASS` must be defined and the application must at least have READ access to it

This socket option is also valid for an AF_INET6 socket that is using IPv4-mapped IPv6 addresses.

This option can be enabled or disabled for a socket with a `setsockopt()`. For UDP and RAW, this option can be enabled or disabled for a single send operation with ancillary data on the `sendmsg()`.

The option value should be in the range 0 - 255. A value of -1 causes the TCP/IP to use the traffic class value specified by policy (if any) or the default of 0.

A `getsockopt()` with this option returns the value set by a `setsockopt()`. If a `setsockopt()` has not been performed, then the stack returns the traffic class value specified by policy (if any) or the default of 0.

Options to provide information about received packets

Use the following options to provide information about received packets:

IPV6_RECVPKTINFO (used with UDP and RAW applications)

The IPV6_RECVPKTINFO socket option allows an application to receive the following pieces of information:

- The destination IP address from the IPv6 header
- The interface index for the interface over which the packet was received

When the IPV6_RECVPKTINFO socket option is enabled, the IP address and interface index are returned as ancillary data on the `recvmsg()` API. The ancillary data level is `IPPROTO_IPV6`. The option name is `IPV6_PKTINFO`. For a detailed explanation of ancillary data, see “Using ancillary data on `sendmsg()` and `recvmsg()`” on page 113.

Restriction: This option can only be enabled or disabled with a `setsockopt()`. `IPV6_RECVPKTINFO` is not valid as ancillary data on `sendmsg()`. A nonzero option value enables the option; a value of 0 disables the option.

A `getsockopt()` with this option returns the value set by a `setsockopt()`. If a `setsockopt()` has not been done, the default value of 0 (disabled) is returned.

IPV6_RECVHOPLIMIT (used with TCP, UDP and RAW applications)

The IPV6_RECVHOPLIMIT socket option allows an application to receive the value of the hop limit field from the IPv6 header. When the IPV6_RECVHOPLIMIT socket option is enabled, the hop limit is returned as ancillary data on the recvmsg() API. The ancillary data level is IPPROTO_IPV6. The option name is IPV6_HOPLIMIT. For a UDP or RAW application, if this option is enabled, the IPV6_HOPLIMIT ancillary data is returned with each recvmsg(). For a TCP application, if this option is enabled, IPV6_HOPLIMIT ancillary data is only returned on recvmsg() when the hop limit value being used has changed. For a detailed explanation of ancillary data, see “Using ancillary data on sendmsg() and recvmsg()” on page 113.

This option can only be enabled or disabled with a setsockopt(). IPV6_RECVHOPLIMIT is not valid as ancillary data on sendmsg(). A nonzero option value enables the option; a value of 0 disables the option.

A getsockopt() with this option returns the value set by a setsockopt(). If a setsockopt() has not been done, the default value of 0 (disabled) is returned.

IPV6_RECVRTHDR (used with UDP and RAW applications)

The IPV6_RECVRTHDR socket option enables the application to receive a routing header.

When the IPV6_RECVRTHDR socket option is enabled, the routing header is returned as ancillary data on the recvmsg() API. Each routing header is returned as one ancillary data object. The ancillary data level is IPPROTO_IPV6. The option name is IPV6_RTHDR. For a detailed explanation of ancillary data, see “Using ancillary data on sendmsg() and recvmsg()” on page 113.

This option can be enabled or disabled only with a setsockopt(). IPV6_RECVRTHDR is not valid as ancillary data on sendmsg(). A nonzero value enables the option; a value of 0 disables the option.

A getsockopt() with this option returns the value set by a setsockopt(). If a setsockopt() has not been performed, then getsockopt() returns a value of 0.

A z/OS UNIX C/C++ application can use the following utilities to process routing headers:

- `inet6_rth_reverse()` - reverse a routing header
- `inet6_rth_segments()` - return number of segments in a routing header
- `inet6_rth_getaddr()` - fetch one address from a routing header

Refer to *z/OS XL C/C++ Run-Time Library Reference* for a description of the above utilities.

A z/OS UNIX Assembler Callable Services application needs to build the options headers explicitly. Refer to *z/OS UNIX System Services Programming: Assembler Callable Services Reference* for information about z/OS UNIX Assembler Callable Services and the data structures defined in the BPXYSOCK macro.

IPV6_RECVHOPOPTS (used with UDP and RAW applications)

The IPV6_RECVHOPOPTS socket option enables the application to receive hop-by-hop options.

When the IPV6_RECVHOPOPTS socket option is enabled, the hop-by-hop options are returned as ancillary data on the recvmsg() API. The ancillary data level is IPPROTO_IPV6. The option name is IPV6_HOPOPTS. For a detailed explanation of ancillary data, see “Using ancillary data on sendmsg() and recvmsg()” on page 113.

This option can be enabled or disabled only with a setsockopt(). IPV6_RECVHOPOPTS is not valid as ancillary data on sendmsg(). A nonzero value enables the option; a value of 0 disables the option.

A getsockopt() with this option returns the value set by a setsockopt(). If a setsockopt() has not been performed, then getsockopt() returns a value of 0.

A z/OS UNIX C/C++ application can use the following utilities to process hop-by-hop options headers:

- inet6_opt_next() - extract the next option from the options header
- inet6_opt_find() - extract an option of a specified type from the header
- inet6_opt_get_val() - retrieve one component of the option content

Refer to *z/OS XL C/C++ Run-Time Library Reference* for a description of the above utilities.

A z/OS UNIX Assembler Callable Services application needs to build the options headers explicitly. Refer to *z/OS UNIX System Services Programming: Assembler Callable Services Reference* for information about z/OS UNIX Assembler Callable Services and the data structures defined in the BPXYSOCK macro.

IPV6_RECVDSTOPTS (used with UDP and RAW applications)

The IPV6_RECVDSTOPTS socket option enables the application to receive destination options.

When the IPV6_RECVDSTOPTS socket option is enabled, the destination options are returned as ancillary data on the recvmsg() API. The application can receive up to two destination options headers (one before a routing header and one after a routing header). Each destination options header is returned as one ancillary data object. The ancillary data level is IPPROTO_IPV6. The option name is IPV6_DSTOPTS. For a detailed explanation of ancillary data, see “Using ancillary data on sendmsg() and recvmsg()” on page 113.

This option can be enabled or disabled only with a setsockopt(). IPV6_RECVDSTOPTS is not valid as ancillary data on sendmsg(). A nonzero value enables the option; a value of 0 disables the option.

A getsockopt() with this option returns the value set by a setsockopt(). If a setsockopt() has not been performed, then getsockopt() returns a value of 0.

A z/OS UNIX C/C++ application can use the following utilities to process destination options headers:

- inet6_opt_next() - extract the next option from the options header
- inet6_opt_find() - extract an option of a specified type from the header
- inet6_opt_get_val() - retrieve one component of the option content

Refer to *z/OS XL C/C++ Run-Time Library Reference* for a description of these utilities.

A z/OS UNIX Assembler Callable Services application needs to build the options headers explicitly. Refer to *z/OS UNIX System Services Programming: Assembler Callable Services Reference* for information about z/OS UNIX Assembler Callable Services and the data structures defined in the BPXYSOCK macro.

IPV6_RECVTCLASS (used with TCP, UDP and RAW applications)

The IPV6_RECVTCLASS socket option enables the application to receive the value of the traffic class field from the IPv6 header.

When the IPV6_RECVTCLASS socket option is enabled, the traffic class is returned as ancillary data on the recvmsg() API. The ancillary data level is IPPROTO_IPV6. The option name is IPV6_TCLASS. For a UDP, or RAW application, if this option is enabled, the IPv6_TCLASS ancillary data is returned with each recvmsg(). For a TCP application, if this option is enabled, IPV6_TCLASS ancillary data is only returned on recvmsg() when the traffic class value being used has changed. For a detailed explanation of ancillary data, see "Using ancillary data on sendmsg() and recvmsg()" on page 113.

This socket option is also valid for an AF_INET6 socket that is using IPv4-mapped IPv6 addresses.

This option can be enabled or disabled only with a setsockopt(). IPV6_RECVTCLASS is not valid as ancillary data on sendmsg(). A nonzero value enables the option; a value of 0 disables the option.

A getsockopt() with this option returns the value set by a setsockopt(). If a setsockopt() has not been performed, then getsockopt() returns a value of 0.

Option to provide checksum processing for RAW applications

Use the following option to provide checksum processing for RAW applications:

IPV6_CHECKSUM (used with RAW applications)

The IPV6_CHECKSUM socket option can be used by a RAW application to enable checksum processing to be done by the TCP/IP protocol stack for packets on a socket. When enabled, the checksum is computed and stored for outbound packets; the checksum is verified for inbound packets. Note that this socket option is not applicable for ICMPv6 RAW sockets because the TCP/IP protocol stack always provides checksum processing for them.

This option can only be enabled or disabled with a setsockopt(). IPV6_CHECKSUM is not valid as ancillary data on sendmsg(). The option value provides the offset into the user data where the checksum field begins. The option value should be an even number between 0 - 65534 inclusive. A value of -1 causes the option to be disabled.

A getsockopt() with this option returns the value set by a setsockopt(). If a setsockopt() has not been done, the value of -1 (disabled) is returned.

Option to provide QoS classification data

Use the following option to provide QoS classification data:

IP_QOS_CLASSIFICATION (used with TCP applications)

This option enables the application to provide QoS classification data. It is a z/OS Communications Server-specific ancillary data type, and is not associated with the IPv6 Advanced Socket API. It can be specified as ancillary data on sendmsg() for AF_INET and AF_INET6 sockets. For AF_INET sockets the level specified should be IPPROTO_IP; for AF_INET6 sockets the level specified should be IPPROTO_IPV6. For a detailed

description of the function, refer to the programming interfaces in the *z/OS Communications Server: IP Programmer's Guide and Reference* for providing classification data to be used in differentiated services policies.

Socket option to support ICMPv6 (IPPROTO_ICMPV6 level)

Table 26. Sockets options at the IPPROTO_ICMPV6 level

Socket options getsockopt() setsockopt()	Assembler Callable Services	C/C++ using Language Environment	REXX	Sockets Extended macro/call
ICMP6_FILTER	N	Y	N	N

Use the following socket option to support ICMPv6 (IPPROTO_ICMPV6 level):

ICMP6_FILTER (used with RAW applications)

The ICMP6_FILTER socket option can be used by a RAW application to filter out ICMPv6 message types that it does not need to receive. There are many more ICMPv6 message types than ICMPv4 message types. ICMPv6 provides function comparable to ICMPv4 plus IGMPv4 and ARPv4 functionality. An application might only be interested in receiving a subset of the messages received for ICMPv6.

This option is enabled or disabled with a setsockopt(). The option value provides a 256-bit array of message types that should be filtered. To disable the option, the setsockopt() should be issued with an option length of 0. This causes the TCP/IP protocol stack's default filter to be in effect.

A getsockopt() with this option returns the value set by a setsockopt(). If a setsockopt() has not been done, the TCP/IP protocol stack's default filter is returned. For more information on default filtering, refer to "ICMP considerations" on page 117.

Table 27 lists the macros that are provided in the Language Environment C/C++ environment to manipulate the filter value.

Table 27. Macros used to manipulate filter value

Macro	Description
void ICMP6_FILTER_SETPASSALL(struct icmp6_filter *);	Specifies that all ICMPv6 messages are passed to the application.
void ICMP6_FILTER_SETBLOCKALL(struct icmp6_filter *);	Specifies that all ICMPv6 messages are blocked from being passed to the application.
void ICMP6_FILTER_SETPASS(int, struct icmp6_filter *);	ICMPv6 messages of type specified in int should be passed to the application.
void ICMP6_FILTER_SETBLOCK(int, struct icmp6_filter *);	ICMPv6 messages of type specified in int should not be passed to the application.
void ICMP6_FILTER_WILLPASS(int, const struct icmp6_filter *);	Returns true if the message type specified in int is passed to the application by the filter pointed to by the second argument.

Table 27. Macros used to manipulate filter value (continued)

Macro	Description
void ICMP6_FILTER_WILLBLOCK(int, const struct icmp6_filter *);	Returns true if the message type specified in int is not passed to the application by the filter pointed to by the second argument.

Using ancillary data on sendmsg() and recvmsg()

The sendmsg() API is similar to other socket APIs, such as send() and write() that allow an application to send data, but also provides the capability of specifying ancillary data. Ancillary data allows applications to pass additional option data to the TCP/IP protocol stack along with the normal data that is sent to the TCP/IP network.

The recvmsg() API is similar to other socket APIs, such as recv() and read(), that allow an application to receive data, but also provides the capability of receiving ancillary data. Ancillary data allows the TCP/IP protocol stack to return additional option data to the application along with the normal data from the TCP/IP network.

These sendmsg() and recvmsg() API extensions are only available to applications using the following socket API libraries:

- z/OS IBM C/C++ sockets with the z/OS Language Environment(R). For more information about these APIs, refer to the *z/OS XL C/C++ Run-Time Library Reference*.
- z/OS UNIX Assembler Callable services socket APIs. For more information about these APIs, refer to *z/OS UNIX System Services Programming: Assembler Callable Services Reference*.

For the sendmsg() and recvmsg() APIs most parameters are passed in a message header input parameter. The mapping for the message header is defined in socket.h for C/C++ and in the BPXYMSGH macro for users of the z/OS UNIX Assembler Callable services. For simplicity, only the C/C++ version of the data structures are shown in the following section:

```
struct msghdr {
    void             *msg_name;           /* optional address      */
    size_t           msg_namelen;        /* size of address       */
    struct           iovec *msg_iov;        /* scatter/gather array  */
    int              msg_iovlen;          /* # elements in msg_iov */
    void             *msg_control;        /* ancillary data        */
    size_t           msg_controllen;      /* ancillary data length */
    int              msg_flags;           /* flags on received msg */
};
```

Notes:

1. The msg_name and msg_namelen parameters are used to specify the destination sockaddr on a sendmsg(). On a recvmsg() the msg_name and msg_namelen parameters are used to return the remote sockaddr to the application.
2. Data to be sent using sendmsg() needs to be described in the msg_iov structure. On recvmsg() the received data is described in the msg_iov structure.
3. The address of the ancillary data is passed in the msg_control field.

4. The length of the ancillary data is passed in msg_controllen. Note that if multiple ancillary data sections are being passed, this length should reflect the total length of ancillary data sections.
5. msg_flags is not applicable for sendmsg().

The msg_control parameter points to the ancillary data. This msg_control pointer points to the following structure (C/C++ example shown below) that describes the ancillary data (also defined in socket.h and BPXYMSGH respectively):

```
struct cmsghdr {
    size_t cmsg_len;      /* data byte count includes hdr */
    int    cmsg_level;    /* originating protocol */
    int    cmsg_type;     /* protocol-specific type */
    /* followed by u_char   cmsg_data[]; */
};
```

Notes:

1. The cmsg_len should be set to the length of the cmsghdr plus the length of all ancillary data that follows immediately after the cmsghdr. This is represented by the commented out cmsg_data field.
2. The cmsg_level should be set to the option level (for example, IPPROTO_IPV6).
3. The cmsg_type should be set to the option name (for example, IPV6_USE_MIN_MTU).

Interactions between socket options and ancillary data

This section describes between socket options and ancillary data, included hop limits.

Understanding hop limit options

The IPv6 header contains a hop limit field that controls the number of hops over which a datagram can be sent before being discarded. This is similar to the TTL field in the IPv4 header. An application can influence the value of the hop limit field using the following options:

- IPV6_UNICAST_HOPS socket option (hop limit value to be used for unicast packets on a socket)
- IPV6_MULTICAST_HOPS socket option (hop limit value to be used for multicast packets on a socket)
- IPV6_HOPLIMIT ancillary data option on sendmsg() (hop limit value to be used for single packet)

The hop limit value can also be influenced by a router advertised hop limit, as well as the globally configured HOPLIMIT parameter value on the IP CONFIG6 statement.

For a unicast packet, the following precedence order is used to determine a packet's hop limit value:

1. If IPV6_HOPLIMIT ancillary data is specified on sendmsg(), use its value.
2. If the IPV6_UNICAST_HOPS socket option is set, use its value.
3. If a router advertised hop limit is known, use its value.
4. If there is a globally configured IPv6 hop limit, use its value.
5. Use the IPv6 default unicast hop limit, 255.

For a multicast packet, the following precedence order is used to determine the packet's hop limit value:

1. If IPV6_HOPLIMIT ancillary data is specified on sendmsg(), use its value.
2. If the IPV6_MULTICAST_HOPS socket option is set, use its value.
3. Use the IPv6 default multicast hop limit, 1.

Understanding options for setting the source address

A UDP or RAW application can influence the setting of the source address with the bind() IPv6 address or with the IPV6_PKTINFO option.

The following precedence order is used to determine the source IP address for a packet:

1. If IPV6_PKTINFO ancillary data is specified on sendmsg() with a nonzero source IP address, use its value. If the IPV6_PKTINFO ancillary data is specified with a length of 0 or with a zero source IP address, go step to 3.
2. If the IPV6_PKTINFO socket option is set and contains a nonzero source IP address, use its value.
3. If the application bound the socket to a specific address, use the Bind address.
4. The TCP/IP protocol stack selects a source address.

Understanding options for specifying the outgoing interface

A UDP or RAW application can influence the outgoing interface for a packet with the IPV6_PKTINFO option, the IPV6_NEXTHOP option, or the IPV6_MULTICAST_IF option. The scope ID field in the send operation's destination sockaddr can also affect the outgoing interface. The options field contains an interface index. The scope ID field contains a zone index.

When responding to a peer, UDP and RAW applications should use the sockaddr_in6 structure which they received, and should not zero out the scope ID field. When sending an unsolicited packet (for example, not responding to one that was received), the scope ID field should be zero, and UDP and RAW applications should use the IPV6_PKTINFO, IPV6_NEXTHOP, or IPV6_MULTICAST_IF options to select the outgoing interfaces.

The following precedence order is used to determine the outgoing interface for a packet:

1. If the send operation specifies a destination sockaddr structure with a scope ID, then the scope ID is used if valid (note that a scope ID should only be provided with a link-local address).
2. If IPV6_PKTINFO ancillary data is specified on sendmsg() with a nonzero interface index, use its value. If the IPV6_PKTINFO ancillary data is specified with a length of 0 or with an interface index of 0, then skip to rule 4.
3. If the IPV6_PKTINFO socket option is set and contains a nonzero interface index, use its value.
4. If this is a multicast packet and the IPV6_MULTICAST_IF socket option is set, use its value.
5. If IPV6_NEXTHOP ancillary data is specified on sendmsg() with a nonzero value, use the stack routing table to determine the interface to the next hop address. If the IPV6_NEXTHOP ancillary data is specified with a length of 0, go to step 7 on page 116.

6. If the IPV6_NEXTHOP socket option is set and contains a nonzero value, use the stack routing table to determine the interface to the next hop address.
7. The TCP/IP protocol stack uses the routing table to determine the interface to the destination IP address.

Why use RAW sockets?

Consider the following factors for RAW sockets use:

- An application (for example, PING) can send and receive ICMPv6 messages.
- An application can send and receive datagrams with an IP protocol that the TCP/IP stack does not support.

The external behavior of IPv6 RAW sockets differs significantly from that of IPv4 RAW sockets, specifically with regards to the following:

- RAW protocol values allowed
- Application visibility of IP headers
- ICMP considerations
- Checksumming data

RAW protocol values

Protocol values 0, 41, 43, 44, 50, 51, 59 and 60 are not allowed because they conflict with the following IPv6 extension header types:

- IPPROTO_HOPOPTS (0)
- IPPROTO_IPV6 (41)
- IPPROTO_ROUTING (43)
- IPPROTO_FRAGMENT (44)
- IPPROTO_ESP (50)
- IPPROTO_AH (51)
- IPPROTO_NONE (59)
- IPPROTO_DSTOPTS (60)

Of the RAW protocol values listed, only the following correspond to well-known IPv4 RAW protocols:

- IPPROTO_ESP (50)
- IPPROTO_AH (51)

Application visibility of IP headers

Applications do not see IP headers of incoming datagrams and cannot provide IP headers with outgoing datagrams.

IPv6 RAW applications can get or set selected IP header information for incoming and outgoing datagrams by way of socket options and ancillary data as follows:

- Applications can set the IPV6_RECVHOPLIMIT socket option in order to get the hop limit for incoming datagrams in ancillary data. By default, this socket option is set to off.
- Applications can set the IPV6_RECVPKTINFO socket option in order to get the destination IP address and interface identifier for incoming datagrams in ancillary data. By default, this socket option is set to off.

- Applications can set the IPV6_RECVRTHDR socket option in order to get the routing header for incoming datagrams in ancillary data. By default, this socket option is set to off.
- Applications can set the IPV6_RECVHOPOPTS socket option in order to get the hop-by-hop options for incoming datagrams in ancillary data. By default, this socket option is set to off.
- Applications can set the IPV6_RECVDSTOPTS socket option in order to get the destination options for incoming datagrams in ancillary data. By default, this socket option is set to off.
- Applications can set the IPV6_RECVTCLASS socket option in order to get the traffic class for incoming datagrams in ancillary data. By default, this socket option is set to off.
- Applications can set the IPV6_UNICAST_HOPS socket option in order to set the hop limit for outgoing unicast datagrams. By default, this socket option is set to off and the configured maximum hop limit or the default hop limit is used.
- Applications can set the IPV6_MULTICAST_HOPS socket option in order to set the hop limit for outgoing multicast datagrams. By default, this socket option is set to off and a hop limit of 1 is used.
- Applications can use the IPV6_HOPLIMIT ancillary data option to set the hop limit for an outgoing datagram.
- Applications can use the IPV6_PKTINFO socket option and ancillary data option to set the source address and interface identifier for outgoing datagrams. By default, the socket option is set to off.
- Applications can use the IPV6_NEXTHOP socket option and ancillary data option to set the next hop address for outgoing datagrams. By default, the socket option is set to off.
- Applications can use the IPV6_RTHDR socket option and ancillary data option to set the routing header for outgoing datagrams. By default, the socket option is set to off.
- Applications can use the IPV6_HOPOPTS socket option and ancillary data option to set the hop-by-hop options for outgoing datagrams. By default, the socket option is set to off.
- Applications can use the IPV6_DSTOPTS socket option and ancillary data option to set the destination options (that get examined by the host at the final destination) for outgoing datagrams. By default, the socket option is set to off.
- Applications can use the IPV6_RTHDRDSTOPTS socket option and ancillary data option to set the destination options (that get examined by every host that appears in the routing header) for outgoing datagrams. By default, the socket option is set to off.
- Applications can use the IPV6_TCLASS socket option and ancillary data option to set the traffic class for outgoing datagrams. By default, the socket option is set to off.

ICMP considerations

IPv6 RAW ICMPv6 applications can set the ICMP6_FILTER socket option to specify which ICMPv6 message types the socket receives. By default, the following message types are blocked (are not received):

- ICMP_ECHO
- ICMP_TSTAMP
- ICMP_IREQ
- ICMP_MASKREQ

- ICMP6_ECHO_REQUEST
- MLD_LISTENER_QUERY
- MLD_LISTENER_REPORT
- MLD_LISTENER_REDUCTION
- ND_ROUTER_SOLICIT
- ND_ROUTER_ADVERT
- ND_NEIGHBOR_SOLICIT
- ND_NEIGHBOR_ADVERT
- ND_REDIRECT

Checksumming data

IPv6 RAW applications can set the IPV6_CHECKSUM socket option in order to have TCP/IP calculate checksums for outgoing datagrams and verify checksums for incoming datagrams. By default, this socket option is set to off.

Part 4. Advanced topics

This section contains the following chapters:

Chapter 10, “Advanced concepts and topics,” on page 121 provides advanced IPv6 protocol information.

Chapter 11, “IPv6 support tables,” on page 131 contains tables with features and applications that support IPv6.

Chapter 10. Advanced concepts and topics

This chapter explains some of the advanced concepts and topics for IPv6 implementation and includes the following sections:

- “Tunneling”
- “Application migration and coexistence overview” on page 125
- “Application migration approaches” on page 127

Tunneling


When IPv6 or IPv6/IPv4 systems are separated from other similar systems that they wish to communicate with by IPv4 networks, then IPv6 packets must be tunneled through the IPv4 network. IPv6 packets are tunneled over IPv4 very simply: the IPv6 packet is encapsulated in an IPv4 datagram, or in other words, a complete IPv4 header is added to the IPv6 packet. The presence of the IPv6 packet within the IPv4 datagram is indicated by a protocol value of 41 in the IPv4 header.

Restriction: z/OS Communications Server cannot function as an endpoint for this type of tunnel.

While there are many tunneling protocols that can be used, all share the following common features and processing characteristics:

- The source tunnel endpoint determines that an IPv6 packet needs to be tunneled over an IPv4 network. This depends on the tunneling protocol that is used. After this decision is made, the source tunnel endpoint adds an IPv4 header to the IPv6 packet. The protocol value in the IPv4 header is set to 41. This indicates that this is an IPv6 over IPv4 tunnel packet. The source and destination addresses in the IPv4 header are set based on the tunneling protocol that is used.
- At the destination tunnel endpoint, the IPv4 layer receives the IPv4 packet (or packets, if the IPv4 datagram was fragmented). The IPv4 layer processes the datagram in the normal way, reassembling fragments if necessary, and records the protocol value of 41 in the IPv4 header. IPv4 security checks are made, and the IPv4 header is removed, leaving the original IPv6 packet. The IPv6 packet is processed as normal.

Figure 19 on page 122 shows a subset of the available tunneling protocols, with descriptions of the more prevalent protocols. Others exist or are in the process of being defined. Select one that is appropriate for your environment.



Tunneling: encapsulate an IPv6 packet in an IPv4 packet and send the IPv4 packet to the other tunnel end-point IPv4 address

Figure 19. Tunneling

Configured tunnels

Configured tunneling refers to IPv6 over IPv4 tunneling, where the IPv4 tunnel endpoint address is determined by configuration information on the encapsulating node. The tunnels can be unidirectional or bidirectional. Bidirectional configured tunnels act similarly as virtual point-to-point links. For each tunnel, the encapsulating node must store the tunnel endpoint address. When an IPv6 packet is transmitted over a tunnel, the tunnel endpoint address configured for that tunnel is used as the destination address for the encapsulating IPv4 header.

Routing information on the encapsulating node usually determines which packets to tunnel. This is typically done by way of a routing table, which directs packets based on their destination address using the prefix mask and match technique.

Configured tunnels can be host-host, host-router, or router-router. Host-host tunnels allow two IPv6/IPv4 nodes to send IPv6 packets directly to one another without going through an intermediate IPv6 router. This can be useful if the applications need to take advantage of IPv6 features that are not available in IPv4.

An IPv6/IPv4 host that is connected to datalinks with no IPv6 routers can use a configured tunnel to reach an IPv6 router. This tunnel allows the host to communicate with the rest of the IPv6 Internet. If the IPv4 address of an IPv6/IPv4 router bordering the IPv6 backbone is known, this can be used as the tunnel endpoint address, and can be used as an IPv6 default route. This default route is used only if a more specific route is not known.

Configured tunnels can also be used between routers, allowing isolated IPv6 networks to be connected by way of an IPv4 backbone. This connectivity can be accomplished by arranging tunnels directly with each IPv6 site to which connectivity is needed, but more typically it is done by arranging a tunnel into a

larger IPv6 routing infrastructure that can guarantee connectivity to all IPv6 end-user site networks. One example of this type of IPv6 routing infrastructure is the 6bone.

When using configured tunnels, a peering relationship must be established between the two IPv6 sites. This requires establishing a technical relationship with the peer and working through the various low-level details of how to configure tunnels between the two sites, including answering questions such as what peering protocol is used (presumably, an IPv6-capable version of BGP4).

Automatic tunnels

Automatic tunnels provide a simple mechanism to establish IPv6 connectivity between isolated dual-stack hosts and/or routers. In automatic tunneling, the IPv4 tunnel endpoint is determined from the IPv4 address embedded in the IPv4-compatible destination address of the IPv6 packet being tunneled. If the destination IPv6 address is IPv4-compatible, then the packet is sent by way of automatic tunneling. If the destination is IPv6-native, the packet cannot be sent by way of automatic tunneling. An IPv6-compatible address is identified by a ::/96 prefix and holds an IPv4 address in the low-order 32 bits. IPv4-compatible addresses are assigned exclusively to nodes that support automatic tunneling. It is globally unique as long as the IPv4 address is not from the private IPv4 address space.

When an IPv6 packet is sent over an automatic tunnel, the IPv6 packet is encapsulated within an IPv4 header as described in “Tunneling” on page 121. The source IPv4 address is an address of the interface the packet is sent over, and the destination IPv4 address is the low-order 32 bits of the IPv6 destination address. The packet is always sent in this form, even if the tunnel endpoint is on an attached link.

Automatic tunneling can be either host-host or router-host. A source host sends an IPv6 packet to an IPv6 router if possible, but that router might not be able to do the same and might have to perform automatic tunneling to the destination host itself. Because of the preference for the use of IPv6 routers rather than automatic tunneling, the tunnel is always as short as possible. However, the tunnel always extends all the way to the destination host. In order to use a tunnel that does not extend all the way to the recipient, another tunneling protocol must be used.

Guidelines: There are several issues to be aware of when using automatic tunnels.

- First, it does not solve the address exhaustion problem of IPv4, as it requires each tunnel endpoint to have an IPv4 address from which the IPv6 compatible address is created.
- Second, the use of IPv4 compatible addresses cause IPv4 addresses to be included in the IPv6 routing table, which in turn can cause a dramatic increase in the size of the IPv6 routing table.

Due to these concerns, use other tunneling protocols, such as 6to4 tunnels, in preference to automatic tunnels.

6to4 addresses

The IANA has permanently assigned one 13-bit IPv6 Top Level Aggregator (TLA) identifier under the IPv6 Format Prefix 001 for the 6to4 scheme. Its numeric value is 0x2002, i.e., it is 2002::/16 when expressed as an IPv6 address prefix.

The format for a 6to4 address is shown in Figure 20:

16 bits	32 bits	16 bits	64 bits
0x0002	V4ADDR	Subnet ID	Interface ID

Figure 20. 6to4 address format

Thus, this prefix has exactly the same format as normal /48 prefixes assigned according to other aggregatable global unicast addresses. It can be abbreviated as 2002:V4ADDR::/48. Within the subscriber site it can be used exactly like any other valid IPv6 prefix, for example, for automated address assignment and discovery for native IPv6 routing, or for the 6over4 mechanism.

6to4 provides a mechanism to allow isolated IPv6 domains, attached to a wide area network with no native IPv6 support, to communicate with other such IPv6 domains with minimal configuration. The idea is to embed IPv4 tunnel addresses into the IPv6 prefixes so that any domain border router can automatically discover tunnel endpoints for outbound IPv6 traffic.

The 6to4 transition mechanism advertises a site's IPv4 tunnel endpoint (to be used for a dynamic tunnel) in a special external routing prefix for that site. When one site tries to reach another site, it discovers the 6to4 tunnel endpoint from a DNS name to address lookup and use a dynamically built tunnel from site to site for communication. The tunnels are transient in that there is no state maintained for them, lasting only as long as a specified transaction uses the path.

A 6to4 site identifies one or more routers to run as a dual-mode stack and to act as a 6to4 router. A globally routable IPv4 address is assigned to the 6to4 router. The 6to4 prefix, which has the 6to4 router's IPv4 address embedded within it, is then advertised by way of the Neighbor Discovery protocol to the 6to4 site, and this prefix is used by hosts within the site to generate a global IPv6 address.

When one IPv6-enabled host at a 6to4 site tries to access an IPv6-enabled host by domain name at another 6to4 site, the DNS returns the IPv6 IP address for that host. The requesting host sends a packet to its nearest router, eventually reaching a site's 6to4 router. When the site's 6to4 router receives the packet and sees that it must send the packet to another site, and the next hop destination prefix is a 2002::/16 prefix, the IPv6 packet is encapsulated as described in "Tunneling" on page 121. The source IPv4 address is the one in the requesting site's 6to4 prefix (which is the IPv4 address of an outgoing interface for one of the site's 6to4 routers) and the destination IPv4 address is the one in the next hop destination 6to4 prefix of the IPv6 packet. When the destination site's 6to4 router receives the IPv4 packet, the IPv4 header is removed, leaving the original IPv6 packet for local forwarding.

6over4 tunnels

The Interface Identifier of an IPv4 interface using 6over4 is the 32-bit IPv4 address of that interface, padded to the left with 0s and is 64 bits in length. Note that the Universal/Local bit is 0, indicating that the Interface Identifier is not globally

unique. When the host has more than one IPv4 address in use on the physical interface concerned, an administrative choice of one of these IPv4 addresses is made.

The IPv6 Link-local address for an IPv4 virtual interface is formed by appending the Interface Identifier, as defined above, to the prefix FE80::/64.

3 bits	45 bits	16 bits	32 bits	32 bits
001	Network	Subnet	0.....0	IPv4 address

Figure 21. 6over4 address format

Site-local and global unicast addresses are generated by prepending a 64-bit prefix to the 6over4 Interface Identifier. These prefixes can be learned in any of the normal ways, for example, as part of stateless address autoconfiguration or by way of manual configuration.

6over4 is a transition mechanism which allows isolated IPv6 hosts, located on a physical link which has no directly connected IPv6 router, to use an IPv4 multicast domain as their virtual local link. A 6over4 host uses an IPv4 address for the interface in the creation of the IPv6 interface ID, placing the 32-bit IPv4 address in the low order bits and padding to the left with 0's for a total of 64 bits. The IPv6 prefix used is the normal IPv6 prefix, and can be manually configured or dynamically learned by way of Stateless Address Autoconfiguration.

Because 6over4 creates a virtual link using IPv4 multicast, at least one IPv6 router using the same method must be connected to the same IPv4 multicast domain if IPv6 routing to other links is required.

When encapsulating the IPv6 packet, the source IP address for the IPv4 packet is an IPv4 address from the sending interface of the 6over4 host. The destination IPv4 address is the low-order 32 bits of the IPv6 address of the next-hop for the packet. Note that the final destination of the packet does not need to be a 6over4 host, although it might be one.

Application migration and coexistence overview

Many IPv6 stacks support both IPv4 and IPv6 interfaces and are capable of receiving and sending native IPv4 and IPv6 packets over the corresponding interfaces. This type of TCP/IP stack is generally referred to as a dual-mode stack IP node. This does not mean that there are two separate TCP/IP stacks running on this type of node. It means that the TCP/IP stack has built-in support for both IPv4 and IPv6. In this document, the term dual-mode stack or IP node is a TCP/IP stack that supports both IPv4 and IPv6 protocols.




Figure 22. Dual-mode stack IP host

For a multihomed dual-mode IP host, it is a likely configuration that the host has both IPv4 and IPv6 interfaces over which requests for host-resident applications are received or sent. Older AF_INET applications are only able to communicate using IPv4 addresses. IPv6-enabled applications that use AF_INET6 sockets can communicate using both IPv4 and IPv6 addresses (on a dual-mode host). AF_INET and AF_INET6 applications are able to communicate with one another, but only using IPv4 addresses.

If the socket libraries on the IPv6-enabled host are updated to support IPv6 sockets (AF_INET6), applications can be IPv6 enabled. When an application on a dual mode stack host is IPv6 enabled, the application is able to communicate with both

IPv4 and IPv6 partners. This is true for both clients and server on a dual-mode stack host.

		Appl. on a dual mode host	
		IPv4-only	IPv6-enabled
IPv4-only partner	✓	✓	
IPv6-only partner			✓

Figure 23. Application communication on a dual-mode host

IPv6-enabling both sockets libraries and applications on dual-mode hosts therefore becomes a migration concern. As soon as IPv6-only hosts are being deployed in a network, applications on those IPv6-only nodes cannot communicate with the IPv4-only applications on the dual mode hosts, unless one of multiple migration technologies are implemented either on intermediate nodes in the network or directly on the dual mode hosts.

Application migration approaches

The ultimate and preferred migration approach for applications that reside on a dual-mode TCP/IP host is to IPv6-enable the applications by migrating them from AF_INET sockets to AF_INET6 sockets.

There are multiple reasons why this approach is not always applicable, such as the following:

- No access to the source code (vendor product, or source no longer available).
- The sockets API implementation does not yet (or never does) support IPv6.
- Resource availability or prioritization dictates a phased IPv6-enabling where not all applications can be available in an IPv6-enabled version at the same point in time where the stack is IPv6-capable.

For those applications that are not or cannot be IPv6 enabled, an alternative migration strategy is needed. The IETF has identified multiple approaches as summarized in draft RFC, *An Overview of the Introduction of IPv6 in the Internet*.

Some of the technologies that are defined by the IETF are supposed to be implemented on intermediate nodes that route traffic between IPv4 and IPv6 network segments. Other technologies are intended for implementation on the dual mode IP nodes themselves.

Translation mechanisms

This section provides an introduction to a few transition mechanisms that can be used when migrating to an IPv6 network.

The key to successful adoption and deployment of IPv6 is the transition from the installed IPv4 base. The goal of all transition strategies is to facilitate the partial and incremental upgrade of hosts, servers, routers, and network infrastructure. There are many possible approaches, and some of the more likely approaches are described below. The transition strategy a company chooses to take varies based on the particular needs of that company.

Several migration issues must be addressed when the backbone routing protocol is IPv4. First, a mechanism is needed to allow communication between islands of IPv6 networks that are interconnected only using the IPv4 backbone. Tunneling of IPv6 packets over the IPv4 network can be used to connect the clouds. Second, end-to-end communication between IPv4 and IPv6 applications must be enabled. Several approaches to accomplish this exist; Application Layer Gateways, NAT-PT, and Bump-in-the-Stack are all possibilities. During the migration phase, it is likely that a combination of one, multiple, or all of these transition mechanisms can be used.

Application Layer Gateways (ALGs) allow an IPv6-only applications to communicate to an IPv4-only peer. Using an ALG, the client connects to the ALG using its native protocol (IPv4 or IPv6) and the ALG connects to the server using the other protocol (IPv6 or IPv4, respectively).

SOCKS gateway

A SOCKS gateway is a method of providing an ALG. The SOCKS64 implementation works as a SOCKS server that relays communication between IPv4 and IPv6 flows. Servers do not require any changes, but client applications (or the stack where the client applications reside) need to be socksified to be able to reach out through a SOCKS64 server to an IPv6-only partner.

Proxy

Protocol translation involves converting IPv4 packets into IPv6 packets and vice versa. This translation typically involves some form of network address translation (NAT) in addition to the protocol translation (PT) function. It might execute in a specialized node that resides between an IPv4 network and an IPv6 network, or it might execute in the host that owns the IPv4 application.

Protocol Translation is useful when devices need to communicate but are not using the same protocol, allowing IPv6-only devices to communicate with IPv4-only devices. However, the following issues make a less-than ideal solution:

- Protocol translation is not foolproof. It is difficult to determine exactly how long to keep the mappings between the real IPv6 address and the locally mapped IPv4 address available. Eventually, an address is going to be reused before all servers have stopped accessing the address.
- Some applications might use the remote IP address as a means of performing a security check. Unless AH or an IPSec tunnel is used, then this method is not foolproof, but it is still done. If the IPv4 address is a locally mapped address, any checks such as this are broken.
- Displays and traces of the remote IP address are meaningless. Today, many applications generate messages, traces, and so on containing the IP address of the remote client.
- All DNS queries for the IPv4-mapped address must flow through the node that performed the NAT function. The DNS resolver or name server at this node, as well as the TCP/IP stack, must maintain a mapping between the IPv4 address and IPv6 address.
- Not all IPv6 protocols have IPv4 equivalents and vice versa. As such, it might not be possible to translate the contents of an IPv4 packet into an equivalent IPv6 packet and vice versa.

Stateless IP/ICMP Translation Algorithm (SIIT)

This algorithm translates between IPv4 and IPv6 packet headers (including ICMP headers) in separate translator boxes in the network without requiring any

per-connection state in those boxes. SIIT can be used as part of a solution that allows IPv6 hosts, which do not have permanently assigned IPv4 addresses, to communicate with IPv4-only hosts.

Network address translation - protocol translation (NAT-PT)

Protocol translation can occur at a specialized node that resides between IPv4 and IPv6 networks. This node is typically referred to as a NAT-PT device because it must translate between the IPv4 and IPv6 addresses, as well as between the IPv4 and IPv6 protocols.

An NAT-PT node plays a similar role to an ALG. Both nodes allow IPv4-only applications to communicate with IPv6-only peers, and both reside in similar places in the network. However, each takes a different approach to accomplish a similar goal.

SOCKS64 is a proxy solution and requires client applications to be updated to use SOCKS64. NAT-PT is not a proxy and requires no changes to either the client or server. Based solely on this, NAT-PT might appear to be a superior solution. However, due to the limitations of NAT-PT and familiarity with SOCKS, it is more likely that SOCKS64 is used to allow IPv4-only applications to communicate with IPv6-only peers.

Chapter 11. IPv6 support tables

This appendix contains the IPv6 support tables and includes the following sections:

- “Supported IPv6 standards”
- “z/OS-specific features”
- “Applications not enabled for IPv6” on page 134

Supported IPv6 standards

Table 28 lists the supported IPv6 standards. RFCs are not implemented in their entirety.

Table 28. Supported IPv6 standards

Standard	RFC or Internet Draft
DNS Extensions to support IP version 6	1886
Path MTU discovery	1981
RIPng for IPv6	2080
An IPv6 Aggregatable Global Unicast Address Format	2374
FTP Extensions for IPv6 and NATs	2428
Internet Protocol, Version 6 (IPv6) Specification	2460
Neighbor discovery for IP Version 6 (IPv6)	2461
IPv6 Stateless Address Autoconfiguration	2462
Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification	2463
Transmission of IPv6 Packets over Ethernet Networks	2464
Multicast Listener Discovery (MLD) for IPv6	2710
IPv6 Router Alert Option	2711
OSPF for IPv6	2740
DNS Extensions to Support IPv6 Address Aggregation and Renumbering	2874
Default Address Selection for Internet Protocol Version 6 (IPv6)	3484
Basic Socket Interface Extensions for IPv6	3493
Internet Protocol Version 6 (IPv6) Addressing Architecture	3513
Advanced Sockets Application Programming Interface (API) for IPv6	3542

z/OS-specific features

The tables in this section summarize z/OS TCP/IP features and the level of support provided in an IPv6 network. In the future, additional features are projected for IPv6 support in subsequent releases of the z/OS Communications Server.

Table 29 lists the link-layer device support.

Table 29. Link-layer device support

Link-layer device support	IPv4 support	IPv6 support	Comments
OSA-Express in QDIO mode	Y	Y	Fast and Gigabit Ethernet support for IPv6 traffic is configured by way of an INTERFACE statement of type IPAQENET6.
CTC	Y	N	None
LCS	Y	N	None
CLAW	Y	N	None
CDLC (3745/3746)	Y	N	None
SNALINK LU0 and LU6.2	Y	N	None
X.25 NPSI	Y	N	None
NSC HyperChannel	Y	N	None
MPC Point-Point	Y	Y	Support is configured by way of an INTERFACE statement of type MPCPTP6.
ATM	Y	N	None
HiperSockets	Y	Y	Support is configured by way of an INTERFACE statement of type IPAQIDIO6 or dynamically configured by way of the IP CONFIG6 DYNAMICXCF statement.
XCF	Y	Y	Support is configured by way of an INTERFACE statement of type MPCPTP6 or dynamically configured by way of the IP CONFIG6 DYNAMICXCF statement.

Table 30 lists virtual IP Addressing support.

Table 30. Virtual IP Addressing support

Virtual IP Addressing support	IPv4 support	IPv6 support	Comments
Virtual Device/Interface Configuration for static VIPA	Y	Y	None

All sysplex functions support IPv6 except for those listed in Table 31.

Table 31. Sysplex support

Sysplex support	IPv4 support	IPv6 support	Comments
Sysplex distributor integration with Cisco MNLB	Y	N	None
Sysplex Wide Security Associations (SWSA)	Y	N	None

Table 32 lists IP routing functions.

Table 32. IP routing functions

IP routing functions	IPv4 support	IPv6 support	Comments
Dynamic Routing - OSPF	Y	Y	None
Dynamic Routing - RIP	Y	Y	None
Static Route Configuration by way of BEGINROUTES statement	Y	Y	None
Static Route Configuration by way of GATEWAY statement	Y	N	None
Multipath Routing Groups	Y	Y	None

Table 33 lists miscellaneous IP/IF-layer functions.

Table 33. Miscellaneous IP/IF-layer functions

Misc. IP/IF-layer functions	IPv4 support	IPv6 support	Comments
Path MTU Discovery	Y	Y	None
Configurable Device or Interface Recovery Interval	Y	Y	None
Link-Layer Address Resolution	Y	Y	None
ARP/Neighbor Cache PURGE Capability	Y	Y	None
Datagram Forwarding Enable/Disable	Y	Y	None
Hipersockets accelerator	Y	N	Support is enabled by way of the IQDIOROUTING parameter on the IPCONFIG statement.
Checksum offload	Y	N	None
Segmentation offload	Y	N	None

Table 34 lists transport-layer functions.

Table 34. Transport-layer functions

Transport-layer functions	IPv4 support	IPv6 support	Comments
Fast Response Cache Accelerator	Y	N	None
Enterprise Extender	Y	Y	IPv6 Enterprise Extender support requires a virtual IP address configured by way of an INTERFACE statement of type VIRTUAL6 and IUTSAMEH configured by way of an INTERFACE statement of type MPCPTP6 or dynamically configured by way of IPCONFIG6 DYNAMICXCF.

Table 34. Transport-layer functions (continued)

Transport-layer functions	IPv4 support	IPv6 support	Comments
Server-BIND Control	Y	Y	None
UDP Checksum Disablement Option	Y	N	None

Table 35 lists network management and accounting functions.

Table 35. Network management and accounting functions

Network management and accounting Functions	IPv4 support	IPv6 support	Comments
SNMP	Y	Y	None
SNMP agent	Y	Y	None
TCP/IP subagent	Y	Y	No IPv6 UDP support
Network SLAPM2 subagent	Y	Y	None
Distributed Protocol Interface	Y	Y	None
OMPROUTE subagent	Y	N	None
Trap forwarder daemon	Y	Y	None
Policy-Based Networking	Y	Y	None
SMF	Y	Y	None
TN3270 subagent	Y	Y	None

Table 36 lists security functions.

Table 36. Security functions

Security functions	IPv4 support	IPv6 support	Comments
IPSec	Y	Y	None
IP filtering	Y	Y	None
IKE daemon	Y	Y	None
NAT traversal	Y	N	None
Network Access Control	Y	Y	None
Stack and Port Access Control	Y	Y	None
Application Transparent TLS	Y	Y	None
Intrusion Detection Services	Y	N	None

Applications not enabled for IPv6

Some applications are not enabled for IPv6. These applications are listed in Table 37.

Table 37. Applications not enabled for IPv6

Server applications	IPv4 support	IPv6 support
SMTPPROC/NJE server	Y	N

Table 37. Applications not enabled for IPv6 (continued)

Server applications	IPv4 support	IPv6 support
Rlogind server	Y	N
MVS Miscellaneous server	Y	N
Popper	Y	N
NDB server	Y	N
MVS LPD server	Y	N
DHCPD server	Y	N
TIMED server	Y	N
NCS LLBD and GLBD servers	Y	N
ONC/RPC MVS portmapper	Y	N
ONC/RPC UNIX portmapper	Y	N
NCROUTE	Y	N
NPF	Y	N
RSVP daemon	Y	N
UNIX named (BIND 4.9.3 based)	Y	N
Client applications		
TSO TELNET client	Y	N
TSO LPR client	Y	N
Command-type applications		
TSO NSLOOKUP	Y	N
UNIX nslookup (BIND 4.9.3-based)	Y	N
UNIX nsupdate (BIND 4.9.3-based)	Y	N
TSO LPRM	Y	N
TSO DIG	Y	N
UNIX dig	Y	N
TSO RPCINFO	Y	N
UNIX rpcinfo	Y	N

Part 5. Appendixes

This section contains the following appendixes:

- Appendix A, “Related protocol specifications (RFCs),” on page 139 contains related protocol specifications (RFCs).
- Appendix B, “Information APARs and technotes,” on page 155 lists information APARs for IP and SNA documents.
- Appendix C, “Accessibility,” on page 159 describes accessibility features to help users with physical disabilities.
- “Notices” on page 161 contains notices and trademarks used in this document.
- “Bibliography” on page 171 contains descriptions of the documents in the z/OS Communications Server library.

Appendix A. Related protocol specifications (RFCs)

This appendix lists the related protocol specifications for TCP/IP. The Internet Protocol suite is still evolving through requests for comments (RFC). New protocols are being designed and implemented by researchers and are brought to the attention of the Internet community in the form of RFCs. Some of these protocols are so useful that they become recommended protocols. That is, all future implementations for TCP/IP are recommended to implement these particular functions or protocols. These become the *de facto* standards, on which the TCP/IP protocol suite is built.

You can request RFCs through electronic mail, from the automated Network Information Center (NIC) mail server, by sending a message to `service@nic.ddn.mil` with a subject line of RFC *nnnn* for text versions or a subject line of RFC *nnnn.PS* for PostScript versions. To request a copy of the RFC index, send a message with a subject line of RFC INDEX.

For more information, contact `nic@nic.ddn.mil` or at:

Government Systems, Inc.
Attn: Network Information Center
14200 Park Meadow Drive
Suite 200
Chantilly, VA 22021

Hard copies of all RFCs are available from the NIC, either individually or by subscription. Online copies are available at the following Web address:
<http://www.rfc-editor.org/rfc.html>.

See "Internet drafts" on page 154 for draft RFCs implemented in this and previous Communications Server releases.

Many features of TCP/IP Services are based on the following RFCs:

RFC	Title and Author
RFC 652	<i>Telnet output carriage-return disposition option</i> D. Crocker
RFC 653	<i>Telnet output horizontal tabstops option</i> D. Crocker
RFC 654	<i>Telnet output horizontal tab disposition option</i> D. Crocker
RFC 655	<i>Telnet output formfeed disposition option</i> D. Crocker
RFC 657	<i>Telnet output vertical tab disposition option</i> D. Crocker
RFC 658	<i>Telnet output linefeed disposition</i> D. Crocker
RFC 698	<i>Telnet extended ASCII option</i> T. Mock
RFC 726	<i>Remote Controlled Transmission and Echoing Telnet option</i> J. Postel, D. Crocker
RFC 727	<i>Telnet logout option</i> M.R. Crispin
RFC 732	<i>Telnet Data Entry Terminal option</i> J.D. Day
RFC 733	<i>Standard for the format of ARPA network text messages</i> D. Crocker, J. Vittal, K.T. Pogran, D.A. Henderson

RFC 734	<i>SUPDUP Protocol</i> M.R. Crispin
RFC 735	<i>Revised Telnet byte macro option</i> D. Crocker, R.H. Gumpertz
RFC 736	<i>Telnet SUPDUP option</i> M.R. Crispin
RFC 749	<i>Telnet SUPDUP—Output option</i> B. Greenberg
RFC 765	<i>File Transfer Protocol specification</i> J. Postel
RFC 768	<i>User Datagram Protocol</i> J. Postel
RFC 779	<i>Telnet send-location option</i> E. Killian
RFC 783	<i>TFTP Protocol (revision 2)</i> K.R. Sollins
RFC 791	<i>Internet Protocol</i> J. Postel
RFC 792	<i>Internet Control Message Protocol</i> J. Postel
RFC 793	<i>Transmission Control Protocol</i> J. Postel
RFC 820	<i>Assigned numbers</i> J. Postel
RFC 821	<i>Simple Mail Transfer Protocol</i> J. Postel
RFC 822	<i>Standard for the format of ARPA Internet text messages</i> D. Crocker
RFC 823	<i>DARPA Internet gateway</i> R. Hinden, A. Sheltzer
RFC 826	<i>Ethernet Address Resolution Protocol: Or converting network protocol addresses to 48.bit Ethernet address for transmission on Ethernet hardware</i> D. Plummer
RFC 854	<i>Telnet Protocol Specification</i> J. Postel, J. Reynolds
RFC 855	<i>Telnet Option Specification</i> J. Postel, J. Reynolds
RFC 856	<i>Telnet Binary Transmission</i> J. Postel, J. Reynolds
RFC 857	<i>Telnet Echo Option</i> J. Postel, J. Reynolds
RFC 858	<i>Telnet Suppress Go Ahead Option</i> J. Postel, J. Reynolds
RFC 859	<i>Telnet Status Option</i> J. Postel, J. Reynolds
RFC 860	<i>Telnet Timing Mark Option</i> J. Postel, J. Reynolds
RFC 861	<i>Telnet Extended Options: List Option</i> J. Postel, J. Reynolds
RFC 862	<i>Echo Protocol</i> J. Postel
RFC 863	<i>Discard Protocol</i> J. Postel
RFC 864	<i>Character Generator Protocol</i> J. Postel
RFC 865	<i>Quote of the Day Protocol</i> J. Postel
RFC 868	<i>Time Protocol</i> J. Postel, K. Harrenstien
RFC 877	<i>Standard for the transmission of IP datagrams over public data networks</i> J.T. Korb
RFC 883	<i>Domain names: Implementation specification</i> P.V. Mockapetris
RFC 884	<i>Telnet terminal type option</i> M. Solomon, E. Wimmers
RFC 885	<i>Telnet end of record option</i> J. Postel
RFC 894	<i>Standard for the transmission of IP datagrams over Ethernet networks</i> C. Hornig
RFC 896	<i>Congestion control in IP/TCP internetworks</i> J. Nagle

RFC 903	<i>Reverse Address Resolution Protocol</i> R. Finlayson, T. Mann, J. Mogul, M. Theimer
RFC 904	<i>Exterior Gateway Protocol formal specification</i> D. Mills
RFC 919	<i>Broadcasting Internet Datagrams</i> J. Mogul
RFC 922	<i>Broadcasting Internet datagrams in the presence of subnets</i> J. Mogul
RFC 927	<i>TACACS user identification Telnet option</i> B.A. Anderson
RFC 933	<i>Output marking Telnet option</i> S. Silverman
RFC 946	<i>Telnet terminal location number option</i> R. Nedved
RFC 950	<i>Internet Standard Subnetting Procedure</i> J. Mogul, J. Postel
RFC 951	<i>Bootstrap Protocol</i> W.J. Croft, J. Gilmore
RFC 952	<i>DoD Internet host table specification</i> K. Harrenstien, M. Stahl, E. Feinler
RFC 959	<i>File Transfer Protocol</i> J. Postel, J.K. Reynolds
RFC 961	<i>Official ARPA-Internet protocols</i> J.K. Reynolds, J. Postel
RFC 974	<i>Mail routing and the domain system</i> C. Partridge
RFC 1001	<i>Protocol standard for a NetBIOS service on a TCP/UDP transport: Concepts and methods</i> NetBios Working Group in the Defense Advanced Research Projects Agency, Internet Activities Board, End-to-End Services Task Force
RFC 1002	<i>Protocol Standard for a NetBIOS service on a TCP/UDP transport: Detailed specifications</i> NetBios Working Group in the Defense Advanced Research Projects Agency, Internet Activities Board, End-to-End Services Task Force
RFC 1006	<i>ISO transport services on top of the TCP: Version 3</i> M.T. Rose, D.E. Cass
RFC 1009	<i>Requirements for Internet gateways</i> R. Braden, J. Postel
RFC 1011	<i>Official Internet protocols</i> J. Reynolds, J. Postel
RFC 1013	<i>X Window System Protocol, version 11: Alpha update April 1987</i> R. Scheifler
RFC 1014	<i>XDR: External Data Representation standard</i> Sun Microsystems
RFC 1027	<i>Using ARP to implement transparent subnet gateways</i> S. Carl-Mitchell, J. Quarterman
RFC 1032	<i>Domain administrators guide</i> M. Stahl
RFC 1033	<i>Domain administrators operations guide</i> M. Lottor
RFC 1034	<i>Domain names—concepts and facilities</i> P.V. Mockapetris
RFC 1035	<i>Domain names—implementation and specification</i> P.V. Mockapetris
RFC 1038	<i>Draft revised IP security option</i> M. St. Johns
RFC 1041	<i>Telnet 3270 regime option</i> Y. Rekhter
RFC 1042	<i>Standard for the transmission of IP datagrams over IEEE 802 networks</i> J. Postel, J. Reynolds
RFC 1043	<i>Telnet Data Entry Terminal option: DODIIS implementation</i> A. Yasuda, T. Thompson

RFC 1044	<i>Internet Protocol on Network System's HYPERchannel: Protocol specification</i> K. Hardwick, J. Lekashman
RFC 1053	<i>Telnet X.3 PAD option</i> S. Levy, T. Jacobson
RFC 1055	<i>Nonstandard for transmission of IP datagrams over serial lines: SLIP</i> J. Romkey
RFC 1057	<i>RPC: Remote Procedure Call Protocol Specification: Version 2</i> Sun Microsystems
RFC 1058	<i>Routing Information Protocol</i> C. Hedrick
RFC 1060	<i>Assigned numbers</i> J. Reynolds, J. Postel
RFC 1067	<i>Simple Network Management Protocol</i> J.D. Case, M. Fedor, M.L. Schoffstall, J. Davin
RFC 1071	<i>Computing the Internet checksum</i> R.T. Braden, D.A. Borman, C. Partridge
RFC 1072	<i>TCP extensions for long-delay paths</i> V. Jacobson, R.T. Braden
RFC 1073	<i>Telnet window size option</i> D. Waitzman
RFC 1079	<i>Telnet terminal speed option</i> C. Hedrick
RFC 1085	<i>ISO presentation services on top of TCP/IP based internets</i> M.T. Rose
RFC 1091	<i>Telnet terminal-type option</i> J. VanBokkelen
RFC 1094	<i>NFS: Network File System Protocol specification</i> Sun Microsystems
RFC 1096	<i>Telnet X display location option</i> G. Marcy
RFC 1101	<i>DNS encoding of network names and other types</i> P. Mockapetris
RFC 1112	<i>Host extensions for IP multicasting</i> S.E. Deering
RFC 1113	<i>Privacy enhancement for Internet electronic mail: Part I — message encipherment and authentication procedures</i> J. Linn
RFC 1118	<i>Hitchhikers Guide to the Internet</i> E. Krol
RFC 1122	<i>Requirements for Internet Hosts—Communication Layers</i> R. Braden, Ed.
RFC 1123	<i>Requirements for Internet Hosts—Application and Support</i> R. Braden, Ed.
RFC 1146	<i>TCP alternate checksum options</i> J. Zweig, C. Partridge
RFC 1155	<i>Structure and identification of management information for TCP/IP-based internets</i> M. Rose, K. McCloghrie
RFC 1156	<i>Management Information Base for network management of TCP/IP-based internets</i> K. McCloghrie, M. Rose
RFC 1157	<i>Simple Network Management Protocol (SNMP)</i> J. Case, M. Fedor, M. Schoffstall, J. Davin
RFC 1158	<i>Management Information Base for network management of TCP/IP-based internets: MIB-II</i> M. Rose
RFC 1166	<i>Internet numbers</i> S. Kirkpatrick, M.K. Stahl, M. Recker
RFC 1179	<i>Line printer daemon protocol</i> L. McLaughlin
RFC 1180	<i>TCP/IP tutorial</i> T. Socolofsky, C. Kale

RFC 1183	<i>New DNS RR Definitions</i> C.F. Everhart, L.A. Mamakos, R. Ullmann, P.V. Mockapetris
RFC 1184	<i>Telnet Linemode Option</i> D. Borman
RFC 1186	<i>MD4 Message Digest Algorithm</i> R.L. Rivest
RFC 1187	<i>Bulk Table Retrieval with the SNMP</i> M. Rose, K. McCloghrie, J. Davin
RFC 1188	<i>Proposed Standard for the Transmission of IP Datagrams over FDDI Networks</i> D. Katz
RFC 1190	<i>Experimental Internet Stream Protocol: Version 2 (ST-II)</i> C. Topolcic
RFC 1191	<i>Path MTU discovery</i> J. Mogul, S. Deering
RFC 1198	<i>FYI on the X window system</i> R. Scheifler
RFC 1207	<i>FYI on Questions and Answers: Answers to commonly asked "experienced Internet user" questions</i> G. Malkin, A. Marine, J. Reynolds
RFC 1208	<i>Glossary of networking terms</i> O. Jacobsen, D. Lynch
RFC 1213	<i>Management Information Base for Network Management of TCP/IP-based internets: MIB-II</i> K. McCloghrie, M.T. Rose
RFC 1215	<i>Convention for defining traps for use with the SNMP</i> M. Rose
RFC 1227	<i>SNMP MUX protocol and MIB</i> M.T. Rose
RFC 1228	<i>SNMP-DPI: Simple Network Management Protocol Distributed Program Interface</i> G. Carpenter, B. Wijnen
RFC 1229	<i>Extensions to the generic-interface MIB</i> K. McCloghrie
RFC 1230	<i>IEEE 802.4 Token Bus MIB</i> K. McCloghrie, R. Fox
RFC 1231	<i>IEEE 802.5 Token Ring MIB</i> K. McCloghrie, R. Fox, E. Decker
RFC 1236	<i>IP to X.121 address mapping for DDN</i> L. Morales, P. Hasse
RFC 1256	<i>ICMP Router Discovery Messages</i> S. Deering, Ed.
RFC 1267	<i>Border Gateway Protocol 3 (BGP-3)</i> K. Lougheed, Y. Rekhter
RFC 1268	<i>Application of the Border Gateway Protocol in the Internet</i> Y. Rekhter, P. Gross
RFC 1269	<i>Definitions of Managed Objects for the Border Gateway Protocol: Version 3</i> S. Willis, J. Burruss
RFC 1270	<i>SNMP Communications Services</i> F. Kastenholz, ed.
RFC 1285	<i>FDDI Management Information Base</i> J. Case
RFC 1315	<i>Management Information Base for Frame Relay DTEs</i> C. Brown, F. Baker, C. Carvalho
RFC 1321	<i>The MD5 Message-Digest Algorithm</i> R. Rivest
RFC 1323	<i>TCP Extensions for High Performance</i> V. Jacobson, R. Braden, D. Borman
RFC 1325	<i>FYI on Questions and Answers: Answers to Commonly Asked "New Internet User" Questions</i> G. Malkin, A. Marine
RFC 1327	<i>Mapping between X.400 (1988)/ISO 10021 and RFC 822</i> S. Hardcastle-Kille

RFC 1340	<i>Assigned Numbers</i> J. Reynolds, J. Postel
RFC 1344	<i>Implications of MIME for Internet Mail Gateways</i> N. Bornstein
RFC 1349	<i>Type of Service in the Internet Protocol Suite</i> P. Almquist
RFC 1350	<i>The TFTP Protocol (Revision 2)</i> K.R. Sollins
RFC 1351	<i>SNMP Administrative Model</i> J. Davin, J. Galvin, K. McCloghrie
RFC 1352	<i>SNMP Security Protocols</i> J. Galvin, K. McCloghrie, J. Davin
RFC 1353	<i>Definitions of Managed Objects for Administration of SNMP Parties</i> K. McCloghrie, J. Davin, J. Galvin
RFC 1354	<i>IP Forwarding Table MIB</i> F. Baker
RFC 1356	<i>Multiprotocol Interconnect on X.25 and ISDN in the Packet Mode</i> A. Malis, D. Robinson, R. Ullmann
RFC 1358	<i>Charter of the Internet Architecture Board (IAB)</i> L. Chapin
RFC 1363	<i>A Proposed Flow Specification</i> C. Partridge
RFC 1368	<i>Definition of Managed Objects for IEEE 802.3 Repeater Devices</i> D. McMaster, K. McCloghrie
RFC 1372	<i>Telnet Remote Flow Control Option</i> C. L. Hedrick, D. Borman
RFC 1374	<i>IP and ARP on HIPPI</i> J. Renwick, A. Nicholson
RFC 1381	<i>SNMP MIB Extension for X.25 LAPB</i> D. Throop, F. Baker
RFC 1382	<i>SNMP MIB Extension for the X.25 Packet Layer</i> D. Throop
RFC 1387	<i>RIP Version 2 Protocol Analysis</i> G. Malkin
RFC 1388	<i>RIP Version 2 Carrying Additional Information</i> G. Malkin
RFC 1389	<i>RIP Version 2 MIB Extensions</i> G. Malkin, F. Baker
RFC 1390	<i>Transmission of IP and ARP over FDDI Networks</i> D. Katz
RFC 1393	<i>Traceroute Using an IP Option</i> G. Malkin
RFC 1398	<i>Definitions of Managed Objects for the Ethernet-Like Interface Types</i> F. Kastenholz
RFC 1408	<i>Telnet Environment Option</i> D. Borman, Ed.
RFC 1413	<i>Identification Protocol</i> M. St. Johns
RFC 1416	<i>Telnet Authentication Option</i> D. Borman, ed.
RFC 1420	<i>SNMP over IPX</i> S. Bostock
RFC 1428	<i>Transition of Internet Mail from Just-Send-8 to 8bit-SMTP/MIME</i> G. Vaudreuil
RFC 1442	<i>Structure of Management Information for version 2 of the Simple Network Management Protocol (SNMPv2)</i> J. Case, K. McCloghrie, M. Rose, S. Waldbusser
RFC 1443	<i>Textual Conventions for version 2 of the Simple Network Management Protocol (SNMPv2)</i> J. Case, K. McCloghrie, M. Rose, S. Waldbusser
RFC 1445	<i>Administrative Model for version 2 of the Simple Network Management Protocol (SNMPv2)</i> J. Galvin, K. McCloghrie
RFC 1447	<i>Party MIB for version 2 of the Simple Network Management Protocol (SNMPv2)</i> K. McCloghrie, J. Galvin

RFC 1448 *Protocol Operations for version 2 of the Simple Network Management Protocol (SNMPv2)* J. Case, K. McCloghrie, M. Rose, S. Waldbusser

RFC 1464 *Using the Domain Name System to Store Arbitrary String Attributes* R. Rosenbaum

RFC 1469 *IP Multicast over Token-Ring Local Area Networks* T. Pusateri

RFC 1483 *Multiprotocol Encapsulation over ATM Adaptation Layer 5* Juha Heinanen

RFC 1497 *BOOTP Vendor Information Extensions* J. Reynolds

RFC 1514 *Host Resources MIB* P. Grillo, S. Waldbusser

RFC 1516 *Definitions of Managed Objects for IEEE 802.3 Repeater Devices* D. McMaster, K. McCloghrie

RFC 1521 *MIME (Multipurpose Internet Mail Extensions) Part One: Mechanisms for Specifying and Describing the Format of Internet Message Bodies* N. Borenstein, N. Freed

RFC 1533 *DHCP Options and BOOTP Vendor Extensions* S. Alexander, R. Droms

RFC 1534 *Interoperation Between DHCP and BOOTP* R. Droms

RFC 1535 *A Security Problem and Proposed Correction With Widely Deployed DNS Software* E. Gavron

RFC 1536 *Common DNS Implementation Errors and Suggested Fixes* A. Kumar, J. Postel, C. Neuman, P. Danzig, S. Miller

RFC 1537 *Common DNS Data File Configuration Errors* P. Beertema

RFC 1540 *Internet Official Protocol Standards* J. Postel

RFC 1541 *Dynamic Host Configuration Protocol* R. Droms

RFC 1542 *Clarifications and Extensions for the Bootstrap Protocol* W. Wimer

RFC 1571 *Telnet Environment Option Interoperability Issues* D. Borman

RFC 1572 *Telnet Environment Option* S. Alexander

RFC 1573 *Evolution of the Interfaces Group of MIB-II* K. McCloghrie, F. Kastenholz

RFC 1577 *Classical IP and ARP over ATM* M. Laubach

RFC 1583 *OSPF Version 2* J. Moy

RFC 1591 *Domain Name System Structure and Delegation* J. Postel

RFC 1592 *Simple Network Management Protocol Distributed Protocol Interface Version 2.0* B. Wijnen, G. Carpenter, K. Curran, A. Sehgal, G. Waters

RFC 1594 *FYI on Questions and Answers— Answers to Commonly Asked "New Internet User" Questions* A. Marine, J. Reynolds, G. Malkin

RFC 1644 *T/TCP — TCP Extensions for Transactions Functional Specification* R. Braden

RFC 1646 *TN3270 Extensions for LUname and Printer Selection* C. Graves, T. Butts, M. Angel

RFC 1647 *TN3270 Enhancements* B. Kelly

RFC 1652	<i>SMTP Service Extension for 8bit-MIMEtransport</i> J. Klensin, N. Freed, M. Rose, E. Stefferud, D. Crocker
RFC 1664	<i>Using the Internet DNS to Distribute RFC1327 Mail Address Mapping Tables</i> C. Allochio, A. Bonito, B. Cole, S. Giordano, R. Hagens
RFC 1693	<i>An Extension to TCP: Partial Order Service</i> T. Connolly, P. Amer, P. Conrad
RFC 1695	<i>Definitions of Managed Objects for ATM Management Version 8.0 using SMIv2</i> M. Ahmed, K. Tesink
RFC 1701	<i>Generic Routing Encapsulation (GRE)</i> S. Hanks, T. Li, D. Farinacci, P. Traina
RFC 1702	<i>Generic Routing Encapsulation over IPv4 networks</i> S. Hanks, T. Li, D. Farinacci, P. Traina
RFC 1706	<i>DNS NSAP Resource Records</i> B. Manning, R. Colella
RFC 1712	<i>DNS Encoding of Geographical Location</i> C. Farrell, M. Schulze, S. Pleitner D. Baldoni
RFC 1713	<i>Tools for DNS debugging</i> A. Romao
RFC 1723	<i>RIP Version 2—Carrying Additional Information</i> G. Malkin
RFC 1752	<i>The Recommendation for the IP Next Generation Protocol</i> S. Bradner, A. Mankin
RFC 1766	<i>Tags for the Identification of Languages</i> H. Alvestrand
RFC 1771	<i>A Border Gateway Protocol 4 (BGP-4)</i> Y. Rekhter, T. Li
RFC 1794	<i>DNS Support for Load Balancing</i> T. Brisco
RFC 1819	<i>Internet Stream Protocol Version 2 (ST2) Protocol Specification—Version ST2+</i> L. Delgrossi, L. Berger Eds.
RFC 1826	<i>IP Authentication Header</i> R. Atkinson
RFC 1828	<i>IP Authentication using Keyed MD5</i> P. Metzger, W. Simpson
RFC 1829	<i>The ESP DES-CBC Transform</i> P. Karn, P. Metzger, W. Simpson
RFC 1830	<i>SMTP Service Extensions for Transmission of Large and Binary MIME Messages</i> G. Vaudreuil
RFC 1831	<i>RPC: Remote Procedure Call Protocol Specification Version 2</i> R. Srinivasan
RFC 1832	<i>XDR: External Data Representation Standard</i> R. Srinivasan
RFC 1833	<i>Binding Protocols for ONC RPC Version 2</i> R. Srinivasan
RFC 1850	<i>OSPF Version 2 Management Information Base</i> F. Baker, R. Coltun
RFC 1854	<i>SMTP Service Extension for Command Pipelining</i> N. Freed
RFC 1869	<i>SMTP Service Extensions</i> J. Klensin, N. Freed, M. Rose, E. Stefferud, D. Crocker
RFC 1870	<i>SMTP Service Extension for Message Size Declaration</i> J. Klensin, N. Freed, K. Moore
RFC 1876	<i>A Means for Expressing Location Information in the Domain Name System</i> C. Davis, P. Vixie, T. Goodwin, I. Dickinson
RFC 1883	<i>Internet Protocol, Version 6 (IPv6) Specification</i> S. Deering, R. Hinden

RFC 1884	<i>IP Version 6 Addressing Architecture</i> R. Hinden, S. Deering, Eds.
RFC 1886	<i>DNS Extensions to support IP version 6</i> S. Thomson, C. Huitema
RFC 1888	<i>OSI NSAPs and IPv6</i> J. Bound, B. Carpenter, D. Harrington, J. Houldsworth, A. Lloyd
RFC 1891	<i>SMTP Service Extension for Delivery Status Notifications</i> K. Moore
RFC 1892	<i>The Multipart/Report Content Type for the Reporting of Mail System Administrative Messages</i> G. Vaudreuil
RFC 1894	<i>An Extensible Message Format for Delivery Status Notifications</i> K. Moore, G. Vaudreuil
RFC 1901	<i>Introduction to Community-based SNMPv2</i> J. Case, K. McCloghrie, M. Rose, S. Waldbusser
RFC 1902	<i>Structure of Management Information for Version 2 of the Simple Network Management Protocol (SNMPv2)</i> J. Case, K. McCloghrie, M. Rose, S. Waldbusser
RFC 1903	<i>Textual Conventions for Version 2 of the Simple Network Management Protocol (SNMPv2)</i> J. Case, K. McCloghrie, M. Rose, S. Waldbusser
RFC 1904	<i>Conformance Statements for Version 2 of the Simple Network Management Protocol (SNMPv2)</i> J. Case, K. McCloghrie, M. Rose, S. Waldbusser
RFC 1905	<i>Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2)</i> J. Case, K. McCloghrie, M. Rose, S. Waldbusser
RFC 1906	<i>Transport Mappings for Version 2 of the Simple Network Management Protocol (SNMPv2)</i> J. Case, K. McCloghrie, M. Rose, S. Waldbusser
RFC 1907	<i>Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2)</i> J. Case, K. McCloghrie, M. Rose, S. Waldbusser
RFC 1908	<i>Coexistence between Version 1 and Version 2 of the Internet-standard Network Management Framework</i> J. Case, K. McCloghrie, M. Rose, S. Waldbusser
RFC 1912	<i>Common DNS Operational and Configuration Errors</i> D. Barr
RFC 1918	<i>Address Allocation for Private Internets</i> Y. Rekhter, B. Moskowitz, D. Karrenberg, G.J. de Groot, E. Lear
RFC 1928	<i>SOCKS Protocol Version 5</i> M. Leech, M. Ganis, Y. Lee, R. Kuris, D. Koblas, L. Jones
RFC 1930	<i>Guidelines for creation, selection, and registration of an Autonomous System (AS)</i> J. Hawkinson, T. Bates
RFC 1939	<i>Post Office Protocol-Version 3</i> J. Myers, M. Rose
RFC 1981	<i>Path MTU Discovery for IP version 6</i> J. McCann, S. Deering, J. Mogul
RFC 1982	<i>Serial Number Arithmetic</i> R. Elz, R. Bush
RFC 1985	<i>SMTP Service Extension for Remote Message Queue Starting</i> J. De Winter
RFC 1995	<i>Incremental Zone Transfer in DNS</i> M. Ohta
RFC 1996	<i>A Mechanism for Prompt Notification of Zone Changes (DNS NOTIFY)</i> P. Vixie

RFC 2010	<i>Operational Criteria for Root Name Servers</i> B. Manning, P. Vixie
RFC 2011	<i>SNMPv2 Management Information Base for the Internet Protocol using SMIv2</i> K. McCloghrie, Ed.
RFC 2012	<i>SNMPv2 Management Information Base for the Transmission Control Protocol using SMIv2</i> K. McCloghrie, Ed.
RFC 2013	<i>SNMPv2 Management Information Base for the User Datagram Protocol using SMIv2</i> K. McCloghrie, Ed.
RFC 2018	<i>TCP Selective Acknowledgement Options</i> M. Mathis, J. Mahdavi, S. Floyd, A. Romanow
RFC 2026	<i>The Internet Standards Process — Revision 3</i> S. Bradner
RFC 2030	<i>Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI</i> D. Mills
RFC 2033	<i>Local Mail Transfer Protocol</i> J. Myers
RFC 2034	<i>SMTP Service Extension for Returning Enhanced Error Codes</i> N. Freed
RFC 2040	<i>The RC5, RC5-CBC, RC-5-CBC-Pad, and RC5-CTS Algorithms</i> R. Baldwin, R. Rivest
RFC 2045	<i>Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies</i> N. Freed, N. Borenstein
RFC 2052	<i>A DNS RR for specifying the location of services (DNS SRV)</i> A. Gulbrandsen, P. Vixie
RFC 2065	<i>Domain Name System Security Extensions</i> D. Eastlake 3rd, C. Kaufman
RFC 2066	<i>TELNET CHARSET Option</i> R. Gellens
RFC 2080	<i>RIPng for IPv6</i> G. Malkin, R. Minnear
RFC 2096	<i>IP Forwarding Table MIB</i> F. Baker
RFC 2104	<i>HMAC: Keyed-Hashing for Message Authentication</i> H. Krawczyk, M. Bellare, R. Canetti
RFC 2119	<i>Keywords for use in RFCs to Indicate Requirement Levels</i> S. Bradner
RFC 2132	<i>DHCP Options and BOOTP Vendor Extensions</i> S. Alexander, R. Droms
RFC 2133	<i>Basic Socket Interface Extensions for IPv6</i> R. Gilligan, S. Thomson, J. Bound, W. Stevens
RFC 2136	<i>Dynamic Updates in the Domain Name System (DNS UPDATE)</i> P. Vixie, Ed., S. Thomson, Y. Rekhter, J. Bound
RFC 2137	<i>Secure Domain Name System Dynamic Update</i> D. Eastlake 3rd
RFC 2163	<i>Using the Internet DNS to Distribute MIXER Conformant Global Address Mapping (MCGAM)</i> C. Allocchio
RFC 2168	<i>Resolution of Uniform Resource Identifiers using the Domain Name System</i> R. Daniel, M. Mealling
RFC 2178	<i>OSPF Version 2</i> J. Moy
RFC 2181	<i>Clarifications to the DNS Specification</i> R. Elz, R. Bush

RFC 2205	<i>Resource ReSerVation Protocol (RSVP)—Version 1 Functional Specification</i> R. Braden, Ed., L. Zhang, S. Berson, S. Herzog, S. Jamin
RFC 2210	<i>The Use of RSVP with IETF Integrated Services</i> J. Wroclawski
RFC 2211	<i>Specification of the Controlled-Load Network Element Service</i> J. Wroclawski
RFC 2212	<i>Specification of Guaranteed Quality of Service</i> S. Shenker, C. Partridge, R. Guerin
RFC 2215	<i>General Characterization Parameters for Integrated Service Network Elements</i> S. Shenker, J. Wroclawski
RFC 2217	<i>Telnet Com Port Control Option</i> G. Clarke
RFC 2219	<i>Use of DNS Aliases for Network Services</i> M. Hamilton, R. Wright
RFC 2228	<i>FTP Security Extensions</i> M. Horowitz, S. Lunt
RFC 2230	<i>Key Exchange Delegation Record for the DNS</i> R. Atkinson
RFC 2233	<i>The Interfaces Group MIB using SMIv2</i> K. McCloghrie, F. Kastenholz
RFC 2240	<i>A Legal Basis for Domain Name Allocation</i> O. Vaughn
RFC 2246	<i>The TLS Protocol Version 1.0</i> T. Dierks, C. Allen
RFC 2251	<i>Lightweight Directory Access Protocol (v3)</i> M. Wahl, T. Howes, S. Kille
RFC 2253	<i>Lightweight Directory Access Protocol (v3): UTF-8 String Representation of Distinguished Names</i> M. Wahl, S. Kille, T. Howes
RFC 2254	<i>The String Representation of LDAP Search Filters</i> T. Howes
RFC 2261	<i>An Architecture for Describing SNMP Management Frameworks</i> D. Harrington, R. Presuhn, B. Wijnen
RFC 2262	<i>Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)</i> J. Case, D. Harrington, R. Presuhn, B. Wijnen
RFC 2271	<i>An Architecture for Describing SNMP Management Frameworks</i> D. Harrington, R. Presuhn, B. Wijnen
RFC 2273	<i>SNMPv3 Applications</i> D. Levi, P. Meyer, B. Stewart
RFC 2274	<i>User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)</i> U. Blumenthal, B. Wijnen
RFC 2275	<i>View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)</i> B. Wijnen, R. Presuhn, K. McCloghrie
RFC 2279	<i>UTF-8, a transformation format of ISO 10646</i> F. Yergeau
RFC 2292	<i>Advanced Sockets API for IPv6</i> W. Stevens, M. Thomas
RFC 2308	<i>Negative Caching of DNS Queries (DNS NCACHE)</i> M. Andrews
RFC 2317	<i>Classless IN-ADDR.ARPA delegation</i> H. Eidnes, G. de Groot, P. Vixie
RFC 2320	<i>Definitions of Managed Objects for Classical IP and ARP Over ATM Using SMIv2 (IPOA-MIB)</i> M. Greene, J. Luciani, K. White, T. Kuo
RFC 2328	<i>OSPF Version 2</i> J. Moy
RFC 2345	<i>Domain Names and Company Name Retrieval</i> J. Klensin, T. Wolf, G. Oglesby

|

RFC 2352	<i>A Convention for Using Legal Names as Domain Names</i> O. Vaughn
RFC 2355	<i>TN3270 Enhancements</i> B. Kelly
RFC 2358	<i>Definitions of Managed Objects for the Ethernet-like Interface Types</i> J. Flick, J. Johnson
RFC 2373	<i>IP Version 6 Addressing Architecture</i> R. Hinden, S. Deering
RFC 2374	<i>An IPv6 Aggregatable Global Unicast Address Format</i> R. Hinden, M. O'Dell, S. Deering
RFC 2375	<i>IPv6 Multicast Address Assignments</i> R. Hinden, S. Deering
RFC 2385	<i>Protection of BGP Sessions via the TCP MD5 Signature Option</i> A. Heffernan
RFC 2389	<i>Feature negotiation mechanism for the File Transfer Protocol</i> P. Hethmon, R. Elz
RFC 2401	<i>Security Architecture for Internet Protocol</i> S. Kent, R. Atkinson
RFC 2402	<i>IP Authentication Header</i> S. Kent, R. Atkinson
RFC 2403	<i>The Use of HMAC-MD5-96 within ESP and AH</i> C. Madson, R. Glenn
RFC 2404	<i>The Use of HMAC-SHA-1-96 within ESP and AH</i> C. Madson, R. Glenn
RFC 2405	<i>The ESP DES-CBC Cipher Algorithm With Explicit IV</i> C. Madson, N. Doraswamy
RFC 2406	<i>IP Encapsulating Security Payload (ESP)</i> S. Kent, R. Atkinson
RFC 2407	<i>The Internet IP Security Domain of Interpretation for ISAKMPD</i> . Piper
RFC 2408	<i>Internet Security Association and Key Management Protocol (ISAKMP)</i> D. Maughan, M. Schertler, M. Schneider, J. Turner
RFC 2409	<i>The Internet Key Exchange (IKE)</i> D. Harkins, D. Carrel
RFC 2410	<i>The NULL Encryption Algorithm and Its Use With IPsec</i> R. Glenn, S. Kent,
RFC 2428	<i>FTP Extensions for IPv6 and NATs</i> M. Allman, S. Ostermann, C. Metz
RFC 2445	<i>Internet Calendaring and Scheduling Core Object Specification (iCalendar)</i> F. Dawson, D. Stenerson
RFC 2459	<i>Internet X.509 Public Key Infrastructure Certificate and CRL Profile</i> R. Housley, W. Ford, W. Polk, D. Solo
RFC 2460	<i>Internet Protocol, Version 6 (IPv6) Specification</i> S. Deering, R. Hinden
RFC 2461	<i>Neighbor Discovery for IP Version 6 (IPv6)</i> T. Narten, E. Nordmark, W. Simpson
RFC 2462	<i>IPv6 Stateless Address Autoconfiguration</i> S. Thomson, T. Narten
RFC 2463	<i>Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification</i> A. Conta, S. Deering
RFC 2464	<i>Transmission of IPv6 Packets over Ethernet Networks</i> M. Crawford
RFC 2466	<i>Management Information Base for IP Version 6: ICMPv6 Group</i> D. Haskin, S. Onishi
RFC 2476	<i>Message Submission</i> R. Gellens, J. Klensin

RFC 2487	<i>SMTP Service Extension for Secure SMTP over TLS</i> P. Hoffman
RFC 2505	<i>Anti-Spam Recommendations for SMTP MTAs</i> G. Lindberg
RFC 2523	<i>Photuris: Extended Schemes and Attributes</i> P. Karn, W. Simpson
RFC 2535	<i>Domain Name System Security Extensions</i> D. Eastlake 3rd
RFC 2538	<i>Storing Certificates in the Domain Name System (DNS)</i> D. Eastlake 3rd, O. Gudmundsson
RFC 2539	<i>Storage of Diffie-Hellman Keys in the Domain Name System (DNS)</i> D. Eastlake 3rd
RFC 2540	<i>Detached Domain Name System (DNS) Information</i> D. Eastlake 3rd
RFC 2554	<i>SMTP Service Extension for Authentication</i> J. Myers
RFC 2570	<i>Introduction to Version 3 of the Internet-standard Network Management Framework</i> J. Case, R. Mundy, D. Partain, B. Stewart
RFC 2571	<i>An Architecture for Describing SNMP Management Frameworks</i> B. Wijnen, D. Harrington, R. Presuhn
RFC 2572	<i>Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)</i> J. Case, D. Harrington, R. Presuhn, B. Wijnen
RFC 2573	<i>SNMP Applications</i> D. Levi, P. Meyer, B. Stewart
RFC 2574	<i>User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)</i> U. Blumenthal, B. Wijnen
RFC 2575	<i>View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)</i> B. Wijnen, R. Presuhn, K. McCloghrie
RFC 2576	<i>Co-Existence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework</i> R. Frye, D. Levi, S. Routhier, B. Wijnen
RFC 2578	<i>Structure of Management Information Version 2 (SMIv2)</i> K. McCloghrie, D. Perkins, J. Schoenwaelder
RFC 2579	<i>Textual Conventions for SMIv2</i> K. McCloghrie, D. Perkins, J. Schoenwaelder
RFC 2580	<i>Conformance Statements for SMIv2</i> K. McCloghrie, D. Perkins, J. Schoenwaelder
RFC 2581	<i>TCP Congestion Control</i> M. Allman, V. Paxson, W. Stevens
RFC 2583	<i>Guidelines for Next Hop Client (NHC) Developers</i> R. Carlson, L. Winkler
RFC 2591	<i>Definitions of Managed Objects for Scheduling Management Operations</i> D. Levi, J. Schoenwaelder
RFC 2625	<i>IP and ARP over Fibre Channel</i> M. Rajagopal, R. Bhagwat, W. Rickard
RFC 2635	<i>Don't SPEW A Set of Guidelines for Mass Unsolicited Mailings and Postings (spam*)</i> S. Hambridge, A. Lunde
RFC 2637	<i>Point-to-Point Tunneling Protocol</i> K. Hamzeh, G. Pall, W. Verthein, J. Taarud, W. Little, G. Zorn
RFC 2640	<i>Internationalization of the File Transfer Protocol</i> B. Curtin

RFC 2665	<i>Definitions of Managed Objects for the Ethernet-like Interface Types</i> J. Flick, J. Johnson
RFC 2671	<i>Extension Mechanisms for DNS (EDNS0)</i> P. Vixie
RFC 2672	<i>Non-Terminal DNS Name Redirection</i> M. Crawford
RFC 2675	<i>IPv6 Jumbograms</i> D. Borman, S. Deering, R. Hinden
RFC 2710	<i>Multicast Listener Discovery (MLD) for IPv6</i> S. Deering, W. Fenner, B. Haberman
RFC 2711	<i>IPv6 Router Alert Option</i> C. Partridge, A. Jackson
RFC 2740	<i>OSPF for IPv6</i> R. Coltun, D. Ferguson, J. Moy
RFC 2753	<i>A Framework for Policy-based Admission Control</i> R. Yavatkar, D. Pendarakis, R. Guerin
RFC 2782	<i>A DNS RR for specifying the location of services (DNS SRV)</i> A. Gubrandsen, P. Vixie, L. Esibov
RFC 2821	<i>Simple Mail Transfer Protocol</i> J. Klensin, Ed.
RFC 2822	<i>Internet Message Format</i> P. Resnick, Ed.
RFC 2840	<i>TELNET KERMIT OPTION</i> J. Altman, F. da Cruz
RFC 2845	<i>Secret Key Transaction Authentication for DNS (TSIG)</i> P. Vixie, O. Gudmundsson, D. Eastlake 3rd, B. Wellington
RFC 2851	<i>Textual Conventions for Internet Network Addresses</i> M. Daniele, B. Haberman, S. Routhier, J. Schoenwaelder
RFC 2852	<i>Deliver By SMTP Service Extension</i> D. Newman
RFC 2874	<i>DNS Extensions to Support IPv6 Address Aggregation and Renumbering</i> M. Crawford, C. Huitema
RFC 2915	<i>The Naming Authority Pointer (NAPTR) DNS Resource Record</i> M. Mealling, R. Daniel
RFC 2920	<i>SMTP Service Extension for Command Pipelining</i> N. Freed
RFC 2930	<i>Secret Key Establishment for DNS (TKEY RR)</i> D. Eastlake, 3rd
RFC 2941	<i>Telnet Authentication Option</i> T. Ts'o, ed., J. Altman
RFC 2942	<i>Telnet Authentication: Kerberos Version 5</i> T. Ts'o
RFC 2946	<i>Telnet Data Encryption Option</i> T. Ts'o
RFC 2952	<i>Telnet Encryption: DES 64 bit Cipher Feedback</i> T. Ts'o
RFC 2953	<i>Telnet Encryption: DES 64 bit Output Feedback</i> T. Ts'o
RFC 2992	<i>Analysis of an Equal-Cost Multi-Path Algorithm</i> C. Hopps
RFC 3019	<i>IP Version 6 Management Information Base for The Multicast Listener Discovery Protocol</i> B. Haberman, R. Worzella
RFC 3060	<i>Policy Core Information Model—Version 1 Specification</i> B. Moore, E. Ellesson, J. Strassner, A. Westerinen
RFC 3152	<i>Delegation of IP6.ARPA</i> R. Bush
RFC 3164	<i>The BSD Syslog Protocol</i> C. Lonwick
RFC 3291	<i>Textual Conventions for Internet Network Addresses</i> M. Daniele, B. Haberman, S. Routhier, J. Schoenwaelder

RFC 3363	<i>Representing Internet Protocol version 6 (IPv6) Addresses in the Domain Name System</i> R. Bush, A. Durand, B. Fink, O. Gudmundsson, T. Hain
RFC 3390	<i>Increasing TCP's Initial Window</i> M. Allman, S. Floyd, C. Partridge
RFC 3410	<i>Introduction and Applicability Statements for Internet-Standard Management Framework</i> J. Case, R. Mundy, D. Partain, B. Stewart
RFC 3411	<i>An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks</i> D. Harrington, R. Presuhn, B. Wijnen
RFC 3412	<i>Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)</i> J. Case, D. Harrington, R. Presuhn, B. Wijnen
RFC 3413	<i>Simple Network Management Protocol (SNMP) Applications</i> D. Levi, P. Meyer, B. Stewart
RFC 3414	<i>User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)</i> U. Blumenthal, B. Wijnen
RFC 3415	<i>View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)</i> B. Wijnen, R. Presuhn, K. McCloghrie
RFC 3419	<i>Textual Conventions for Transport Addresses</i> M. Daniele, J. Schoenwaelder
RFC 3484	<i>Default Address Selection for Internet Protocol version 6 (IPv6)</i> R. Draves
RFC 3493	<i>Basic Socket Interface Extensions for IPv6</i> R. Gilligan, S. Thomson, J. Bound, J. McCann, W. Stevens
RFC 3513	<i>Internet Protocol Version 6 (IPv6) Addressing Architecture</i> R. Hinden, S. Deering
RFC 3526	<i>More Modular Exponential (MODP) Diffie-Hellman groups for Internet Key Exchange (IKE)</i> T. Kivinen, M. Kojo
RFC 3542	<i>Advanced Sockets Application Programming Interface (API) for IPv6</i> W. Richard Stevens, M. Thomas, E. Nordmark, T. Jinmei
RFC 3584	<i>Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework</i> R. Frye, D. Levi, S. Routhier, B. Wijnen
RFC 3602	<i>The AES-CBC Cipher Algorithm and Its Use with IPsec</i> S. Frankel, R. Glenn, S. Kelly
RFC 3629	<i>UTF-8, a transformation format of ISO 10646</i> R. Kermode, C. Vicisano
RFC 3658	<i>Delegation Signer (DS) Resource Record (RR)</i> O. Gudmundsson
RFC 3715	<i>IPsec-Network Address Translation (NAT) Compatibility Requirements</i> B. Aboba, W. Dixon
RFC 3947	<i>Negotiation of NAT-Traversal in the IKE</i> T. Kivinen, B. Swander, A. Huttunen, V. Volpe
RFC 3948	<i>UDP Encapsulation of IPsec ESP Packets</i> A. Huttunen, B. Swander, V. Volpe, L. DiBurro, M. Stenberg

Internet drafts

Internet drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Other groups may also distribute working documents as Internet drafts. You can see Internet drafts at <http://www.ietf.org/ID.html>.

Several areas of IPv6 implementation include elements of the following Internet drafts and are subject to change during the RFC review process.

Draft Title and Author

draft-bivens-sasp-02

Server/Application State Protocol v1 A. Bivens

draft-ietf-ipngwg-icmp-v3-07

Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification A. Conta, S. Deering

draft-ietf-ipsec-esp-v3-10

IP Encapsulating Security Payload (ESP) S. Kent

draft-ietf-ipsec-rfc2402bis-11

IP Authentication Header S. Kent

draft-ietf-ipsec-rfc2401bis-06

Security Architecture for the Internet Protocol S. Kent, K. Seo

draft-ietf-ospf-ospfv3-auth-07

Authentication/Confidentiality for OSPFv3 M. Gupta, N. Melam

Appendix B. Information APARs and technotes

This appendix lists information APARs for IP and SNA documents.

Notes:

1. Information APARs contain updates to previous editions of the documents listed in Table 38 and Table 39 on page 156. Documents updated for V1R7 are complete except for the updates contained in the information APARs that might be issued after V1R7 documents went to press.
2. Information APARs are predefined for z/OS V1R7 Communications Server and might not contain updates.
3. Information APARs for z/OS documents are in the document called *z/OS and z/OS.e DOC APAR and PTF ++HOLD Documentation*, which can be found at http://publibz.boulder.ibm.com:80/cgi-bin/bookmgr_OS390/BOOKS/ZIDOCMST/CCONTENTS.

Information APARs for IP documents

Table 38 lists information APARs for IP documents. For releases V1R7 and later, updates are available as technotes, which can be found at <http://www.ibm.com/support/docview.wss?uid=swg21178966>.

Table 38. IP information APARs for z/OS Communications Server

Title	V1R6	V1R5	V1R4
New Function Summary (both IP and SNA)	II13824		
Quick Reference (both IP and SNA)	II13831		II13246
IP and SNA Codes	II13842		II13254
IP API Guide	II13844	II13577	II13255 II13790
IP CICS Sockets Guide		II13578	II13257
IP Configuration Guide	II13826	II13568	II13244 II13541 II13652 II13646
IP Configuration Reference	II13827	II13569 II13789	II13245 II13521 II13647 II13739
IP Diagnosis	II13836	II13571	II13249 II13493
IP Messages Volume 1	II13838	II13572	II13624 II13250
IP Messages Volume 2	II13839	II13573	II13251
IP Messages Volume 3	II13840	II13574	II13252
IP Messages Volume 4	II13841	II13575	II13253 II13628
IP Migration		II13566	II13242 II13738

Table 38. IP information APARs for z/OS Communications Server (continued)

Title	V1R6	V1R5	V1R4
IP Network and Application Design Guide	II13825	II13567	II13243
IP Network Print Facility			
IP Programmer's Reference	II13843	II13581	II13256
IP User's Guide and Commands	II13832	II13570	II13247
IP System Admin Commands	II13833	II13580	II13248 II13792

Information APARs for SNA documents

Table 39 lists information APARs for SNA documents. For releases V1R7 and later, updates are available as technotes, which can be found at <http://www.ibm.com/support/docview.wss?uid=swg21178966>.

Table 39. SNA information APARs for z/OS Communications Server

Title	V1R6	V1R5	V1R4
New Function Summary (both IP and SNA)	II13824		
Quick Reference (both IP and SNA)	II13831		II13246
IP and SNA Codes	II13842		II13254
SNA Customization	II13857	II13560	II13240
SNA Diagnosis		II13558	II13236 II13735
SNA Diagnosis, Vol. 1: Techniques and Procedures	II13852		
SNA Diagnosis, Vol. 2: FFST Dumps and the VIT	II13853		
SNA Messages	II13854	II13559	II13238 II13736
SNA Network Implementation Guide	II13849	II13555	II13234 II13733
SNA Operation	II13851	II13557	II13237
SNA Migration		II13554	II13233 II13732
SNA Programming	II13858		II13241
SNA Resource Definition Reference	II13850	II13556	II13235 II13734
SNA Data Areas, Vol. 1 and 2			II13239
SNA Data Areas, 1	II13855		
SNA Data Areas, 2	II13856		

Other information APARs

Table 40 on page 157 lists information APARs not related to documents.

Table 40. Non-document information APARs

Content	Number
Index to APARs that list recommended VTAM maintenance	II11220
Index to APARs that list trace and dump requests for VTAM problems	II13202
Index of Communication Server IP information APARs	II12028
MPC and CTC	II01501
Collecting TCPIP CTRACEs	II12014
CSM for VTAM	II13442
CSM for TCP/IP	II13951
DLUR/DLUS for z/OS V1R2, V1R4, and V1R5	II12986, II13456, and II13783
DOCUMENTATION REQUIRED FOR OSA/2, OSA EXPRESS AND OSA QDIO	II13016
DYNAMIC VIPA (BIND)	II13215
DNS — common problems and solutions	II13453
Enterprise Extender	II12223
FTPing doc to z/OS Support	II12030
FTP problems	II12079
Generic resources	II10986
HPR	II10953
iQDIO	II13142
LPR problems	II12022
MNPS	II10370
NCROUTE problems	II12025
OMPROUTE	II12026
PASCAL API	II11814
Performance	II11710 II11711 II11712
Resolver	II13398 II13399 II13452
Socket API	II11996 II12020
SMTP problems	II12023
SNMP	II13477 II13478
SYSLOGD howto	II12021
TCPIP connection states	II12449
Telnet	II11574 II13135
TN3270 TELNET SSL common problems	II13369

Appendix C. Accessibility

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully. The major accessibility features in z/OS enable users to:

- Use assistive technologies such as screen readers and screen magnifier software
- Operate specific or equivalent features using only the keyboard
- Customize display attributes such as color, contrast, and font size

Using assistive technologies

Assistive technology products, such as screen readers, function with the user interfaces found in z/OS. Consult the assistive technology documentation for specific information when using such products to access z/OS interfaces.

Keyboard navigation of the user interface

Users can access z/OS user interfaces using TSO/E or ISPF. Refer to *z/OS TSO/E Primer*, *z/OS TSO/E User's Guide*, and *z/OS ISPF User's Guide Vol I* for information about accessing TSO/E and ISPF interfaces. These guides describe how to use TSO/E and ISPF, including the use of keyboard shortcuts or function keys (PF keys). Each guide includes the default settings for the PF keys and explains how to modify their functions.

z/OS information

z/OS information is accessible using screen readers with the BookServer/Library Server versions of z/OS books in the Internet library at:

www.ibm.com/servers/eserver/zseries/zos/bkserv/

Notices

IBM may not offer all of the products, services, or features discussed in this document. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

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

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

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

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

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

Site Counsel
IBM Corporation
P.O. Box 12195
3039 Cornwallis Road
Research Triangle Park, North Carolina 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, IBM International Program 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.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

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

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

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application

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

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_. All rights reserved.

IBM is required to include the following statements in order to distribute portions of this document and the software described herein to which contributions have been made by The University of California. Portions herein © Copyright 1979, 1980, 1983, 1986, Regents of the University of California. Reproduced by permission. Portions herein were developed at the Electrical Engineering and Computer Sciences Department at the Berkeley campus of the University of California under the auspices of the Regents of the University of California.

Portions of this publication relating to RPC are Copyright © Sun Microsystems, Inc., 1988, 1989.

Some portions of this publication relating to X Window System** are Copyright © 1987, 1988 by Digital Equipment Corporation, Maynard, Massachusetts, and the Massachusetts Institute Of Technology, Cambridge, Massachusetts. All Rights Reserved.

Some portions of this publication relating to X Window System are Copyright © 1986, 1987, 1988 by Hewlett-Packard Corporation.

Permission to use, copy, modify, and distribute the M.I.T., Digital Equipment Corporation, and Hewlett-Packard Corporation portions of this software and its documentation for any purpose without fee is hereby granted, provided that the above copyright notice appears in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the names of M.I.T., Digital, and Hewlett-Packard not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. M.I.T., Digital, and Hewlett-Packard make no representation about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

Copyright © 1983, 1995-1997 Eric P. Allman

Copyright © 1988, 1993 The Regents of the University of California. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement:

This product includes software developed by the University of California, Berkeley and its contributors.
4. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This software program contains code, and/or derivatives or modifications of code originating from the software program "Popper." Popper is Copyright ©1989-1991 The Regents of the University of California, All Rights Reserved. Popper was created by Austin Shelton, Information Systems and Technology, University of California, Berkeley.

Permission from the Regents of the University of California to use, copy, modify, and distribute the "Popper" software contained herein for any purpose, without fee, and without a written agreement is hereby granted, provided that the above copyright notice and this paragraph and the following two paragraphs appear in all copies. **HOWEVER, ADDITIONAL PERMISSIONS MAY BE NECESSARY FROM OTHER PERSONS OR ENTITIES, TO USE DERIVATIVES OR MODIFICATIONS OF POPPER.**

IN NO EVENT SHALL THE UNIVERSITY OF CALIFORNIA BE LIABLE TO ANY PARTY FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, INCLUDING LOST PROFITS, ARISING OUT OF THE USE OF THE POPPER SOFTWARE, OR ITS DERIVATIVES OR MODIFICATIONS, AND ITS DOCUMENTATION, EVEN IF THE UNIVERSITY OF CALIFORNIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

THE UNIVERSITY OF CALIFORNIA SPECIFICALLY DISCLAIMS ANY WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE POPPER SOFTWARE PROVIDED HEREUNDER IS ON AN "AS IS" BASIS, AND THE UNIVERSITY OF CALIFORNIA HAS NO OBLIGATIONS TO PROVIDE MAINTENANCE, SUPPORT, UPDATES, ENHANCEMENTS, OR MODIFICATIONS.

Copyright © 1983 The Regents of the University of California. All rights reserved.

Redistribution and use in source and binary forms are permitted provided that the above copyright notice and this paragraph are duplicated in all such forms and that any documentation, advertising materials, and other materials related to such distribution and use acknowledge that the software was developed by the University of California, Berkeley. The name of the University may not be used to endorse or promote products derived from this software without specific prior written permission. THIS SOFTWARE IS PROVIDED ``AS IS'' AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

Copyright © 1991, 1993 The Regents of the University of California. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement:
This product includes software developed by the University of California, Berkeley and its contributors.
4. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Copyright © 1990 by the Massachusetts Institute of Technology

Export of this software from the United States of America may require a specific license from the United States Government. It is the responsibility of any person or organization contemplating export to obtain such a license before exporting.

WITHIN THAT CONSTRAINT, permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of M.I.T. not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. Furthermore

if you modify this software you must label your software as modified software and not distribute it in such a fashion that it might be confused with the original M.I.T. software. M.I.T. makes no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

Copyright © 1998 by the FundsXpress, INC. All rights reserved.

Export of this software from the United States of America may require a specific license from the United States Government. It is the responsibility of any person or organization contemplating export to obtain such a license before exporting.

WITHIN THAT CONSTRAINT, permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of FundsXpress not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. FundsXpress makes no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

THIS SOFTWARE IS PROVIDED ``AS IS'' AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

Copyright © 1999, 2000 Internet Software Consortium.

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS" AND INTERNET SOFTWARE CONSORTIUM DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL INTERNET SOFTWARE CONSORTIUM BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Copyright © 1995-1998 Eric Young (eay@cryptsoft.com) All rights reserved.

This package is an SSL implementation written by Eric Young (eay@cryptsoft.com). The implementation was written so as to conform with Netscape's SSL.

This library is free for commercial and non-commercial use as long as the following conditions are adhered to. The following conditions apply to all code found in this distribution, be it the RC4, RSA, lhash, DES, etc., code; not just the SSL code. The SSL documentation included with this distribution is covered by the same copyright terms except that the holder is Tim Hudson (tjh@cryptsoft.com).

Copyright remains Eric Young's, and as such any Copyright notices in the code are not to be removed. If this package is used in a product, Eric Young should be

given attribution as the author of the parts of the library used. This can be in the form of a textual message at program startup or in documentation (online or textual) provided with the package.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement: "This product includes cryptographic software written by Eric Young (eay@cryptsoft.com)". The word 'cryptographic' can be left out if the routines from the library being used are not cryptographic related.
4. If you include any Windows specific code (or a derivative thereof) from the apps directory (application code) you must include acknowledgement:
"This product includes software written by Tim Hudson (tjh@cryptsoft.com)"

THIS SOFTWARE IS PROVIDED BY ERIC YOUNG ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The license and distribution terms for any publicly available version or derivative of this code cannot be changed. i.e. this code cannot simply be copied and put under another distribution license [including the GNU Public License.]

This product includes cryptographic software written by Eric Young.

Copyright © 1999, 2000 Internet Software Consortium.

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS" AND INTERNET SOFTWARE CONSORTIUM DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL INTERNET SOFTWARE CONSORTIUM BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Copyright © 2004 IBM Corporation and its licensors, including Sendmail, Inc., and the Regents of the University of California. All rights reserved.

Copyright © 1999,2000,2001 Compaq Computer Corporation

Copyright © 1999,2000,2001 Hewlett-Packard Company

Copyright © 1999,2000,2001 IBM Corporation

Copyright © 1999,2000,2001 Hummingbird Communications Ltd.

Copyright © 1999,2000,2001 Silicon Graphics, Inc.

Copyright © 1999,2000,2001 Sun Microsystems, Inc.

Copyright © 1999,2000,2001 The Open Group

All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, provided that the above copyright notice(s) and this permission notice appear in all copies of the Software and that both the above copyright notice(s) and this permission notice appear in supporting documentation.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR HOLDERS INCLUDED IN THIS NOTICE BE LIABLE FOR ANY CLAIM, OR ANY SPECIAL INDIRECT OR CONSEQUENTIAL DAMAGES, OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Except as contained in this notice, the name of a copyright holder shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization of the copyright holder.

X Window System is a trademark of The Open Group.

If you are viewing this information softcopy, photographs and color illustrations may not appear.

Trademarks

The following terms are trademarks of the IBM Corporation in the United States or other countries or both:

Advanced Peer-to-Peer Networking	MVS/XA
AFS	NetView
AD/Cycle	Network Station
AIX	Nways
AIX/ESA	OS/2
AnyNet	OS/390
APL2	Parallel Sysplex
AS/400	pSeries
BookManager	RACF
C/370	Redbooks
CICS	RETAIN
CICS/ESA	RISC System/6000
C Set ++	RMF
DB2	RS/6000
DFSMSdfp	S/370
DFSMShsm	S/390
DPI	S/390 Parallel Enterprise Server
ESCON	SecureWay
eServer	SiteCheck
ES/9000	SQL/DS
ES/9370	System/360
FFST	System/370
FICON	System/390
First Failure Support Technology	System z
GDDM	System z9
IBM	Tivoli
ibm.com	VM/ESA
IBMLink	VSE/ESA
IMS	VTAM
IMS/ESA	WebSphere
HiperSockets	z9
Language Environment	z/Architecture
Multiprise	z/OS
MVS	z/VM
MVS/DFP	z9
MVS/ESA	zSeries
MVS/SP	400

The following terms are trademarks of other companies:

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

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

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

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

Intel is a registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

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

Bibliography

z/OS Communications Server information

This section contains descriptions of the documents in the z/OS Communications Server library.

z/OS Communications Server documentation is available:

- Online at the z/OS Internet Library web page at <http://www.ibm.com/servers/eserver/zseries/zos/bkserv>
- In softcopy on CD-ROM collections. See "Softcopy information" on page xix.

z/OS Communications Server library

z/OS Communications Server documents are available on the CD-ROM accompanying z/OS (SK3T-4269 or SK3T-4307). Unlicensed documents can be viewed at the z/OS Internet library site.

Updates to documents are available on RETAIN® and in information APARs (info APARs). See Appendix B, "Information APARs and technotes," on page 155 for a list of the documents and the info APARs associated with them.

Info APARs for z/OS documents are in the document called *z/OS and z/OS.e DOC APAR and PTF ++HOLD Documentation* which can be found at http://publibz.boulder.ibm.com:80/cgi-bin/bookmgr_OS390/BOOKS/ZIDOCMST/CCONTENTS.

Planning

Title	Number	Description
<i>z/OS Communications Server: New Function Summary</i>	GC31-8771	This document is intended to help you plan for new IP for SNA function, whether you are migrating from a previous version or installing z/OS for the first time. It summarizes what is new in the release and identifies the suggested and required modifications needed to use the enhanced functions.
<i>z/OS Communications Server: IPv6 Network and Application Design Guide</i>	SC31-8885	This document is a high-level introduction to IPv6. It describes concepts of z/OS Communications Server's support of IPv6, coexistence with IPv4, and migration issues.

Resource definition, configuration, and tuning

Title	Number	Description
<i>z/OS Communications Server: IP Configuration Guide</i>	SC31-8775	This document describes the major concepts involved in understanding and configuring an IP network. Familiarity with the z/OS operating system, IP protocols, z/OS UNIX System Services, and IBM Time Sharing Option (TSO) is recommended. Use this document in conjunction with the <i>z/OS Communications Server: IP Configuration Reference</i> .

Title	Number	Description
<i>z/OS Communications Server: IP Configuration Reference</i>	SC31-8776	<p>This document presents information for people who want to administer and maintain IP. Use this document in conjunction with the <i>z/OS Communications Server: IP Configuration Guide</i>. The information in this document includes:</p> <ul style="list-style-type: none"> • TCP/IP configuration data sets • Configuration statements • Translation tables • SMF records • Protocol number and port assignments
<i>z/OS Communications Server: SNA Network Implementation Guide</i>	SC31-8777	<p>This document presents the major concepts involved in implementing an SNA network. Use this document in conjunction with the <i>z/OS Communications Server: SNA Resource Definition Reference</i>.</p>
<i>z/OS Communications Server: SNA Resource Definition Reference</i>	SC31-8778	<p>This document describes each SNA definition statement, start option, and macroinstruction for user tables. It also describes NCP definition statements that affect SNA. Use this document in conjunction with the <i>z/OS Communications Server: SNA Network Implementation Guide</i>.</p>
<i>z/OS Communications Server: SNA Resource Definition Samples</i>	SC31-8836	<p>This document contains sample definitions to help you implement SNA functions in your networks, and includes sample major node definitions.</p>
<i>z/OS Communications Server: IP Network Print Facility</i>	SC31-8833	<p>This document is for system programmers and network administrators who need to prepare their network to route SNA, JES2, or JES3 printer output to remote printers using TCP/IP Services.</p>

Operation

Title	Number	Description
<i>z/OS Communications Server: IP User's Guide and Commands</i>	SC31-8780	<p>This document describes how to use TCP/IP applications. It contains requests that allow a user to log on to a remote host using Telnet, transfer data sets using FTP, send and receive electronic mail, print on remote printers, and authenticate network users.</p>
<i>z/OS Communications Server: IP System Administrator's Commands</i>	SC31-8781	<p>This document describes the functions and commands helpful in configuring or monitoring your system. It contains system administrator's commands, such as TSO NETSTAT, PING, TRACERTE and their UNIX counterparts. It also includes TSO and MVS commands commonly used during the IP configuration process.</p>
<i>z/OS Communications Server: SNA Operation</i>	SC31-8779	<p>This document serves as a reference for programmers and operators requiring detailed information about specific operator commands.</p>
<i>z/OS Communications Server: Quick Reference</i>	SX75-0124	<p>This document contains essential information about SNA and IP commands.</p>

Customization

Title	Number	Description
<i>z/OS Communications Server: SNA Customization</i>	SC31-6854	<p>This document enables you to customize SNA, and includes the following:</p> <ul style="list-style-type: none"> • Communication network management (CNM) routing table • Logon-interpret routine requirements • Logon manager installation-wide exit routine for the CLU search exit • TSO/SNA installation-wide exit routines • SNA installation-wide exit routines

Writing application programs

Title	Number	Description
<i>z/OS Communications Server: IP Sockets Application Programming Interface Guide and Reference</i>	SC31-8788	<p>This document describes the syntax and semantics of program source code necessary to write your own application programming interface (API) into TCP/IP. You can use this interface as the communication base for writing your own client or server application. You can also use this document to adapt your existing applications to communicate with each other using sockets over TCP/IP.</p>
<i>z/OS Communications Server: IP CICS Sockets Guide</i>	SC31-8807	<p>This document is for programmers who want to set up, write application programs for, and diagnose problems with the socket interface for CICS using z/OS TCP/IP.</p>
<i>z/OS Communications Server: IP IMS Sockets Guide</i>	SC31-8830	<p>This document is for programmers who want application programs that use the IMS TCP/IP application development services provided by IBM's TCP/IP Services.</p>
<i>z/OS Communications Server: IP Programmer's Guide and Reference</i>	SC31-8787	<p>This document describes the syntax and semantics of a set of high-level application functions that you can use to program your own applications in a TCP/IP environment. These functions provide support for application facilities, such as user authentication, distributed databases, distributed processing, network management, and device sharing. Familiarity with the z/OS operating system, TCP/IP protocols, and IBM Time Sharing Option (TSO) is recommended.</p>
<i>z/OS Communications Server: SNA Programming</i>	SC31-8829	<p>This document describes how to use SNA macroinstructions to send data to and receive data from (1) a terminal in either the same or a different domain, or (2) another application program in either the same or a different domain.</p>
<i>z/OS Communications Server: SNA Programmer's LU 6.2 Guide</i>	SC31-8811	<p>This document describes how to use the SNA LU 6.2 application programming interface for host application programs. This document applies to programs that use only LU 6.2 sessions or that use LU 6.2 sessions along with other session types. (Only LU 6.2 sessions are covered in this document.)</p>
<i>z/OS Communications Server: SNA Programmer's LU 6.2 Reference</i>	SC31-8810	<p>This document provides reference material for the SNA LU 6.2 programming interface for host application programs.</p>
<i>z/OS Communications Server: CSM Guide</i>	SC31-8808	<p>This document describes how applications use the communications storage manager.</p>

Title	Number	Description
<i>z/OS Communications Server: CMIP Services and Topology Agent Guide</i>	SC31-8828	This document describes the Common Management Information Protocol (CMIP) programming interface for application programmers to use in coding CMIP application programs. The document provides guide and reference information about CMIP services and the SNA topology agent.

Diagnosis

Title	Number	Description
<i>z/OS Communications Server: IP Diagnosis Guide</i>	GC31-8782	This document explains how to diagnose TCP/IP problems and how to determine whether a specific problem is in the TCP/IP product code. It explains how to gather information for and describe problems to the IBM Software Support Center.
<i>z/OS Communications Server: SNA Diagnosis Vol 1, Techniques and Procedures and z/OS Communications Server: SNA Diagnosis Vol 2, FFST Dumps and the VIT</i>	GC31-6850 GC31-6851	These documents help you identify an SNA problem, classify it, and collect information about it before you call the IBM Support Center. The information collected includes traces, dumps, and other problem documentation.
<i>z/OS Communications Server: SNA Data Areas Volume 1 and z/OS Communications Server: SNA Data Areas Volume 2</i>	GC31-6852 GC31-6853	These documents describe SNA data areas and can be used to read an SNA dump. They are intended for IBM programming service representatives and customer personnel who are diagnosing problems with SNA.

Messages and codes

Title	Number	Description
<i>z/OS Communications Server: SNA Messages</i>	SC31-8790	This document describes the ELM, IKT, IST, ISU, IUT, IVT, and USS messages. Other information in this document includes: <ul style="list-style-type: none"> • Command and RU types in SNA messages • Node and ID types in SNA messages • Supplemental message-related information
<i>z/OS Communications Server: IP Messages Volume 1 (EZA)</i>	SC31-8783	This volume contains TCP/IP messages beginning with EZA.
<i>z/OS Communications Server: IP Messages Volume 2 (EZB, EZD)</i>	SC31-8784	This volume contains TCP/IP messages beginning with EZB or EZD.
<i>z/OS Communications Server: IP Messages Volume 3 (EZY)</i>	SC31-8785	This volume contains TCP/IP messages beginning with EZY.
<i>z/OS Communications Server: IP Messages Volume 4 (EZZ, SNM)</i>	SC31-8786	This volume contains TCP/IP messages beginning with EZZ and SNM.
<i>z/OS Communications Server: IP and SNA Codes</i>	SC31-8791	This document describes codes and other information that appear in z/OS Communications Server messages.

APPC Application Suite

Title	Number	Description
<i>z/OS Communications Server: APPC Application Suite User's Guide</i>	SC31-8809	This document describes the end-user interface (concepts, commands, and messages) for the AFTP, ANAME, and APING facilities of the APPC application suite. Although its primary audience is the end user, administrators and application programmers may also find it useful.

Title	Number	Description
<i>z/OS Communications Server: APPC Application Suite Administration</i>	SC31-8835	This document contains the information that administrators need to configure the APPC application suite and to manage the APING, ANAME, AFTP, and A3270 servers.
<i>z/OS Communications Server: APPC Application Suite Programming</i>	SC31-8834	This document provides the information application programmers need to add the functions of the AFTP and ANAME APIs to their application programs.

Index

A

accessibility 159
address assignment 64
address autoconfiguration 5
address resolution, in IPv6 32
address states
 deprecated 20
 preferred 20
 tentative 19
 unavailable 20
addressing 9
Advanced socket APIs 99
AF_INET socket applications
 and dual-mode stack 45
 and IPv4-only stack 44
AF_INET6 socket applications
 and dual-mode stack 45
 and IPv4-only stack 44
 and IPv6-only stack 45
AF_INET6 support, enabling 55
aggregatable global addresses, unicast 13
ALG 43
ancillary data 100, 113, 114
APIs 71
APIs, advanced 99
Application Layer Gateway (ALG) 43
 and z/OS CS TCP/IP 44
ARP, in IPv4 27
authentication, with IPv6 OSPF 25
autoconfiguration
 stateful 5
 stateless 5, 33
 guidelines 64
 steps 34
automatic tunnels 123
automation impacts
 due to message changes 51
 due to netstat changes 51
autonomous flag, in router advertisement 29

B

Basic socket API extensions for IPv6 75
 address conversion functions 83
 Address families 75
 address testing macros 83
 Design considerations 75
 interface identification 84
 name and address resolution functions 77
 name translation 77
 Protocol families 75
 socket options 84
 special IP addresses 76
BPXPRMxx
 and enabling IPv6 support 54
 CINET IPv4-only sample 54
 CINET IPv4/IPv6 dual-mode sample 55
 IPv4-only sample 54
 IPv4/IPv6 dual-mode sample 54
broadcast 27

C

C sockets 71
checksum processing for RAW applications 111
CICS sockets 72
coexistence overview, application 125
Common INET
 and AF_INET6 support 46
 configuring 48
 considerations 46
COMMONSEARCH statement, in resolver setup file 57
Communications Server for z/OS, online information xxi
configured tunnels 122

D

DAD 32
data stream, including IP addresses 95
data tracing 62
default address selection 36
default destination address selection 36
 rules 36
default source address selection
 rules 38
deprecated, address state 20
DHCPv6 5
diagnosing problems 62
disability 159
Distributed Protocol Interface (DPI) 134
DNS
 and VIPAs, guidelines 66
DNS definitions, updating 66
DNS, guidelines 66
DNS, online information xxii
dual-mode stack 40
 and IPv4 application 42
 and IPv6 application 40
dual-mode stack support 7
Duplicate Address Detection (DAD) 32
 and loopback addresses 32
 and VIPA 32, 35
 how to disable 32
 processing steps 32
Dynamic routing protocols 24

E

exits 58
extension headers 21

F

fragmentation 21
 support 6, 22
FTP exits 58

G

getaddrinfo 77
gethostbyaddr 82

gethostbyname 77
getnameinfo 82
getservbyname 77
getservbyport 82

H

header format 4
hierarchical addressing 4
hierarchical addressing and routing infrastructure 4
host names, defining 66

I

IBM Software Support Center, contacting xiv
ICMP considerations 117
ICMPv4 messages 27
ICMPv6 26
 message types 27
IGMP, in IPv4 27
IMS sockets 72
INET
 and dual-mode IPv4/IPv6 stack 46
 and IPv4-only stack 45
inetd 58
 configuration file 58
information APARs for IP-related documents 155
information APARs for non-document information 156
information APARs for SNA-related documents 156
interface ID
 defining for physical interfaces 64
interface identifiers, in IPv6 unicast address 16
Internet, finding z/OS information online xxi
IP addresses, impermanence 94
IP header format 4
IPAQENET6 interface type 63
IPCS 62
IPPROTO_IPV6 level 100
IPv4 TCP server program 95
IPv4-mapped IPv6 address 15
IPv6
 address space 4
 applications not enabled 134
 Basic socket API extensions 75
 expanded routing and addressing 4
 supported standards 131
 z/OS-specific features 131
IPv6 address
 anycast 19
 assigned to a node 19
 categories 12
 model 11
 multicast 17
 states 19
 textual representation 9
 types 11
 global unicast aggregatable 11
 link-local unicast 11
 loopback 11
 multicast 11
 site-local unicast 11
 unspecified 11
 unicast 13
IPv6 address space 11
IPv6 and IPv4 characteristics, comparison 6
IPv6 header 4

IPv6 header (*continued*)
 header options 4
IPv6 packet header 99
IPv6 prefix
 textual representation 10
IPv6, enabling applications for 89

K

keyboard 159

L

license, patent, and copyright information 161
link-layer device support 131
local use address, unicast 14
LookAt message retrieval tool xxii
loopback address
 and DAD 32
loopback address, unicast 15

M

message retrieval tool, LookAt xxii
migration and coexistence overview, application 125
migration approaches 127
migration overview, application 125
MLD 27
MPCPTP6 interface type 63
MTU discovery, options 101
multicast 17
 groups
 group ID 18
 scope 17
multicast and IPv6, using 94
Multicast Listener Discovery
 listener function 28
 query message 28
 router function 28
Multicast Listener Discovery (MLD) 27
multicasting 27
multipath routes, considerations 26

N

NAT 44
NAT-PT 129
native TCP/IP sockets 72
Neighbor Discovery (ND) 5, 27, 28
 Address Resolution 28, 32
 Duplicate Address Detection (DAD) 28
 parameter discovery 5
 prefix discovery 5
 Reachability Detection 28
neighbor node interaction, protocol 5
neighbor unreachability detection 33
Netstat 60
Network address translation (NAT) 44
network prefix 65
Network SLAPM2 subagent 134

O

OMPROUTE 24
OMPROUTE subagent 134

OMPROUTE, guidelines 67
on-link flag, in router advertisement 29
options, support 4
orexecd 58
orshd 58
otelnetd 58
outgoing interface, specifying 115

P

packet header, controlling the content 99
packet tracing 62
packets, controlling sending 103
Pascal API 72
path MTU discovery 22
Ping 62
Policy Agent 59
preferred, address state 20
prefix information options
 autonomous flag 29
 on-link flag 29
proxy 128

Q

QoS
 and flow label 7
QoS classification data 111
QoS policies 59

R

RAW applications, checksum processing 111
RAW sockets 116
received packets 108
redirect messages 6
redirect processing
 IGNOREREDIRECT on IPCONFIG6 31
Resolver 56
 and DNS 58
resolver API processing 92
Resolver configuration
 files 56
 search orders 56
REXX sockets 72
RFC (request for comments)
 accessing online xxi
 list of 139
route selection 26
router advertisements 28
 prefix information options 29
 autonomous flag 29
 on-link flag 29
router discovery 23
routing 22
 and VARY TCPIP,,OBEYFILE command 26
routing infrastructure 4

S

scope 11
 multicast 17
scope zone 12
scope zone index 12
scope zones 11

security considerations 53
shortcut keys 159
SIIT 128
site-local addresses 64
SMF records 59
SNA application access
 and EE 52
 and TN3270 52
SNMP, agent 134
SNMP, subagents 134
 DPI 134
 Network SLAPM2 134
 OMPROUTE 134
 TCP/IP 134
socket APIs 71
 and IPv6 support 71
socket APIs, advanced 99
socket options and ancillary data, interactions 114
sockets extended call instruction API 72
sockets extended macro API 72
SOCKS 43, 128
SOCKS server 43
SOCKS64 43, 128
source address, options 115
source VIPA, guidelines 66
SOURCEVIPA, for IPv6 35
special IPv6 addresses 93
static routes
 BEGINROUTES profile statement 22
 GATEWAY profile statement 22
static routes, guidelines 67
support tables 131
sysplex support 132

T

takeover function, interface 34
tasks
 implementing the resolver functions
 steps 57
TCP server program enabled for IPv6 96
TCP/IP
 online information xxi
 protocol specifications 139
TCP/IP, subagent 134
tentative, address state 19
textual representation
 IPv6 addresses 9
 IPv6 prefixes 10
TN3270
 and SNA access 52
Traceroute 40, 62
tracing
 data 62
 packet 62
trademark information 169
translation mechanisms 127
 NAT-PT 129
 proxy 128
 SIIT 128
 SOCKS 128
Trap forwarder daemon 134
tunneling
 6over4 tunnels 124
 6to4 addresses 123
 automatic tunnels 123
 configured tunnels 122

tunneling (*continued*)

 overview 121

Tunneling 121

 and z/OS/CS 40

U

unavailable, address state 20

unicast 13

unspecified address, unicast 15

user exits 58

V

VARY TCPIP,,OBEYFILE , and autoconfiguration 34

VARY TCPIP,,OBEYFILE command considerations, in router

 advertisements 31

VIPA

 and Duplicate Address Detection(DAD) 32, 35

 and prefixes 35

 and source address selection 38

 how to get addresses 35

 interface identifier

 guidelines 65

 network prefix

 guidelines 65

 static

 guidelines 64

VTAM, online information xxi

Z

z/OS, documentation library listing 171

z/OS, listing of documentation available 155

zone index 12

Communicating Your Comments to IBM

If you especially like or dislike anything about this document, please use one of the methods listed below to send your comments to IBM. Whichever method you choose, make sure you send your name, address, and telephone number if you would like a reply.

Feel free to comment on specific errors or omissions, accuracy, organization, subject matter, or completeness of this document. However, the comments you send should pertain to only the information in this manual and the way in which the information is presented. To request additional publications, or to ask questions or make comments about the functions of IBM products or systems, you should talk to your IBM representative or to your IBM authorized remarketer.

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.

Please send your comments to us in either of the following ways:

- If you prefer to send comments by FAX, use this number: 1+919-254-4028
- If you prefer to send comments electronically, use this address:
 - comsvrcf@us.ibm.com.
- If you prefer to send comments by post, use this address:

International Business Machines Corporation
Attn: z/OS Communications Server Information Development
P.O. Box 12195, 3039 Cornwallis Road
Department AKCA, Building 501
Research Triangle Park, North Carolina 27709-2195

Make sure to include the following in your note:

- Title and publication number of this document
- Page number or topic to which your comment applies.



Program Number: 5694-A01 and 5655-G52

Printed in USA

SC31-8885-04



Spine information:



z/OS Communications Server

**z/OS V1R8.0 Comm Srv: IPv6 Network and Appl
Design Guide**

Version 1
Release 8